

Evaluating location predictors with extensive Wi-Fi mobility data

Libo Song*, David Kotz*, Ravi Jain†, Xiaoning He†

*{lsong, dfk}@cs.dartmouth.edu, Department of Computer Science, Dartmouth College, Hanover, NH, USA
 †{jain, xiaoning}@docomolabs-usa.com, NTT DoCoMo Labs, Palo Alto, CA, USA

Abstract—Location is an important feature for many applications, and wireless networks can better serve their clients by anticipating client mobility. As a result, many location predictors have been proposed in the literature, though few have been evaluated with empirical evidence. This poster reports on the results of the first extensive empirical evaluation of location predictors, using a two-year trace of the mobility patterns of over 6,000 users on Dartmouth College’s campus-wide Wi-Fi wireless network.

We implemented and compared the prediction accuracy of several location predictors drawn from two major families of domain-independent predictors, namely Markov-based and compression-based predictors. We found that low-order Markov predictors fared as well or better than the more complex and more space-consuming compression-based predictors. Predictors of both families fail to make a prediction when the recent context has not been previously seen. To overcome this drawback, we added a simple fallback feature to each predictor and found that it significantly enhanced its accuracy in exchange for modest effort.

I. INTRODUCTION

A fundamental problem in mobile computing and wireless networks is the ability to track and predict the location of mobile devices. An accurate location predictor can significantly improve the performance or reliability of wireless network protocols, the wireless network infrastructure itself, and many applications in pervasive computing. These improvements lead to a better user experience, to a more cost-effective infrastructure, or both.

Location prediction has been proposed in many areas of wireless cellular networks as a means of enhancing performance, including better mobility management, improved assignment of cells to location areas, more efficient paging, and call admission control.

To the best of our knowledge, no other researchers have evaluated location predictors with extensive mobility data from real users.

In this poster we compare the most significant domain-independent predictors using a large set of user mobility data collected at Dartmouth College. In this data set, we recorded for two years the sequence of wireless cells (Wi-Fi access points) frequented by more than 6000 users.

We found that the simple Markov predictors performed as well or better than the more complicated LZ predictors, with smaller data structures. We also found that many predictors often fail to make any prediction, but our simple fallback technique provides a prediction and improves overall accuracy.

II. BACKGROUND

A. Domain-independent predictors

We consider here only domain-independent predictors. We are interested in *on-line predictors*, which examine the history so far, extract the current context, and predict the next location. Once the next location is known, the history is now one location longer, and the predictor updates its internal tables in preparation for the next prediction.

We implemented two families of domain-independent predictors, Order- k Markov predictors and LZ-based predictors. For detail of these predictors, please see a survey paper by Cheng, Jain, and van den Berg [1].

B. Accuracy metric

During an on-line scan of the location history, the predictor is given a chance to predict each location. There are three possible outcomes for this prediction, when compared to the actual location: correct, incorrect and “no prediction”.

We define the *accuracy* of a predictor for a particular user to be the fraction of locations for which the predictor correctly identified the next move. Thus, “no prediction” is counted as an incorrect prediction.

C. Entropy metric

We believe that there are some intrinsic characteristics of a trace that ultimately determine its predictability and hence the performance of different predictors. In the results section, we compare the accuracy metric with the entropy metric, for each user.

The general entropy does not help much for its predictability. However the conditional entropy is intuitively related.

D. Data collection

We have been monitoring usage on the Wi-Fi network at Dartmouth College, since installation began in April 2001. Today there are 543 access points providing 11 Mbps coverage to the entire campus. As a result we recorded the history of location changes for each of 6000 wireless cards, many for a period of months or years.

For more information about Dartmouth’s network and our data collection, see our previous study [2].




Fig. 1. Prediction accuracy for a sample user

III. RESULTS

We evaluated the location predictors using our Wi-Fi mobility data.

As we step through the locations in a single trace, attempting to predict each location in sequence, we can examine the accuracy up to that point in the trace. Figure 1 shows how the accuracy metric varies over the course of one user’s history, using the $O(1)$ Markov predictor. Ultimately, for comparison purposes, we define the accuracy over the entire trace, the rightmost value in this plot.

Of course, we have several thousand users and the predictor was more successful on some traces than on others. We use the accuracy of predictors in cumulative distribution function (CDF) curves to compare with each other.

In this study we examine only those user with long histories (1000+ moves). We compare the best Markov predictors with the best LZ predictors in Figure 2. It is difficult to distinguish the LZ family, which all seem to have performed equally well. The $O(2)$ Markov predictor, with fallback to $O(1)$ whenever it had no prediction, was the best overall. It is striking that the extra complexity, and the theoretical aesthetics, of the LZ predictors apparently gave them no advantage.

We include an “Optimal” curve in Figure 2, as a simple upper bound on the performance of history-based location predictors. In our definition, the “optimal” predictor can accurately predict the next location, except when the current location has never been seen before. Although it should be possible to define a tighter, more meaningful upper bound for domain-independent predictors like those we consider here, it seems clear that there is room for better location-prediction algorithms in the future. It may be that some user traces are simply less predictable than others. In Figure 3 we compare the entropy of each user, based on the probability table built by the $O(1)$ Markov predictor, with the accuracy of the $O(1)$ predictor. The correlation is striking, and indeed the correlation coefficient is -0.95 (a coefficient of 1.0 or -1.0 represents perfect correlation). This strong correlation indicates that some

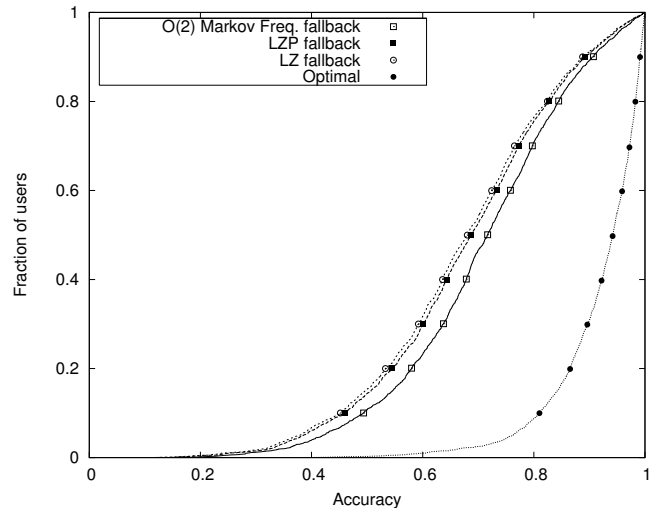


Fig. 2. The best predictors, compared




Fig. 3. Correlating accuracy with entropy

users with high entropy are doomed to poor predictability.

Acknowledgements

The Dartmouth authors thank NTT DoCoMo USA Labs, for their funding of this research, and Cisco Systems, for their funding of the data-collection effort.

REFERENCES

- [1] Christine Cheng, Ravi Jain, and Eric van den Berg, “Location prediction algorithms for mobile wireless systems,” in *Handbook of Wireless Internet*, M. Illyas and B. Furht, Eds. CRC Press, 2003.
- [2] David Kotz and Kobby Essien, “Analysis of a campus-wide wireless network,” in *Proceedings of the Eighth Annual International Conference on Mobile Computing and Networking*, September 2002, pp. 107–118, Revised and corrected as Dartmouth CS Technical Report TR2002-432.