Blue-Fi: Enhancing Wi-Fi Performance

Mobile and Adaptive Systems
Alex Borés and Dror Dayan

Department of Computer Science
University College London

Outline

1. Introduction
2. Predicting Wi-Fi Availability
3. Bluetooth discovery
4. Special Bluetooth Devices
5. Collaborative prediction
6. Evaluation
7. Further Improvements
8. Related work
9. Summary

Introduction(1)

- **Mobile devices** increasingly equipped with Wi-Fi, Bluetooth and cellular interfaces.
- These interfaces have widely **different characteristics** in terms of throughput, range and power.
- **Wi-Fi** is the **least power efficient** in idle state and incurs a high overhead when scanning for new networks.
- We should **avoid scanning** whenever possible.
- Efficiently detect Wi-Fi availability **without switching it on**.

Introduction(2)

- Blue-Fi predicts Wi-Fi availability by using **Bluetooth contact-patterns** and **cell-tower information**.
- Most users tend to **find the same set** of Bluetooth devices and cell-towers.
- Examples: Bluetooth mouse, printer at work, shop-owner's mobile device...
- Bluetooth **drawbacks**: lower range, time-consuming discovery process, mobile devices.
- Leverage the **complementary properties** of Bluetooth and cell-tower to improve the prediction.

Predicting Wi-Fi Availability(1)

- Learning process: each mobile device **periodically logs** all the network signals.
- The question is: **how accurately** can a Bluetooth device or a cell-tower **predict** the Wi-Fi availability?
- Bluetooth devices **may be mobile** (phones, notebooks...) so they cannot be used to predict fixed Wi-Fi access points.
- **Predictability**: only consider a bluetooth / cell-tower to be **reliable** if **most of the log entries** in which it appears **contain at least one AP**.

Predicting Wi-Fi Availability(2)

- Two metrics to evaluate predictions: **coverage** and **accuracy**.
- Bluetooth based Prediction (high accuracy, low coverage).
- Cell-tower based Prediction (high coverage, low accuracy).
- **Hybrid** based prediction.
  1. Discover **Bluetooth** devices.
  2. If reliable devices found, Wi-Fi Connectivity **available**.
  3. If not, discover **cell-towers**.
  4. If reliable cell-towers found, Wi-Fi Connectivity **available**.
  5. If not, Wi-Fi Connectivity **unavailable**.

Predicting Wi-Fi availability(3)

- How to calculate the **prediction reliability threshold**.
- Prediction matrix where **s** indicates the **presence and absence** of Wi-Fi availability and **p** indicates cases when Blue-Fi **predicts the availability** of Wi-Fi.
- Calculate energy waste **as a function of the threshold** and find the threshold which **minimizes energy waste**.

	p	$\bar{p}$
s	**1. Probe for Wi-Fi network when there is Wi-Fi availability (p_1)** ✓	**2. Use the cellular interface in the presence of Wi-Fi (p_2)** ✗
$\bar{s}$	**3. Waste energy to probe for Wi-Fi networks (p_3)** ✗	**4. Use the cellular interface because there is no Wi-Fi availability (p_4)** ✓

Bluetooth discovery

- Bluetooth discovery takes over **10 seconds**. We may **disable it** when not needed (user stationary).
- Cell tower signal strength values have **low variance** when the device is stationary.
- A **cell-tower fingerprint** is a set of tuples containing the cell-tower identifier and signal strength.
- Calculate **Euclidean distance** between two cell-tower fingerprints.
- The threshold of Euclidean distance for inferring stationary periods is calculated using **proximity of a stationary Bluetooth device** or being **connected to the same AP**.

Special Bluetooth Devices

- **Two classes** of Bluetooth devices: landmark and mobile.
- **Landmark** or stationary (mouse, keyboard, printer...).
- **Mobile** (Bluetooth headsets, iPod...). **Not reliable** indicators and should be removed from the logs.
- How do we **identify** them?
  - **Correlating** APs and cell-towers, which are stationary, with the Bluetooth devices.
  - Capture the **variation** in the different locations at which the Bluetooth device was sighted.
  - For every Bluetooth device in the log, extract the **list of Wi-Fi and cell-tower signatures** that were **co-sighted** along with it.

Collaborative Prediction(1)

- P2P sharing: Get information from devices which are currently connected to some Aps.
  - No need to use the cellular connection.
  - Accurate information: throughput, latency.
  - Easy to deploy.
- Global sharing: Get information from previously connected clients.
  - AP clients upload Wi-Fi usage to a central server.
  - Upload entries: Timestamp, {Bluetooth devices}, {Cell Towers}, {Wi-Fi network, {Characteristics}}.
  - Timestamps
    - Dynamic Bluetooth devices need to be within a time interval.
    - Cell towers and landmarks are stationary.
  - Expensive to maintain.

Collaborative Prediction(2)

- Increase in coverage.
- Reaching 100%.
- What about accuracy?

User	$\tau = 0.8$	p2p	Global
U1	59.6%	68.3%	79.4%
U2	57.8%	72.2%	84.5%
U3	51.1%	78.4%	86.3%
U4	73.4%	78.8%	89.1%
U5	74.2%	81.8%	90.1%
U6	50.2%	66.5%	88.1%
U7	58.8%	71.1%	90.1%
U8	61.2%	69.4%	83.3%
U9	44.6%	65.1%	80.8%
U10	60.2%	70.8%	76.4%
U11	60.1%	72.2%	77.3%
U12	64.5%	70.9%	73.2%

- **Is P2P Safe?**
  - Intrusion: Can we trust Bluetooth devices?
  - Usage Pattern: Learning user habits.
  - Industrial Espionage: can be used in order to learn about business profile.

Evaluation

- 12 volunteers Berkley students and professionals.
- Duration: 2 weeks.
- 32.1% to 68% of the time volunteers were connected to preferred networks.
- 49.6 – 77.2% of the time Bluetooth devices where around.

Power Consumption

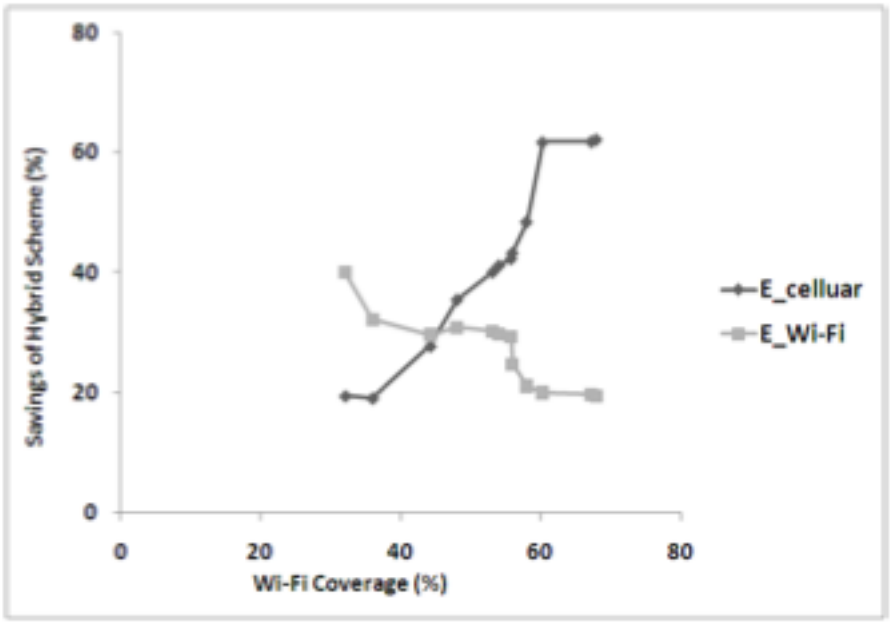

- **Cellular 3G, GPRS (E_Cellular):** Blue-Fi saves more energy when Wi-Fi coverage gets Higher.
- **Wi-Fi when available (E_Wi-Fi):** Blue-Fi saves more energy when Wi-Fi coverage gets lower.

- **Large file sizes -> less Wi-Fi scans:** lower Blue-Fi contribution.
- **Battery life increase?**

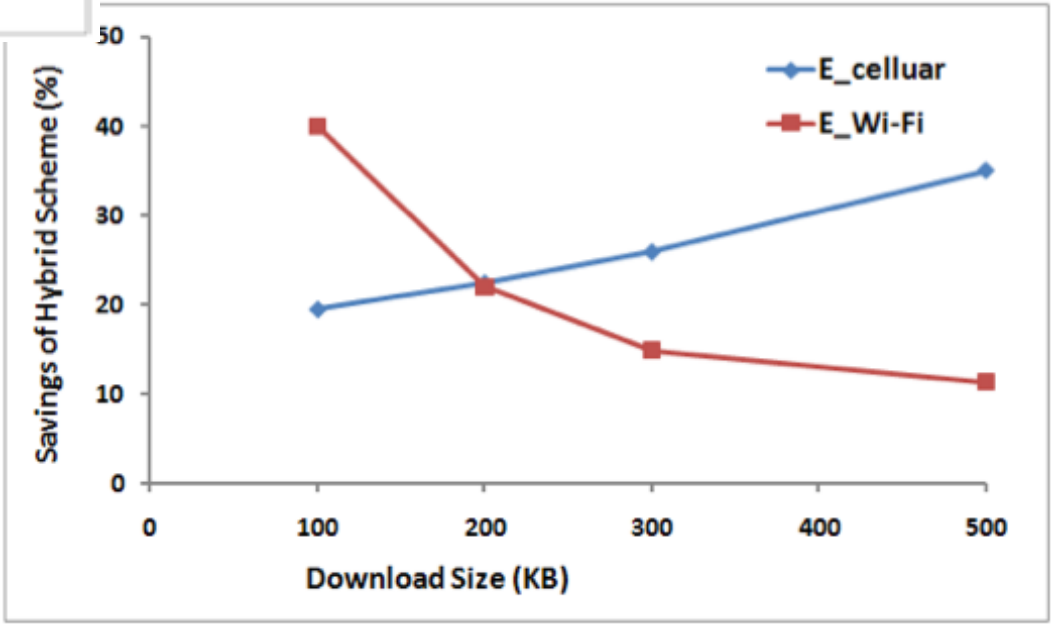

Accuracy vs Coverage

- High threshold -> better Accuracy.
- Low threshold -> better coverage.

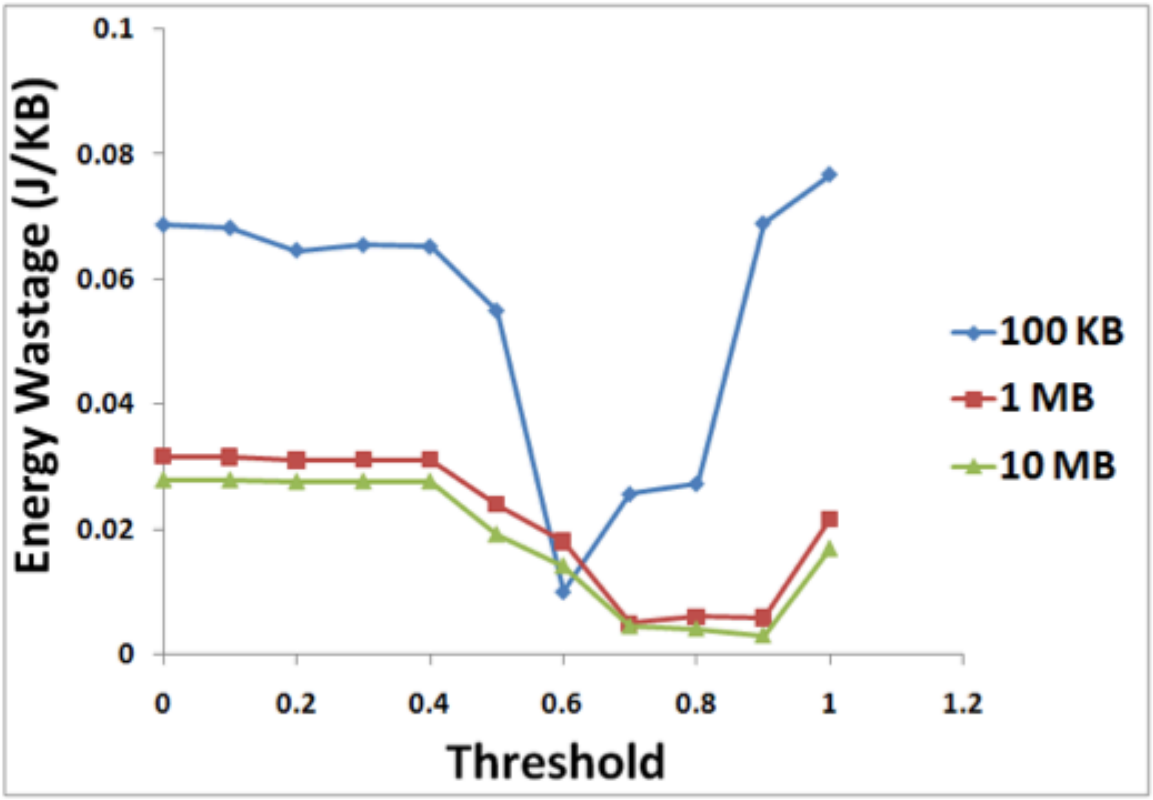

Periodic Discovery

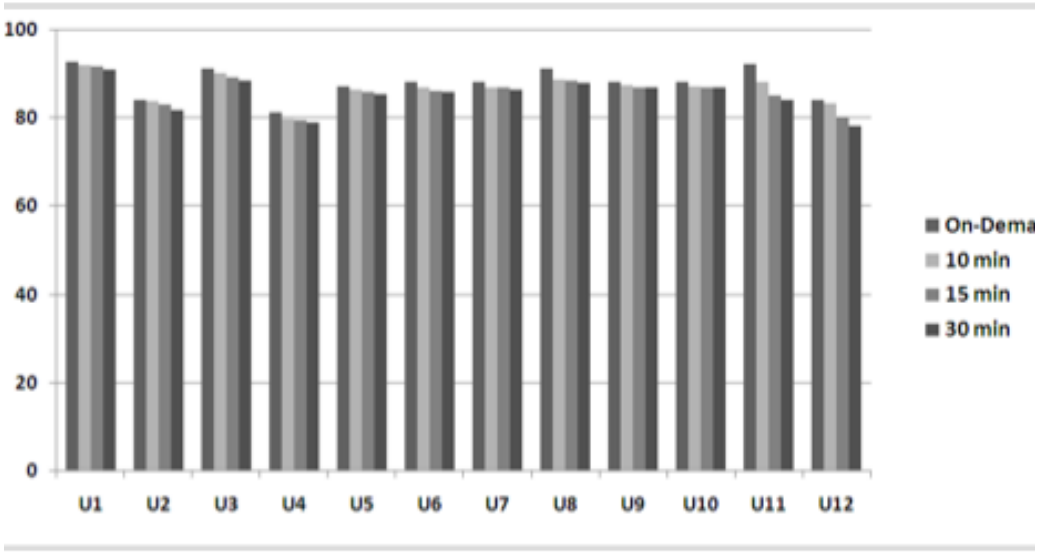

Reduction in Accuracy due to lower periodic Bluetooth discovery rate.

Reduction in Coverage due to lower periodic Bluetooth discovery rate.

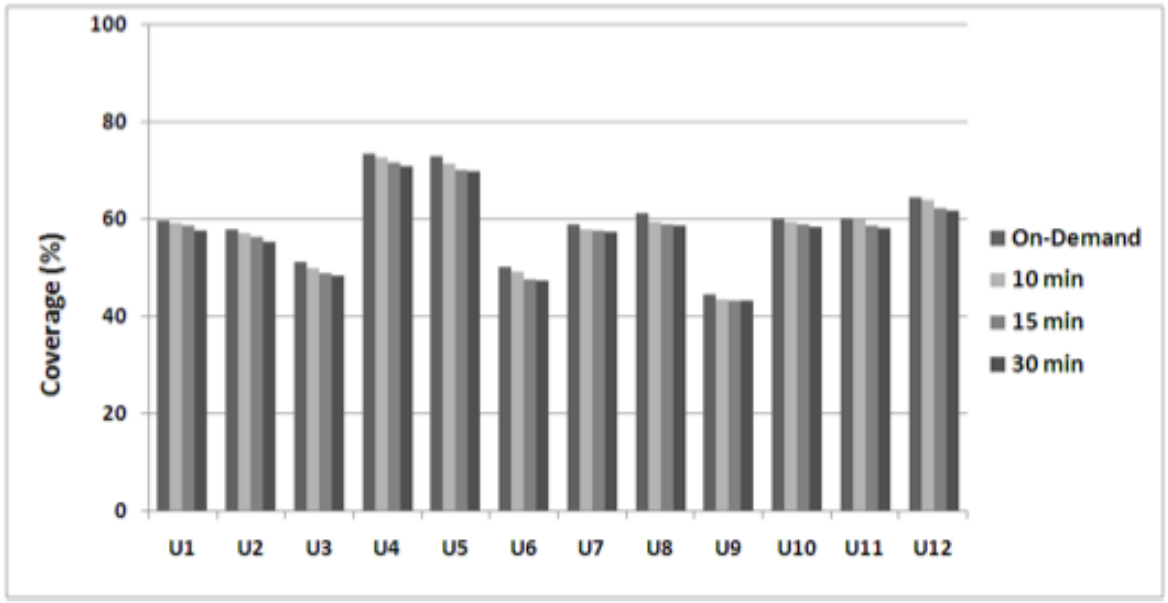

Further Improvements

- Browser authentication: Avoid connecting to redirecting AP's.
  - How to check for connectivity?
    - **Response size:** Ask for a large file, check response size.
      - What if the file was moved or the file server is redirecting?
    - **Secure HTTP:** Ask for unsecured page, check if received page is secure.
      - Identifies only 68.4% of authentication pages.
- Multi-hop sharing
  - Use multiple Bluetooth devices.
  - Higher number of hops -> Higher coverage -> Lower accuracy.
  - 234% increase of detected Bluetooth devices.

Related Work(1)

- BlueDust – Spread Bluetooth landmarks around.
  - Simple and cheap deployment. e.g. Bluetooth dongles.
- $Predict_{BT}$ – Learn prediction patterns, and construct rules based on patterns.
  - e.g. " Wi-Fi is available ten minutes after you spot the Bluetooth device b".
- Wake-on-Wireless – dedicated radio transmitter, expensive to deploy.
- On-Demand-Paging – Requires APs with both Wi-Fi and Bluetooth connectivity.
- Cell2Notify – Wake on VOIP call from dedicated servers. Expensive to deploy.

Related Work(2)

- War-driving – Drive while sniffing packets, similar to Google street, very expensive.
- Context-for-Wireless – Use only cell towers, high coverage.
- Intel Place Lab – Better positioning Wi-Fi hotspots.
- CoolSpots – Bluetooth access points..

Summary

- Main contributions: Minimizing Wi-Fi scanning.
  - Identify landmarks.
  - Identify being stationary.
- Strengths:
  - Easy to deploy.
  - Low energy consumption.
- Weaknesses:
  - Evaluation scope.
  - P2P security.
  - Central server complexity.
  - Sensing being stationary.
- How much power does it save during the day?
- Can it work when cellular connectivity is not available? (Underground).
- Using other sensors.