

Modeling Steady-state and Transient Behaviors of User Mobility: Formulation, Analysis, and Application

Jong-Kwon Lee^{*}
Ubiquitous Computing Laboratory (UCL)
IBM Korea
Seoul, Korea
jkwlee@kr.ibm.com

Jennifer C. Hou
Department of Computer Science
University of Illinois
Urbana, IL 61801
jhou@cs.uiuc.edu

ABSTRACT

Recent studies on mobility modeling have focused on characterizing user mobility from real traces of wireless LANs (WLANs) and creating mobility models based on such characterization. However, most of the work does not study how user mobility is correlated in time at different time scales. For example, the future APs with which a user will be associated are predicted without the knowledge of *when* the association will take place and for how long. In this paper, we build a mathematical model for characterizing both steady-state and transient behaviors of user mobility in WLANs. Specifically, we model user mobility by a semi-Markov process, and obtain the transition probability matrix and the sojourn time distribution from the *association history* of WLAN users available at Dartmouth college [21]. With the steady-state characterization of user mobility in WLANs, we can estimate the long-term wireless network usage among different access points. By comparing the steady-state distributions of semi-Markov models built based on trace data collected at *different* time scales, we are able to characterize the *degree of correlation* in time and location.

We also perform a transient behavior analysis of the semi-Markov process (that characterizes user mobility), and devise a timed location prediction algorithm that accurately predicts the future locations of users — *both* the future access points they will associate themselves with and the association duration. We demonstrate the utility of timed location prediction, by showing how it can be utilized to predict the distribution of future user locations with the time information figured in, and redistributing loads among neighboring APs. An improvement of 80% (in terms of load balance) is observed in a wide spectrum of traffic loads in the simulation.

^{*}Jong-Kwon Lee was previously a visiting scholar at the Department of Computer Science, University of Illinois at Urbana-Champaign.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiHoc'06, May 22–25, 2006, Florence, Italy.
Copyright 2006 ACM 1-59593-368-9/06/0005 ...\$5.00.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication; I.6.5 [Model Development]: Modeling methodologies

General Terms

Algorithms, Performance, Measurement

Keywords

WLAN, load balancing algorithm, location prediction, mobility model, semi-Markov process

1. INTRODUCTION

Thanks to the proliferation of wireless-enabled, portable laptops and PDAs, cost-effective deployment of access points, and availability of the license exempt bands and appropriate standards, IEEE 802.11 based wireless Local Area Networks (WLANs) have become widely available at universities and corporations to provide wireless Internet access. These wireless-enabled devices (and applications built on top of them) have, however, aggravated the competition for limited bandwidth of the underlying wireless infrastructure. This is particularly true for devices operating in the unlicensed frequency band such as IEEE 802.11-enabled devices [8]. Without adequate resource provisioning, wireless-enabled devices may soon become the victims of their own success. A rescue is really a thorough understanding of how wireless resources are spatially and temporally used, based on which intelligent management and provisioning policies can be deployed. As the limited frequency band for wireless communication is essentially a spatial resource defined by the location, the availability of user mobility information will enable advanced management and control of wireless networks. The wireless infrastructure will likely sustain, even in the presence of constantly increasing offered loads.

Mobility and network usage of users in WLANs have been characterized by several studies for campus or corporate WLANs [16, 3, 9, 20]. In particular, most of the work on real-trace-based mobility models focuses on extracting features that characterize mobility parameters and creating parameterized mobility models for simulating wireless networks. (We will provide a succinct summary in Section 2.) In spite of all the pioneering work, there have been few studies on how user mobility is correlated in time at different time scales. In particular, among the few research efforts that leverage real WLAN traces to empirically predict user

locations [18, 6], the future APs with which a user will be associated are predicted without knowing *when* the association will take place and for how long. This somewhat limits the extent to which these models can be used for network planning and resource management.

In this paper, we aim to bridge the gap, and build a mathematical model for characterizing both steady-state and transient behaviors of user mobility in WLANs. Specifically, we model user mobility by a semi-Markov process, with the renewal points embedded at the time instants when users associate themselves with a *new* access point. We obtain the transition probability matrix and the sojourn time distribution from the *association history* of WLAN users available at Dartmouth college [21]. With the steady-state characterization of user mobility in WLANs, we can estimate the long-term wireless network usage among different access points. Using this methodology, we build different semi-Markov models based on the trace data collected at *different* time scales and for different user groups (where users are grouped based on their home locations). By comparing the steady-state distributions of these models (with respect to a similarity measure, the *cosine distance* [17]), we are able to characterize the *degree of correlation* in time and location.

The second part of our work focuses on the transient behavior analysis of the semi-Markov process (that characterizes user mobility). Based on the model, we devise a timed location prediction algorithm that accurately predicts the future locations of users — *both* the future access points they will associate themselves with and the association duration. In contrast, the majority of existing location prediction algorithms predict only what the next location a user will be associated with. We also demonstrate the utility of timed location prediction, by showing how it can be utilized to predict the distribution of future user locations with the time information figured in, and redistributing loads among neighboring APs.

There are several important findings in this study. First, we show that the user mobility on campus is correlated in time on a daily, weekly or monthly basis. The cosine distance measurement confirms that the level of *temporal* similarity in user mobility often exceeds 0.9 (with 1 characterizing exactly the same pattern). Second, in the course of characterizing user mobility from the association patterns recorded in the trace data, we found that users frequently make repetitive re-associations between two or three neighboring APs. This phenomenon is known as a *ping-pong* transition in cell-based wireless networks such as cellular networks. As a matter of fact, the ping-pong transitions occur so often that they affect the *accuracy* of the transition probability between APs and the residence time characteristics at each AP. After eliminating the ping-pong transitions from the traces, the mean residence time thus calculated at several APs becomes larger, and reflects better the actual user mobility behavior. Third, we show that the ping-pong phenomenon can actually be exploited to balance loads across APs in WLANs. In the simulation study (carried out based on the user mobility model obtained from the traces), we show that a much balanced load can be achieved with the use of the proposed timed location prediction algorithm. An improvement of 80% (in terms of load balance) is observed in a wide spectrum of traffic loads.

The remainder of the paper is organized as follows. In Section 2, we give a succinct summary of existing work that

pertains to our work, and motivate the need for our work. In Section 3, we present the preliminaries of semi-Markov processes and how we model user mobility with such processes. Following that, we discuss in Section 4 the steady-state characteristics of user mobility based on the Dartmouth traces, and present in Section 5 the transient behavior of user mobility and its use in timed location prediction. In Section 5, we demonstrate how timed location prediction can be used to balance loads among APs. An evaluation study based on the Dartmouth traces is also given there. Finally, we conclude the paper with a list of future research agendas in Section 7.

2. RELATED WORK

There have been several mobility modeling studies that leverage traces collected from public or campus-wide wireless LANs (WLANs). Earlier studies focus on characterizing how wireless LANs are utilized and the behaviors of users therein [19, 16, 1, 3, 9]. Tang and Baker [19] focused on traces collected in a university building of 74 users over a period of 12 weeks. They analyzed the overall user behavior, the overall network traffic/load characteristics, and the per-user traffic characteristics. Kotz and Essien [16] characterized the usage of a university campus WLAN with 1706 users and spanning 161 buildings with a total of 476 access points. They focused on studying large-scale characteristics of the campus, such as aggregate network usage activities, aggregate traffic patterns and access point behaviors. Balachandran *et al.* [1] examined the usage of a wireless network in a large auditorium during a three-day conference. They used SNMP to poll each of their four APs every minute, and collected the traces. As the study was performed in a conference setting, the user behavior is homogeneous, with clients following the conference schedule. Balazinska and Castro [3] traced 1366 corporate users on 117 APs over four weeks. They analyzed user and load distribution across access points. They also introduced the notion of *persistence* and *prevalence* as two metrics for user mobility, where persistence reflects the session durations while prevalence reflects the frequency with which users visit various locations. They found that the probability distributions of both measures follow power laws. Henderson *et al.* [9] analyzed the Dartmouth trace with more than 550 access points and 7000 users over seventeen weeks, and examined geographic mobility characteristics for different wireless devices such as laptops, PDAs, and VoIP devices.

Another thrust of mobility modeling aims to create mobility models for simulating mobile wireless networks. In the random walk or random waypoint model [5], a user selects a random direction or destination and a random speed to reach that destination from a specified range $[V_{min}, V_{max}]$, and moves in compliance with the selected parameters. After the user reaches the destination, it pauses for some random time (again drawn from the uniform distribution) before it repeats the process. Although these random models are simple and scalable, they have been shown to exhibit unrealistic behaviors [13]. The work reported in [13] attempts to make the models more realistic by incorporating random waypoints and physical obstacle models.

More recent studies that leverage real traces to create mobility models for use in simulating wireless mobile networks include those reported in [10, 4, 20, 11, 12, 14, 15]. (Note that the work reported in [11, 12, 14, 15] used the

same source of mobility traces collected at Dartmouth College.) Both Helmy *et al.* [10, 4] and Tuduca *et al.* [20] developed more realistic mobility models for simulating ad hoc networks. The former is based on actual movement patterns of users on a university campus, while the latter uses real-life mobility characteristics extracted from WLAN traces to generate mobility scenarios for simulation. Jain *et al.* [11] derived a model that characterizes the time and the sequence of locations at which user devices register, and showed that the time features indicate heavy-tailed distributions. In their another paper [12], they also derived an empirical model, called *Model T*, for characterizing spatial registration patterns of mobile users. They showed that the user registration patterns exhibit hierarchy, and that WLAN APs can be clustered based on registration patterns.

Kim and Kotz [14, 15] investigated, with the use of *discrete Fourier transform*, the time intervals between which user mobility repeats itself. They counted the number of user visits to APs and found that the interval between which the user behavior is strongly correlated was one day. Based on this observation, they aggregated multiple days of hourly visits to characterize the hourly variations over a day. Then, they clustered APs based on their peak hour and developed a model that (rather than modeling the movements of individual users) characterizes wireless network usage in terms of the arrival rate at each AP.

Our work differs from the aforementioned efforts in that we aim to develop a mathematical mobility model for characterizing both the *long-term* and *transient* behaviors of user mobility. We also demonstrate how this model can be leveraged to predict future user locations – both the future APs users will associate themselves with and the association duration. The work that comes closest to ours with respect to empirically predicting user locations based on real traces of WLANs is reported in [18, 6]. However, they focused on evaluating how well future user locations can be predicted, without knowing *when* the movement to the next location will take place. In this paper, we extend (with the use of the proposed semi-Markov model for user mobility) location prediction in WLANs along the temporal domain.

3. SEMI-MARKOV MODEL FOR USER MOBILITY

We will model user mobility by a *semi-Markov process*. The reason for using semi-Markov processes (rather than continuous-time Markov chains) for modeling user mobility is because the *sojourn* time during which a user is associated with an AP in a campus-wide WLAN has been shown in [3, 11] *not* to follow the exponential distribution, but instead a heavy-tailed distribution. A semi-Markov process allows for arbitrary distributed sojourn times and can be viewed as a process with an embedded Markov chain, where the embedded points are the time instants when a user associates itself with a new AP.

Specifically, we consider a Markov renewal process $\{(X_n, T_n) : n \geq 0\}$ with a discrete state space $S = \{1, \dots, m\}$, where T_n is the time instant of the n -th transition ($T_0 = 0$) and $X_n \in S$ is the state at the n -th transition. For simplicity, we assume the renewal process is *time homogeneous* during the period in which the mobility model is built. Then, the associated *time homogeneous semi-Markov kernel* Q is defined

by:

$$\begin{aligned} Q_{ij}(t) &= \Pr(X_{n+1} = j, T_{n+1} - T_n \leq t | X_n = i) \\ &= p_{ij} H_{ij}(t), \end{aligned} \quad (1)$$

where $p_{ij} = \lim_{t \rightarrow \infty} Q_{ij}(t) = \Pr(X_{n+1} = j | X_n = i)$ is the state transition probability between states i and j , $P = [p_{ij}]$ is the transition probability matrix of the embedded Markov chain, and $H_{ij}(t) = \Pr(T_{n+1} - T_n \leq t | X_{n+1} = j, X_n = i)$ is the sojourn time distribution in state i when the next state is j . Let $D_i(t) \triangleq \Pr(T_{n+1} - T_n \leq t | X_n = i)$ denote the probability distribution of the sojourn time in state i regardless of the next state. Then, the distribution of the *residence* time, $D_i(t)$, during which the user is associated with AP i before his/her next transition takes place can be expressed as

$$D_i(t) = \sum_{j=1}^m Q_{ij}(t). \quad (2)$$

Now we define the homogeneous semi-Markov process as $Z = (Z_t, t \in \mathcal{R}_0^+)$, with the transient distributions

$$\begin{aligned} \phi_{ij}(t) &\triangleq \Pr(Z_t = j | Z_0 = i), \\ &= (1 - D_i(t))\delta_{ij} + \sum_{l=1}^m \int_0^t \phi_{lj}(t - \tau) dQ_{il}(\tau) \\ &= (1 - D_i(t))\delta_{ij} + \sum_{l=1}^m \int_0^t \dot{Q}_{il}(\tau) \phi_{lj}(t - \tau) d\tau, \end{aligned} \quad (3)$$


where δ_{ij} represents the Kronecker δ function defined by

$$\delta_{ij} = \begin{cases} 0 & \text{for } i \neq j \\ 1 & \text{for } i = j. \end{cases} \quad (4)$$


4. CHARACTERIZATION OF USER MOBILITY IN WLANS

We characterize user mobility in WLANs using the above semi-Markov process with the transition probability matrix P and the residence time distribution $D_i(t)$. The transition probability matrix P is a $m \times m$ matrix, where m is the number of locations (i.e. the number of APs). The (i, j) -th entry of P , p_{ij} , is the probability that a user will be associated with AP j in the next transition given that he/she is currently associated with AP i . $D_i(t)$ represents the distribution of the *residence* time. In addition, we define the $m \times 1$ mean residence time vector $\bar{D} = [\bar{d}_i]$, where \bar{d}_i is the mean value for $D_i(t)$. We will use the terms, *state* and *AP*, interchangeably in subsequent sections.


To empirically derive P and $D_i(t)$, we use the wireless network traces collected at Dartmouth College between November 1, 2003 and June 30, 2004. Specifically, we use the *syslog* data for the association patterns of mobile users (i.e., when a user (re)associates with an AP). Each *syslog* message contains a timestamp in seconds, the client's MAC address, the AP name, and the event type. From these *syslog* messages, the mobility of each user is extracted in the form of a series of two tuples (AP name, the timestamp when the association with this AP occurs). The "OFF" state is introduced to represent the fact that the user departs from the WLAN. For ease of expressions, we convert AP names into numeric states: 1 for OFF state, 2-155 for APs in *Academic* buildings, 156-228 for APs in *Administrative* buildings, 229-250




(a) Steady-state transition probability distribution (log-scale)



(b) Mean residence time at each AP



(c) Steady-state distribution of users over APs (including OFF state; log-scale)



(d) Steady-state distribution of users over APs (w/o OFF state; linear-scale)

Figure 1: Steady-state characteristics of user mobility.

for APs in *Athlete* buildings, 251-300 for APs in *Library* buildings, 301-514 for APs in *Residence* buildings, and 515-560 for APs in *Social* buildings. The transition probability matrices and mean residence time vectors can then be empirically calculated based on traces collected in intervals of different time scales and/or among different user groups, where the user group is defined with respect to the home locations of users.

4.1 Steady-state probability distribution of users over APs

Given the transition probability matrix P , we can derive the *steady-state transition probability* $\tilde{\pi} = [\tilde{\pi}_1, \dots, \tilde{\pi}_m]$ by solving the following equations:

$$\tilde{\pi} = \tilde{\pi}P, \quad \sum_{i=1}^m \tilde{\pi}_i = 1. \quad (5)$$

Then, with the mean residence time vector $\bar{D} = [\bar{d}_i]$, we can characterize user mobility by calculating the *steady-state user distribution* $\pi = [\pi_i]$ as follows:

$$\pi_i = \frac{\bar{d}_i \tilde{\pi}_i}{\sum_{j=1}^m \bar{d}_j \tilde{\pi}_j}. \quad (6)$$

The steady-state distribution π is the probability distribution of users over APs *at any time instant*, and is hence the *long-term average* distribution of users over APs.

In order to compare mobility behaviors to see if they are different or similar, we use the *cosine distance* as a similarity measure. The cosine distance has been widely used as

a pattern similarity measure [17]. In the current problem setting, the cosine distance $\text{sim}(\vec{p}, \vec{q})$ is defined as

$$\text{sim}(\vec{p}, \vec{q}) = \frac{\sum_{i=1}^m p_i q_i}{\left(\sqrt{\sum_{i=1}^m p_i^2}\right) \left(\sqrt{\sum_{i=1}^m q_i^2}\right)} = \frac{\vec{p} \cdot \vec{q}}{|\vec{p}| |\vec{q}|}, \quad (7)$$

where $\vec{p} = [p_i]$ and $\vec{q} = [q_i]$ are the steady-state distributions of two mobility models under comparison, and p_i and q_i are the probabilities that a user is associated with AP i under the two models, respectively. Note that $\text{sim}(\vec{p}, \vec{q})$ ranges in $[0, 1]$, with $\text{sim}(\vec{p}, \vec{q}) = 1$ indicating that the two mobility behaviors are identical as far as the long-term user distribution is concerned.

Figure 1(a)-1(b) show, respectively, the steady-state transition probability distribution and the mean residence time at each AP, obtained from the semi-Markov model derived with the trace of 3567 active users during the period of November 1, 2003 through February 29, 2004. Figure 1(c) is the steady-state distribution, π , of users over APs obtained from the results of Figures 1(a) and 1(b). It can be observed that users are in the OFF state (AP 1) most of the time. This is because most users disconnect from the network at night and will not return until the next morning. Although the mean residence time in the OFF state is not the largest, the product of the mean residence time and the transition probability into the OFF state is much higher than that in any other states. In contrast, a few states that have a larger mean residence time than the OFF state actually have very low steady-state transition probabilities. To give a clear picture of the steady state distribution of users over only *actual*




Figure 2: Examples of ping-pong transitions

APs, Figure 1(d) gives the zoom-in view of the distribution without the the OFF state (AP 1).

4.2 Ping-pong phenomena in user mobility

In the course of characterizing user mobility from the association patterns in the trace data, we found that sometimes users experience frequent re-associations between two or three APs in a short period of time. This phenomenon is known as a *ping-pong* transition in cell-based wireless networks such as cellular networks. In this section, we measure how often the ping-pong phenomena occurs and its effect on characterizing user mobility.

Definition of, and method for identifying, ping-pong transitions

Ping-pong transitions occur when a user is geographically located among several neighboring cells (APs). We define the *ping-pong transition* as follows:

- (i) For APs i and j , if at any time a user makes a sequence of transitions $i \rightarrow j \rightarrow i \rightarrow j$, then these transitions are ping-pong transitions.
- (ii) For APs i , j and k , if at any time a user makes a sequence of transitions $i \rightarrow j \rightarrow k \rightarrow i$, then these transitions are ping-pong transitions.

Figure 2 depicts the corresponding scenarios in which ping-pong transitions described in cases (i) and (ii) may occur. The above definition errs on the pessimistic side, in the sense that it does not distinguish a transition that results from real movement from a ping-pong transition. We may extend the definition by taking into account the residence time information to better distinguish the two types of transitions.

Once a ping-pong transition is identified in the association patterns of a user, we cluster the states incurred in the ping-pong transition into an *aggregate state (AS)*. For example, i and j form an aggregate state in case (i) ($AS = \{i, j\}$), while i , j , and k form an aggregate state in case (ii) ($AS = \{i, j, k\}$). The size of an AS is at most three according to the above definition of ping-pong transitions.

Frequency of ping-pong transitions

Figure 3 gives the *cumulative* distribution of users against their ping-pong ratios, where the *ping-pong ratio* of a user is defined as the ratio of the number of ping-pong transitions over the total number of transitions made by that user. The average ping-pong ratio is 0.36 and the median is 0.30. This means that approximately half of the users experience ping-pong transitions over 30% of all their transitions. As ping-pong transitions occur quite often, their effect on modeling user mobility should not be ignored.





Figure 3: Cumulative distribution of users against their ping-pong ratios. (For 3567 active users during the period between November 1, 2003 and February 29, 2004.)


Impact of ping-pong transitions on user mobility characteristics

With the frequent occurrence of ping-pong transitions, the transition probability matrix and the residence time at each AP derived under the *original* association patterns may not represent the exact mobility behavior of users. To investigate the impact of ping-pong transitions on user mobility characteristics, we eliminate ping-pong transitions from the original association patterns of users by identifying each ping-pong transition and replacing it with the *dominant* AP with which the user has been associated most of the time among the APs in the aggregate state (AS). For example, in the scenario described in Figure 2, if a sequence of AP transitions $a \rightarrow 1 \rightarrow 4 \rightarrow 1 \rightarrow 4 \rightarrow b$ appears in the association patterns of a user, then the ping-pong transition $1 \rightarrow 4 \rightarrow 1 \rightarrow 4$ is identified with the aggregate state $AS = \{1, 4\}$. If AP 1 is the dominant AP with which the user has associated most of the time in the $AS = \{1, 4\}$, then the original sequence of AP transitions is replaced by the sequence of $a \rightarrow 1 \rightarrow b$.


Figure 4 gives the steady-state transition probability, the mean residence time at each AP, and the steady-state distribution of users over APs, after eliminating the ping-pong transitions from the association patterns of users. As compared to the results given in Figure 1, the steady-state distribution of users over APs is similar to that derived with the original user association patterns (cosine distance = 0.998). However, the mean residence times at several APs become longer. We believe these APs are the dominant APs with which users have mostly associated. Similar results have also been reported in [9], where the authors attempted to eliminate the effect of the ping-pong phenomenon by using the notion of the session diameter. Figure 5 further quantifies the increase in the mean residence time when the ping-pong transitions are eliminated. The mean and median residence times obtained from the original user association patterns are approximately 51 minutes and 33 minutes, respectively, while those obtained after the ping-pong transitions are eliminated are approximately 108 minutes and 73 minutes, respectively. This result implies that several previous studies based on real WLAN traces [3, 20, 11] may




(a) Steady-state transition probability distribution (log-scale)



(b) Mean residence time at each AP



(c) Steady-state distribution of users over APs (including OFF state; log-scale)



(d) Steady-state distribution of users over APs (w/o OFF state; linear-scale)

Figure 4: Steady-state characteristics of user mobility *after eliminating ping-pong transitions*.

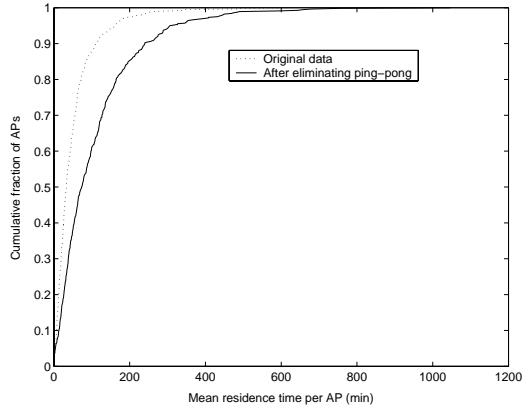


Figure 5: Change in the mean residence time at APs.

have under-estimated residence times and/or over-estimated transition frequencies. We will henceforth use trace data with the *corrected* association patterns of users in which ping-pong transitions are eliminated.

4.3 Similarity of user mobility in time intervals of different scales

In order to investigate whether or not, and to what extent, mobility behaviors correlate in time, we calculate the steady-state distributions of users over APs based on traces

collected in time intervals of different scales. Both the transition probability matrix P and the residence time distribution $D_i(t)$ can be constructed from the association patterns of users collected in a month or a week, and the steady-state distribution of users over APs can then be calculated following the steps in Section 4.1. For example, for traces collected in an 8-month period (32 weeks between November 2, 2003 and June 12, 2004), we can derive 8 different steady-state distributions of users over APs (based on monthly traces), and compare them with the *cosine distance* as described in Section 4.1.

Table 1 gives the cosine distance between steady-state distributions derived based on monthly traces collected in a period of 8 months (32 weeks) between November 2, 2003 and June 12, 2004. It can be observed that the closer in time the two monthly traces, the more similar the corresponding two steady-state distributions, though there exist some exceptions. For example, the cosine distance becomes much smaller when **m6** is involved. This is because **m6** corresponds to the period between March 21, 2004 and April 17, 2004 that includes the week of *spring break*. This administrative event affects the user behavior and hence user mobility dramatically.

Table 2 gives the similarity measure between the steady-state distributions derived based on *weekly* traces collected in a 14-week period between February 1, 2004 and May 8, 2004. The result confirms that user mobility correlates on a weekly basis as well. Moreover, the closer in time the two weekly traces, the more similar the corresponding

Table 1: Cosine distance between steady state distributions derived based on traces collected on a *monthly* basis.

	m1	m2	m3	m4	m5	m6	m7	m8
m1	1	0.9015	0.8336	0.7913	0.7936	0.4679	0.7698	0.7548
m2	-	1	0.8425	0.7779	0.8023	0.4661	0.7525	0.7331
m3	-	-	1	0.8935	0.8501	0.6085	0.7142	0.7275
m4	-	-	-	1	0.9166	0.7234	0.7241	0.7251
m5	-	-	-	-	1	0.7282	0.7388	0.7357
m6	-	-	-	-	-	1	0.4849	0.4859
m7	-	-	-	-	-	-	1	0.9259
m8	-	-	-	-	-	-	-	1

Table 2: Cosine distance between steady state distributions derived based on traces collected on a *weekly* basis.

	w1	w2	w3	w4	w5	w6	w7	w8	w9	w10	w11	w12	w13	w14
w1	1	0.8784	0.9128	0.8674	0.8992	0.8222	0.6126	0.3459	0.5566	0.5425	0.5555	0.7471	0.7394	0.6893
w2	-	1	0.8945	0.8883	0.8641	0.8556	0.6007	0.3136	0.6461	0.6354	0.6311	0.7484	0.7026	0.6688
w3	-	-	1	0.9023	0.9178	0.8498	0.6318	0.3287	0.5745	0.5532	0.5642	0.7655	0.7698	0.7170
w4	-	-	-	1	0.9059	0.8904	0.6079	0.3609	0.5761	0.6062	0.6214	0.7406	0.7037	0.6651
w5	-	-	-	-	1	0.8967	0.6613	0.3445	0.5056	0.5482	0.5668	0.7899	0.7740	0.7439
w6	-	-	-	-	-	1	0.6674	0.3714	0.5460	0.5914	0.6058	0.7583	0.6589	0.6256
w7	-	-	-	-	-	-	1	0.4363	0.3416	0.3860	0.3602	0.4840	0.4688	0.4650
w8	-	-	-	-	-	-	-	1	0.4673	0.4724	0.4616	0.3646	0.2595	0.2177
w9	-	-	-	-	-	-	-	-	1	0.7342	0.8071	0.5331	0.4224	0.3518
w10	-	-	-	-	-	-	-	-	-	1	0.8233	0.6382	0.4534	0.4132
w11	-	-	-	-	-	-	-	-	-	-	1	0.6694	0.4854	0.4214
w12	-	-	-	-	-	-	-	-	-	-	-	1	0.8061	0.7887
w13	-	-	-	-	-	-	-	-	-	-	-	-	1	0.8877
w14	-	-	-	-	-	-	-	-	-	-	-	-	-	1

two steady-state distributions. It is again confirmed that user mobility during the week of *spring break* (**w8**) exhibits different behaviors from other weeks. In addition, several weeks before and after the spring break week, i.e., **w7**, **w9**, **w10**, and **w11**, have somewhat lower similarities to other weeks. This result again shows how administrative events on campus may affect user mobility.

As the mobility data collected from the daily trace is insufficient to construct the transition probability matrix and characterize the mobility behavior, we construct the mobility models at the daily scale by gathering all the data on each day of a week (Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday) during a 8-week period between January 18, 2004 and March 13, 2004, and using them to derive the steady-state distribution at the daily scale. Table 3 gives the cosine distance between steady-state distributions derived based on daily traces. Again the results indicates the close similarity of user mobility between each day of a week, particularly between weekdays. As expected, the similarity in user mobility between a weekday and a weekend day is a little lower.

4.4 Unsimilarity of user mobility for different user groups

To investigate whether different groups of users exhibit similar or different mobility behaviors, we construct mobility models separately on a per-group basis, where the group is defined with respect to the home locations of users and a home location is defined as the AP with which a user is associated most of the time. Since there are too many APs in the traces, we further aggregate all the APs in the same building into one group. As a result, there are seven user groups, each of which corresponds to the following building categories — *Academic*, *Administrative*, *Athlete*, *Library*, *Residence*, *So-*

cial, and *NoHome*. We classify a user to a home location, if he/she is associated with that location more than 50% of his/her total on-time (during which the user is associated with an AP). If a user is not associated with any AP with more than 50% of his/her on-time, he/she is classified into the *NoHome* group.

Figure 6 shows the steady-state distribution of users over APs for each of the seven user groups. It is clear that each group shows dramatically *different* mobility behaviors. Another interesting finding is that users tend to stay or move around their home locations.

4.5 Effect of user sample sizes on steady-state distributions

To investigate the effect of the user sample size on the representativeness of the mobility model thus devised, we vary the number of active users in a given trace *by selecting only a subset of active users as samples in that trace*. Only the data of active users is used to update the transition probability matrix and the residence time for a given period. If a user does not make any association during the period, he/she is not counted as an active user. For each given number, n , of active users, we retrieve randomly the association patterns of n active users from the trace, and construct the mobility model accordingly. For each number n , we construct two mobility models, one based on the trace collected in the period between November 1, 2003 and December 31, 2003, and the other in the period between January 1, 2004 and February 29, 2004. We randomly selected k sets of user samples for each experiment with 100, 200, 400, 800, 1600, and 3200 active users, where $k = 5$ in the case of 100-400 active users; $= 4$ in the case of 800 active users; $= 2$ in the case of 1600 active users; and $= 1$ in the case of 3200 active users. This is due to the limitation of the total number of users available

Table 3: Cosine distance between steady state distributions derived based on traces collected on a *daily* basis.

	Sun	Mon	Tue	Wed	Thu	Fri	Sat
Sun	1	0.8255	0.7982	0.8099	0.7302	0.6438	0.6922
Mon	-	1	0.8711	0.8290	0.8456	0.7376	0.7084
Tue	-	-	1	0.8984	0.8363	0.7521	0.7428
Wed	-	-	-	1	0.8402	0.7380	0.7045
Thu	-	-	-	-	1	0.7900	0.7422
Fri	-	-	-	-	-	1	0.7245
Sat	-	-	-	-	-	-	1

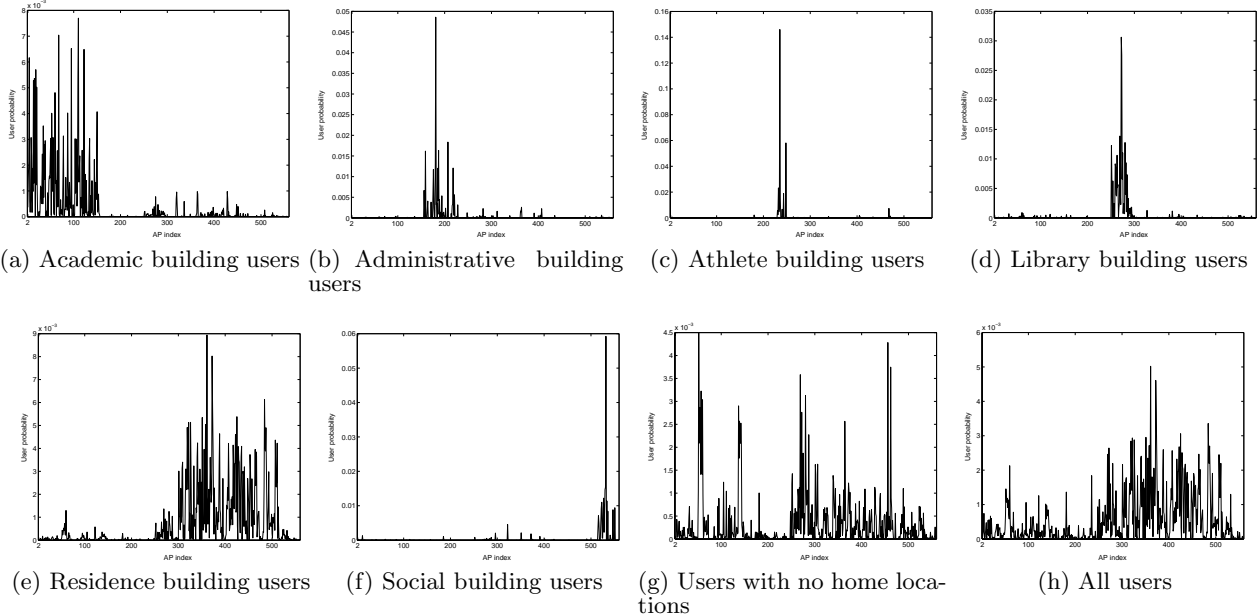


Figure 6: Steady-state distribution of users over APs for different user groups.

in the traces, which is around four thousand active users. We ensure that users selected as samples in the trace with a smaller sample size are selected again in the same trace with a larger user sample size.

Figure 7 gives the results of the cosine distance between the two mobility models (constructed based on the two traces). As expected, the degree of similarity becomes higher as the user sample size increases. Moreover, the variance of the cosine distance is reduced as the sample size n increases. In fact, more samples allow us to obtain more stable statistical estimates for the parameters of interest. An important conclusion that can be drawn from the observation is that, rather than using all the WLAN users in the entire traces, it suffices to using a (reasonably large, ≥ 1600) subset of user samples in building representative mobility models.

5. TRANSIENT BEHAVIOR OF USER MOBILITY AND ITS USE IN TIMED LOCATION PREDICTION

In this section, we carry out a transient behavior analysis of the semi-Markov process. Based on the results derived, we devise a timed location prediction algorithm that predicts the future locations of users — both the future access points they will associate themselves with and the associa-

tion duration. We also demonstrate the utility of the timed location prediction algorithm, by showing how it can be utilized to predict the distribution of future user locations and redistributing loads among neighboring APs.

5.1 Transient behavior of the semi-Markov model

In order to obtain the transient behavior of the semi-Markov model, we have to compute the transient distribution, $\phi_{ij}(t)$, in the evolution equations (Eq. (3)) in Section 3. A numerical solution has already been proposed in the literature [7]. Specifically, the evolution equations can be re-written for the discrete-time homogeneous semi-Markov process as

$$\phi_{ij}(k) = d_{ij}(k) + \sum_{l=1}^m \sum_{\tau=1}^k v_{il}(\tau) \phi_{lj}(k - \tau), \quad (8)$$

where $d_{ij}(k) = (1 - D_i(kh))\delta_{ij}$, $v_{ij}(k) = h\dot{Q}_{ij}(kh)$, and h is the discretization step. Since $Q_{ij}(t)$ is not given in a closed form, but is obtained in the form of histogram distribution from experimental data, we need to further approximate $\dot{Q}_{ij}(kh)$ in the expression of v_{ij} as

$$v_{ij}(k) = \begin{cases} \tilde{Q}_{ij}(h) & \text{for } k = 1 \\ \tilde{Q}_{ij}(kh) - \tilde{Q}_{ij}((k-1)h) & \text{for } k > 1, \end{cases} \quad (9)$$




Figure 7: Cosine distance between two mobility models constructed based on two traces, one for the period between November 1, 2003 and December 31, 2003, and the other for the period between January 1, 2004 and February 29, 2004.

where $\tilde{Q}_{ij}(kh)$ is the empirical distribution of $Q_{ij}(t)$, obtained from the trace data.

5.2 Timed location prediction based on the devised semi-Markov model

Once a semi-Markov mobility model has been constructed based on the association patterns of a user, we can calculate the future transient distribution $\phi_{ij}(k) = \Pr(Z_k = j | Z_0 = i)$ by the discrete-time evolution equations given in Eq. (8) in Section 5.1. With $\phi_{ij}(k)$'s, we can estimate the future locations of a mobile user, given his/her current location.

Specifically, we assume that a WLAN management server is responsible for collecting *on-line* the user mobility information from all the APs and predicting the location of *each* user periodically with a pre-configured interval. Given the past association history of a user <association time, location(AP)>, the transient behavior of the user (i.e. $\phi_{ij}(k)$) can be computed through the following procedure:

Procedure 1: Whenever there is a transition from AP i to j ,

Step 1. Update $P = [p_{ij}]$ and $\tilde{H}_{ij}(kh)$ for $k = 1, 2, \dots, K$.

Step 2. Compute $\tilde{Q}_{ij}(kh) = p_{ij}\tilde{H}_{ij}(kh)$.

Step 3. Obtain $\tilde{D}_i(kh) = \sum_j \tilde{Q}_{ij}(kh)$.

Step 4. Compute $\phi_{ij}(k)$ for $k = 1, 2, \dots, K$ from Eq. (8).

Note that since $\tilde{H}_{ij}(kh)$, $\tilde{Q}_{ij}(kh)$, and $\tilde{D}_i(kh)$ are all the cumulative distribution functions discretized by the discretization step h , K should be set to infinity in the ideal case. In this paper, however, we limit the number of bins so that all the data samples above Kh are aggregated into a single bin. Both the values of K and h determine the “accuracy” of cumulative distribution functions obtained from the experimental data. We will elaborate on the rule with which we determine these values in Section 5.3.

Note that the $\phi_{ij}(k)$ gives the probability that a user is in state j after kh time from the time instant when he/she just made a transition (from somewhere) to state i . However, to

predict the location of a user at every k' time steps, we have to estimate the probability distribution $\hat{\phi}_{ij}(k', s) = \Pr(Z_{s+k'} = j | Z_0 = i, t_{soj} = s)$, i.e., the probability that a user is in state j after $k'h$ time, given that the current state is i and the user has stayed in the current state for time t_{soj} . Note that, as the semi-Markov model does not assume an exponential sojourn time distribution, $\hat{\phi}_{ij}(k', s) \neq \phi_{ij}(k')$. Instead,

$$\hat{\phi}_{ij}(k', s) = \frac{\Pr(Z_{s+k'} = j, t_{soj} = s | Z_0 = i)}{\Pr(t_{soj} = s | Z_0 = i)}. \quad (10)$$

The denominator of Eq. (10) is the probability that the user has been in state i for time s without transitioning to other states, i.e.,

$$\Pr(t_{soj} = s | Z_0 = i) = \Pr(T_{n+1} - T_n > s | X_n = i) = 1 - D_i(s). \quad (11)$$

The numerator of Eq. (10) can be derived in a similar fashion to the evolution equation in Eq. (8), i.e.,

$$\Pr(Z_{s+k'} = j, t_{soj} = s | Z_0 = i) = d_{ij}(s + k') + \sum_{l=1}^m \sum_{\tau=s+1}^{s+k'} v_{il}(\tau) \phi_{lj}(s + k' - \tau). \quad (12)$$

Plugging Eqs. (11) and (12) together into Eq. (10), we obtain

$$\hat{\phi}_{ij}(k', s) = \frac{d_{ij}(s + k') + \sum_{l=1}^m \sum_{\tau=s+1}^{s+k'} v_{il}(\tau) \phi_{lj}(s + k' - \tau)}{[1 - D_i(s)]}. \quad (13)$$

Note that as all the values required in Eq. (13) have been computed in **Procedure 1**, the WLAN management server can readily calculate $\hat{\phi}_{ij}(k', s)$.

The complexity in computing the values in **Procedure 1** when a transition from i to j occurs is $O(K^2m)$. Since K can be set to a reasonably small value without aggravating the prediction accuracy too much (see Section 5.3), this computational overhead is not a significant problem. On the other hand, the computational overhead in obtaining $\hat{\phi}_{ij}(k', s)$ for all j is $O(k'm^2)$ once the values in **Procedure 1** have been computed. If there are N users in a WLAN, the computational overhead at each prediction interval becomes $O(Nk'm^2)$. Although N and m can become several hundreds or thousands, we believe that the overall computational overhead can be handled with today's computing power.

5.3 Evaluation of the timed location prediction algorithm

To evaluate the performance of the proposed timed location prediction algorithm and to investigate the effects of different parameters on prediction accuracy, we have carried out experiments with the use of the WLAN mobility data available at Dartmouth College [21]. We compare 1-location prediction with 2-locations prediction. In 1-location prediction, the location with the highest probability among the corresponding $\hat{\phi}_{ij}(k', s)$ of Eq. (13) is chosen as the future location at the next prediction time. In 2-locations prediction, the location with the next highest probability among the corresponding $\hat{\phi}_{ij}(k', s)$ is also chosen as the second potential location at the next prediction time.

Unlike existing location predictors that do not consider the timing information (e.g., [18, 6]), the proposed algorithm has to specify the time intervals at which predictions

Table 4: Prediction accuracy under various combinations of parameters (median/mean). (1) 1-location prediction (2) 2-locations prediction.

T_p (sec)		$h = 300$ sec	$h = 600$ sec
600	$K = 6$	(1) 0.8377/0.7724 (2) 0.9000/0.8644	(1) 0.8000/0.7311 (2) 0.9000/0.8426
	$K = 12$	(1) 0.8377/0.7661 (2) 0.9129/0.8579	(1) 0.8000/0.7294 (2) 0.9083/0.8459
1800	$K = 6$	(1) 0.6962/0.6778 (2) 0.8500/0.8204	(1) 0.6962/0.6727 (2) 0.8500/0.8189
	$K = 12$	(1) 0.7071/0.6790 (2) 0.8500/0.8223	(1) 0.6923/0.6676 (2) 0.8571/0.8173

are made. Thus, the parameters for the proposed prediction algorithm are (1) h : the discretization step in time; (2) K : the number of bins in the discretized empirical cumulative probability distributions $\tilde{H}_{ij}(kh)$, $\tilde{Q}_{ij}(kh)$, and $\tilde{D}_i(kh)$; and (3) T_p : the period at which predictions are made. To determine the value of K , note that Kh is the maximum time value above which the data samples are aggregated into a single bin. If Kh is less than the prediction interval T_p , we cannot compute $\phi_{ij}(k)$ in Eq. (8) and thus $\hat{\phi}_{ij}(k')$ in Eq. (13). As a result, the value of K has to be constrained by the following rule:

$$K \geq \frac{T_p}{h}. \quad (14)$$

To calculate prediction accuracy, two points are in order. Some users are disconnected from the network for long periods of time, and are in the OFF state for several consecutive prediction intervals. As this may lead to over-estimating the prediction accuracy, we ignore predictions made in those consecutive OFF intervals. That is, only predictions made in non-OFF states are considered *valid*. In addition, if, for any user, the number of *valid* predictions are less than 20% of the total N_{pred} predictions, the prediction results for those users are not counted toward calculation of the prediction accuracy.

Table 4 gives the prediction accuracy of the proposed prediction algorithm under various combinations of the three parameters. A total of 20 predictions are made for each user in each set of experiments. With $T_p = 600$ seconds and $T_p = 1800$ seconds, respectively, this corresponds to an overall intervals of 200 minutes (= 3 hours and 20 minutes) and 10 hours, respectively. The start time of the prediction is Tuesday, Feb 3, 2004 at 14:00 UTC (9:00 EST). For each combination of three parameters h , K , and T_p , the median and mean values of the accuracy of 1-location and 2-locations prediction are shown in the table. Comparing the results of $h = 300$ seconds with those of $h = 600$ seconds for the same values of K and T_p , we observe that higher prediction accuracy is achieved with a smaller discretization step h . This is expected because a smaller discretization step implies more refined characterization of the distribution functions. On the other hand, given the same discretization step h , no notable discrepancy is observed when $T_p = 600$ seconds, although the prediction accuracy for $T_p = 1800$ seconds becomes higher with a larger value of K . Since K can significantly affect the computational overhead of the prediction algorithm, it is desirable to use as small values K as possible. The results of Table 4 indicate that the value of K can be chosen as small as possible, provided that it still satisfies the constraint of Eq. (14).




Figure 8: Cumulative fraction of users against the prediction accuracy with $h = 600$, $K = 12$, and $T_p = 1800$

Figure 8 shows the cumulative fraction of users against prediction accuracy under the combination of $h = 600$, $K = 12$, and $T_p = 1800$. A total of 50 predictions are made, so that the location prediction is made spanning almost one day (25 hours). In addition to 1-location prediction ((a) in Fig. 8) and 2-location prediction ((b) in Fig. 8), we consider one more variation: (c) the second candidate with the next highest probability *among those in the same AS (Aggregate State)* is included as the location with the highest probability. The prediction accuracy for 1-location prediction with the highest probability (case (a)) is reasonable (mean = 0.6908), but if we include the location with the second highest probability as another potential future location in the next time interval, the accuracy is much improved — the mean accuracy of case (b) and case (c) is 0.8406 and 0.8129, respectively. (This was already observed in Table 4.) Among the two variations of 2-locations prediction, choosing the two locations with the highest probabilities (case (b)) has slightly better performance than choosing the second location in the same AS as the location with the highest probability (case (c)).

6. APPLICATION OF TIMED LOCATION PREDICTION — MOBILITY-AWARE LOAD BALANCING

To demonstrate the utility of timed location prediction, we apply the proposed timed location prediction algorithm to balance the loads among APs and devise a mobility-aware load balancing algorithm, where the load at each AP refers to the number of users at that AP in the WLAN. The extension to considering the amount of traffic as the load is quite straightforward, if the traffic characteristics of individual users are given.

6.1 Overview of proposed load balancing algorithm

As discussed in Section 4.2, ping-pong transitions refer to the phenomenon that a user is geographically located among the coverage areas of several APs and re-associates itself with two or three neighboring APs in the same aggregate state. In practice, these transitions occur quite fre-

quently, and lead to the under-estimate of residence times. On the flip side, one can utilize this seemingly undesirable phenomenon for the purpose of redistributing user loads across APs in WLANs. Specifically, APs in the same AS can take turns to serve a user with an acceptable SNR. If a user initially associates with AP i and AP i is overloaded, the user can be instrumented to switch to a lightly-loaded AP j in the same AS.

The idea of load balance can be realized by exploiting the timed location prediction algorithm as follows. A $1 \times m$ bit vector AC is used to control the association of users to APs. Initially, $AC(i) = 1$ for all AP i . Whenever an AP is predicted to be overloaded, the bit corresponding to the AP in the bit vector AC is set to 0. For example, $AC(i) = 0$ ($AC(i) = 1$) if the time location prediction algorithm estimates that the number of users to be associated with AP i in the next interval is greater than (equal to or less than) a pre-configured threshold L . Thus, the AC bit vector is updated in every prediction interval. The value of $AC(i)$ can also be reset to 1 if the actual load at the AP becomes lower than L .

Now when a user moves to another location and attempts to associate itself with an AP in the vicinity of its new location, instead of associating with the AP with the largest signal strength, the user first obtains (in, for example, the AP beacon messages) the bit vector AC and checks the corresponding AC bit. If this bit is 0, the user avoids associating with that AP but instead attempts for another AP with a reasonably large signal strength. Only when there exist no alternative APs with acceptable signal strengths or when the association attempts with alternative APs fail, will the user be allowed to associate itself with the overloaded AP.

An important point that is worthy of mentioning is that the key in the proposed load balancing algorithm is to put the *ping-pong* phenomenon to good use. Thus, the history of ping-pong transitions has to be updated for each user despite the load balancing algorithm. As described above, we achieve this by setting the AC bit to 0 *only when* the load at an AP is greater than a certain threshold L . Thus in lightly loaded conditions users can still make ping-pong transitions between a number of APs.

6.2 Simulation results

To assess the performance of the proposed mobility-aware load balancing algorithm in its capability of redistributing users across APs, we applied the basic idea of the load balancing in the simulation with real traces of 786 users. With the steady-state distribution obtained from the semi-Markov model of each user, we randomly determine the location of each user to obtain the initial distribution of users across APs. We also identify ping-pong transitions and the associated Aggregate States for each user from the same trace data. Then, we apply the proposed load balancing algorithm to obtain the new distribution of users across APs. Since the load is distributed within the same AS for each user, the physical location of each user is not changed.

To quantitatively evaluate the performance, we use the performance metric, the *balance index* β , that is defined as

$$\beta = \left(\sum_{i=1}^m L_i \right)^2 / \left(m \sum_{i=1}^m L_i^2 \right), \quad (15)$$

where m is the number of APs over which the load is redistributed, and L_i is the load at AP i . The above definition is

similar to that defined in [2]. The balance index β is 1 when all APs have exactly the same load, and becomes closer to $1/m$ when the APs are heavily unbalanced.


Figure 9 shows the user distribution before and after applying the proposed load balancing algorithm. Figure 9(a) corresponds to the trace that contains 277 active users and 509 users in the OFF state. In this case, the balance index of the original distribution is $\beta_{org} = 0.1808$, with the maximum load of 9 users at AP 361, while the balance index of the balanced distribution is $\beta_{bal} = 0.3279$, with the maximum load of 3 at AP 373. To investigate the performance under more heavily loaded cases, we artificially reduce the number of users in the OFF state to 201, producing 585 active users (Figure 9(b)).¹ In this case, the original distribution has $\beta_{org} = 0.2846$, with the maximum load of 12 at AP 361, while the balanced distribution has $\beta_{bal} = 0.5064$, with the maximum load of 4 at AP 275. In both cases, an 80% improvement results after applying the proposed load balancing algorithm.

7. CONCLUSION


In this paper, we present a semi-Markov model for characterizing user mobility both in the temporal and spatial domains in wireless LANs. With this semi-Markov model, we will be able to capture and analyze both the steady-state and transient behaviors of user mobility. In particular, from the steady-state analysis of the mobility model, we can identify the long-term mobility characteristics, such as the steady-state user distribution over APs. On the other hand, the transient analysis of the semi-Markov model enables us to devise a timed location prediction algorithm that can predict the future location of a user — *both* the future access points they will associate themselves with and the association duration. We also demonstrate the utility of the proposed timed location prediction algorithm by showing how it can be exploited, along with the ping-pong phenomenon, to re-balance user loads among APs. An 80% improvement (in terms of the balance index) has been reported in the simulation study.

There are several refinements that can be made for the timed location prediction algorithm. For example, the algorithm can be further improved by (i) building individual semi-Markov models for each time of the day, each day of the week, or each week of the month; and (ii) varying the rules for the transition probability matrix and the sojourn time distributions and giving more weights to recent mobility trace data. Another research avenue is to prototype the proposed load balance algorithm and empirically study the performance. We are in the process of leveraging open-source Linksys WRT54G Wireless-G routers as APs and customizing its kernel sources with the proposed load balance algorithm. These Linksys routers are inexpensive (less than \$60 per piece), support IEEE 802.11b/g, and run Linux kernel 2.4 with its source codes [22] available under GPL. This allows us to customize the kernel sources and embed in it the Kismet 802.11 layer-2 wireless monitoring software [23], as well as other required monitoring functions including controlled channel scanning, neighborhood discovery, and selective capturing of management, control and data frames.

¹We have carried out experiments with the other mixture of active/inactive users, and observed similar performance trends.



(a) 277 active and 509 inactive users



(b) 585 active and 201 inactive users

Figure 9: User distribution before and after load balancing

The Ethernet interface at the Linksys router allows APs to send user association patterns to the WLAN management server and the server to send the bit vector AC. The empirical results, along with the systems prototyping details, will be reported in a fore-coming manuscript.

8. ACKNOWLEDGMENTS

We gratefully acknowledge the use of wireless data from the CRAWDAD archive at Dartmouth College.

9. REFERENCES

- [1] Anand Balachandran, Geoffrey M. Voelker, Paramvir Bahl, and P. Venkat Rangan. Characterizing user behavior and network performance in a public wireless LAN. In *Proceedings of IEEE WMCSA '02*, 2002.
- [2] Anand Balachandran, Paramvir Bahl, and Geoffrey M. Voelker. Hot-spot congestion relief in public-area wireless networks. In *Proceedings of IEEE WMCSA '02*, 2002.
- [3] Magdalena Balazinska and Paul Castro. Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network. In *Proceedings of MobiSys'03*, San Francisco, CA, May 2003.
- [4] D. Bhattacharjee, A. Rao, C. Shah, M. Shah, and A. Helmy. Empirical modeling of campus-wide pedestrian mobility: observations on the USC campus. In *Proceedings of IEEE VTC'04*, Los Angeles, CA, September 2004.
- [5] T. Camp, J. Boleng, and V. Davies. Mobility models for ad hoc network simulations. *Wireless Communication and Mobile Computing*, 2(5):483–502, 2002.
- [6] F. Chinchilla, M. Lindsey, and M. Papadopouli. Analysis of wireless information locality and association patterns in a campus. In *Proceedings of IEEE INFOCOM'04*, Hong Kong, China, March 2004.
- [7] Gianfranco Corradi, Jacques Janssen, and Raimondo Manca. Numerical treatment of homogeneous semi-Markov processes in transient case—a straightforward approach. *Methodology and Computing in Applied Probability*, 3:233–246, 2004.
- [8] N. Golmie. Interference in the 2.4 GHz band. In *Proceedings of International Conference on Applications and Services in Wireless Networks*, pages 187–199, 2001.
- [9] Tristan Henderson, David Kotz, and Ilya Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of ACM MobiCom'04*, Philadelphia, PA, 2004.
- [10] Wei-je Hsu, Kashyap Merchant, Haw-wei Shu, Chih-hsin Hsu, and Ahmed Helmy. Preference-based mobility model and the case for congestion relief in WLANs using Ad hoc Networks. In *Proceedings of IEEE VTC'04*, Los Angeles, CA, September 2004.
- [11] Ravi Jain, Anupama Shivaprasad, Dan Lelescu, and Xiaoning He. Towards a model for user mobility and registration patterns. *Mobile Computing and Communications Review*, 8(4):59–62, Oct. 2004.
- [12] Ravi Jain, Dan Lelescu, and Mahadevan Balakrishnan. Model T: an empirical model for user registration patterns in a campus wireless LAN. In *Proceedings of ACM MobiCom'05*, Cologne, Germany, August 2005.
- [13] A. Jardosh, E. Belding-Royer, K. Almeroth, and S. Suri. Towards realistic mobility models for ad hoc networks. In *Proceedings of ACM MobiCom'03*, September 2005.
- [14] Minkyung Kim and David Kotz. Classifying the mobility of users and the popularity of access points. In *Proceedings of Int. Workshop on Location- and Context-Awareness (LoCA '05)*, May 2005.
- [15] Minkyung Kim and David Kotz. Modeling users' mobility among WiFi access points. In *Proceedings of WiTMeMo'05*, June 2005.
- [16] David Kotz and Kobby Essien. Analysis of a campus-wide wireless network. In *Proceedings of ACM MobiCom'02*, September 2002, pages 107–118.
- [17] Gerard Salton and Chris Buckley. Term weighting approaches in automatic text retrieval. Technical Report TR87-881, Department of Computer Science, Cornell University, November 1987. Available at <http://techreports.library.cornell.edu:8081/Dienst/UI/1.0/Display/cul.cs/TR87-881/>.
- [18] Libo Song, David Kotz, Ravi Jain, and Xiaoning He. Evaluating location predictors with extensive Wi-Fi mobility data. In *Proceedings of IEEE INFOCOM'04*, Hong Kong, China, March 2004.
- [19] Diane Tang and Mary Baker. Analysis of a local-area wireless network. In *Proceedings of ACM MobiCom'00*, August 2000.
- [20] C. Tudeuce and T. Gross. A mobility model based on WLAN traces and its validation. In *Proceedings of IEEE INFOCOM'05*, Miami, FL, March 2005.
- [21] A community resource for archiving wireless data at Dartmouth. <http://crawdad.cs.dartmouth.edu/data.php>.
- [22] Linksys WRT54G Wireless-G broadband router source codes. <http://www.linksys.com/support/gpl.asp>.
- [23] Kismet. <http://www.kismetwireless.net>.