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

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

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

Don Towsley
University of Massachusetts Amherst
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 and validate a simple mixed queueing network model of user mobility among access points in a campus network. We identify two classes of users, an open and a closed class, corresponding to mobile users that visit the network for a short time before departure, and users that are always resident in the network during the observation period. 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 mixed 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 *mixed queueing network models* of user mobility among access points (APs), consisting of users in an *open* and a *closed* class. Users in the open class arrive according to a random process, move from AP to AP, and depart the network. Users of this class might be laptop users, leaving the network after being served in public hot spots; here, each new arrival to the campus network is treated as a new, independent customer, considerably simplifying the computation of performance metrics. Users in the closed class are of a fixed population, circulating among APs but never leaving the network. These customers could be users carrying their smart phones that are always connected to campus APs, or users whose laptops are similarly always connected.

The question we address is the following: can such a mixed 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 and average path length) 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 with the performance actually observed in the traces, our findings here are that such a simple model of mobility can indeed be used to accurately predict a number of performance measures of interest. We also illustrate the application of our model in system-level performance and dimensioning analyses.

The remainder of this paper is structured as follows. Section II describes the traces we use, and how we pre-process them. Section III presents our proposed queueing network model, which is validated in Section IV. We show an application of our model for network dimensioning in Section V. Related work is discussed in Section VI and Section VII concludes this paper.

II. The Traces

There are several publicly available traces of long term user activity in wireless LANs (WLANs) [10] [6] [3]. As we are interested in modeling user-level mobility among APs in larger (e.g., campus-level) wireless networks, we seek traces that contain information of user movements in a large network (both in terms of the number of APs and the user population) over a long period of time. The trace we use to construct our model, and against which we will validate model predictions, is the Dartmouth trace [6], which records wireless user activity for a 17-week period, from 11/2/2003 to 2/28/2004.

A. Trace Description

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 authenticate, associate, roam, disassociate, or deauthenticate. We find, 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 cannot be determined 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. In response to 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 estimate a user's departure time by that user's absence in a subsequent poll, as discussed below.

B. Trace Preprocessing

To circumvent the problem of diurnal user behavior (people's daytime and nighttime behaviors are different), we only consider user activity during those periods of time when the university is most active. Hence, we extracted traces from 9 AM to 5 PM of each day (as will be discussed below), and removed all weekend, holiday, and inter-session periods as well. The processed trace contains 544 APs across 6 different types of buildings (as listed in Table I), with 5,715 distinct MAC addresses.

1) Departure Length Threshold: 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) and lasts until the user disassociates from all network APs.

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 later SNMP reports. There are three possible explanations for this:

- The user left the network and later returned.
- The user was in motion, leaving one AP and then later associating with another AP.
- An SNMP update was missing or lost.

Without explicit disassociations, it is difficult to determine 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.

We introduce a departure length threshold, T_d, such that if the user does not appear in an SNMP report for an amount of time greater than 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 $\delta > T_d$ (with no SNMP reports of that user during the intervening δ) 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 10 minutes or less.

2) The Observation Period: Since we are interested in the period of time that the campus network is most active, we examine the trace during the time that there are a relatively large number of users in the network, and the network is relatively stable and stationary. Figure 2 plots the average weekday user arrival rate to the network at different times of day, averaged over the entire measurement period. We note that the user arrival rate to the network increases sharply between 6 AM to 9 AM, and remains relatively stable until 5 PM, and then slowly decreases till 6 AM the next morning. Thus, we chose to model user activity during the weekday hours of 9 to 5, as discussed above.

We also found a not-insignificant fraction of users who were present in the network at 9 AM and remained in the network after 5 PM (i.e., have their wireless devices always in the connected-mode). We thus divide network users into two groups: those present all day (9 AM - 5 PM) and those that first arrive and depart during the day. We refer to users in the first group as being in the "*closed class*", and refer to the second group of users - those that come and go - as the "*open class*" of users. For each day, we computed the population of this closed class, and found that it was relatively stable over the entire measurement period. On average, the population of this closed class is, $N = 441$.

III. The Model

We model the campus wireless network of APs as a mixed network of infinite server (i.e., $\cdot/G/\infty$) queues, where each AP is represented by an infinite server queue. The network is mixed in that it serves two classes of users: a closed class and an open class. The *closed class* consists of N users that always remain in the system; users in the *open class* arrive according to a Poisson process and can depart the system. Since each AP is modeled as an infinite server queue, each user (regardless of class) is served immediately (there being an infinite number of servers) and independently of the other users[1].

Before discussing the details of our model, let us first introduce the key parameters and the notation that we will use in our model ($1 \le i, j \le 544$):

- U: total number of users at steady state
- M: total number of APs on campus
- N: total number of users in closed class
- U_i: number of users associated with AP_i
- $1/\mu_i$: the expected user residence 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 an open class user moving from AP_i to AP_j
- v_i: fraction of time of a closed class user visits AP_i

We will refer to the *open class* as class 0, and the *closed class* as class 1. Since each AP is modeled as an infinite server queue, arriving customers of both classes are served immediately and independently. Hence, we can treat the network as a combination of two independent networks.

A. Open Class

We will model the exogenous arrivals of new users to each AP as a Poisson process, and hence model each AP of this open class as an $M/G/\infty$ queue (i.e., an infinite server queue with Poisson arrivals and general service times).

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

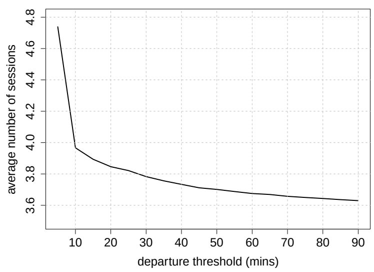

Fig. 1. Average number of sessions for various departure length threshold.

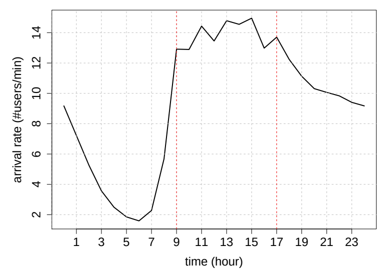

Fig. 2. Average user arrival rate to the network.

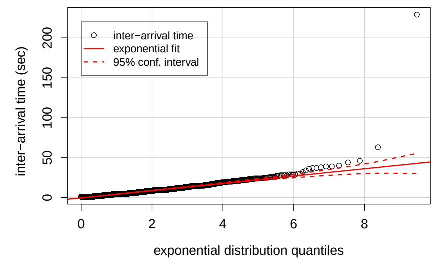

Fig. 3. Q-Q plot of inter-arrival time distribution against exponential distribution.

Since the aggregation of Poisson processes is still a Poisson process, then the aggregate arrivals to the campus can also be characterized as a Poisson process. A natural question is how closely the actual exogenous arrival process is to a Poisson process, with its exponential inter-arrival times. To investigate this question, we consider how well the empirical user inter-arrival time distribution between 9 AM and 5 PM matches an exponential distribution. We look at each set of daily arrivals to the campus network. As a standard statistical measure in regression analysis, we use R^2 (the square correlation, $0 \leq R^2 \leq 1$) to quantify how well the empirical inter-arrival times for a day are fitted by the corresponding exponential distributions [7].

We find that the average R^2 of all examined days is 0.9664. That is, on average, 96.64% of the data variation in our daily traces is explained by a corresponding exponential distribution. This high value of R^2 supports our assumption that daily user inter-arrival times of the campus network are exponentially distributed. Figure 3 is the quantile-to-quantile plot (Q-Q plot) of empirical inter-arrival times against an exponential distribution for the worst fitted day, with $R^2 = 0.81$, throughout the entire day, also showing 95% confidence interval of the estimated regression line.

Given this relatively good fit between the observed inter-arrival times and the exponential distribution, we model the exogenous arrivals to AP_i as a Poisson process with rate γ_i. We then have the aggregate arrival rate to AP_i as:

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

where p_{ji} denotes the probability of a user's moving from AP_j to AP_i. The probability that a user leaves the system is $p_{i0} = 1 - \sum_{j=1}^{M} p_{ij}$. Let $\pi_0(\vec{u}) = P(U_{01} = u_1, \ldots, U_{0M} = u_M)$ denote the joint steady state AP occupancy probability distribution, where $u_i = 0, 1, \ldots$ and $1 \leq i \leq M$. The marginal AP_i occupancy distribution is $P(U_{0i} = u_i) = e^{-\rho_{0i}} \frac{\rho_{0i}^{u_i}}{u_i!}$, and the joint steady state AP occupancy probability distribution hence has the following product form [4]:

$$\begin{aligned} \pi_0(\vec{u}) &= P(U_{01} = u_1, \ldots, U_{0M} = u_M) \\ &= \prod_{i=1}^{M} \frac{\rho_{0i}^{u_i} e^{-\rho_{0i}}}{u_i!}, u_i \geq 0; 1 \leq i \leq M \end{aligned} \tag{2}$$

B. Closed Class

As discussed previously, since each AP in the network is modeled as a $\cdot/G/\infty$ queue, user behavior of this closed class is independent of user behavior in the open class. We can, therefore, model the AP occupancy distribution of this closed class as a binomial distribution (since users of this class always circulate among APs and never leave the network), and the joint distribution is given as a multinomial distribution.

As we are only interested in the marginal statistics of each AP, we only present the marginal distribution of the user occupancy at AP_i as $P(U_{1i} = u_i) = \binom{N}{u_i} v_i^{u_i} (1 - v_i)^{N - u_i}$. Note that v_i, the probability that a closed class user visits AP_i, can be obtained directly from the trace, and N is the average number (to the closest integer) of always-active users over the entire observation period.

C. Mixed Queueing Network

Integrating Section III-A and III-B, our proposed mixed queueing network mobility model is the combination of the above open and closed classes of users, and the user occupancy distribution of AP_i is simply the convolution of distributions of the closed and the open network: $U_i = U_{0i} + U_{1i}$.

In this work, we are investigating the performance of a large scale campus network where the population in the closed network is large, and the probability of finding a user at AP_i is relatively small. Hence, we can approximate the binomial distribution $b(N, v_i)$ by a Poisson distribution with parameter ρ_{1i} such that $\rho_{1i} = N \times v_i$ (as suggested in [7], we have $\max\{v_i\} = 0.017 < 0.1$ and $N = 441 \geq 100$).

Hence, the convolution of two Poisson distributions leads to the marginal occupancy distribution of AP_i in the mixed network with $\rho_i = \rho_{0i} + \rho_{1i}$:

$$\begin{aligned} P\{U_i = u_i\} &\approx \sum_{k=0}^{u_i} e^{-\rho_{0i}} \frac{\rho_{0i}^k}{k!} e^{-\rho_{1i}} \frac{\rho_{1i}^{u_i - k}}{(u_i - k)!} \\ &= \frac{e^{-(\rho_{0i} + \rho_{1i})}}{u_i!} (\rho_{0i} + \rho_{1i})^{u_i} = e^{-\rho_i} \frac{\rho_i^{u_i}}{u_i!} \end{aligned} \tag{3}$$

With this simple expression for the user occupancy distribution of users in both the open and the closed class, in the following section, we will investigate how closely the predictions from our model match empirically observed results.

IV. Model Validation

We validate our model against the empirical trace data by considering the following metrics: AP occupancy distribution,

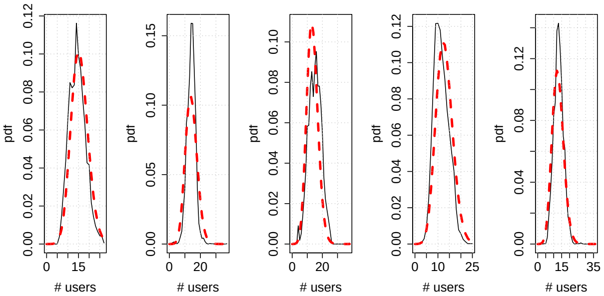

Fig. 4. Occupancy distribution of the most heavily loaded 5 APs.

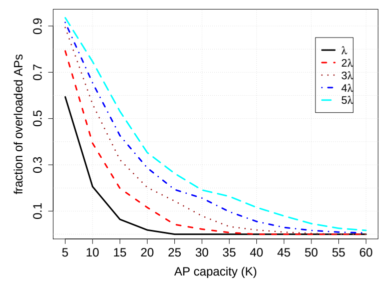

Fig. 5. Scale up arrival rates.

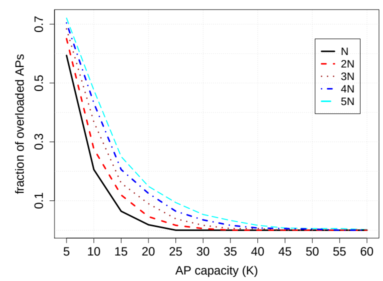

Fig. 6. Scale up closed population.

mean user sojourn time (i.e., a user's session time in the system), and the average number of transitions of a user during a session.

A. AP Occupancy Distribution

We first consider how well the model-predicted AP occupancy distribution matches the empirically-observed occupancy distribution. We observe that the most heavily loaded APs are in residential buildings, followed by academic buildings. Figure 4 shows the user occupancy distributions of the five most heavily loaded APs on campus. In each plot, the dashed line is the result predicted by the model (with load, $\rho_i \approx \lambda_i/\mu_i + Nv_i$ at AP_i), while the solid line is the empirical population distribution. We note a good match between the model predictions and the empirical 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 K-S test is used to determine whether a hypothesized distribution (i.e., predictions from our mixed $\cdot/G/\infty$ queueing model) matches the empirical distribution, and is not sensitive to the binning of our data (in our case, the number of users), as in the Chi-square test [7].

In our study, we set the significance level of K-S tests to 0.05 (i.e., a 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 AP occupancy. Again, we note a good match between the model-predicted and empirically-observed results. The overall accuracy of predictions of user population distribution reported by K-S tests is 93.57%.

AP Type	# passed K-S test	total No. APs	Ratio
Residential	207	211	98.1%
Academic	131	152	86.18%
Administrative	68	69	98.55 %
Social	44	44	100%
Library	40	49	81.63%
Athletic	19	19	100%
Total	509	544	93.57 %

TABLE I
NUMBER OF MODEL-PREDICTED AP OCCUPANCY DISTRIBUTIONS THAT PASSED K-S TEST.

B. Mean Sojourn Time and Average Path Length

We next analytically compare the mean sojourn time (i.e., the duration of a user's session length), and the average path length (i.e., the number of transitions that a user makes before leaving the network) predictions from our model against those of the empirical data.

We first complete the entries in the transition matrix related to the additional state, 0, a state that models users leaving the network. We add $p_{00} = 0$, and $p_{0i} = \gamma_i / \sum_{j=1}^{M} \gamma_j$, for $i = 1, \ldots, M$, as the fraction of exogenous arrivals to each AP, to the transition matrix.

Let $\mathcal{M} = \{AP_1 \ldots AP_M\}$ be the set of states in which user transitions result in their remaining in the network, and AP_0 be the exit state, such that $|\mathcal{M}| = M$. The transition probability matrix P is of the canonical form such that the submatrix $Q = P\{p_{ij}\}, i, j \in \mathcal{M}$, governs the transitions of a user that moves from one AP to another AP in the network; the submatrix $R = P\{p_{k0}\}, k \in \mathcal{M}$ concerns a user's departure (i.e., transition from the network to state 0) [8].

We now use above notations to derive the expected user sojourn time and the average path length.

1) Mean Sojourn Time: Let T_i be the time user spends in system given that he/she is currently at AP_i ($i \in \mathcal{M}$), including the period of time staying at AP_i. We then have

$$\mathbb{E}[T_i] = \frac{1}{\mu_i} + \sum_{j \in M} p_{ij} \mathbb{E}[T_j] \quad (4)$$

Define the diagonal matrix $D = diag(1/\mu_1, \ldots, 1/\mu_M)$, and $T = \{\mathbb{E}[T_i]\}$. Then (4) can be expressed as $T = \{\mathbb{E}[T_i]\} = D + QT$. Note that the inverse of $(I - Q)$ exists [8], and thus the mean system stay time can be computed and represented as $T = (I - Q)^{-1} D$.

As we have computed the exogenous arrival rate to AP_i, γ_i, from the trace, for each user's first arriving to the campus network, he/she has a probability $p_{0i} = \gamma_i / \sum_{j=1}^{M} \gamma_j$ of associating with AP_i. Let S be the user sojourn time, hence the mean user sojourn time of a user is $\mathbb{E}S = \sum_{i \in \mathcal{M}} p_{0i} \mathbb{E}[T_i]$. The mean sojourn time observed in the *empirical data* is 2.23 hours, and the corresponding prediction from our *analytical model* is $\mathbb{E}S = 2.36$ hours.

2) Average Path Length: The average path length can be easily derived using our analysis above and setting the expected stay time at each AP to 1 (i.e., setting diagonal matrix $D = I$). The average path length observed in the *empirical data* is 2.07 transitions, and the corresponding prediction from our *analytical model* is 2.10 transitions.

In summary, our model predicts an average path length of 2.1, which is very close to the empirical value of 2.07. Recall that we also found that the predicted mean sojourn time matches the empirical mean sojourn time well, with only 7.8 minutes difference in the sojourn times, where the mean sojourn time is longer than 2 hours.

V. Applications and Network Dimensioning

Our proposed model can now be used to analyze the performance and dimensioning wireless networks. Suppose each AP has a capacity to serve K users at a time with a guaranteed quality service, we then say that AP_i is *overloaded* if $P(U_i > K) > 0.01$. That is, an AP would function properly if 99% of time the number of users associated with it is smaller than its capacity K. Note that in our model, both open and closed class users contribute to the load, ρ_i, of AP_i ($\rho_i \approx \rho_{0i} + \rho_{1i} = \frac{\lambda_i}{\mu_i} + Nv_i$). We assume that the mobility of mobile users (i.e., $1/\mu_i$ and v_i) does not change in our study.

We first look at the case when the exogenous arrivals to the network increase. In such scenario, what is the fraction of APs that will become overloaded? Figure 5 shows the fraction of overloaded APs for different AP capacities (from 5 to 60) under different levels of arrival rates. The solid line is the load and population computed from the trace; if we seek to have a stable campus network with fewer than 5% overloaded APs when the campus population (of the open class) increases five-fold, then the AP capacity should be tripled.

Secondly, we investigate the case where additional smart phone users are introduced to the campus network given that the exogenous arrival rate to the network APs remains constant. Figure 6 shows the fraction of overloaded APs with respect to different AP capacities at different scales of closed class population being introduced, and the solid line is the load and population of the trace used. Similarly, if we hope to have a stable campus network with fewer than 5% overloaded APs when the closed class population increases five-fold, a doubling of AP capacity from 15 users to 30 users will allow the campus network running more smoothly.

VI. Related Work

There are several works on modeling mobility in cellular networks. Kim and Choi [9] developed a mobility model of cell phone users, but focused on calculating the call handoff rate and loss probabilities in a cell for an assumed exogenous arrival and inter-cell probabilistic mobility model. Ashtiani et al. [2] characterized spatial traffic distribution of a fixed number of active users in a closed network to obtain user location density in cellular networks. Both works made additional assumptions about cell dwell time and call holding time with no supporting field data. Ghosh et al. [5] examined traces of specific types of public Wi-Fi hotspots. They modeled the number of users and their stay time at each hotspot, but did not consider user mobility.

To our knowledge, this paper presents the first analytical model with a simple queueing model of mobility with empirical validation to predict various network and user-level measures in a simple yet efficient manner.

VII. Conclusion

In this paper, we proposed a simple mixed queueing network model of mobility with infinite server ($\cdot/G/\infty$) queues as APs on campus. We divide mobile users into two groups, the open class and the closed class. In such a network, users in the open class arrive to the network according to a Poisson process, move from AP to AP, and depart the network; users in the closed class are of a fixed population, circulating among APs, and they always remain active. We show that our model accurately predicts the AP occupancy distribution, the average number of AP transitions a user makes, and the mean sojourn time (of open class users) compared to results from empirical data. We also show that our model can be used for network dimensioning, answering "what if" questions, such as how user performance changes as the number of users increase, and the amount of capacity that must be deployed to maintain user-perceived performance within a specified range.

As our mobility model, with its many simplifying assumptions, has shown its use of predicting a set of performance metrics accurately, we are also interested in exploring its capability of capturing more detailed user behaviors. We are currently investigating this as future work.

Acknowledgment

This research was sponsored by US Army Research laboratory and the UK 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 authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. This material is also based upon work supported by the National Science Foundation under Grant No. CNS-1040781.

References

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

[2] F. Ashtiani, J. A. Salehi, and M. R. Aref. Mobility modeling and analytical solution for spatial traffic distribution in wireless multimedia networks. *IEEE JSAC*, 21:1699–1709, 2003.

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

[4] F. Baskett, K. M. Chandy, R. R. Muntz, and F. G. Palacios. Open, closed, and mixed networks of queues with different classes of customers. *J. ACM*, 22:248–260, April 1975.

[5] A. Ghosh, R. Jana, V. Ramaswami, J. Rowland, and N. Shankaranarayanan. Modeling and characterization of large-scale wi-fi traffic in public hot-spots. In *IEEE Infocom*, pages 2921–2929, 2011.

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

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

[8] J. Kemeny and L. Snell. *Finite Markov Chains*. Van Nostrand, 1960.

[9] K. Kim and H. 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.

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