Modeling Cellular User Mobility using a Leap Graph

Wei Dong† Nick Duffield‡ Zihui Ge‡ Seungjoon Lee‡ Jeffrey Pang‡
† The University of Texas at Austin ‡ AT&T Labs – Research

Abstract. User mobility prediction can enable a mobile service provider to optimize the use of its network resources, e.g., through coordinated selection of base stations and intelligent content prefetching. In this paper, we study how to perform mobility prediction by leveraging the base station level location information readily available to a service provider. However, identifying real movements from *handovers* between base stations is non-trivial, because they can occur without actual user movement (e.g., due to signal fluctuation). To address this challenge, we introduce the *leap graph*, where an edge (or a *leap*) corresponds to actual user mobility. We present the properties of leap based mobility and demonstrate how it yields a mobility trace more suitable for mobility prediction. We evaluate mobility prediction on the leap graph using a Markov model based approach. We show that prediction using model can substantially improve the performance of content prefetching and base station selection during handover.

1 Introduction

Mobile network providers have a strong desire to optimize network resources due to the scarcity of radio frequency spectrum and the rapidly increasing bandwidth demands of mobile users. The ability to predict short-term user mobility can be useful in optimizing these resources. For example, at the network layer, accurate prediction can inform the choice of basestation(s) used to communicate with a mobile device. At the application level, different delivery strategies in the network based on expected movement (e.g., prefetching) could improve both the user experience and network efficiency.

A recent body of work has examined user mobility prediction from data collected on user devices, e.g., by using GPS or Wi-Fi associations readings [10,15,16,19]. However, the spatial information most relevant from a provider's perspective would be the cellular basestations that mobile devices associate with. This data is readily and ubiquitously available to mobile operators without requiring additional instrumentation of devices, and does not pose the coverage, energy-consumption, and privacy concerns of GPS and Wi-Fi association based techniques. Thus, this paper takes a novel provider-centric approach: we study how to perform mobility prediction by using the base station-level location information readily available to a cellular service provider. This data reports the *active set* of basestations with which a given mobile device is currently associated, and, in particular, a data record is generated for each *soft handover* event that changes a user's active set.

Despite the advantages of handover traces, there are a number of challenges that make it non-trivial to use handover data directly as a mobility trace. First, while it seems natural to use the active set to define a user's location, fine granularity of the coverage intersections is unachievable or unreliable due to the dynamic nature of radio environment. Using the active set to define location can also suffer from the state-explosion problem, since a mobile device may see any combination of tens of sectors in densely covered regions. Secondly, not all handovers happen due to user mobility. Other causes include radio signal and workload fluctuations. In these cases, it is not obvious how to distinguish fluctuations from real user mobility.

To address these challenges, we introduce a *leap graph*, where an edge (or a *leap*) between two sectors denotes that moving from one sector to the other requires actual user mobility. A leap consists of two or more sector-level transitions. We use a data driven approach to identify sectors that overlap, and find leaps in the handover data between non-overlapping sectors. We then design procedures to effectively leverage the transitions that are not leaps and fully extract the leap information. The resulting leap graph differs significantly from the direct handover graph in terms of the number of state changes and degree. It effectively reduces fluctuations, yielding a mobility trace more suitable for mobility prediction. We study mobility prediction on the leap graph using a Markov-based approach. We also show the performance of our mobility prediction in two example applications: prefetching and handover optimization. Using a month of handover data from a cellular service provider, we show that our approach can improve content prefetching hit-rate to 84%, compared with 40% for a popularity-based approach. We also show that our approach can potentially reduce the number of handovers by 38% on average.

The rest of the paper is organized as follows: Section 2 introduces background on cellular handovers and discusses the challenges of using handover data for mobility prediction in detail; Section 3 describes our approach to extract the leap based mobility; in Section 4 we study properties of leap traces and evaluate leap based mobility prediction; we study prediction performance in real applications in Section 5; we review related works in Section 6 and conclude in Section 7.

2 Background and Challenges

2.1 Soft Handover and Active set

To maintain a data connection in a UMTS cellular network, each mobile device connects to several cell sectors when it is actively sending or receiving. A cell sector is defined by an antenna on the base station and the frequency that it transmits in. There are typically 1–3 sectors pointing in each of three directions on each macrocell base station. The set of sectors to which a mobile device is connected is called the *active set*. The size of active set typically varies from 1 to 4 sectors depending on the quality of the radio channel and the load on the base stations. In a UMTS network, any or all of these cell sectors may transmit to the device at once, depending on the radio technology used. Most modern devices use HSPA technology and only receive data from a single *serving sector* in the active set at a time although this serving sector can change very quickly.

The process of adding or removing sectors from the active set is called *soft handover* and is controlled by the radio network. A sector is added if its signal strength is greater than a threshold and the sector has not already admitted the maximum number of connections, while a sector is removed if its signal strength falls below another threshold [17]. Hence, the active set typically contains the sectors with the highest signal strength with respect to the mobile device. Since signal strength falls off with the square of the distance from the antenna, the active set cells are usually close to the mobile device in geographic space as well. We leverage this fact to use soft handover traces to predict a device's mobility.

```
17:46:59.296  S1 S2
17:46:59.976  S2
17:47:00.936  S1 S2
17:59:41.395  S3 S2
17:59:43.195  S2
17:59:43.875  S3 S2
17:59:46.995  S2
17:59:48.355  S3 S2
18:00:35.194  S4 S5
18:04:09.481  S6
```

Fig. 1. Example records of timestamp and active set for a stationary device

```
13:22:32.012  U0
13:22:47.795  U1
13:22:56.088  U2
13:23:57.005  U1
13:24:56.118  U3
13:24:59.625  U4
13:25:38.340  U5
13:25:38.775  U6
13:25:40.593  U3
13:36:38.473  U7
```

Fig. 2. Example raw trace. Each pair of adjacent sectors overlap. In addition, `U3`, `U4`, `U5`, and `U6` mutually overlap.

```
Trace 1  13:22:32.012  U0
         13:22:56.088  U2
         13:24:56.118  U3

Trace 2  13:22:47.795  U1
         13:24:59.625  U4
         13:36:38.473  U7

Trace 3  13:23:57.005  U1
         13:24:59.625  U4
         13:36:38.473  U7
```

Fig. 3. 3-hop leap traces for Figure 2

2.2 Challenges

As described in Section 1, modeling user mobility through handover traces offers many advantages. Yet there are several rather unique challenges with this approach. The first challenge is on how to define users' location. Different cell sectors have overlapping coverage areas, and a mobile user is located within the intersect of the coverage areas of all the sectors in the active set. It seems natural to use the active set to define the user's location as it may offer high precision. However, it turns out that the radio environment dynamics and sectors' workload variability can lead to significant fluctuation in the active set, making the fine granularity of the coverage intersects unachievable or unreliable. In addition, the combinatorial nature of the active set can potentially create a state-explosion problem in densely covered regions where tens of sectors are visible to a mobile device. Another approach is that we cluster/partition the geo-space into regions and take the union of sectors in the region to define the location. However, we lose the precision with this approach. In this work, we choose to use the serving cell in the active set as the representative for location as we find it achieving a good balance between precision and accuracy.

Another major challenge of examining the handover trace is to identify real user mobility from handovers due to radio signal fluctuations. To understand this aspect, we performed a controlled experiment where a stationary phone owned by a cellular provider is set up to transmit data packets periodically. We obtained the handover logs shown in Figure 1. We observe that handovers occur even when a user is stationary. We also see a diverse set of sectors in the active sets. While we present more details later in Section 4.2, it is clear that these handovers are inherently different from the ones induced by user movement and hence present noise for mobility modeling. Signal strength triangulation does not help in identifying which handover is due to real user movements as signal strength at a single location can vary a lot [14].

Before describing our solution, we examine the limitations of two heuristics:

Loop detection and elimination: Stationary users are much more likely to alternate among a small set of sectors than mobile users. Hence, one simple approach to eliminate non-mobility handovers would be to remove trace segments between repeated occurrences of the same serving sector. However, not all stationary traces manifest a loop, and thus this approach does not eliminate superfluous handovers. Moreover it is possible that a users comes back to the same location after making real movements in a short period of time, in which case the real movements will be discarded.

Low-pass filter: Low-pass smoothing is a principled approach to suppress membership fluctuations in the active set among near-by sectors. For example, we can pass the the sectors in each consecutive active set through a queue. If a sector is already in queue, we move it to the tail of queue. When the queue is full, we evict the oldest member and produce a "smoothed" trace using the eviction sequence. However, the number of sectors visible to a user varies, making it difficult to determine a single fixed queue size. When we apply this approach to real traces with a fixed queue size, we find that it admits superfluous handovers as mobility-induced and misses true user movement.

3 Mobility and Leaps

Individual changes in the active set do not in themselves indicate whether a handover was due to user mobility. Thus in trying to infer mobility from the handovers, we try to eliminate handovers involving changes whose interpretation is ambiguous, and focus instead on minimal groups of successive handovers which together likely indicate mobility. The boundaries of these groups will be termed a *leap*, and a set of successive adjacent leaps together constitute a *leap trace*. These are constructed in a two step procedure.

Step One: Identifying Overlapping Sectors
Informally, two sectors overlap if a handover can take place between them. Although overlap could in principle be inferred from auxiliary sector configuration data (such as sector antenna locations, directions and powers), this would be a complex task in general. Instead we designate two sectors s_i and s_j as overlapping (written $s_i \sim s_j$) if at least one of the following two criteria holds: (i) Configurational: s_i and s_j are based at the same cell tower; (ii) Empirical: s_i and s_j appear within an active set reported in the handover trace during a specified time period (3 weeks in our evaluation). We considered alternate ways of determining overlap (e.g., considering serving sector transitions), and they yielded marginal performance differences and thus are not discussed further.

Step Two: Creating Leap Traces
We partition the handover trace by user, then further extract the sequence of serving sectors reported for each such user. We call each such sequence a *raw trace*. A *segment* is a maximal ordered subset of a raw trace in which the time between handovers does not exceed a specified timeout value. A *leap* is a pair $s_i s_j$, $i < j$ of sectors within a segment such that $s_i \not\sim s_j$ but $s_i \sim s_k \, \forall i < k < j$. A *leap trace* on a segment $\{s_1, s_2, \ldots, s_m\}$ is a maximal set of some number ℓ of adjacent leaps $s_{i(1)} s_{i(2)}, s_{i(2)} s_{i(3)}, \ldots, s_{i(\ell)} s_{i(\ell+1)}$. A first leap trace constructed by finding a leap with initial node $s_{i(1)} = s_1$ and then using that leap's final node as the initial node for the next leap, and so on until reaching the end of raw segment. As many as $m - 2$ further leap traces may be constructed from the segment by the same procedure, taking each s_k, $k = 2, \ldots, m-1$, as the initial node $s_{i(1)}$ of the initial leap. But in order to avoid double counting of leaps, we stipulate that if a leap trace starts to repeat leap segments already identified in trace from a previous starting sector, we ignore the remaining trace after including at most some number n of further leaps. The n is determined by the mobility modeling requirement. For example, $n = 1$ for first order Markov model and $n = 2$ for second order Markov model.

To illustrate, we show an example raw trace in Figure 2, where all adjacent sectors overlap (e.g., `U0`$\sim$`U1`, `U1`$\sim$`U2`, etc.). In addition, `U3`, `U4`, `U5`, and `U6` overlap with each other. Starting from `U0`, we can get Trace 1 in Figure 3. There can be multiple leap traces from the same raw trace. For example, starting from `U1`, we can get Trace 2 in Figure 3. In fact, the number of different leap traces can be exponentially large to the length of raw trace. Moreover, leap traces from different starting sectors may become identical after a few leaps since they are derived from the same raw trace. In this case we only keep the useful information and discard the repeated part as described above.

The leap trace ignores handovers due to signal fluctuations or user movements in small areas, while focusing on longer trips. This is sufficient for our target applications because we focus on improving handovers or prefetching in larger areas. We study application specific performance in section 4.

4 Properties of Leap Traces and Leap-based Mobility Prediction

4.1 Data Set

We use anonymized event logs collected from several RNCs (Radio Network Controllers) in a major U.S. cellular operator in December 2011. These RNCs control a significant fraction of the base stations in a large U.S. city. The logs record soft handover events, i.e., additions and removals from each device's active set. Each log entry has a timestamp, and devices are anonymously identified by an irreversible hash of the device's IMSI, which is unique per SIM card. All device and subscriber identifiers are anonymized to protect privacy without affecting the usefulness of our analysis. Furthermore, the data set does not permit reversing the anonymization or re-identification of subscribers. We use the data of the whole month (the first 3 weeks are considered known and used for training, the last week is considered unknown for testing purpose) and include all users from the trace. The logs recorded 67 million soft handover events for 413K users distributed over 5K sectors. The logs are generated only for active devices transmitting data, but not for idle devices. In our evaluation, if two subsequent soft handover records for a given device is apart by more than 30 minutes, we assume that the device has been idle, and start a new mobility segment using the latter record. Our data set contains proprietary information and cannot be made public.

4.2 Characteristics of Leap Traces

We extract the leap traces from the above data set. In this subsection, we present a high-level characterization of the aggregate leap traces.

We first compare the length of raw trace segments and leap segments. Figure 4(a) plots their CDF. We can see that leap segments are much shorter than raw segments. Specifically, over 80% of the raw segments generate no leap at all, indicating limited or no user mobility. In contrast, around 20% of raw segments contain only 1 active set report. This result illustrates that many soft handovers in raw traces either do not involve serving sector transition or the transition happens between close-by sectors and are likely not due to user mobility. This highlights the importance of our approach in separating the different causes for handovers. In Figure 4(b), we compare the inter-leap time and inter-handover time. We observe that the inter-leap time is much longer than inter-handover time. Specifically, the median of inter-leap time is 636 seconds, while for inter-handover time the number is 2.8 seconds.

We define the **leap graph** as the graph of sectors in which the edges represent the presence of a leap transition in any of the leap traces. Similarly, we also consider graphs

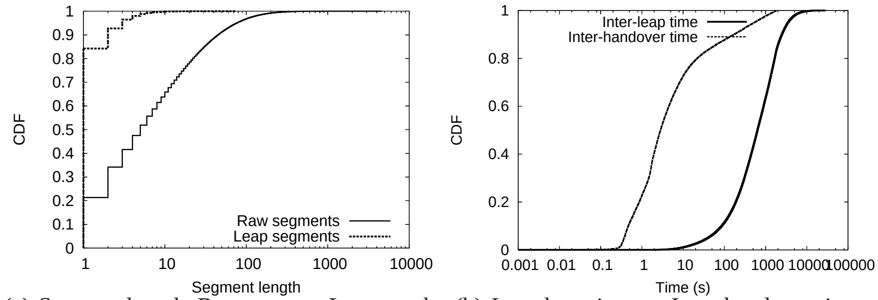

(a) Segment length: Raw trace vs. Leap graph (b) Inter-leap time vs. Inter-handover time

Fig. 4. Segment Characteristics

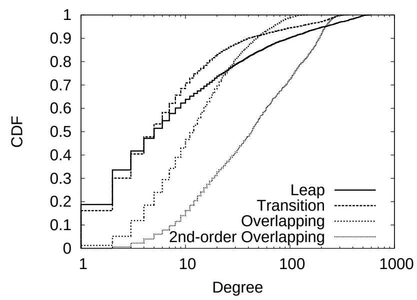

Fig. 5. Degree distribution comparison

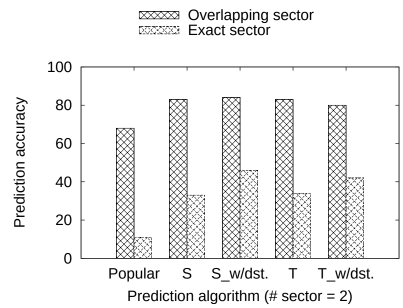

Fig. 6. Prediction accuracy. S and T denotes second- and third- order Markov models, respectively.

obtained from serving sector changes, from overlapping sets, and from second-order overlapping sets[1]. In Figure 5, we compare their degree distributions. We first observe that the size of overlapping set can be quite large (e.g., more than 20 for around 30% of cases), while the second-order overlapping set is even larger. Compared to the second-order overlapping set, the degree of leap graph is significantly low (e.g., 10 or less for more than 60% of cases). For many sectors, it is even smaller than the number of serving sectors they can transition to (marked as "Transition"), which suggests that people follow similar patterns (e.g., along a highway) and only move out of a region with very limited choices of ways. While fewer choices can make the mobility prediction using the leap graph easier, we also observe in the tail part that the leap degrees may approach the size of second-order overlapping, indicating that areas with dense cellular coverage tend to have more mesh-like transportation paths (e.g., downtown areas) – posing a challenge for mobility prediction. A very small fraction of sectors have larger leap degree than the degree of second-order overlapping, that is because a leap may happen between two sectors more than two hops away, *e.g.*, when a user stops using his phone for a short period of time while he is still moving.

4.3 Mobility Prediction on Leap Graph

In this subsection we study mobility prediction on the leap based graph. We adopt Markov-based approaches for prediction, as it has been proven effective in the liter-

[1] u is in the second-order overlapping set of s if s and u are not overlapping, and there exists t that overlaps with s and u.

ature [12, 19]. We consider two variants: second-order Markov model using one leap as a state and third-order Markov model using two consecutive leaps as a state. Using a higher-order model allows us to make predictions based on not only the user's current location but also the recent path trajectory. We further consider a variant where we assume the knowledge of the destination of a segment, to understand how such additional information can help with mobility prediction. This is motivated by the observation that many people have highly predictable daily routine and that the destination may be projected simply based on the time of day [8, 9]. In the data trace, we estimate the probability of state s being the next state as $P(s|d, o)$ where d is the destination (the last sector in trace segment) and o is the current state.

We train our prediction models with the first 3 weeks' data and evaluate them with the leap segments extracted from the last week's data. Given the current state o, we predict m next sectors using the m highest-probability leaps, while we vary m from 1 to 3. We evaluate how often we can correctly predict the next leap. We adopt two accuracy measures: (1) the predicted sector exactly matches the actual sector in the testing data, and (2) the predicted sector is in the overlapping set of the actual sector. We also count how often we cannot make a prediction and report the result. To form a base for comparison, we also employ a naive scheme (denoted by "Popular"), where among all possible next leaps, we pick m leaps with the most transitions.

In Figure 6, we compare the prediction accuracy of the popularity based approach, second- and third-order Markov models (with and without destination information). We select m=2 and report both accuracy measures. In the figure, the accuracy of predicting any one among the overlapping set is significantly higher than predicting the exact match, which is well expected. We make three observations. First, the Markov models significantly outperform the simple popularity-based approach. Specifically, when using the overlapping set of the actual sector, the accuracy of the popularity-based approach is 68%, while the accuracy of Markov models is 80% or higher. Second, the knowledge of destination information can further improve the prediction accuracy of the exact sector (e.g., from 33% to 46% in second-order Markov model) but helps little when we use the overlapping set of the predicted sector to measure accuracy. Finally we observe that the accuracy gain from using longer history in the third-order Markov models is marginal. On the other hand, the probability of being able to predict a sector is lowest with the third-order Markov model with destination information (62%). This is because training data is often unavailable for consecutive leap transitions with a particular destination. In contrast, the popularity based approach, due to its simplicity, can make a prediction for 99.5% of the cases, and the second-order Markov model without destination information can make a prediction for 98.4% of the cases. In practice, we can start with as much information as possible and fall back to less demanding settings if needed [16].

5 Applications

There are many potential applications of future sector prediction. In this section we focus on two example applications, namely prefetching and handover optimization, and quantify the application specific performance of our prediction schemes.

5.1 Prefetching

In this application scenario, we prefetch user requested content to a predicted future cell tower, such that the user can retrieve the content upon entering the range of the cell tower. The content a user is going to request is often predictable [5, 18]. Users can also request for prefetching since it saves time for them.

Once a prefetching request is made, a prediction is made based on the user's current mobility history. One complication here is that at a given sector in a raw trace, there could be multiple different leap traces leading to it as well as multiple different leap traces "leaping over" it, e.g., depending on which sector we start with in the raw trace. We lose information by considering only leap trace and ignoring the rest. In this paper we combine the predictions made with these different leap traces using the following simple heuristic (assuming second order Markov model).

We extract all possible leap traces and from them we get the different ending leaps. Let $L_{ij} = (s_i, s_j)$ be one of the ending leaps, where s_i and s_j are sectors. Let $P(S|L_{ij})$ be the probability vector predicted using L_{ij}, where S is a vector of potential future sectors. The final prediction is then computed as $P(S) =<< P(S|L_{ij}) >_i>_j$ ($<>_i$ means taking average over all i).

For prefetching, the criteria for a good prediction is that the cell tower of the predicted sector become within reach later. In our evaluation we consider a prediction correct if the sectors on the predicted cell tower appears in the active set in the future.

Due to space limit, we only present herein the result using second order Markov model without assuming destination knowledge. We make a prediction in the middle of the segment, and we choose the segments that have more than 3 leaps after the prediction is made to ensure that the user remains active after our prediction point. Then we prefetch the content to the cell tower of the top m predicted sectors. We vary m from 1 to 3. We run 10k tests and record how many times the prefetched content become available to the user after the prediction. We find that we can make a prediction for 99.6% of the times, which is slightly higher than the leap-based case (98.4%) because we effectively combine the predictions of different ending leaps. Out of the predictions we make, the accuracy is 84.7% when $m = 1$, suggesting 84.7% of the times the prefetched content becomes available to the user. Increasing m to two and three increase the number to 91.3% and 94.4%, respectively. We also find that 97% of the time the content becomes available within an hour. In comparison, prefetching to the cell tower at the most popular leap achieves lower than 40% accuracy.

5.2 Handover optimization

Next we use future sector prediction to optimize handovers. The idea is based on the predicted leap, we can suggest which sector to hand over to, such that we reduce the total number of handovers. Detailed simulation of handovers is not trivial, as it requires a detailed modeling of signal strength variations, traffic load, load changes, etc. In this paper we only consider an idealized scenario to demonstrate the potential gain we can achieve. Specifically, after each leap we make a prediction of the next leap. Based on the prediction we first rank the sectors in the active set by giving preference to sectors that overlap with the predicted sector. Then we break ties using physical distance to the predicted sector (the closer the better). We use the highest ranked sector as the suggested handover target. To evaluate how many handovers we can potentially save, we count how many consecutive future handovers in the real trace have the suggested sector in the active set before the next leap actually happens, as these handovers can potentially be replaced by one handover to the suggested sector. In our evaluation we only consider the traces that have at least 2 leaps. We predict the second leap based on the first one using second-order Markov model without assuming destination information. We run the test for 10k times and we find that less than 0.1% of the times the suggested sector does not appear in future handovers, which means the prediction is wrong and may cause extra handovers; 59% of the times we can save at least one handover by using the suggested sector, and 32% of the times we can save three or more. On average, our handover optimization reduces the handover count by 38%.

6 Related Work

Despite the plethora of work in mobility modeling, we believe that this is the first work to address the unique challenges of mobility prediction using cellular handover traces and to present an approach that works well on real data. We survey a key selection of prior work here.

An important body of work focused on predicting locations as defined by Wi-Fi associations [10, 15, 16, 19]. Much of this work focused on evaluating the effectiveness of well known location predictors such as Markov-models [19], compression-based predictors [16], and CDF predictors [15]. Others have focused more on modeling [10]. In contrast to cellular handover traces, which include many handovers that occur when a user is stationary, changes in Wi-Fi associations represent real movement in most cases. Thus, our work is unique in addressing the inherent challenge of ambiguity in predicting future cell sectors in cellular traces. Nonetheless, we build upon the same principled predictors, such as Markov-models.

There have been previous proposals on location prediction in cellular networks [1,3, 4, 11, 13]. However, they either make unrealistic assumptions (e.g., that basestations are mobile [13] or a prefect sector structure [4]), require information that is not typically available to a cellular operator (e.g., reports of device location and velocity [1, 11]), or are designed for different purposes (e.g., to limit cell updates [3] and paging [6]). As a result, all of these approaches have only been evaluated on synthetic data. Our approach is the first to be evaluated on real cellular handover traces.

Finally, there have been studies on the predictability of wireless attributes, such as cellular connectivity [7], Wi-Fi connectivity [12], and commute routes [2]. The predictors for them use similar data to our work, but are orthogonal in design and purpose.

7 Conclusion

In this paper we introduced a novel leap based approach to extract user mobilities from soft handover data, which is readily available but also contains significant fluctuations even for a stationary device due to signal strength change or load balancing. Our study showed that our approach can effectively reduce fluctuations in the raw handover data while maintaining real user mobility pattern. In our experiments, we demonstrated significant gain in prediction accuracy using our leap based approach, and performance improvement in two example applications that we considered. While our approach is provider-centric, a service provider can potentially make location prediction for a user available to selected applications on the user's device, so that the applications can provide a better service to the user. In the future, we plan to apply our approach to real world location based services to find further application specific optimizations and see its benefit in real systems.

References

1. I. F. Akyildiz and W. Wang. The predictive user mobility profile framework for wireless multimedia networks. *IEEE/ACM Trans. Netw.*, 12(6):1021–1035, Dec. 2004.
2. R. A. Becker, R. Caceres, K. Hanson, J. M. Loh, S. Urbanek, A. Varshavsky, and C. Volinsky. Route classification using cellular handoff patterns. UbiComp '11, pages 123–132, New York, NY, USA, 2011. ACM.
3. A. Bhattacharya and S. K. Das. Lezi-update: an information-theoretic approach to track mobile users in pcs networks. MobiCom '99, pages 1–12, New York, NY, USA, 1999. ACM.
4. R. Chellappa, A. Jennings, and N. Shenoy. The sectorized mobility prediction algorithm for wireless networks. In *in Proc. ICT*, 2003.
5. X. Chen and X. Zhang. A popularity-based prediction model for web prefetching. *Computer*, 36(3):63 – 70, mar 2003.
6. S. Das, , S. K. Das, and S. K. Sen. Adaptive location prediction strategies based on a hierarchical network model in cellular mobile environment. *The Computer Journal*, 42:473–486, 1996.
7. P. Deshpande, A. Kashyap, C. Sung, and S. R. Das. Predictive methods for improved vehicular wifi access. MobiSys '09, pages 263–276, New York, NY, USA, 2009. ACM.
8. M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, June 2008.
9. S. Isaacman, R. A. Becker, R. Cceres, S. G. Kobourov, J. Rowland, and A. Varshavsky. A tale of two cities. In *HOTMOBILE'10*, pages 19–24, 2010.
10. M. Kim and D. Kotz. Extracting a mobility model from real user traces. In *In Proceedings of IEEE INFOCOM*, 2006.
11. B. Liang and Z. J. Haas. Predictive distance-based mobility management for multidimensional pcs networks. *IEEE/ACM Trans. Netw.*, 11(5):718–732, Oct. 2003.
12. A. J. Nicholson and B. D. Noble. Breadcrumbs: forecasting mobile connectivity. MobiCom '08, pages 46–57, New York, NY, USA, 2008. ACM.
13. P. N. Pathirana, A. V. Savkin, and S. Jha. Mobility modelling and trajectory prediction for cellular networks with mobile base stations. MobiHoc '03, pages 213–221, New York, NY, USA, 2003. ACM.
14. A. Schulman, V. Navda, R. Ramjee, N. Spring, P. Deshpande, C. Grunewald, K. Jain, and V. N. Padmanabhan. Bartendr: a practical approach to energy-aware cellular data scheduling. MobiCom '10, pages 85–96, New York, NY, USA, 2010. ACM.
15. L. Song, U. Deshpande, U. C. Kozat, D. Kotz, and R. Jain. Predictability of wlan mobility and its effects on bandwidth provisioning. In *INFOCOM*. IEEE, 2006.
16. L. Song, D. Kotz, R. Jain, and X. He. Evaluating location predictors with extensive wi-fi mobility data. In *In Proceedings of INFOCOM*, pages 1414–1424, 2004.
17. S.-F. Su. *The UMTS Air-Interface in RF Engineering*. McGraw Hill, 2007.
18. Z. Su, Q. Yang, and H.-J. Zhang. A prediction system for multimedia pre-fetching in internet. MULTIMEDIA '00, pages 3–11, New York, NY, USA, 2000. ACM.
19. J. Yoon, B. D. Noble, and M. Liu. Building realistic mobility models from coarse-grained traces. In *in Proc. MobiSys*, pages 936–5983. ACM Press, 2006.