Predicting Wi-Fi Availability using Bluetooth and Cellular Signals

Ganesh Ananthanarayanan and Ion Stoica
University of California, Berkeley
{ganesha, istoica} @ cs.berkeley.edu

ABSTRACT

With the various network interfaces of a mobile device having complementary characteristics in terms of power consumption, throughput and range, efforts are on to combine their positives. We aim to reduce the idle power consumption of Wi-Fi by using a combination of bluetooth and cellular data for localization, learning and prediction. We show that the combination of bluetooth and cellular data offer a good mix of coarse-grained and fine-grained accuracy as well as coverage.

Categories and Subject Descriptors

D.4.8 [**Operating Systems**]: Performance

General Terms

Measurement, Performance

Keywords

Bluetooth contact-patterns, Wi-Fi prediction

1. INTRODUCTION

Mobile devices are increasingly equipped with multiple network interfaces which have widely different, often complementary, characteristics (see Table 1). These devices would like to intelligently leverage the interfaces to maximize throughput and minimize the power consumption. In particular, a device would prefer to use the Wi-Fi network (whenever available) to exploit its high throughput and transfer power efficiency, but not waste the energy to probe for its availability. In this paper, we aim to address the following problem − *Can we predict the availability of a Wi-Fi network without switching it on, using just the bluetooth and cellular data?*

	Throughput	Transfer (J/MB)	Idle (W)	Scan (W)	Range (m)
Cellular	few 100 kbps	100	**0**	**0**	**500**
Wi-Fi	**11-54 Mbps**	**5**	0.77	1.29	100
Bluetooth	700 kbps	0.1	**0.01**	**0.22**	**10**

Table 1: Various interface characteristics.

SIGCOMM'08, August 17–22, 2008, Seattle, Washington, USA.
ACM 978-1-60558-175-0/08/08.

We present a simple scheme that uses a combination of bluetooth contact-patterns [1] and cellular information to localize a user's context that will predict the availability of Wi-Fi. The main observation behind our scheme is that a user is likely to see the almost same set of bluetooth devices and cell-towers while at home, work, shops and even commute. To predict Wi-Fi availability, we use bluetooth signals − which by virtue of their low range would provide accurate and fine-grained context − in conjunction to cell-tower signals − which would provide wider coverage but a coarse-grained context. To the best of our knowledge, this is the first application of bluetooth contact-patterns to context localization. While Context-for-Wireless [2] uses cell-tower data for predicting Wi-Fi availability, we show how taking advantage of the complementary characteristics of the bluetooth data can significantly improve the prediction accuracy. Combining bluetooth and cellular data results in an average coverage of 94% and an average accuracy of approximately 84%, a 47% improvement in accuracy over the pure cell-tower based prediction and a 57% improvement in coverage over the pure bluetooth based prediction scheme. On our sample workload, the combined prediction scheme is 47.5% more energy-efficient as compared to data transfers using only the cellular interface. The corresponding savings are 26% and 32.2% for cell-tower based prediction and bluetooth-based prediction schemes respectively. Such Wi-Fi predictions would be particularly useful for background tasks on the phone like synchronizing a user's mail box.

2. WI-FI PREDICTION SCHEMES

In this section, we present the Wi-Fi prediction schemes. All the schemes have a learning phase when devices periodically log all the network signals in a log L. The log entries are of the form (*Timestamp*, {*Bluetooth Devices*}, {*Cell Towers*}, {*Wi-Fi beacons*}). The goal of the learning phase is to identify the set of bluetooth devices, $predict_{BT}$ and cell-towers, $predict_{cell}$, whose presence indicates the availability of Wi-Fi connectivity.

Bluetooth based Prediction: Let W_P be the set of Wi-Fi networks to which a device can connect, and set_{BT} the set of bluetooth devices in the log L. In the learning phase, for each Bluetooth device, $BT_i \, \epsilon \, set_{BT}$, we calculate *reliability* as $n(L, BT_i, W_P)/n(L, BT_i)$, where $n(L, BT_i, W_P)$ is the number of entries in L when BT_i was present and at least one of the networks in W_P was present, and $n(L, BT_i)$ is the number of entries in L when BT_i was present. If *reliability* is greater than a threshold τ, we add BT_i to $predict_{BT}$. After we populate $predict_{BT}$ in the learning phase, we predict Wi-

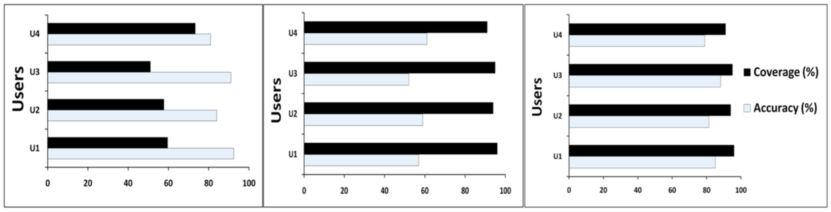

Figure 1: Accuracy and Coverage for Bluetooth, Cell-Tower and Hybrid Prediction

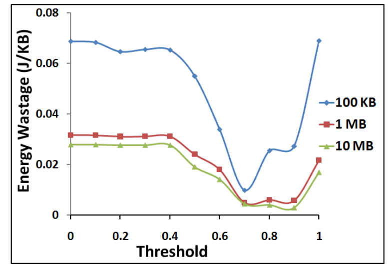

Figure 2: Expected Energy Wastage per KB for different file transfers

Fi availability if any of the bluetooth devices in $predict_{BT}$ is currently nearby.

Cell-Tower based Prediction: This scheme uses only the cell-tower data to predict Wi-Fi availability [2]. The algorithm to obtain $predict_{cell}$ is exactly the same as above with Bluetooth devices replaced by cell-towers.

Hybrid Prediction: The Hybrid scheme predicts Wi-Fi availability using both bluetooth contact-patterns and cell-tower data. It first discovers the bluetooth devices that are currently nearby and checks if any of them are present in $predict_{BT}$. If none of the discovered devices are present in $predict_{BT}$ and set_{BT}, it resorts to cell-tower based prediction. This combines the advantage of both schemes.

3. THRESHOLD SELECTION

To choose the threshold, τ, we use a simple model for computing the expected energy wastage as a function of τ. Besides τ, this model takes as input the energy consumed in probing for Wi-Fi networks, and energy consumed per unit for transferring data using the cellular and Wi-Fi interfaces. Figure 2 plots the expected energy wastage per KB versus τ for transfers of 100 KB, 1 MB and 10 MB respectively. We pick $\tau = 0.7$ as it works well across wide range of file transfers. In our experiments we use identical thresholds for populating $predict_{BT}$ and $predict_{cell}$.

4. EVALUATION

Four volunteers were given PDAs that logged signals every two minutes over a period of 2-3 weeks. Learning was done on one half of the data and the other half was used for evaluation. We used two metrics for evaluation:

- *Accuracy:* What fraction of the predictions of Wi-Fi availability are correct?
- *Coverage:* What fraction of Wi-Fi availability is predicted?

We measure the *Accuracy* and *Coverage* of using bluetooth and cellular data individually, highlight their complementary properties and demonstrate the benefits of using them in conjunction. Due to the limited range of bluetooth, the accuracy of the prediction is high (87%), but the coverage is low (60%). Cell-tower based prediction results in high coverage of 94% but an accuracy of 57% (Figure 1). The Hybrid Scheme combines the advantage of both and shows a coverage of 94% and an accuracy of 84% (see Figure 1), a 47% improvement in accuracy over the pure cell-tower based prediction and a 57% improvement in coverage over pure bluetooth based prediction scheme. We believe this to be an encouraging validation for the usage of bluetooth contact-patterns for localizing contexts.

We also tested the energy savings of our prediction using a simulated sample workload that downloaded 100 KB once in every 15 mins for 10 hours (see Figure 3). The hybrid scheme is 47.5% more energy efficient than $E_{cellular}$ where the device uses only the cellular interface for its data transfer, and 37.5% more energy efficient than E_{Wi-Fi} where the device scans and uses Wi-Fi if available and the cellular interface otherwise. The corresponding savings are 26% and 32.2% for cell-tower based and bluetooth-based prediction schemes respectively.

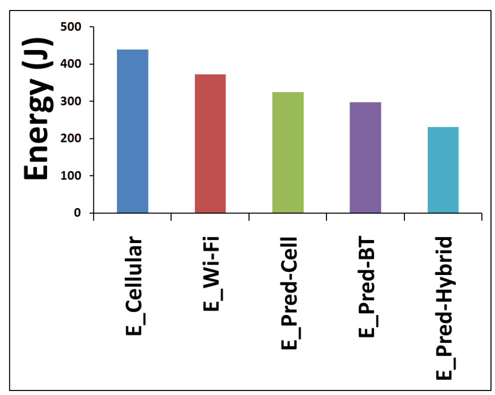

Figure 3: Energy Consumption for Data Transfer

5. FUTURE WORK

We argued for the usage of bluetooth contact-patterns for localizing contexts for Wi-Fi prediction. While our preliminary results are promising, much more remains to be done. We intend to test the performance of our schemes on a larger scale. Clustering algorithms could weed out the user's personal accessories (e.g., Bluetooth headphones) and identify static devices (e.g., Bluetooth mouse at work) for sharing. Sharing bluetooth and cell-tower data in a peer-to-peer and global fashion is part of future exploration.

6. REFERENCES

[1] A. Natarajan et al. Understanding Urban Interactions from Bluetooth Phone Contact Traces. In *PAM*, Apr. 2007.

[2] A. Rahmati and L. Zhong. Context-for-Wireless: Context-Sensitive Energy-Efficient Wireless Data Transfer. In *MobiSys*, June 2007.