Ubiquitous and Mobile Computing CS 525M: An Empirical Approach to Smartphone Energy Level Prediction

Qian He (Steve)

Computer Science Dept.

Worcester Polytechnic Institute (WPI)

Overview

- Introduction
- Related Work
- Data Collection
- Energy Prediction
- Successful Execution Prediction
- Contribution

Some background knowledge

- Authors: ***Earl Oliver*** and ***Srinivasan Keshav*** from David R. Cheriton School of Computer Science, ***University of Waterloo***, Waterloo, ON, Canada
- ***BlackBerry***, made by Research In Motion Limited (RIM), a Canadian company
- Delay Tolerant Network applications
  - DARPA
  - NASA

Introduction

- The energy density of smartphone, batteries has grown at a comparably insignificant rate.

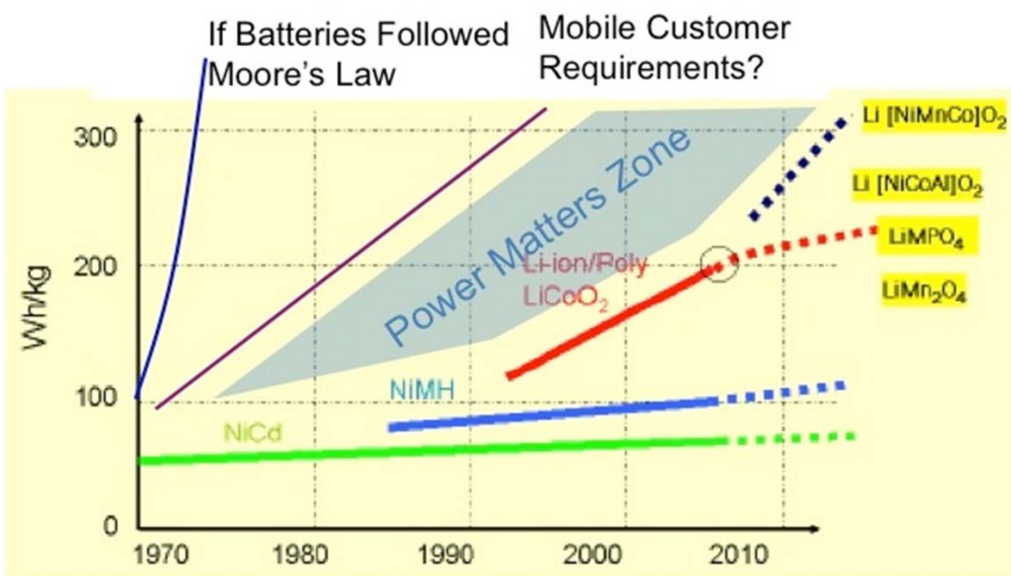

Source: Avicenne

http://www.kk.org/thetechnium/archives/2009/07/was_moores_law.php

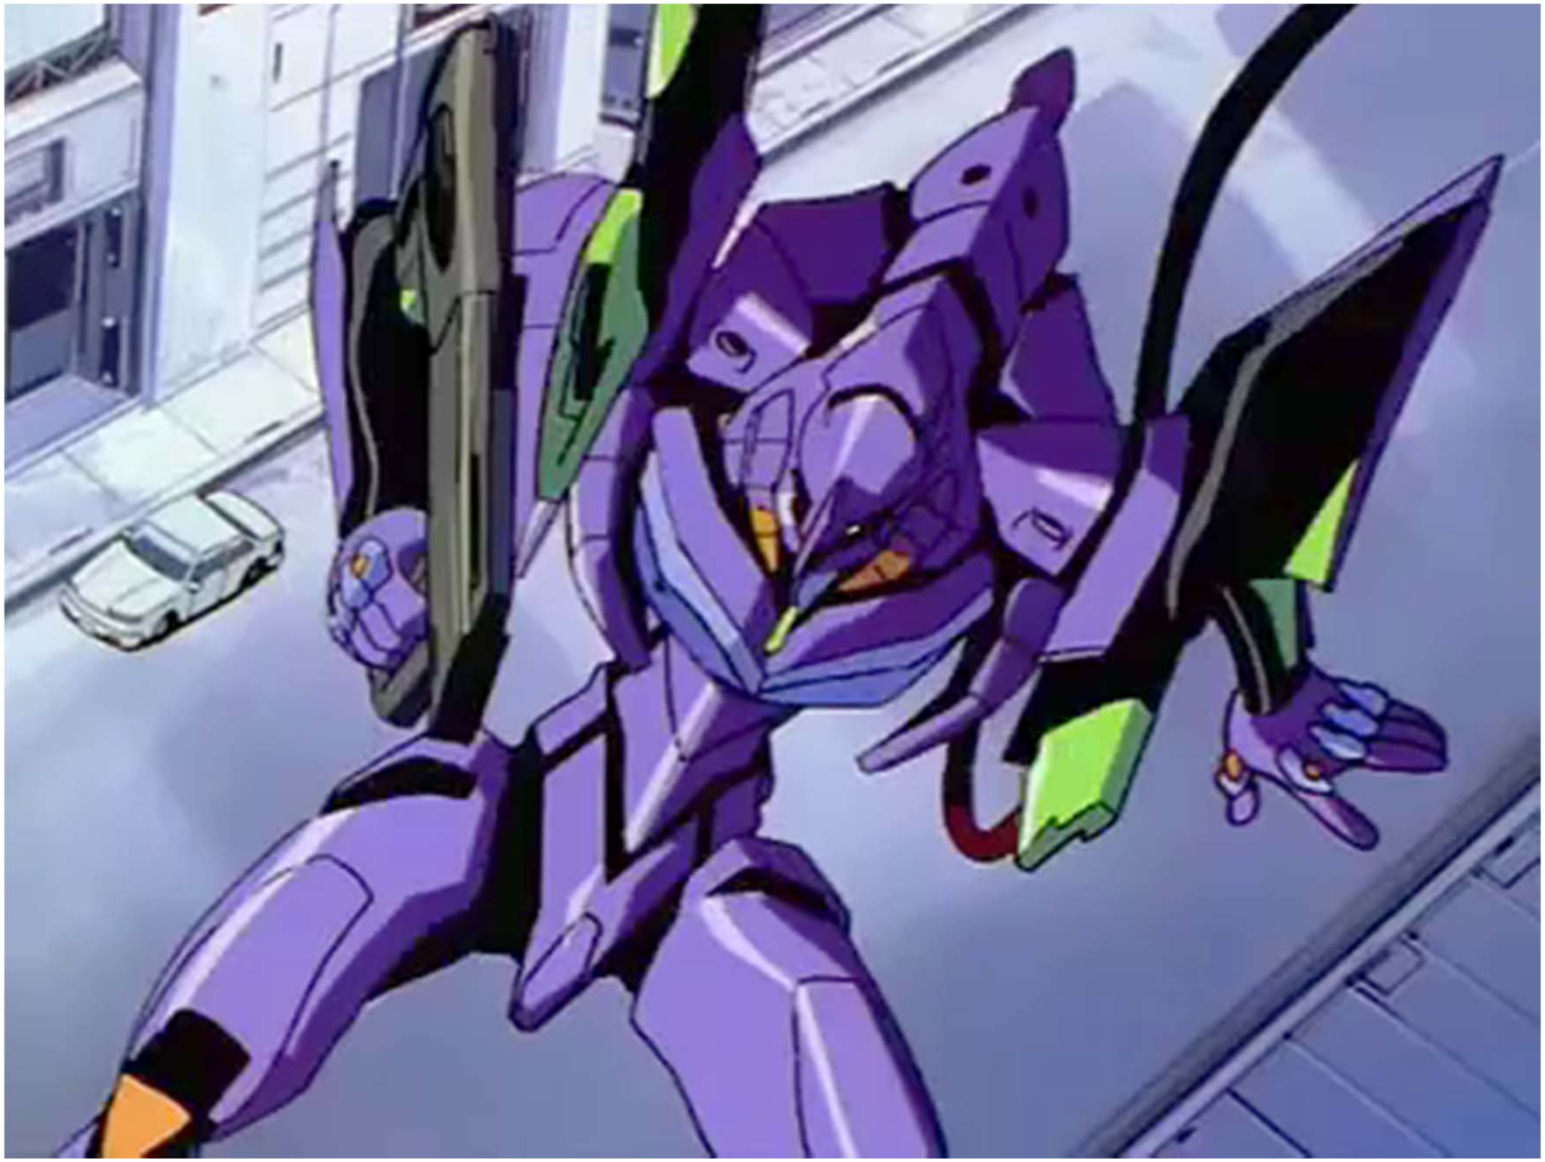

Neon Genesis Evangelion, Episode 3

Introduction

- The authors want to help a class of applications that:
  - must operate with **LITTLE** or **NO** user-intervention
  - must operate over **LONG** durations, and
  - consume **LARGE** amounts of energy

  - e.g. pedometer, GPS tracker, and delay tolerant applications

Related Work

- Energy consumption has been widely studied, ***little is known*** about the charging characteristics of users.
- ***Android*** (33) and ***Windows Mobile*** (222) were studied, but BlackBerry was not. (How about iPhone?)
- Energy model for estimating energy consumption ***in an application*** was studied, but model for concurrent apps was not.
- ***OS***-directed power management was studied.
- The relation between a user's ***location*** and when the user will charge their device was analyzed.

Data Collection

- ***<u>Logger</u>***
  - an event-driven BlackBerry application that runs continuously in the background of a device
    - Backlight activity
    - Idle counter
    - Charging activity
    - Battery level
    - Soft shutdown
    - Device type
  - augmented by another logger developed by a major BlackBerry software developer, and can upload the data to ***<u>the company's servers</u>*** each week

Data Collection: Logger

- ***Backlight Activity***
  - records the time that the backlight turns ***ON*** and ***OFF*** with OS callbacks.
- ***Idle Counter***
  - records the time when the backlight is ***ON***, but the user ***is NOT interacting*** with the device.
- ***Charging Activity***
  - records when a device is ***plugged*** and ***unplugged*** from an external power source
- ***Battery Level***
  - records the battery level ***every 10 minutes***

Data Collection: Logger

(cont.)

- ***Soft Shutdown***
  - records when a user powers off a device when the ***battery is low***, plug it in, unplug it at a later time, ands power on at a full charge.
- ***Device Type***
  - records the device type and OS version.

Data Collection: Challenges

- **Volatile File Systems**
  - Buffer in memory, and then upload data each night
- **Energy Constrains**
  - Polling or even-driven (callback)?
- **Third-party application Intervention**
  - Statistical abnormality detection, e.g. spyware
- **Non-linear Time**
  - Time Sync

Data Collection: Result

- Over *6 months*
- Over *20,100 smartphone users* from *23 time zones*
- 1150 years of cumulative interaction and energy consumption behavior
- Approximately 15 years of suspicious data from 213 users was discarded due to the reasons previously discussed
- BlackBerry device types released since early 2006
  - BlackBerry Pearl 8100+

Energy Prediction: Characteristics

- Char
- Char
- Batt
  - Pa… …y or we…
- Char
  - Th… rep… …gged fro…

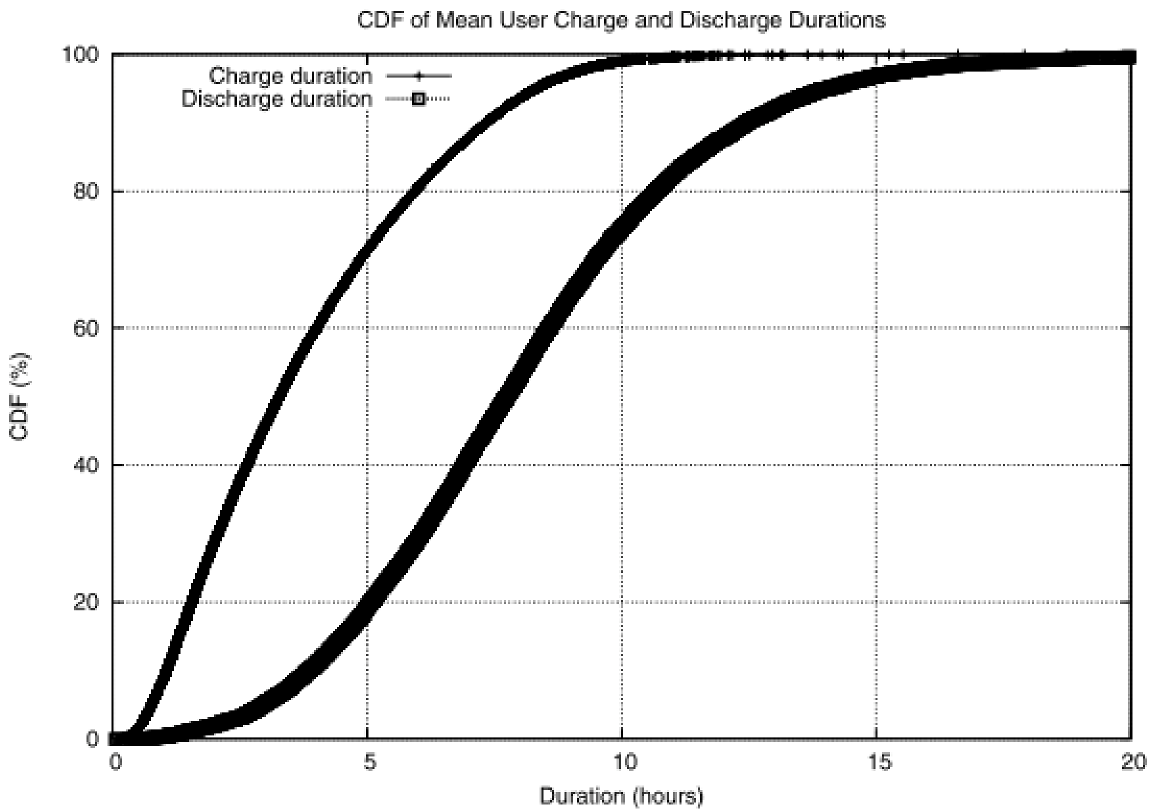

Figure 2. CDF of participants' mean charge/discharge duration.

Energy Prediction: Algorithm

- Divide each week into 336 discrete 30-minute buckets
- Knowledge we have:
  - The ***δ*** vector contain the mean charge/discharge durations for cycles initiated during a specific bucket
    - $\boldsymbol{\delta_{charge}}$
    - $\boldsymbol{\delta_{discharge}}$
  - The ***ρ*** vector contains the mean charge/discharge rate for each bucket.
    - $\boldsymbol{\rho_{charge}}$
    - $\boldsymbol{\rho_{discharge}}$

Energy Prediction: Algorithm

(cont.)

- Required input:
  - the current charge/discharge state: γ
  - the time that the current charge/discharge cycle began: t_{last}
  - the current time: t_{curr}
  - the current battery level: b_{curr}
  - the desired prediction time: t_{pred}
- Required input:
  - the predicted battery level: b_{pred}

Energy Prediction: Predictor

- **Individual Predictor**
  - prediction made on an individual basis
- **Simple Predictor**
  - prediction made on the entire participant population
- **Device Predictor**
  - prediction made on the population with the same device type
- **Clustered Predictor (explain later)**
  - prediction made on the population within the same cluster

Energy Prediction: Predictor Performance Comparison

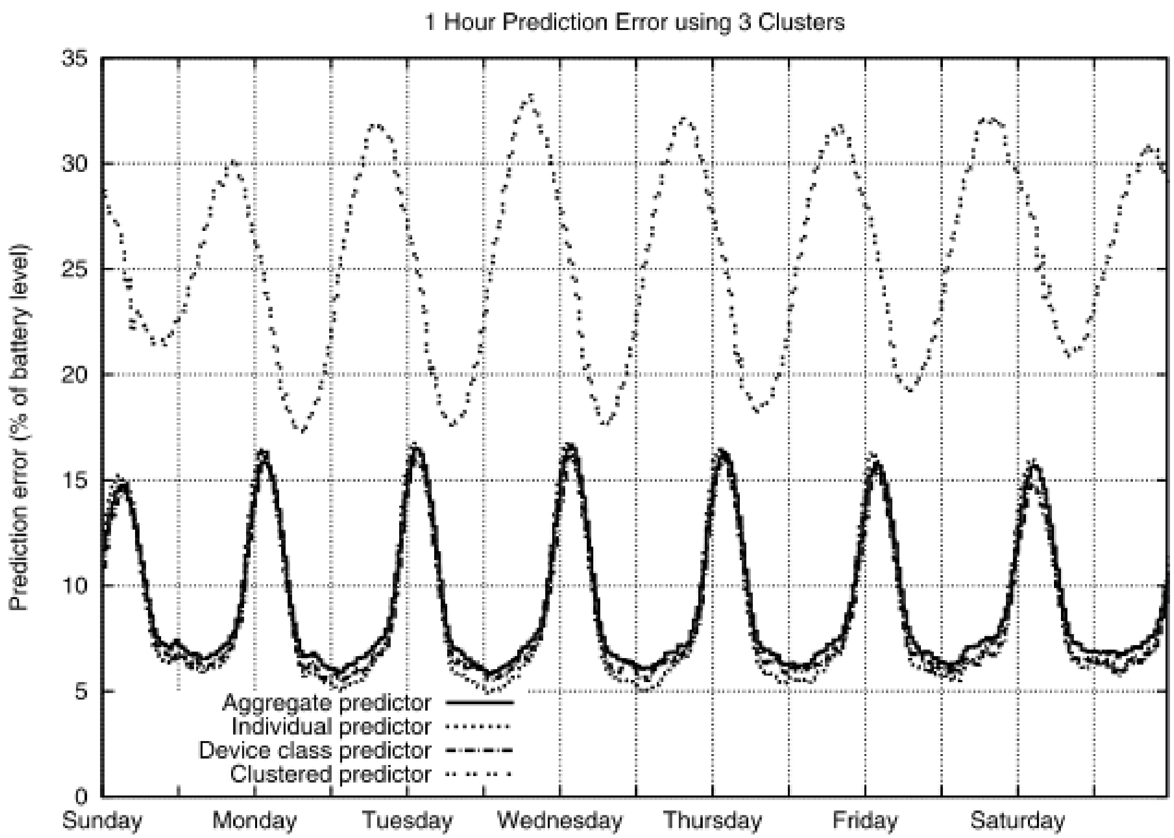

(a) One hour absolute prediction error.

Energy Prediction: Predictor Performance Comparison

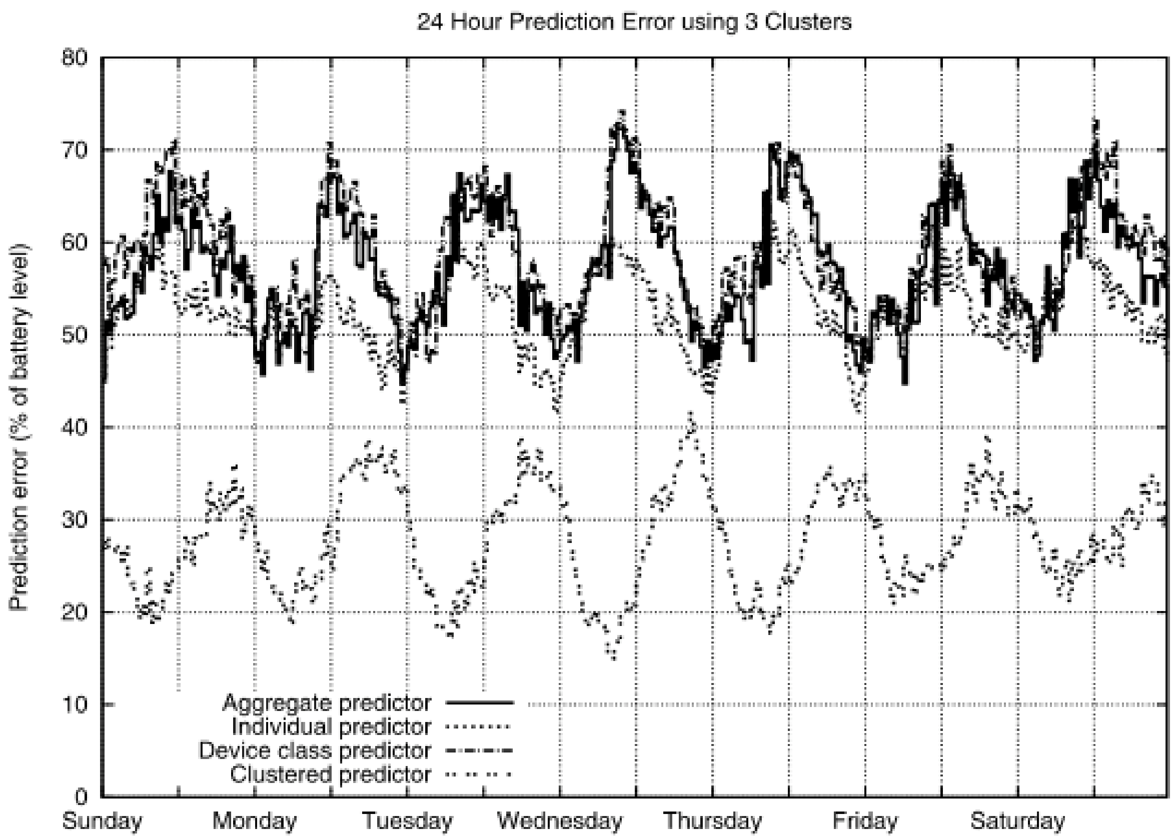

(b) Mean absolute one day prediction error.

Energy Prediction: Predictor Performance Comparison

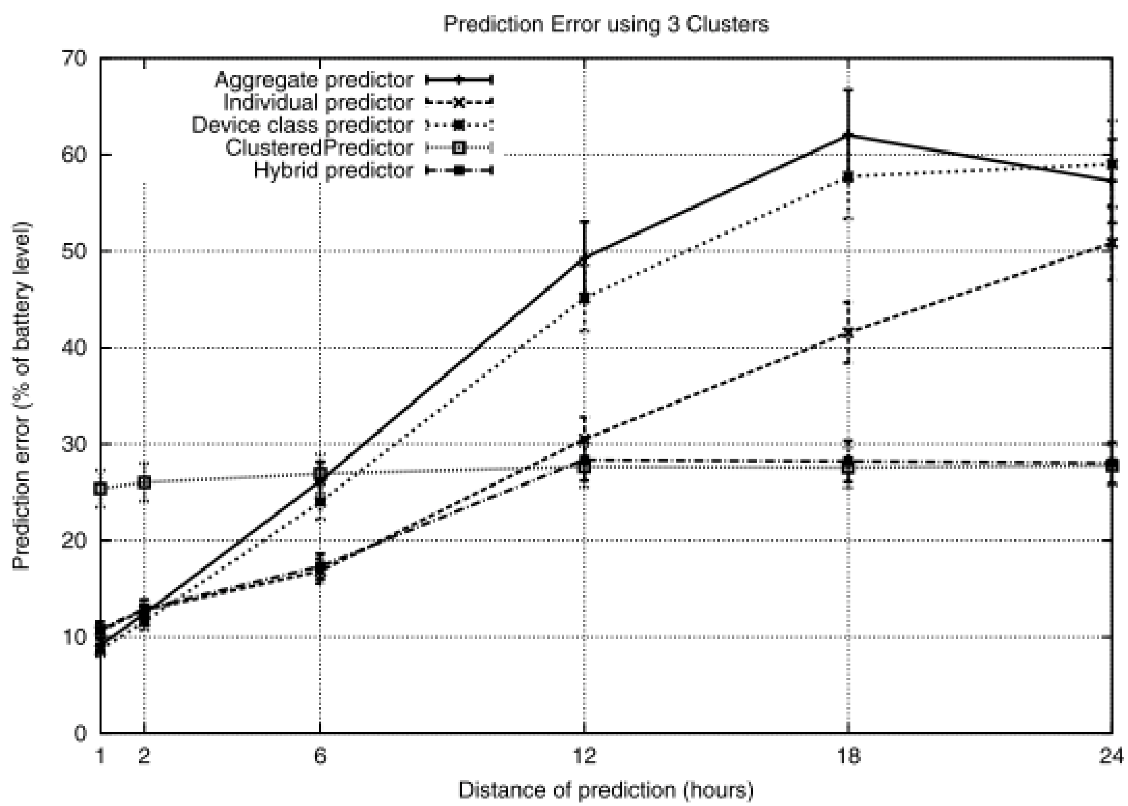

Energy Prediction: User Classification

- For each user, predict 6 b_{pred} with t_{pred} = {1, 2, 6, 12, 18, 24} hours
- Cluster the users with attribute tuple (b_1, b_2, b_6, b_{12}, b_{18}, b_{24})
  - Clustering Algorithm: K-means
    - K = 3 is the best among [2, 6]
  - Distance function: Euclidean
  - Validation: 3-fold cross classification

Energy Prediction: User Classification Result

- **Opportunistic chargers (approximately 63%)**
  - the ***most common*** type of smartphone users
  - the ***most aggressive*** energy consumers (nearly ***4.8%*** per hour)
  - ***frequent***, ***short charge*** durations during the hours of 8am to 5pm

Energy Prediction: User Classification Result

- **Light consumers (approximately 20%)**
  - the ***lowest*** energy discharge rate
  - ***longer*** charge duration
  - discharge their devices over a longer duration
  - allow their battery to drop to its lowest level before initiating a charge (on average, ***34%***)
  - the ***lowest*** discharge rate, but "***surprisingly***" (?) maintain the lowest mean battery level of **56.0%**

Energy Prediction: User Classification Result

- **Nighttime chargers (17%)**
  - these users initiate a charge (probably) before going to bed (***10pm to 11pm***)
  - their mean charge duration is ***significantly higher*** than the other two groups (during the night)
  - maintain a mean battery level of ***72.5%***
  - initiate a charge at an average battery level of ***56%***

Charge Duration

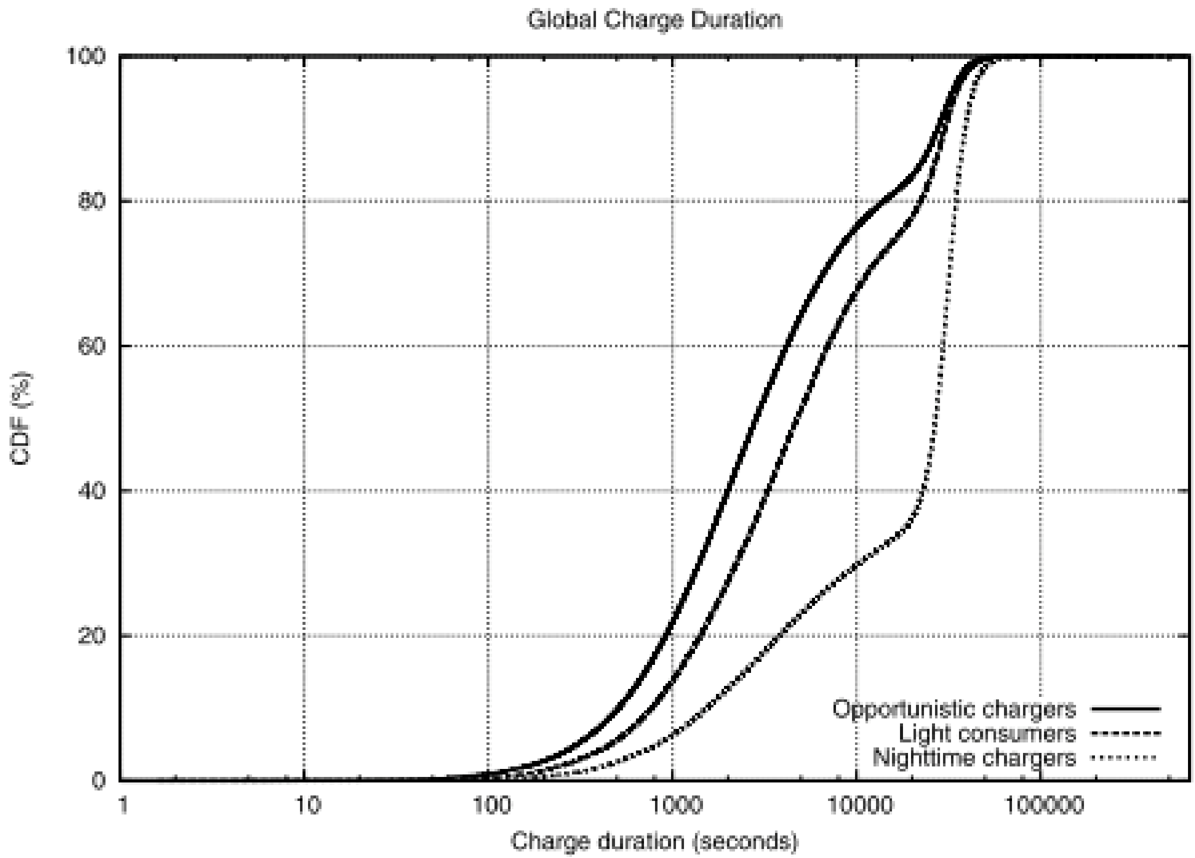

(a) CDF of charge duration.

Discharge Rate

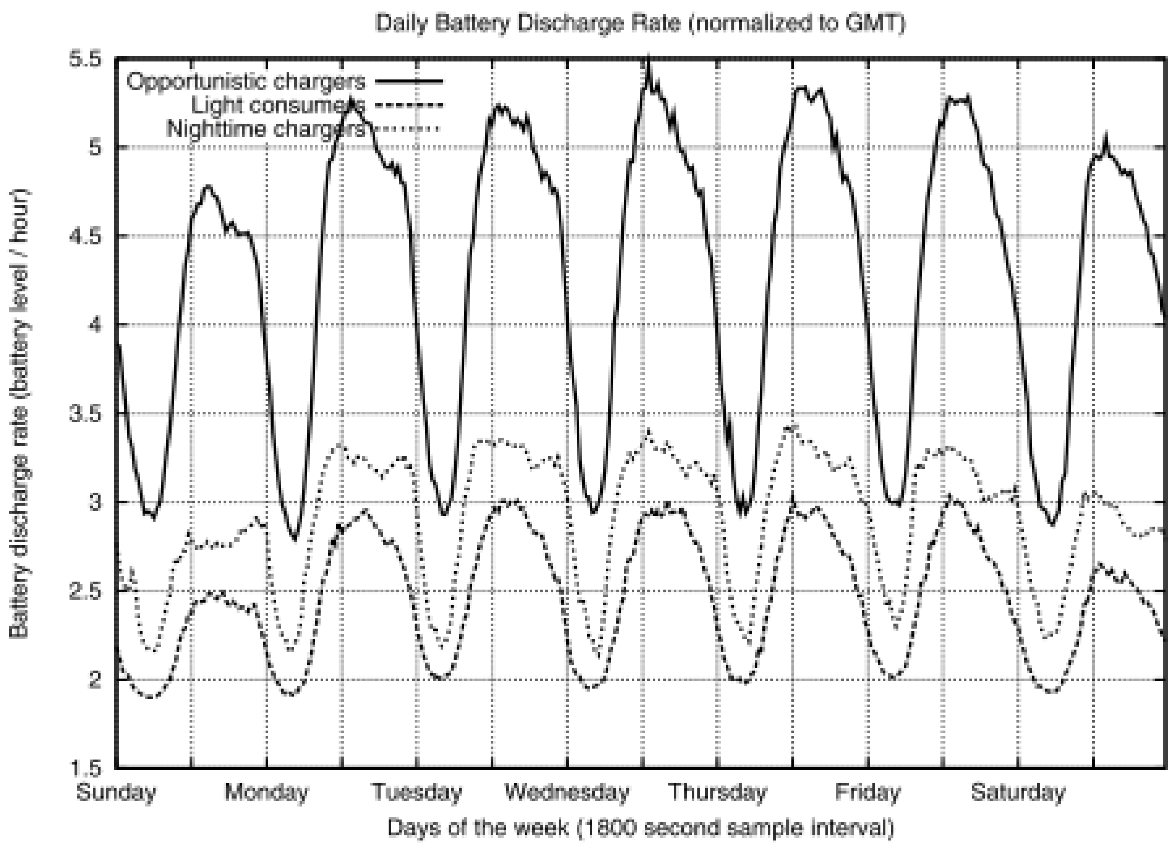

(a) Discharge rate over the week.

Charge Initiation Time

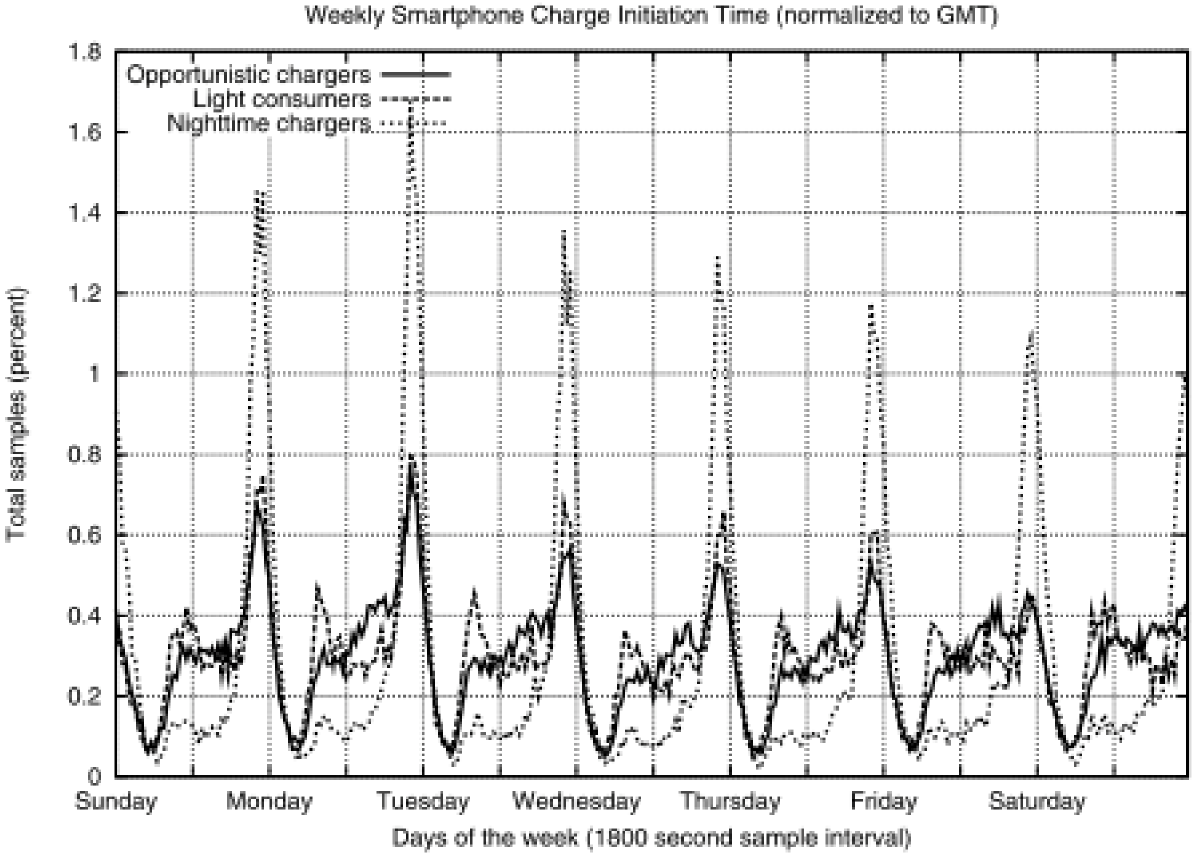

(c) ~~PDF~~ of charge initiation time.

Mean Battery Level

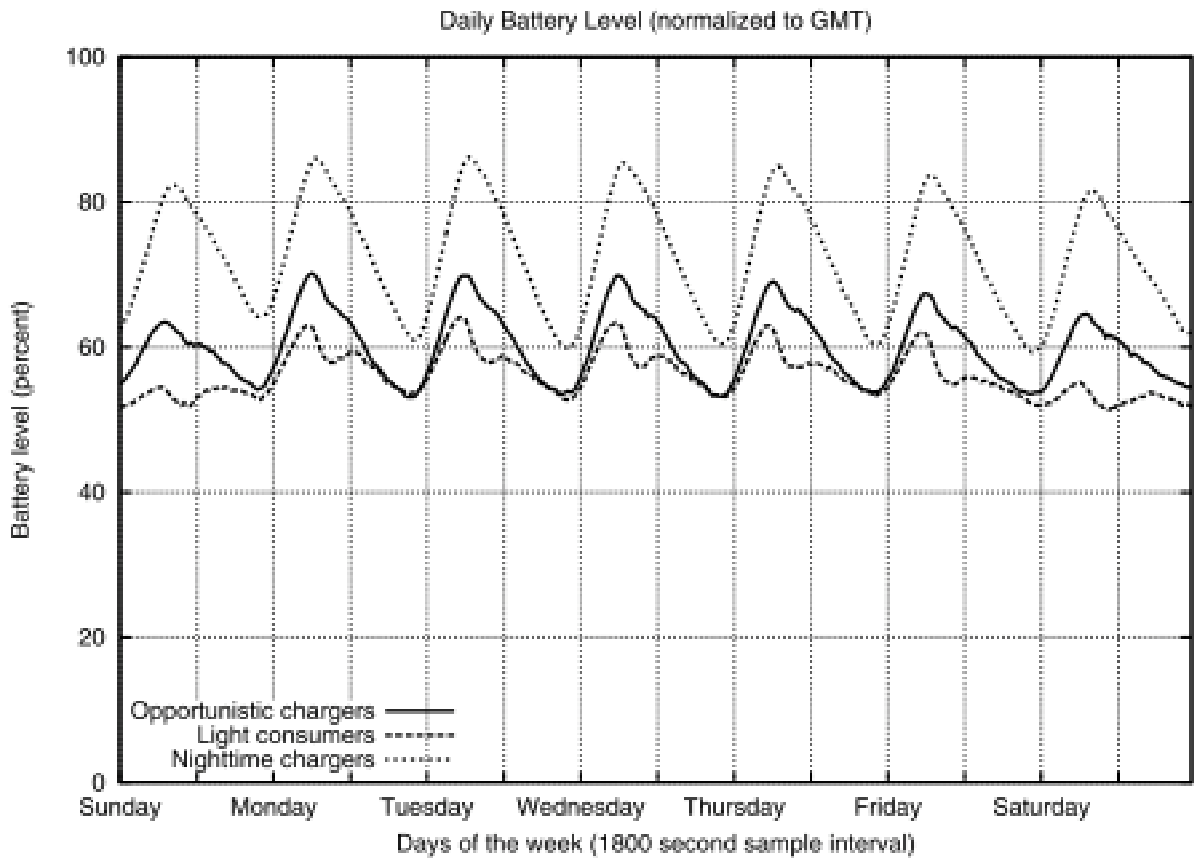

(b) Mean battery level over the week.

Successful Execution Prediction (EET, the tool for this)

- Cut the timeline with 5-minute interval
- Notate the start time of each 5-minute interval as ***$time_i$***
- **Then, f(*$time_i$*, *app_j*)** =
  - ***1*** (if the application ***app_j*** invoked at time ***t_i*** successfully executed for the next ***30 minutes, a.k.a. "bucket"***)
  - ***0*** (otherwise)

Successful Execution Prediction: App Example

- A ***delay-tolerant network (DTN)*** application
  - scans for neighboring Bluetooth devices ***every 1 min***
  - connects to a server on the Internet over WiFi and uploads 100 KB of memory-resident data on ***every second scan***
  - reads a 100 MB file from flash memory and uploads the contents to the server ***once per hour***
  - ***then (one per hour)***, downloads 100 MB of data from the server and subsequently writes the data to persistent flash memory
- The average consumption rate is ***569mW***

Successful Execution Prediction: App Example

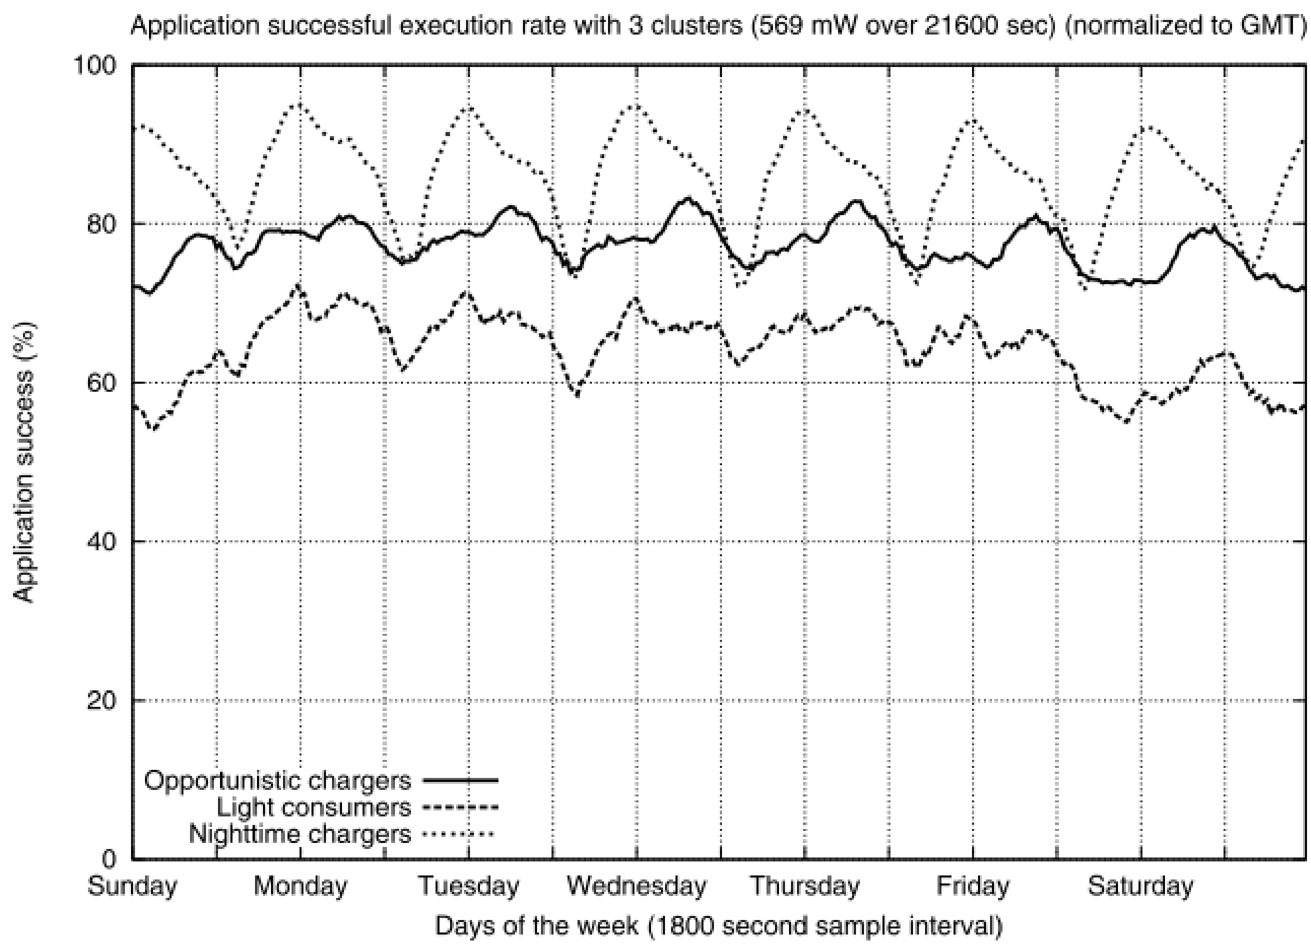

(a) Six hour execution duration.

Successful Execution Prediction: App Example

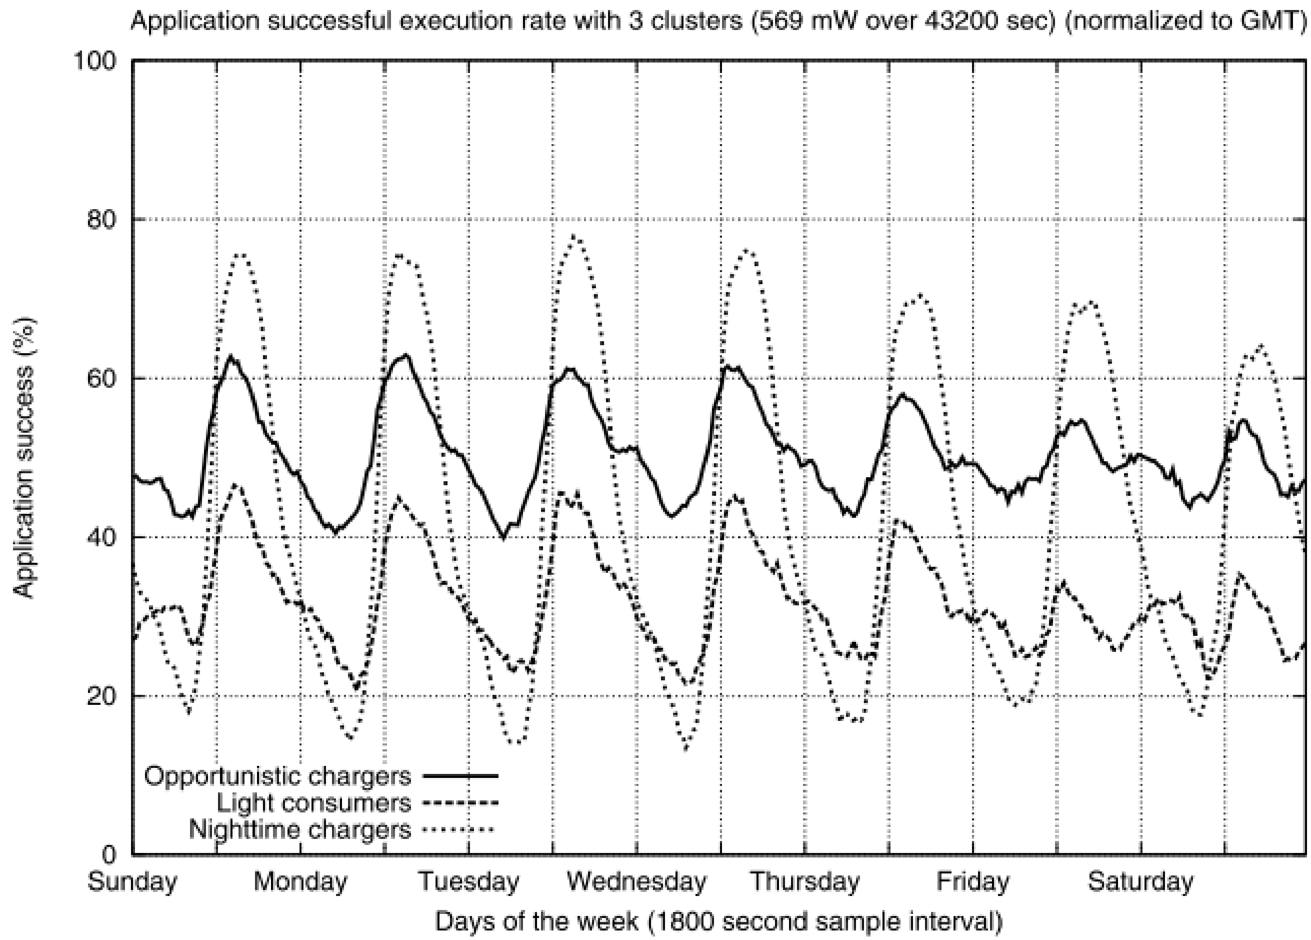

(b) Twelve hour execution duration.

Energy Management Oracle (EMO)

- Input:
  - Same as EET
- Output:
  - ***1***, if the application can safely execute the operation.
  - ***0***, if the operation will result in the depletion of the battery.
- Use Energy Prediction
  - Hybrid Predictor

EMO vs. EET (modified version)

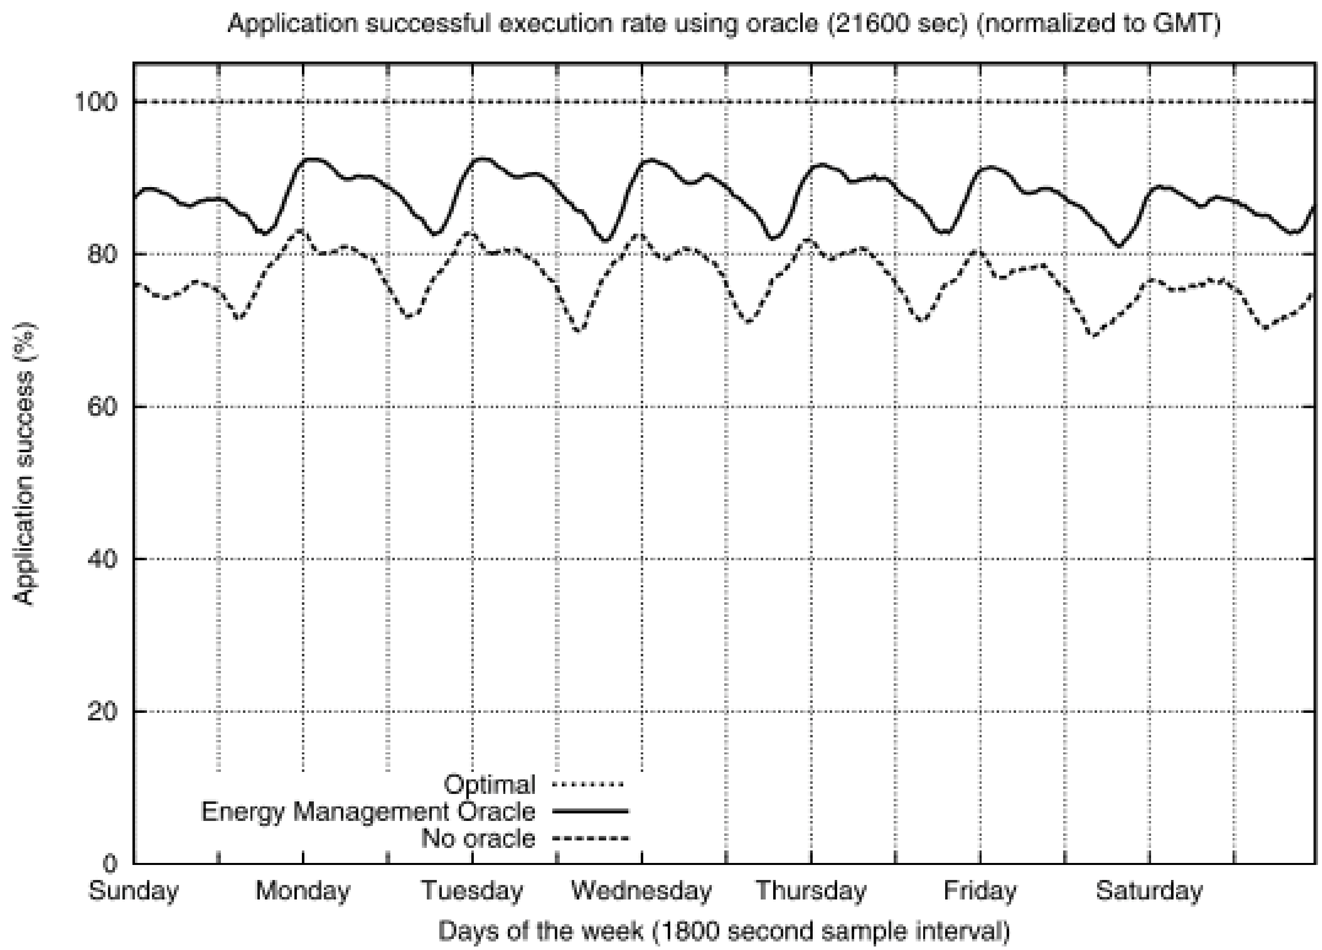

(a) Six hour execution duration using EMO.

EMO vs. EET (modified version)

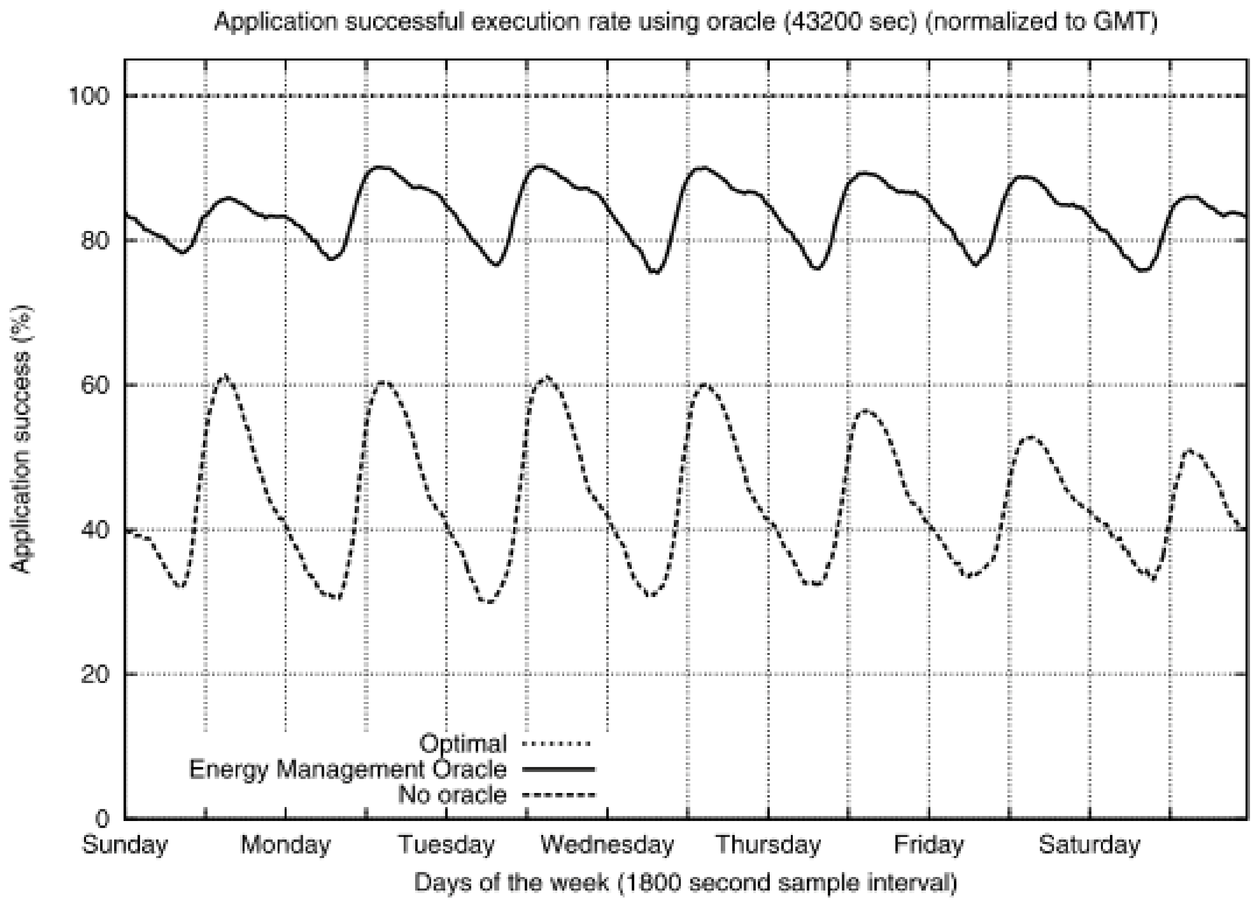

(b) Twelve hour execution duration using EMO.

Contributions

1. The authors built ***a dataset*** containing the smartphone usage and energy consumption characteristics of 20,100 BlackBerry smartphone users.
2. They exploited the dataset and built the ***Energy Emulation Toolkit (EET)***. Developers can test their apps' energy consumption behavior against existing energy traces.
3. They classified users into one of three groups according to their unique energy consumption characteristics, they demonstrated that energy level ***can be predicted***.

Q & A

THANKS