

MobiCom Poster Abstract: Bandwidth Reservation using WLAN Handoff Prediction

Libo Song[†], Udayan Deshpande[†], Ulas Kozat[‡], David Kotz[†], Ravi Jain[‡]

[†] {lsong, udayan, dfk}@cs.dartmouth.edu Dartmouth College, Hanover, NH 03755

[‡] kozat@docomolabs-usa.com, ravi.jain@acm.org DoCoMo USA Labs, San Jose, CA 95110

I. Introduction

Many network services may be improved or enabled by successful predictions of users' future mobility. The success of predictions depend on how much accuracy can be achieved on real data and on the sensitivity of particular applications to this achievable accuracy. We investigate these issues for the case of advanced bandwidth reservation using real WLAN traces collected on the Dartmouth College campus [1].

In our system model, we envision a wireless network in which users associate with one access point (AP) at a time, and can change their associations from one AP to another as needed to remain connected. Such reassociation can be caused by a roaming user or the changes in connection quality. This sequence of handoffs for each user, indicating the time and AP of each move (or "OFF" when the user leaves the network), is their movement history. We assume that there is a centralized or distributed mechanism in place for (i) collecting the history, (ii) performing online predictions, and (iii) distributing prediction results to relevant application agents.

In the case study, we focus on VoIP as the application of interest; we evaluate mobility predictors for advanced bandwidth reservation to maintain VoIP service quality after handoffs. We measure the performance using application-specific *call drop rate* and *call block rate* metrics. The results show that intelligent prediction can lead to significant reductions in the rate at which active calls are dropped due to handoffs with marginal increments in the rate at which new calls are blocked.

II. Predictors

We designed a CDF time predictor that produces the probability that the time of the next move is less than a given value. It computes the observed cumulative distribution function (CDF) of the historic values, and using the CDF to measure the probability of a given value appearing in the distribution.

Consider a history H of values v_1, v_2, \dots, v_n . Suppose V is the random variate, which outputs the actual values in H , and P is its distribution. The CDF predictor computes the observed CDF function \hat{P} of V

from the histogram, $\hat{P}(V < v) = \frac{1}{n} \sum_{i=1}^n I(v_i < v)$, where I is the indicator function. In a similar fashion, we can compute the probability of values occurring in range $a \leq V < b$, by simply computing $\hat{P}(a \leq V < b) \approx (\hat{P}(V < b) - \hat{P}(V < a))$.

We combine the CDF time predictor with location predictors from our previous work [3] for integrated location and time predictions; we call it the *MarkovCDF Predictor*. This integrated predictor outputs a vector of probabilities, one for each AP that the user will move to within a certain time threshold. Therefore, we can make reservations according to the probabilities of predicted locations and times. These predictors build their internal tables on per-user basis, but it is equally possible to build aggregate tables from all users' movement histories. We name them the *MarkovCDF Individual* predictor and the *MarkovCDF Aggregate* predictor, respectively.

We introduce a simple "straw-man" predictor, the *Neighbor Graph Predictor*, to compare with our MarkovCDF predictor. Using users' current neighbor locations as the prediction is an obvious way to predict future locations. Mishra et al. [2] present an algorithm to dynamically build a user's neighbor graph to cache context for fast handoffs.

III. Case study

In our Case Study, we simulate a wireless network that is capable of supporting roaming telephone users. When a user has an ongoing call and moves from one access point to another, we refer to that call as a *hand-off call*. When a user initiates a call, we refer to that call as a *new call*. All calls require dedicated bandwidth at their current AP. If the AP lacks the bandwidth for a new or handoff call, the call fails: a failed handoff call is a "call drop" and a failed new call is a "call block". The literature often assumes that call drops are much more frustrating to users than call blocks, so the goal of mobility prediction in this particular application is to reserve bandwidth, in advance of handoffs, to reduce call drops at the expense of a small increase in call blocks. Specifically, we define the drop rate $DR = \frac{N(\text{dropped calls})}{N(\text{attempted call handoffs})}$,




Figure 1: The ratio (Base DR /Reserved DR) for all the predictors with training. Higher ratio is better.

and the block rate $BR = \frac{N(\text{blocked calls})}{N(\text{attempted calls})}$, where $N(x)$ is the number of x .

We evaluated the three predictors described above using two months real mobility traces collected on our campus-wide wireless network; these include 545 APs and 6,181 users. We use the first months traces to train our MarkovCDF predictors and build the neighbor graph for the Neighbor Graph predictor. We also used exponentially distributed call duration and inter-call time to simulate the voice calls.

We implemented four reservation schemes:

All probabilities, non-normalized: We reserve bandwidth proportional to the probabilities returned. The sum of the probabilities is not necessarily 1 for the MarkovCDF predictors.

All probabilities, normalized: We reserve bandwidth proportional to the normalized probabilities returned by the predictor. The normalization makes the sum of the probabilities to be 1.

Top-3 normalized probabilities: We make reservations at the three most probable APs, proportional to the normalized probabilities returned by the predictor. The three normalized probabilities sum to 1.

Top-3 probabilities, 1BU: We make reservations of 1 bandwidth unit (BU) at each of the three most probable APs.

Figure 1 shows the ratios of Base DR and Reserved DR , where the Base DR is the drop-rate without using any reservation scheme, and the Reserved DR is the drop-rate with reservations. Figure 2 shows the ratios of Reserved BR and Base BR , where the Base BR is the block-rate without using any reservation scheme, and the Reserved BR is the block-rate with reservations.

The Neighbor Graph Predictor reduces the drop-




Figure 2: The ratio (Reserved BR /Base BR) for all the predictors with training. Lower ratio is better.

rate the most, while it increases the call block-rate the most. Although the low drop-rate of the Neighbor Graph Predictor is desirable, it does so by wasting valuable bandwidth resources.

We observe that using the “Top-3 normalized probabilities” scheme is the best of the four resource reservation schemes for the MarkovCDF predictor, because it concentrates resources at the top three APs that the user is most likely to roam to. The “Top-3 probabilities with 1 BU” scheme reserves excessive resources and causes an increase in the block rate. In the case of the MarkovCDF Individual predictor, the “all probabilities normalized” reservation scheme performs similar to the “Top-3 normalized” scheme because, it is likely that the sum of the top three probabilities is equal to or close to 1. In case of the MarkovCDF predictors, the “all probabilities non-normalized” scheme does not improve the DR greatly neither does it worsen the BR greatly, because the probabilities returned by the predictor are low.

References

- [1] D. Kotz and K. Essien. Analysis of a campus-wide wireless network. *Wireless Networks*, 11:115–133, 2005.
- [2] A. Mishra, M. ho Shin, and W. A. Arbaugh. Context caching using neighbor graphs for fast hand-offs in a wireless network. In *Proc. IEEE INFOCOM 2004*, pages 351–361, March 2004.
- [3] L. Song, D. Kotz, R. Jain, and X. He. Evaluating location predictors with extensive Wi-Fi mobility data. In *Proc. IEEE INFOCOM 2004*, pages 1414 – 1424, March 2004.