12-2011

Location-Aware Traffic Management on Mobile Phones

Sarath Krishna Mandava
University of Arkansas, Fayetteville

LOCATION-AWARE TRAFFIC MANAGEMENT ON MOBILE PHONES

LOCATION-AWARE TRAFFIC MANAGEMENT ON MOBILE PHONES

A thesis submitted in partial fulfillment

of the requirements for the degree of

Master of Computer Science and Computer Engineering

By

Sarath Krishna Mandava

Jawaharlal Nehru Technological University

Bachelor in Computer Science & Engineering, 2009

December 2011

University of Arkansas

ABSTRACT

The growing number of mobile phone users is a primary cause of congestion in cellular networks. Therefore, cellular network providers have turned to expensive and differentiated data plans. Unfortunately, as the number of smartphone users keeps increasing, changing data plans only provides a temporary solution. A more permanent solution is offloading 3G traffic to networks in orthogonal frequency bands. One such plausible network is open Wi-Fi, which is free by definition. As Wi-Fi networks become ubiquitous, there are several areas where there is simultaneous Wi-Fi and 3G coverage.

In this thesis, we study the feasibility of offloading 3G traffic to Wi-Fi networks. First, we design a custom tool for the Android phone, which helps us collect data on CPU usage, GPS coordinates, applications running on the platform, and traffic generated by the smartphone. With the help of initial data collected from the tool, we quantify the amount and characteristics of traffic that users generate from smartphones. Next, using the data we show that at several locations offloading a considerable amount of data is possible from 3G to Wi-Fi. Our study can lead to the design of multiradio systems that prevent traffic overload on 3G networks.

This thesis is approved for recommendation

to the Graduate Council

Thesis Director:

____

Dr. Nilanjan Banerjee

Thesis committee:

____

Dr. Craig W. Thompson

____

Dr. Dale R. Thompson

____

Dr. Wing-Ning Li

THESIS DUPLICATION RELEASE

I hereby authorize the University of Arkansas Libraries to duplicate this Thesis when needed for research and/or scholarship.

Agreed ____

Sarath Krishna Mandava

Refused ____

Sarath Krishna Mandava

ACKNOWLEDGEMENTS

I thank are due to the staff of the University of Arkansas Graduate School for all of their help with theses. It would be impossible to make it through the semester without their help. Secondly, I thank Falaki [et al][17] for their Systemsens tool which was helpful in the project. Finally, I thank the faculty and staff at the University of Arkansas for their commitment to the University and to the students.

TABLE OF CONTENTS

LIST OF FIGURES AND TABLES

Chapter 1: Introduction

Cellular networks such as 3G or EVDO are popular modes of mobile data connectivity. Among the various data standards, 3G is popular due to its moderate bandwidth and wide coverage. However, due to increasing user demand, cellular networks are overloaded and congested. As a result, the end-users suffer from reduced bandwidth, missed voice calls and unreliable coverage, particularly in metropolitan areas. To mitigate the above problem, cellular providers have resorted to expensive and often differentiated data plans. While the solution is plausible in the short term, it is unlikely to scale in the long term. Expensive data plans causes irritation among users.

Another interesting alternative solution to the problem is augmentation using radios in other frequency bands. For instance, Wi-Fi networks are becoming more and more ubiquitous. Open Wi-Fi is prevalent in residential, downtown, and metropolitan regions. Open Wi-Fi by definition is free and provides much higher bandwidth than 3G and can be used as a complimentary radio for 3G access. However, given that the range of Wi-Fi is short, its coverage is sparse. Hence, it is unclear whether augmenting 3G with Wi-Fi is practical. To this end, in this thesis we ask the following question: *At popular places for individuals, is offloading from 3G to Wi-Fi plausible?*

There are several open research questions that need to be tackled to address the above question. First, it is unknown what traffic patterns are most common among smartphone users. The amount and success of traffic offload is a function of the type of data generated from the device. For example, applications such as VoIP that require near ubiquitous connectivity are not good candidates for offload. However, applications that require download or dissemination of bulk data can be used for offload since Wi-Fi provides larger bandwidth for short periods of time.

Second, given the locations at which users spend the majority of their time, it is desirable to know how much data can be offloaded to balance the complexity of protocols that facilitate such an offload.

To address these challenges, we perform a systematic two step study in this thesis. First, we design and implement a custom tool for the Android platform that collects data. We modify the SystemSens application and use it to collect location data, CPU usage data, data on applications resident in the Android memory at any time, and the amount of traffic transmitted through the wireless interfaces. We deploy the tool on two smartphones and collect data over three weeks. We use the collected data to quantify the type of traffic generated from applications running on the Android platform. Further, we use GPS coordinates to understand where the data is generated by the user. Our analysis shows that amount of data transmitted by users is small and is not continuous and hence is a good candidates for offload.

Next, we study what share of data from 3G can be offloaded to Wi-Fi. To answer this question quantitatively, we perform extensive experiments in Chicago and Fayetteville. We calculate the amount of data that can be offloaded from 3G to Wi-Fi at popular user locations. From the experiments, we observe that about 11% of data can be offloaded in downtown Chicago and 5% can be offloaded in downtown Fayetteville, indicating that augmentation works better in cities like Chicago, where the density of open Wi-Fi access is higher. We validate our results by performing experiments at different speeds --- vehicular Internet access and access on foot. We observe that vehicular Internet access faces more disruptions and data loss in Chicago and almost nonexistent in Fayetteville. But, the quality of access (in terms of network speeds) is better in downtown of Fayetteville as compared to downtown Chicago in spite of its low Access Point density.

Overall, under different network conditions and locations, we show that offloading data from 3G to Wi-Fi can lead to substantial bandwidth savings. Our study can lead to the design of multiradio systems that intelligently determine when radios like Wi-Fi can be used to offload traffic from cellular radio networks like 3G.

Chapter 2: Related Work

Our research builds on related work in several areas of networking, energy efficiency, vehicular networking, and multi-radio systems. Here, we compare and contrast our work with the literature most relevant to our work.

Section 1: Wireless and Mobile Networking

Several studies in the past have analyzed the speeds and the availability of 3G and Wi-Fi [2, 16-26]. These studies focused on three core areas. First, there are several papers in the literature that study the performance of Wi-Fi in a mobile scenario. These studies show that the performance of Wi-Fi is reduced in a mobile scenario due to several reasons including interference, multipath fading, and physical obstructions. Second, there are several research papers on trying to improve the performance of Wi-Fi in a mobile scenario. These include aggregating bandwidth across multiple access points [30-33], coordination using multiple access points [30], and improving the performance of the join process of Wi-Fi, namely the association and DHCP [33]. Third, there are recent studies on the performance of cellular data connections in a mobile scenario. Much of data has been collected and analyzed to show that the performance of 3G is also affected adversely by the prevalence of coverage holes, and co-channel interference from other cell phones in the vicinity. Moreover, recently, cellular providers have restored to differentiated data plans and bandwidth caps to circumvent the limited bandwidth issue. The goal of our project is complimentary to the above research. Specifically, we in our project, we are studying the feasibility of augmenting 3G networks with Wi-Fi. The study can act as a reference to build multi-radio systems that uses 3G and Wi-Fi to offload 3G traffic onto Wi-Fi.

Section 2: Vehicular Testbeds

Our work has relevance in the area of vehicular testbeds. Vehicular testbeds act as an appropriate evaluation platform for 3G,Wi-Fi systems. Among the vehicular testbeds, Dieselnet [33] is the most prominent one, where a 3G augmented Wi-Fi system can be built and deployed. Moreover, there are other mobile evaluation platforms, including simulation and emulation based where multi-radio systems as suggested by our study can be built. As future work, it might be useful to further augment our study by extending the feasibility study on these mobile testbeds.

Section 3: Multi-radio Systems

Our work builds on a large corpus of work on multi-radio systems. Multiple radios have been used in the wireless and mobile systems community for several purposes. First, energy efficiency is a primary goal of multi-radio systems. For example, energy hungry radios like Wi-Fi are coupled with low power radios like Zigbee to facilitate low power consumption of the entire system. Wi-Fi has been augmented with the cellular voice 2G radio to reduce the energy consumption of VoIP calls [34]. Second, multiradio systems have been used to increase coverage of wireless networks. For example, short range radios like Zigbee are augmented with long range radios like the 900 MHz radios to provide better coverage [35]. Recent efforts on trying to mitigate the increasing bandwidth demand of 3G involve trying to offload some of the traffic to Wi-Fi or WiMax. Efforts involve measurement studies that show the amount of offload possible. In line with these research attempts, our goal is to determine whether 3G can be augmented with Wi-Fi. However, we take a different approach to the problem. We first deploy a tool on Android phones that collects of data including location, applications resident in the phone memory, and the amount of traffic transmitted from the wireless interfaces. After thorough analysis using the collected data, we determine the speeds of the two networks are different locations. Assuming that the user spends a certain fraction of time at a particular location, we then determine the amount of data that can be offloaded from 3G to Wi-Fi.

Chapter 3: System Architecture

The proposed system consists of two components, a mobile client, and a server. The client end of our system is a data and statistics collection tool called ***"SystemSens+"***. The collected statistics are transferred over to the server over a 3G or Wi-Fi connection. Here, we describe in detail, the different components of the mobile client and the data collection tool. The tool is modified from the SystemSens tool, and facilitates data collection of traffic measurements that is key to our project.

Section 1: Client setup

The designed client runs on an Android device equipped with Wi-Fi, 3G, and GPS radios. It has several modules that run parallel and as a background process collect data on GPS locations, Wi-Fi access points, proximate Bluetooth devices, battery statistics, and traffic generated over the network interfaces.

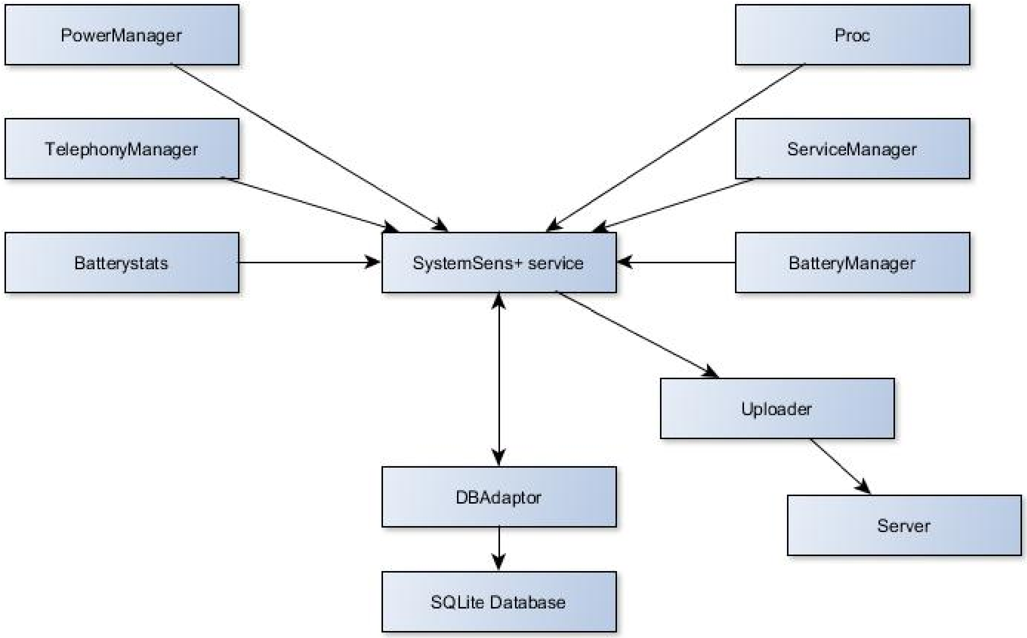

Figure 1: Overall System architecture

Section 2: Mobile Statistics Tool

The SystemSens+ tool installed on the mobile client runs as a background service when the application is launched. Hence, the service is continuously functional and collecting data and statistics and offloading them to our backend server. This service collects 3G and Wi-Fi data usage, the Wi-Fi status, call states, screen states and residual battery capacity. The service also provides an optional user interface that can be used to view the collected data.

The tool uses the following Android SDK APIs for data collection.

Section 2.1: PowerManager.

The Power manager provides control over the power state of the device. It provides wavelocks (used to control power states of the device) to switch the screen and keyboard to various power states and reboot. The continuous usage of the API corresponding to the power manager can lead to significant power draw of the entire device.

Section 2.2: SQLiteDatabase.

SQLite is the on-device database for temporary data storage on the Android devices. It provides a sql style API for storing and retrieving relational data. It provides intuitive APIs in the form of methods to create, delete, and update tables on the database. It is critical for our monitoring tool due to two reasons. First, it provides an efficient interface for accessing and storing data locally on the phone---use of files is inefficient given the rate of data access and storage for our application. Second, it provides data backing. If the data collected is not continuously transferred over to the server due to some disruptions in the internet access, it can lead to a significant data loss. Storing critical data locally on the devices is always a safe option.

Section 2.3: TelephonyManager

The telephony manager provides access to information about the telephony services on the device. The SystemSens+ tool uses this service to determine the call states of the phone.

Specifically it records the following call states: CALL_STATE_OFFHOOK, CALL_STATE_RINGING and CALL_STATE_IDLE.

Section 2.4: ServiceManager

The service manager provides methods to get a reference to a service, add a new service, and retrieve an existing service. The SystemSens+ tool uses this class to get battery information and usage statistics.

Section 2.5: Batterystats

The batterystats provides access to battery usage statistics, information on wavelocks, processes, packages, and services. SystemSens+ application uses it to get GPS readings, battery usage and process statistics.

Section 2.6: /proc

The API to the /proc helps read information from the /proc file system. In the SystemSens+ tool, the API is used to read the content of the file, /proc/net/dev. This file consists of a row for each network interface. Each row in turn contains the number of bytes and packets that were transmitted and received over that network interface. This information is encoded as a JSON object and inserted into the SQLite database.

Section 2.7: BatteryManager

The BatteryManager contains the strings and constants used for **ACTION_BATTERY_CHANGED** intent. This intent is used to collect statistics on the status of the battery and the health of the battery.

Section 3: Server Setup

The mobile client collects data on several important statistics, stores it locally in a SQLite database and transmits it to a server at a regular interval. A PHP script on the server acquires this data using POST methods. The format of the POST messages used in our implementation is JSON. However, XML is another acceptable format. The PHP script reads this data and writes it to a file on the server.

Chapter 4: Experimental setup, Testbed and Implementation Details

Experimental setup: We have performed experiments in two cities, Chicago and Fayetteville both on foot and in the local transits. Chicago represents a dense urban area while Fayetteville is a semi-urban locality. We have collected data from four Android users. Our results represent data collected over a period of 2.5 months for one user in both Chicago and Fayetteville and 6, 15 and 20 days for other users in Fayetteville respectively. Data was collected once every two minutes and offloaded to a central server both at peak hours (10:00 AM to 4:00 PM) and off-peak hours (8:00 PM to 1:00 AM).

Testbed : Our testbeds were the local railway transits (Red Line and Blue Line) in Chicago. Coming to Fayetteville we used the local buses (Green, Maple Hill and Tan routes).

Section 1: Phone Statistics collection

As described in the previous chapter, we collect data from Android phones on GPS coordinates, data transmitted and received over 3G and Wi-Fi, and battery statistics. Additionally, screen states and the status of the battery was collected once every two minutes and stored in the local database.

GPS statistics

While GPS radios on most mobile phones offer high accuracy (order of 10m if there are sufficient number of satellite fixes), it can be a large drain on the battery. For example, if the GPS unit is used continuously for localization, the phone unit may not last for more than a few hours. Hence, it is critical to exploit the optimal tradeoff between localization accuracy and energy utilization of the phone. Our application polls for a GPS reading once every two minutes and the data is stored in the SQLite database of the phone. The goal behind collecting GPS data is to correlate traffic generated over 3G and Wi-Fi with locations of the user. Specifically, we want to understand if there are places where users transfer more data than others. Such a location based network analysis also helps us determine network availability at different locations.

Data usage statistics

Collection of 3G and Wi-Fi usage

On the Android platform, each radio sends data over a separate network port. While the name of the specific port is phone dependent, on the Samsung phone ***pdp0*** is tied to the 3G interface. Similarly, the port bound to the Wi-Fi radio is ***eth0***. The SystemSens+ service collects the number of packets and bytes sent and received using the Wi-Fi and 3G device every 2 minutes using the ***/proc/net/dev*** interface file.

In addition to the raw data transferred over the Wi-Fi interface, we collect several pieces of data. Specifically, the SystemSens+ also tracks states of the Wi-Fi radio. A message handler and a Broadcast receiver are used for this purpose. These results are event triggered. Whenever there is a change in the state, the data is recorded, stored in the local database, and finally transferred over to the central server.

The various states in Wi-Fi the SystemSens+ service can track are the following**:**

(1) *WIFI_STATE_DISABLED*: This state occurs when the Wi-Fi is disabled.

(2) WIFI_STATE_DISABLING: This state occurs when the Wi-Fi is disabling.

(3) WIFI_STATE_ENABLED: This state occurs when the Wi-Fi is enabled.

(4) WIFI_STATE_CONNECTED: This state occurs when the Wi-Fi is connected.

(5) WIFI_STATE_UNKNOWN: This state occurs when the Wi-Fi state is unknown.

(6) WIFI_STATE_SCANNING: This state occurs when the Wi-Fi state is scanning.

Process statistics (currently running applications):

In addition to collecting data on GPS locations and data transmitted/received through network interfaces, we also collect data on processes resident in the phone's memory. The goal is twofold. First, as future work, we plan to correlate the amount of traffic generated by different applications. While at this point, there is no easy way of calculating the data usage of individual third party applications, we plan to use regression analysis to infer this statistics. Second, the number of applications resident in memory helps us fine tune our monitoring tool to use the optimal amount of system resources such as processing power and memory. We presently collect the currently running processes along with their user-id and timestamp. The processes can be system processes or user processes. The process list also helps us understand usage patterns of mobile phone users. For example, if the ***android.at&t*****music** process runs frequently we can infer that the user listens to music most of the time. If the ***android.browser*** runs frequently the user uses the web frequently.

Battery Statistics

We have discussed the battery statistics collected by our monitoring tool in the previous chapter. Here, we concentrate on the battery states that are also collected as part of the statistics collected by our monitoring tool. Specifically, we collect data on battery status, level, voltage, temperature provided by various labels.

The health of a battery: The following labels refer to the various states regarding the health of a battery.

- ***battery_info_health_unknown***
- ***battery_info_health_good***
- ***battery_info_health_overheat***
- ***battery_info_health_dead***
- ***battery_info_health_over_voltage***
- ***battery_info_health_unspecified_failure***

Here the application checks whether the battery health is good, overheated, dead, over voltage or any other unspecified failure.

Call states

The various call states of the phone are also recorded by the application

CALL_STATE_OFFHOOK: This indicates that the device call state is off-hook

CALL_STATE_RINGING: This indicates that the device call state is ringing.

CALL_STATE_IDLE: This indicates that there is no call activity

We use ***TelephonyManager*** package in android sdk to track these states. These statistics helps us track the phone activity of the user.

Chapter 5: Evaluation

In our evaluation we answer the following key questions using data collected from the four Android users**.**

- How does the data usage over 3G and Wi-Fi vary at different locations and different times of the day?
- What is the period of time a user uses 3G and Wi-Fi at different locations?
- What is the bandwidth of 3G and Wi-Fi at different locations and different times of the day?

While answering these questions, we also present results on the number of recharges that users perform during the period of data collection. Further, we present results on the number of memory resident processes while the user is actively using his mobile device.

How diverse are the user visited locations?

We have collected GPS statistics from four test users. For the first user the location coordinates varied a significantly, indicating that the user travelled during the period when the data was collected. On the other hand, the other two users remained in their town. For the fourth user the data variation was small indicating that he remained in a relatively small area.

The data for the first user was a combination of Fayetteville and Chicago. There are several sources of inaccuracy in the location data. First, location data collected indoors are inaccurate since the receiver does not get a direct line of sight with the satellites. Additionally, we found that the location coordinates are inaccurate when the user moves through tunnels and urban canyons and railway tunnels.

A feature that we found in our experiments was the inaccuracy in GPS readings when the user was traveling in a vehicle at high speed. For example, we found that the GPS readings were inaccurate up to 35m when the vehicle was driven at a speed close to 60mph. We claim that several techniques including Software-based receivers or differential GPS units can circumvent the above errors.

What network states was the Android device in?

We collected statistics on the states of the Wi-Fi network card is in. These include scanning, obtaining DHCP leases, associated, and data transfer. The data collected is event triggered and hence it is collected whenever the network state changes. A change in a network state also occurs when the user connects or disconnects from a network or when he switches from one network to another.

These statistics help us to determine the number of times a user used Wi-Fi or switched from one Wi-Fi network to another. Moreover, it helps us understand how robust the Wi-Fi association process is? Figure 2 shows the average of number of times the user is in a particular network state for four days.

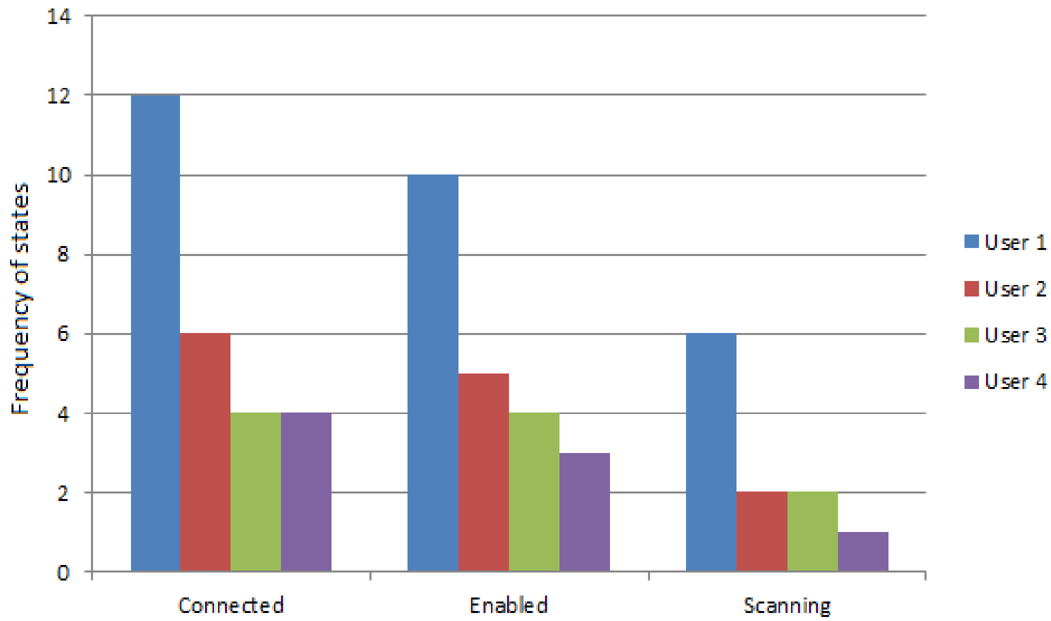

Figure 2: Average number of times the user is in a particular network state for four days

How many processes does an Android phone run?

We used the process service counter to collect data on the number of processes resident in the phone's memory. The Figure 3 below shows the number of processes per user over four intervals.

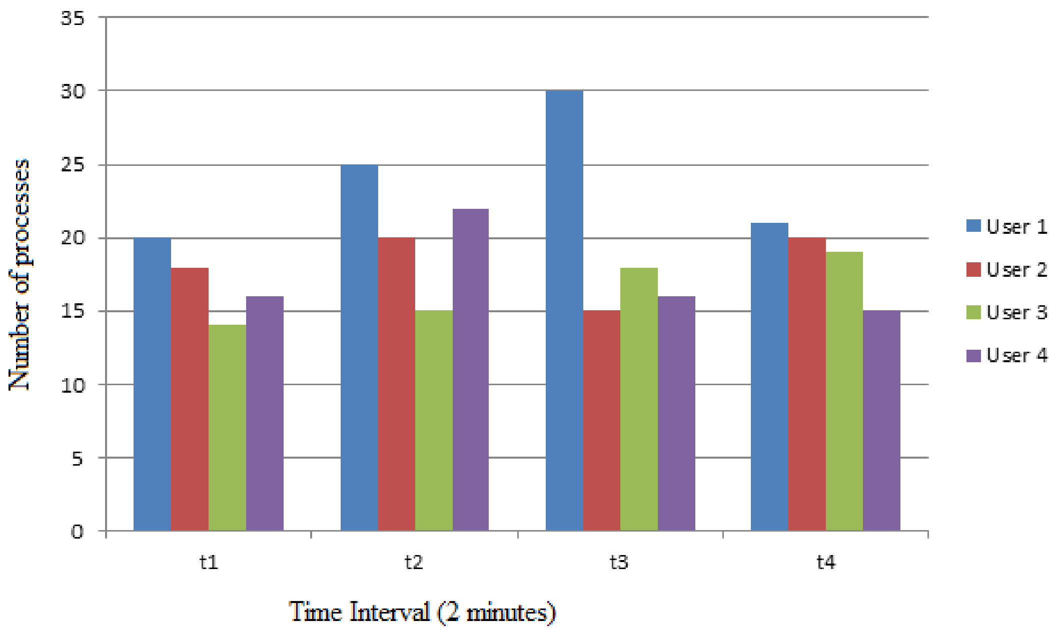

Figure 3: Number of process running during a time interval

Clearly, the number of processes resident in memory is large. Hence, it is critical that our logging tool consumes a minimal amount of energy, to be minimally intrusive. Further analysis of data shows that ***com.google.android.youtube*** was the most frequently used process, indicating that the users spend a large fraction of time watching streaming videos. This also illustrates the need for offload 3G data to networks such as Wi-Fi. There are two reasons for the above. First, the network speed of 3G is lower than Wi-Fi, as we illustrate later. Thus, watching videos over a Wi-Fi connection would lead to better user experience than 3G. Second, with differentiated caps such as 2GB/month, it limits the number of videos that a user can watch over youtube. Hence, it calls for offloading data from 3G to radios such as Wi-Fi. The process list in the above figure is a mix of user and system initiated processes. Moreover, several of the system processes are background processes.

How often do users recharge their phones?

We are interested in the number of times a user recharges his phone. This is key to understanding whether it is feasible to use energy hungry radios such as Wi-Fi for data transfer. Similarly, it also provides information on whether our service can offload data to the server often---an optimization to our monitoring service could be to upload data to the server only when the device is plugged in. Figure 4 shows that the number of recharges that each user performed during the course of four days was of the order of 3-4. This points to the fact that it is feasible to use energy hungry radios like Wi-Fi for data transfer when 3G is either unavailable or expensive.

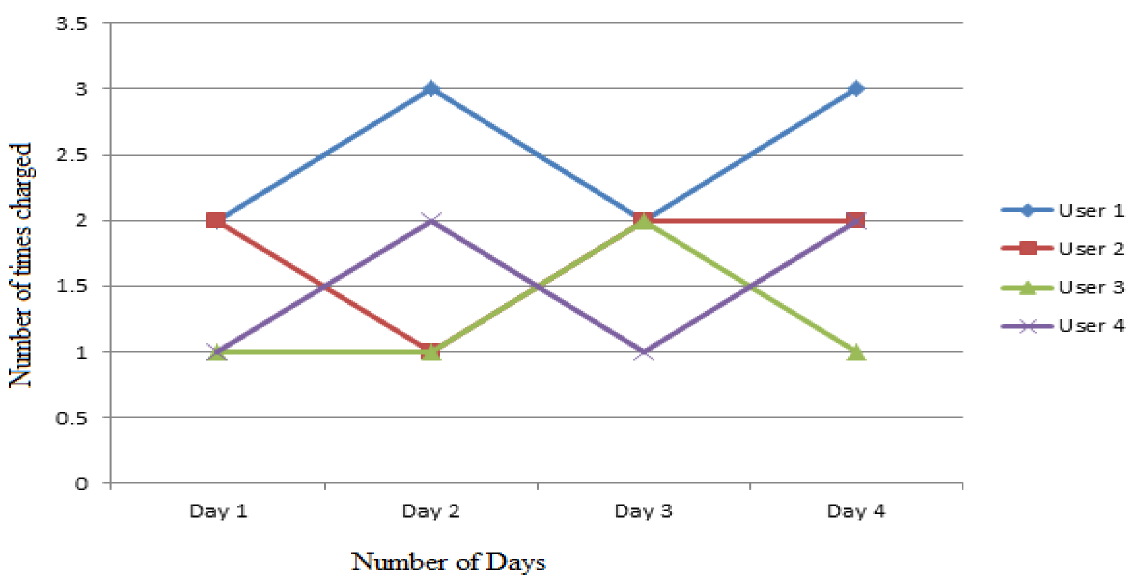

Figure 4: Number of times the users charged their phones

Another important profile that we measured was the energy consumption of the SystemSens+ application. Our measurements show that the application did not affect the number of times a user charged. Since it is built on top of the SystemSens+ service, it is highly optimized to consume the minimal amount of energy throughout the day. Another orthogonal feature deficiency that we found on the Android phone was the battery level indicator. We have observed that the battery discharges slowly when the indicator level is high and quickly when the indicator level is low. However, in the ideal scenario the discharge rates should be uniform.

How much data was transmitted/received over 3G?

In this section, we present results from data collected in Chicago and Fayetteville on the amount of data transmitted/received over the 3G radio. We collected statistics on the number of packets and number of bytes transmitted/received for upstream and downstream 3G traffic. For the Samsung phone we used the ***pdp0*** interface for data collection. Specifically, we used the ***getnetDev()*** function call that parses the file in /proc/net/dev in to obtain the results from the pdp0 Interface. The results are returned in the form of JSONobject.

The usage of 3G of various users during peak and off-peak hours is shown in the graphs below.

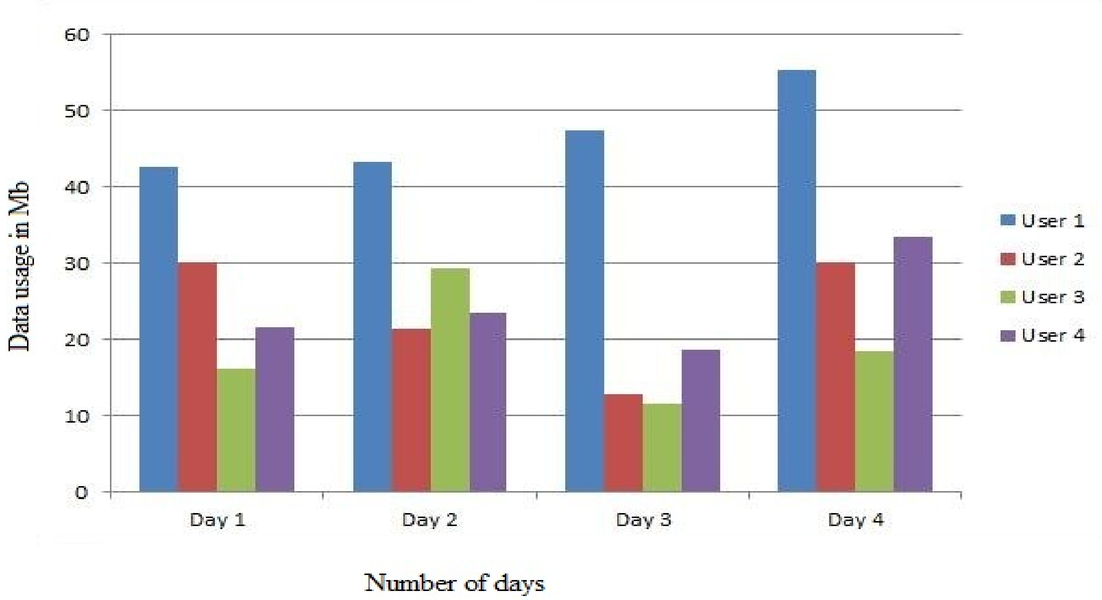

Figure 5: 3G Usage during peak hours (10:00 AM to 4:00 PM)

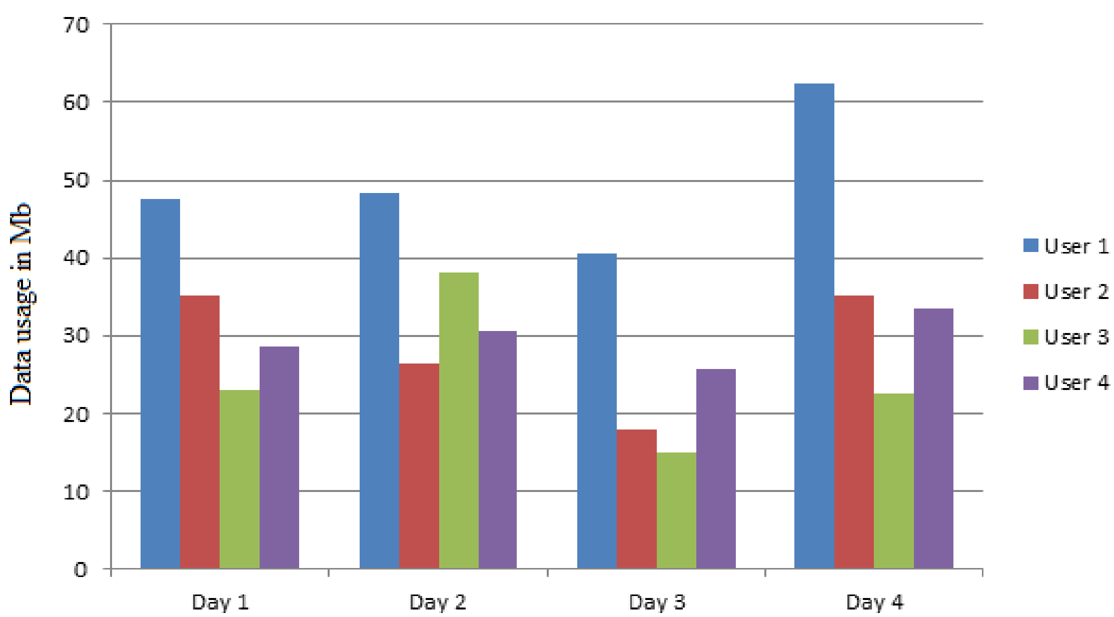

Figure 6: 3G Usage during off-peak hours (8:00 PM to 1:00 AM)

There are several conclusions that we can make through these graphs. First, the amount of data transferred by various users varies. While certain users transmit/receive 40 MBs of data, others transmit fairly low amounts of data. Moreover, the usage level varies depending on the day and peak or off-peak hours. A counterintuitive result that we find from the graphs is that the amount of data transferred during peak hours is less than non-peak hours for 3G except when the user was on a trip. We conjecture that the reason for the above observation is the more frequent use of Wi-Fi during peak hours. Additionally, we find the total amount of data transferred during a day is less than 50 MBs in most cases. This points to the fact that the users are cognizant of the bandwidth caps and try to use the 3G radio as minimally as possible.

How does the time of use of Wi-Fi and 3G compare?

Table 1: 3G and Wi-Fi usage of User1

Place	3G	Wi-Fi
Residence	15min	4.3hrs
Office	32min	3.1hrs
Downtown	4.1hrs	32min

Table 2: 3G and Wi-Fi usage of User2

User 2	3G	Wi-Fi
Residence	26min	3.3hrs
Office	21min	3.1hrs

Table 3: 3G and Wi-Fi usage of User3

User 3	3G	Wi-Fi
Residence	12min	2.3hrs
Office	6min	3.4hrs

Table 4: 3G and Wi-Fi usage of User4

User 4	3G	Wi-Fi
Residence	6min	3.23hrs
Office	3min	2.6hrs

What is the usage of Wi-Fi?

The SystemSens+ service collects the amount of packets and bytes of both upstream and downstream of the Wi-Fi network interface (eth0) along with the data and time for every two minute interval. The results show that users preferred to use Wi-Fi when an access point was near.

Experiments in downtown of Chicago show that there are at least ten Wi-Fi access points at any location. Out of these at least three were open to connection. Due to this high density of Wi-Fi access points, they are also used for indoor localization where GPS does not work.

A comparison of the data usage through Wi-Fi of various users is shown in the graphs below. It was interesting to note that three of the four users were static while using Wi-Fi.

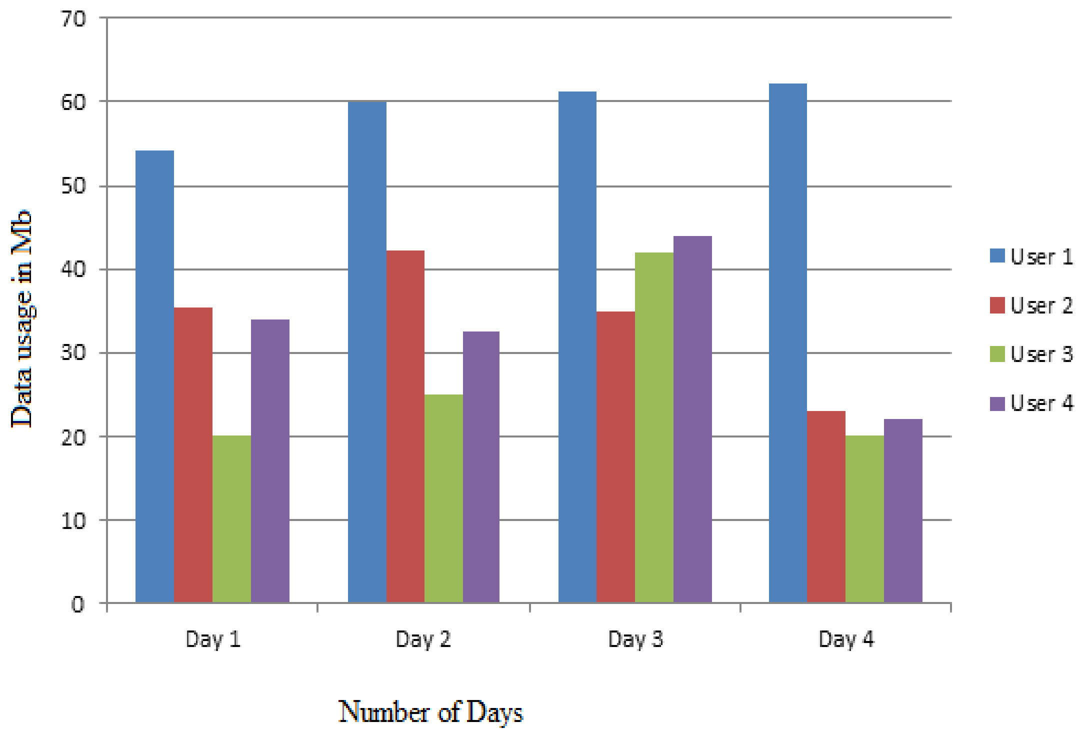

Figure 7 Wi-Fi data usage at peak hours

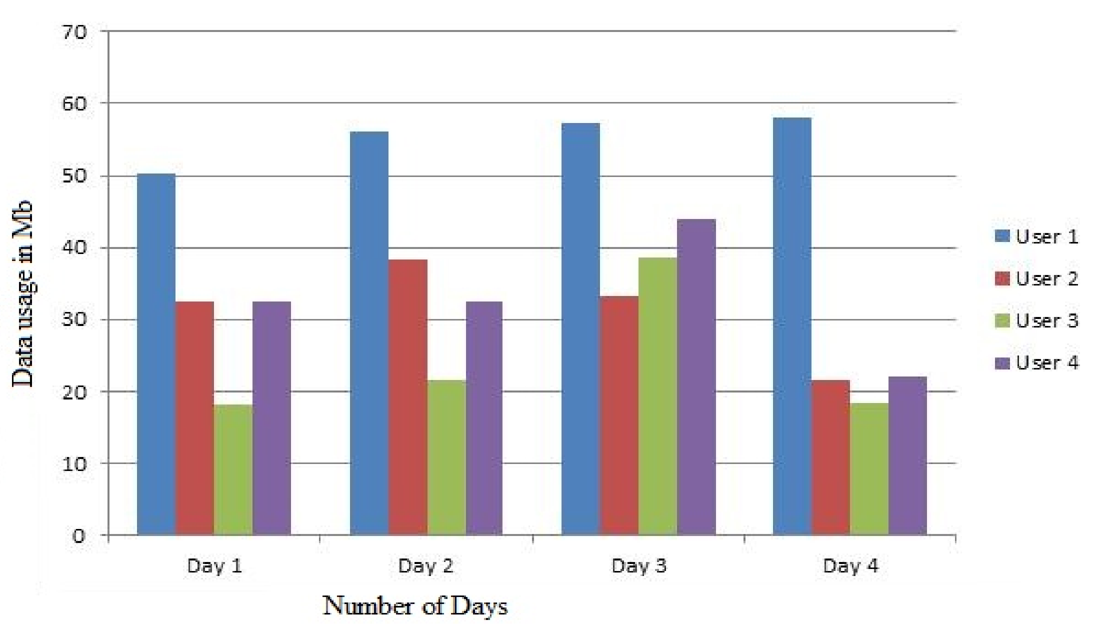

Figure 8 Wi-Fi data usage at off-peak hours

Recently some of the Wi-Fi routers are being fitted with directional antennas. They allow Wi-Fi access only in a particular direction. These might be a disadvantage in our case although they were built in order to address the privacy issues.

Also, there are various factors which affect the Wi-Fi network speed. The network speed is reduced if the access point has peak number of users. Different access points have their peak number of users at various times in a day. For example, access points at a cafe experience the peak during evenings. In order to avoid this problem we have collected the readings from the same access point at various times in a day. Also, the access points that are stable have been identified where the peak hours are not significant. Most of these stable access points were in hotels or in residential areas.

How does the bandwidth of 3G and Wi-Fi compare?

In this section, we compare the download speeds for 3G and Wi-Fi. To perform these experiments, we used inetwork tool. The goal is to understand whether substituting 3G with Wi-Fi at different locations is feasible for applications that require high bandwidth and speed. In a nutshell, our analysis shows that Wi-Fi has double the speed of 3G on an average. So, Wi-Fi would be a preferred solution over 3G if we are interested in applications that require high network bandwidth, such as watching videos.

In downtown, we have considered experiments at locations that have ideal conditions for Wi-Fi i.e., the mobile device is near an open access point and the signal strength is strong. In other locations the Wi-Fi speed was sometimes less than that of 3G particularly when the signal strength was poor.

Broadly speaking, top-down, virtually integrated approach is offered by 3G for delivering wireless Internet access. It represents a natural evolution of the business models of existing mobile providers. The providers of this service have already invested billions of dollars in purchasing licenses to support the advanced data services and equipment makers have been gearing up to produce the base stations and handsets for wide-scale deployment of 3G services.

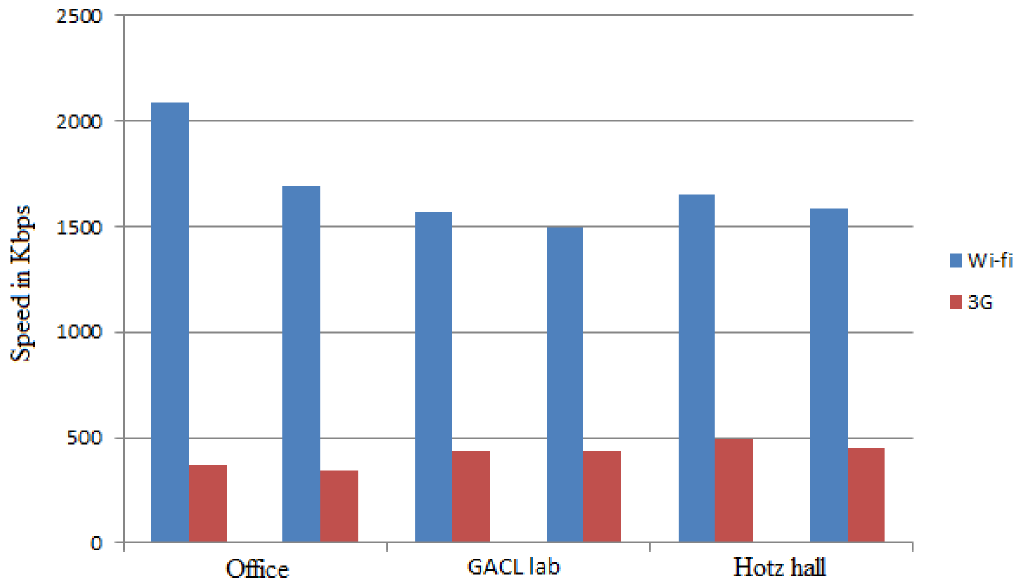

Figure 9 3G vs Wi-Fi speed at Workplace at peak hours

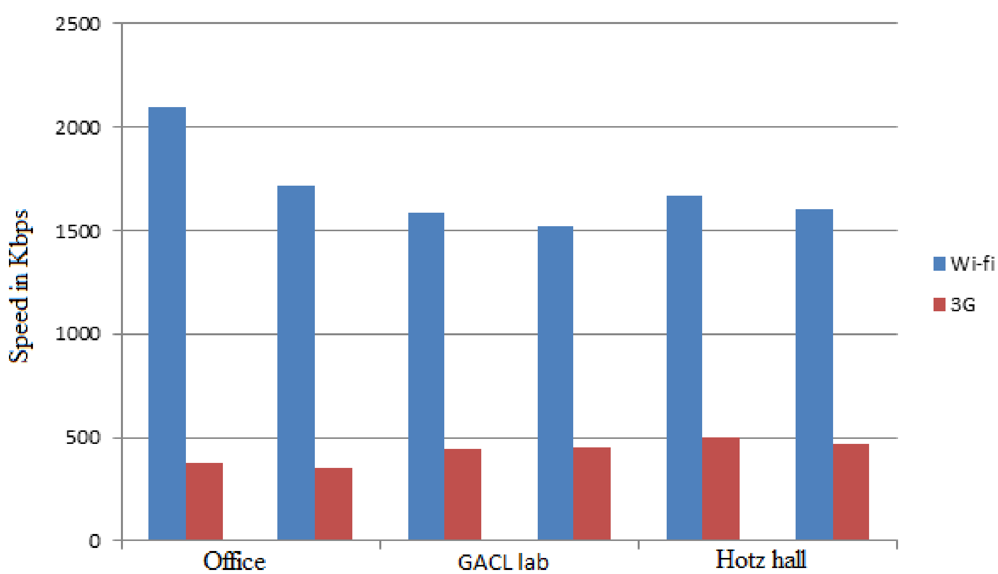

Figure 10 3G vs Wi-Fi speed at Workplace at off-peak hours

In the Figure 9 and Figure 10, we compare the speed of the Wi-Fi and 3G connections at workplaces like offices and labs in the University of Arkansas. Measurements were performed for every two hours during peak and off-peak hours at three locations (My office, GACL lab and Hotz hall). Clearly speeds varied significantly and the speed of Wi-Fi was almost twice that of 3G both during peak and off-peak hours. Also, the speeds during peak hours were less compared to off-peak hours. This difference can be attributed to the larger number of users sharing the medium during peak hours.

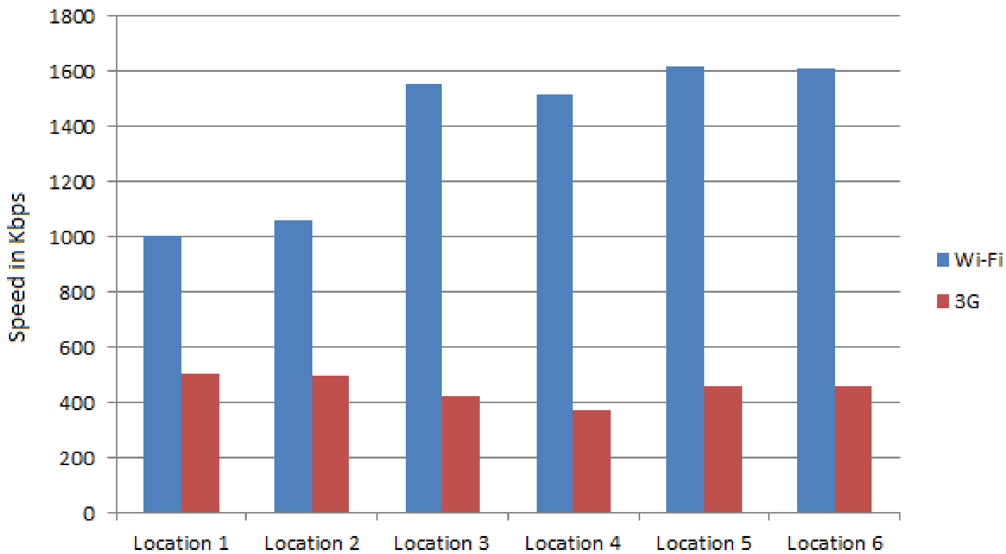

Figure 11: 3G vs Wi-Fi speed in downtown Fayetteville at peak hours

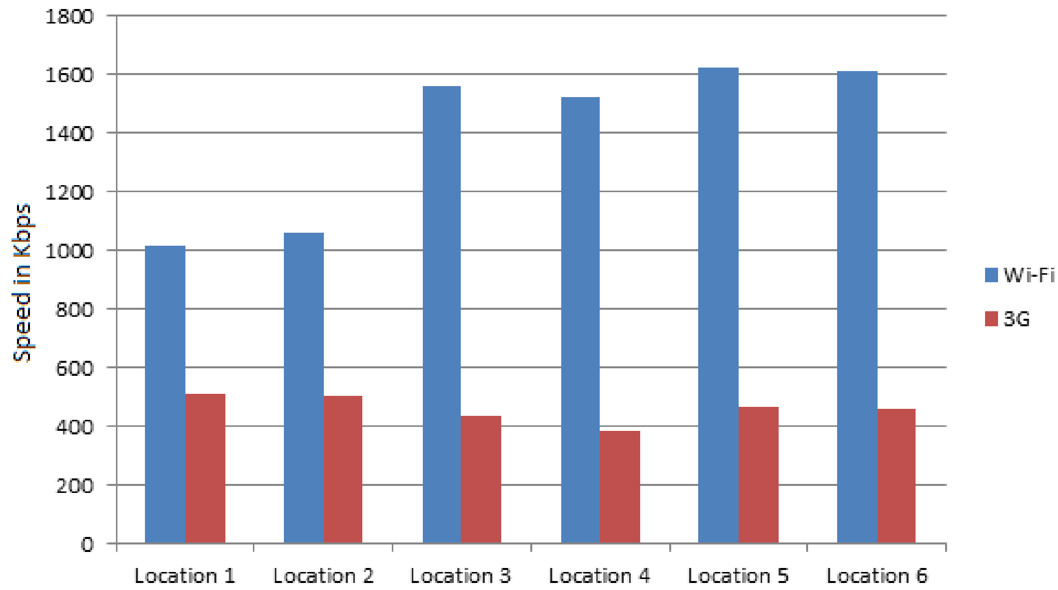

Figure 12 3G vs Wi-Fi speed at downtown Fayetteville at off-peak hours

Figure 11 and Figure 12 compare the speeds of Wi-Fi and 3G downloads at various locations in downtown Fayetteville recorded for every 2 hours when the user was walking. The downtown could be considered as a place where the user spends time in a cafeteria --- popular locations for users to spend time working. Again we find that there is variation in the speeds of Wi-Fi and 3G --- moreover, the variance is larger for Wi-Fi as compared to 3G. This points to the known fact that long range radios are immune to multipath effects and hence the instantaneous bandwidth can show larger variance in Wi-Fi compared to 3G.

We further compare the speeds of Wi-Fi and 3G recorded for every two hours at residential areas during the peak and off-peak hours on foot.

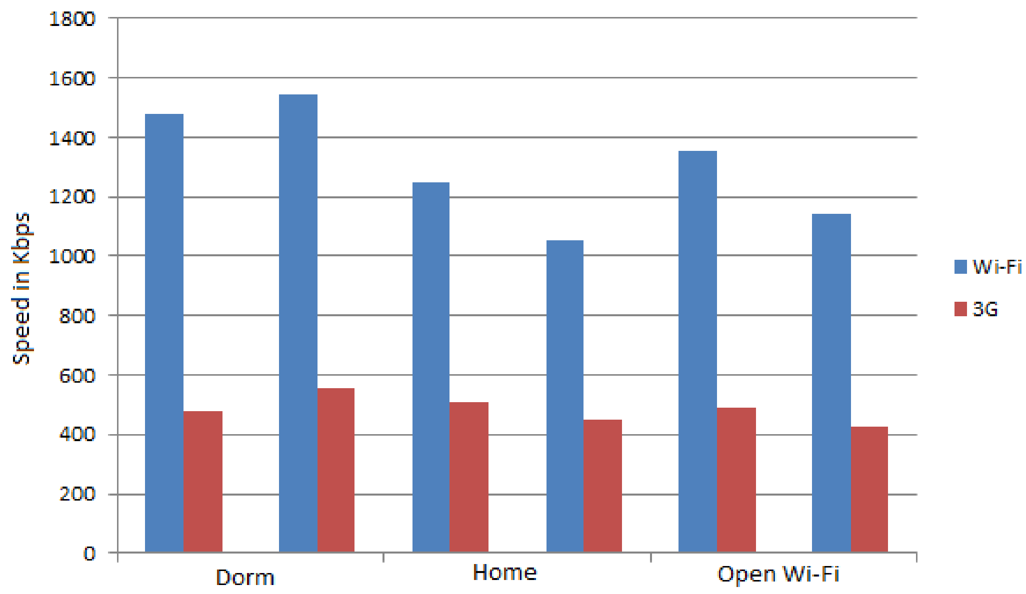

Figure 13 Wi-Fi vs 3G speed at Residential area of Fayetteville at peak hours

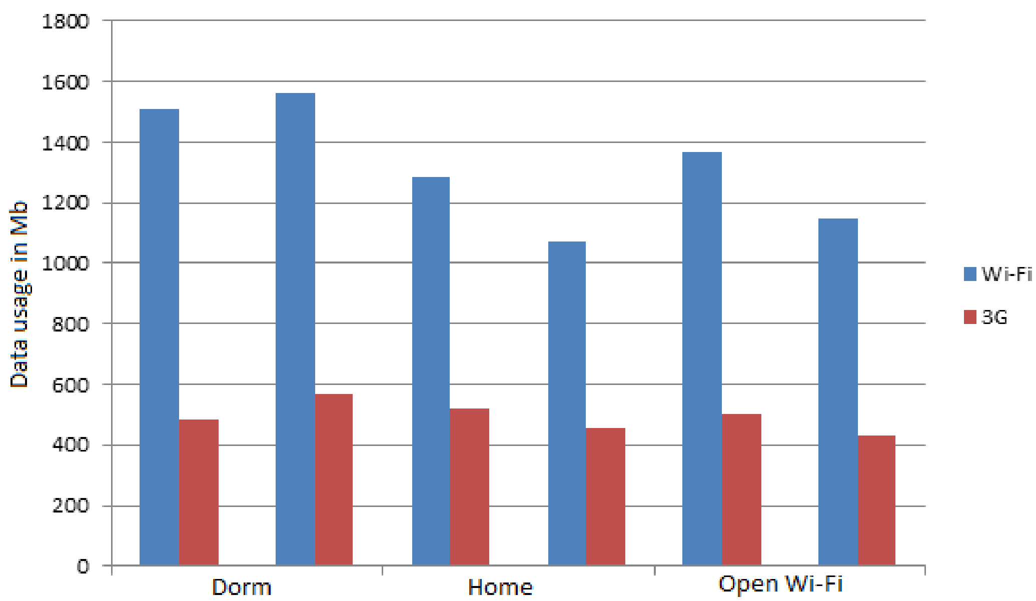

Figure 14 Wi-Fi vs 3G speed at Residential area of Fayetteville at off-peak hours

Figure 13 and Figure 14 clearly show that Wi-Fi speeds are way more than 3G although Wi-Fi availability is meager. While we performed experiments on 3G and Wi-Fi connectivity in downtown Chicago, it was difficult to stay connected to Wi-Fi at vehicular speeds.

Can 3G be augmented with Wi-Fi?

The goal of this section is to determine what percentage of 3G can be substituted with Wi-Fi at a given location. Experiments were conducted on foot for four days. In downtown Chicago, the user used Wi-Fi 9% of the time (Figure 15). And, 34.1Mb of data was transmitted through Wi-Fi and 266.4 Mb was transmitted through 3G. Given speeds of Wi-Fi and 3G, if we assume that 3G is used 100% of the time, from the above results we observe that at least 11% of the total data (300.5 Mb) on 3G can be offloaded to Wi-Fi. Note that this is a conservative estimate, since the user may be using Wi-Fi only on demand. If we have a system that does this offloading automatically, a larger amount of data can be transferred over Wi-Fi.

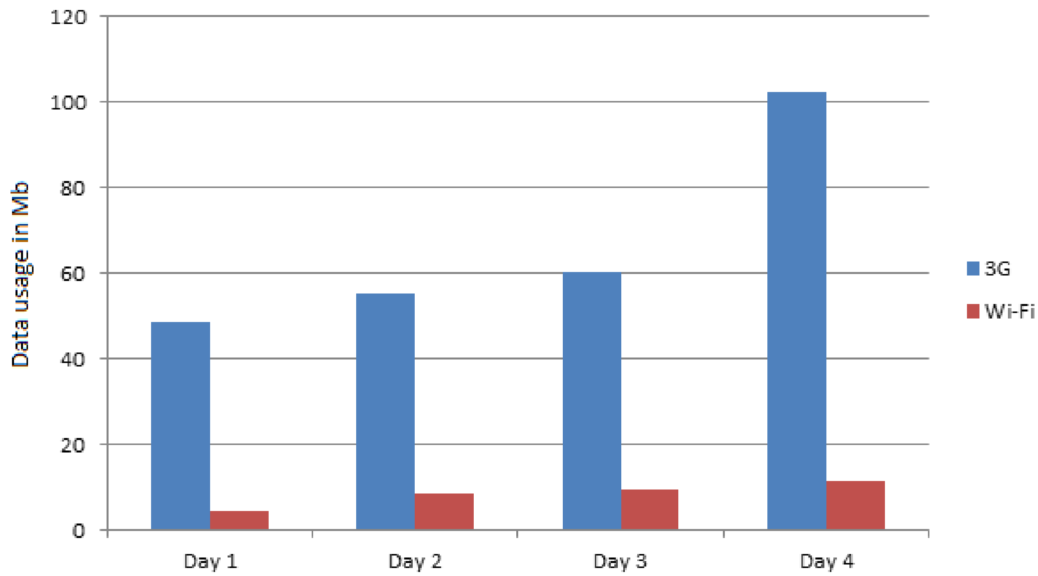

Figure 15: Data usage of 3G and Wi-Fi in Downtown Chicago

We also measured data transferred over Wi-Fi and 3G at various locations on foot in the downtown Fayetteville for four days. We found that the user used 3G for transferring data 96% of the time (Figure 16). However, given that about 4% of the time Wi-Fi was used, about 5.6Mb (5%) of the total data (109.355Mb) transfer can be offloaded to connections over the Wi-Fi radio.

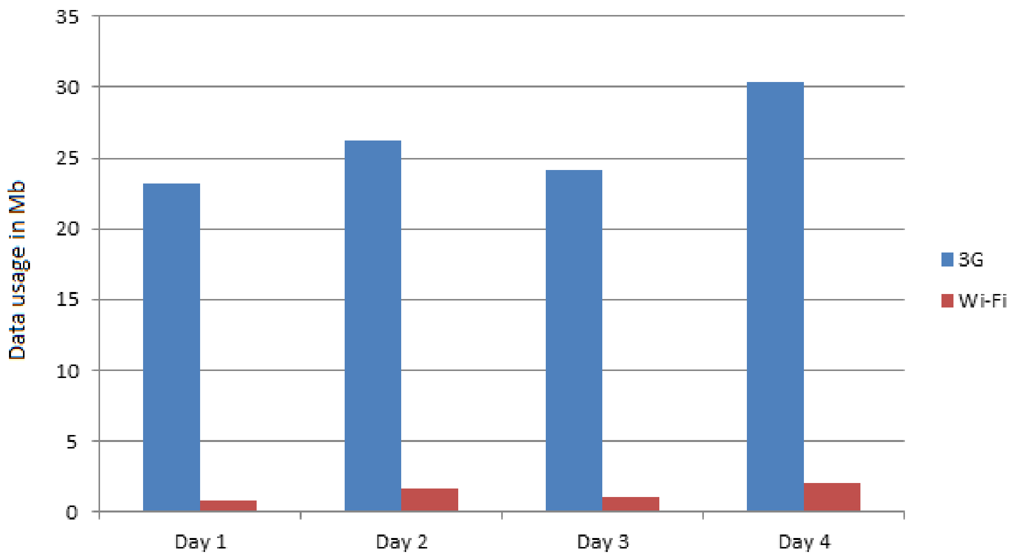

Figure 16 Data usage of 3G and Wi-Fi in downtown Fayetteville

Challenges with augmenting 3G with Wi-Fi

While it might be practically feasible to augment 3G with Wi-Fi in several places, there are several challenges that need to be encountered. When the user is mobile in semi-urban areas, he sees very few open access points. Most of the access points have access control restrictions or are secured. Hence, finding access points that are open and free is challenging.

The second challenge is to develop an architecture that can provide seamless transition from Wi-Fi to 3G. There are several options to build such architecture.

- Tightly couple integration, where a WLAN emulates functions of a 3G radio access network. It is treated as a 3G access network from the point view of the 3G core network

- In loosely coupled Integration, access point connects the gateway which in turn connects to the Internet. A mobile station contacts 3G towers via an access point. Mobile IP and authentication, authorization and accounting (AAA) services are implemented by the gateway to interact with the 3G's home servers. This service helps enable the exchange accounting information and billing information between a 3G network and a WLAN network.

- A hybrid coupled integration schema differentiates the data paths according to the type of the traffic and can accommodate traffic from WLAN efficiently with guaranteed seamless mobility.

Chapter 6: Conclusion

This thesis demonstrates the feasibility of augmenting 3G networks with Wi-Fi. Specifically, we show through experimentation and measurements that there are several locations where 3G and Wi-Fi connectivity co-exist. Moreover, the speed of Wi-Fi is twice that of 3G at both peak and non-peak hours. Hence, on mobile systems, there are several opportunities where traffic from 3G can be shifted over to Wi-Fi. Specifically, we show that 11% of data in Chicago and 5% of the data in Fayetteville can be offloaded to Wi-Fi. This could act as a solution to ever increasing congestion on 3G networks and bandwidth caps imposed by cellular providers like AT&T and Verizon. Additionally, we present the design of a general purpose tool that can be used for collecting usability statistics on Android devices. The tool can be used for several measurement studies related to the Android phones in the future.

However, there are several challenges that need to be solved before such augmentation becomes a reality. First, Wi-Fi is inherently unreliable. While there are 21,000 public hotspots in the US alone, Wi-Fi is a short-range radio and is prone to losses due to reflection, multipath fading, and obstacles. Hence, key to augmenting 3G with Wi-Fi lies is making Wi-Fi more robust.

.

VIII. References

[1] Andres Sevtsuk, Carlo Ratti, "*iSPOTS. How Wireless Technology is Changing Life on the MIT Campus*", CUPUM, 2005.

[2] Theodore Hillestad and Dr. Rajit Gadh "*An Overview of Wi-Fi "Hotspot" Pricing Models*", UCLA-WINMEC, 2003.

[3] Dan Saugstrup and Anders Henten "*3G Standards: the battle between WCDMA and CDMA2000*", Emerald Group Publishing Limited, 2006.

[4] Justin Manweiler and Romit Roy Choudhury "*Avoiding the Rush Hours: Wi-Fi Energy Management via Traffic Isolation*", MobiSys 2011.

[5] Vinoth Gunasekaran and Fotios C. Harmantzis "*Towards a Wi-Fi Ecosystem: Technology Integration and Emerging Service Models*", Telecommunications Policy, 2008.

[6] Keon Jang, Mongnam Han, Soohyun Cho "*3G and 3.5G Wireless Network Performance Measured from Moving Cars and High-Speed Trains*", MICNET 2009.

[7] Anmol Sheth Srinivasan Seshan and David Wetherall "*Geo-fencing: Confining Wi-Fi Coverage to Physical Boundaries*", Pervasive 2009.

[8] G. Anastasi, M. Conti, E. Gregori, A. Passarella "*802.11 Power-Saving Mode for Mobile Computing in Wi-Fi hotspots: Limitations, Enhancements and Open Issues*", Journal Wireless Networks archive, 2008

[9] Papadopoulos Homer "*Mobile data services usage - a methodological research approach*", Communications of the IBIMA, 2008.

[10] Mads Bødker, Greg Gimpel, Jonas Hedman "*The User Experience of Smart Phones: A Consumption Values Approach*", Eighth Global Mobility Roundtable, 2009.

[11] Long-Term Evolution (LTE): The vision beyond 3G White paper by Nortel.

[12] Yang Xiao, Kin K. Leung, Yi Pan, Xiaojiang Du "*Architecture, Mobility Management, and Quality of Service for Integrated 3G and WLAN Networks*". Wireless Communications and Mobile Computing, 2005.

[13] Aaron Carroll and Gernot Heiser "*An Analysis of Power Consumption in a Smartphone*", USENIX, 2010.

[14] William Lehr,Lee W. McKnight, "*Wireless Internet Access: 3G vs. WiFi?*", Technical report, MIT, 2002

[16] Hossein Falaki, Dimitrios Lymberopolous, Ratul Mahajan, Srikanth Kandula, Deborah Estrin, "*A First Look at Traffic on Smartphones*", IMC, 2010.

[17] Hossein Falaki, Ratul Mahajan, Srikanth Kandula, Dimitrios Lymberopolous, Ramesh Govindan, Deborah Estrin, "*Diversity in Smartphone Usage*", MobiSys, 2010

[18] Aruna Balasubramanian, Ratul Mahajan, Arun Venkataramani, "*Augmenting Mobile 3G Using WiFi*", MobiSys, 2010.

[19] Ratul Mahajan, John Zahorjan, and Brian Zill, "*Understanding WiFi-based Connectivity From Moving Vehicles*", IMC, 2007.

[20] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. "*Link-level measurements from an 802.11b mesh network*", SIGCOMM, 2004.

[21] Y. Cheng, J. Bellardo, P. Benko, A. C. Snoeren, G. M. Voelker, and S. Savage. "*Jigsaw: Solving the puzzle of enterprise 802.11 analysis*", SIGCOMM, 2006.

[22] R. Gass, J. Scott, and C. Diot. "*Measurements of in-motion 802.11 networking*", WMSCA workshop, 2006.

[23] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. "*Understanding congestion in IEEE 802.11b wireless networks*", IMC, 2005.

[24] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. "*Analyzing the MAC-level behavior of wireless networks in the wild*", SIGCOMM, 2006.

[25] J. Ott and D. Kutscher, "Drive-thru Internet: IEEE 802.11b for automobile users", INFOCOM, 2004.

[26] K. Papagiannaki, M. Yarvis, and W. S. Conner. "*Experimental characterization of home wireless networks and design implications*", INFOCOM, 2006.

[27] UMass DieselNet. http://prisms.cs.umass.edu/dome/

[28] City-wide Wi-Fi rolls out in UK. http://news.bbc.co.uk/2/hi/technology/4578114.stm

[29] Cities unleash free Wi-Fi. http://www.wired.com/gadgets/wireless/news/2005/10/68999.

[30] Aruna Balasubramanian, Ratul Mahajan, Arun Venkataramani, Brian Levine, and John Zahorjan, "*Interactive WiFi Connectivity for Moving Vehicles*", SIGCOMM, 2008.

[31] Srikanth Kandula, Kate Ching-Ju Lin, Tural Badirkhanli and Dina Katabi, "*FatVAP: Aggregating AP Backhaul Bandwidth*", NSDI, 2008

[32] Anthony J. Nicholson, Scott Wolchok and Brian D. Noble, "*Juggler: Virtual Networks for Fun and Profit*", IEEE Transactions on Mobile Computing, 2009.

[33] Jakob Eriksson, Hari Balakrishnan, Samuel Madden, "*Cabernet: Vehicular Content Delivery Using WiFi*", Mobicom, 2008.

[34] Yuvraj Agarwal, Ranveer Chandra, Alec Wolman, Paramvir Bahl, Kevin Chin, and Rajesh Gupta, "*Wireless wakeups revisited: energy management for voip over wi-fi smartphones*", Mobisys, 2007.

[35] H. Soroush, N. Banerjee, B. N. Levine, M. D. Corner, "*Patching Mobile WiFi Networks*", Technical Report, UMass-Amherst, 2009.