

Enabling Energy-Aware Collaborative Mobile Data Offloading for Smartphones

Aaron Yi Ding*, Bo Han†, Yu Xiao‡, Pan Hui§¶, Aravind Srinivasan||, Markku Kojo* and Sasu Tarkoma*

*Department of Computer Science, University of Helsinki, FI-00014, Helsinki, Finland

†Department of Computer Science, University of Maryland, College Park, MD 20742, USA

‡Department of Computer Science and Engineering, Aalto University, Espoo, Finland

§Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

¶Telekom Innovation Laboratories, Ernst-Reuter-Platz 7, 10587 Berlin, Germany

||Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, USA

Abstract—Searching for mobile data offloading solutions has been topical in recent years. In this paper, we present a *collaborative* WiFi-based mobile data offloading architecture - Metropolitan Advanced Delivery Network (MADNet), targeting at improving the energy efficiency for *smartphones*. According to our measurements, WiFi-based mobile data offloading for moving smartphones is challenging due to the limitation of WiFi antennas deployed on existing smartphones and the short contact duration with WiFi APs. Moreover, our study shows that the number of open-accessible WiFi APs is very limited for smartphones in metropolitan areas, which significantly affects the offloading opportunities for previous schemes that use only open APs. To address these problems, MADNet intelligently aggregates the collaborative power of cellular operators, WiFi service providers and end-users. We design an *energy-aware* algorithm for energy-constrained devices to assist the offloading decision. Our design enables smartphones to select the most energy efficient WiFi AP for offloading. The experimental evaluation of our prototype on smartphone (Nokia N900) demonstrates that we are able to achieve more than 80% energy saving. Our measurement results also show that MADNet can tolerate minor errors in localization, mobility prediction, and offloading capacity estimation.

I. INTRODUCTION

3G cellular networks are currently overloaded with data traffic generated by various bandwidth-hungry smartphone applications (e.g., mobile TV), especially in metropolitan areas [9]. Although several cellular operators have already upgraded their networks to LTE (4G) for higher capacity, the traffic demand from end-users also continues to increase. Mobile data offloading may relieve this problem by using complementary communication technologies (considering the increasing capacity of WiFi) to deliver traffic originally planned for transmission over cellular networks.

Although WiFi has been shown to be promising for mobile data offloading in Wiffler [2] and by Lee et al. [13], there are still several challenging issues when offloading mobile traffic for *smartphones*. First, Wiffler [2] is designed for PCs (e.g., netbooks) on vehicular networks without considering the offloading energy consumption. Lee et al. [13] evaluate the energy saving of offloading through a trace-driven simulation with several simplified assumptions, but how to harvest the energy gain of mobile data offloading *in practice* is still an open problem. Second, through our extensive war-driving

and war-walking measurements using smartphones in three cities of US and Europe, we found that the number of open-accessible WiFi access points (APs) is very limited. Therefore the schemes using only open APs as in Wiffler [2] may not be enough. Finally, the goal of previous work is to increase the amount of delay-tolerant data traffic that can be offloaded to WiFi networks. However, delay-tolerant applications generate only a small amount of data traffic, compared to streaming applications [16].

To address the above challenges, we propose MADNet, a collaborative mobile data offloading architecture for smartphones. The main design principle of MADNet is to extend smartphone battery life. According to our measurements, transferring the same amount of data may consume more energy on a low throughput WiFi network than transferring over a high speed 3G access. If a scheme only aims to increase the offloaded traffic to WiFi networks without considering the energy consumption on smartphones, it may drain the battery much faster than using only 3G networks. Furthermore, due to the limitation of WiFi antennas deployed on existing smartphones, offloading mobile data traffic for smartphones is much more challenging than that for PCs. To compensate such existing restrictions, a dedicated scheme is needed.

The above observations motivate the design of MADNet that harvests the collaboration power across cellular operators, WiFi providers, and end-users to achieve energy-aware mobile data offloading for smartphones. In this paper, we make the following contributions:

- We investigate and compare the performance of WiFi Internet access for both netbooks and smartphones in metropolitan areas. Our observations offer valuable insights for future design and inspire us to exploit the predictable nature of streaming content via data prefetching.
- We design an efficient and deployable *energy-aware* algorithm to assist the offloading decision. Our algorithm, integrated with the collaborative MADNet architecture, is able to tolerate minor errors of input parameters such as user location, mobility prediction, and estimation of system offloading capacity. We also utilize content prefetching to compensate the impact of hardware limitations on smartphones.

- Our prototype implementation confirms the feasibility of MADNet to be deployable in existing environments. By enabling smartphones to select the most energy efficient WiFi AP for offloading, our field experiments further verify that MADNet can achieve notable 80% energy saving.

The rest of the paper is organized as follows. Section II provides a thorough characterization study of WiFi Internet access in metropolitan areas. We present the MADNet architecture in Section III and evaluate our implementation in Section IV. Section V covers several challenging issues. We discuss the related work in Section VI and conclude in Section VII.

II. METROPOLITAN WiFi NETWORKS

As WiFi-based mobile data offloading is the major focus of MADNet, the goal of our measurement study is to answer the following two questions: (1) *Is it feasible to offload mobile data traffic to WiFi networks, given the current deployment and availability of WiFi APs?* (2) *Is the network performance of WiFi-based Internet access for moving smartphones good enough to offload mobile traffic?*

A. Accessible WiFi APs for Smartphones

We performed field studies mainly in three cities, Berlin in Germany, Chicago and Baltimore in the US, using a tool called 3G-WiFi. It consists of two threads: the first one measures TCP uplink and downlink throughput and latency to a reference server over 3G networks; the second one scans neighboring WiFi APs and then performs the same measurements as over the 3G networks for all open-accessible APs. 3G-WiFi conducts these tests over 3G and WiFi networks periodically with 10-second intervals. We used Nokia N900 smartphones for most of the measurements and experiments in this paper. N900 default OS, Maemo 5, is an open source Linux distribution (2.6.28 kernel). Its WiFi chipset is Texas Instruments WL1251. The full experimental results are available in a technical report [10].

We chose four areas in Berlin for war-driving: two neighborhoods near Nokia Siemens Networks and Schloss Charlottenburg, and along Kurfürstendamm avenue and Unter den Linden boulevard (two of the most popular avenues in Berlin). In Chicago, the war-walking was around Michigan Avenue. We also conducted war-driving in the downtown area of Baltimore. We summarize the statistics of detected APs in these three cities in Table I. *Accessible APs for offloading* are those that allow us to test at least one of the above three metrics, TCP uplink or downlink throughput, or latency. Nicholson et al. [17] reported that about 40% of APs detected in three neighborhoods of Chicago were open in 2006 and Wiffler [2] also found that around 40% APs encountered by 20 public transit vehicles in Amherst, MA were not encrypted. In contrast with these results, the percentage of APs without encryption (i.e., open APs) that we detected is very low, less than 10% in Berlin and less than 20% in Chicago and Baltimore. One of the possible reasons may be that more and

TABLE I
STATISTICS OF DETECTED APs IN BERLIN, CHICAGO AND BALTIMORE.

City Speed	Berlin driving	Chicago walking	Baltimore driving
Detected APs	4421	4588	3418
APs without encryption	351 (7.9%)	775 (16.9%)	621 (18.2%)
APs granting IP addresses	12 (0.27%)	18 (0.39%)	6 (0.18%)
Accessible APs for offloading	0 (0.0%)	7 (0.15%)	1 (0.03%)

more people care about the security of their home networks, especially in big cities.¹

Most of these ‘open’ APs are *not* accessible, as they may apply MAC address filter or web-based authentication for access control. The percentage of accessible APs is extremely low, smaller than 1%, which makes the offloading solutions using only open APs challenging. To detect and measure as many APs as possible, we set the timeout of DHCP messages and TCP connection setup to be 0.1 seconds. These small values may affect the measurement results [8]. Thus, we also conducted war-driving in a neighborhood of College Park, MD with three different timers (1, 2 and 3s) for DHCP discovery messages and TCP connection setup. The driving speed was ~ 20 – 30 km/h. The results (available in our technical report [10]) show that even in a small college town, the percentage of open APs is less than 20%. Moreover, there is almost no open-accessible APs in that neighborhood. Therefore, *active participation from end-users (by sharing their home APs) is a key enabling factor for mobile data offloading.*²

Through the identifiable patterns of ESSID (Extended Service Set Identifier) from the collected trace of AP information in Berlin, we found that there are a large number of APs (~ 1000) deployed/distributed by cellular operators (see details in [10]). We also identified 616 WiFi APs produced by 2Wire, Inc. and distributed by AT&T to home users from the Chicago trace, and 334 WiFi APs produced by Westell or Abocom and distributed by Verizon from the Baltimore trace. These numbers are lower bounds because users may change the ESSID after they get the APs from their providers.

B. WiFi-based Internet Access for Moving Smartphones

We conducted experiments at both walking and driving speeds to investigate the performance of WiFi-based Internet access for moving smartphones.

1) *Experimental Setup:* We measured the amount of transferred data, duration of TCP connections and average TCP throughput from two servers, a remote server and a local server running on a laptop, to our moving smartphone. As shown in Figure 1, the laptop was connected to a 100M Ethernet

¹We confirmed that in the US most of the home APs distributed by AT&T U-verse and Verizon FiOS are encrypted by default (with their technical support of customer services). The same is true for Deutsche Telekom T-Home in Germany.

²End-users may also be WiFi providers in MADNet, and we discuss the security issues of WiFi sharing in Section V.




Fig. 2. Map for the walking and driving experiments. The locations of WiFi APs are marked.




Fig. 1. The experimental setup.

port of the AP, which was connected to a campus network. The smartphone was associated with the AP through its WiFi interface. The remote server was located in an industrial research lab in Europe. The experimental setup was similar to that in Hadaller et al. [8]. The major differences are that we used a smartphone as the client, instead of a laptop, and experimented with both walking and driving speeds. We did all the experiments in an apartment neighborhood in College Park, MD and show the map for these experiments in Figure 2.

2) *Moving Smartphones at Walking Speed:* We conducted a group of outdoor experiments for moving smartphones at walking speed, passing by the home AP. The AP is around 36 meters away from the road. To guarantee that the duration of experiment is longer than the actual AP association time of the smartphone, we walked from one location to another along the road for the selected trail. Both locations are out of the coverage area of the AP. We performed experiments for both directions along the road, to the east and west. We measured three settings for this scenario: remote server with DHCP, local server with DHCP, and local server with static IP address configuration. For each setting, we repeated the experiments 10 times.

We plot the mean TCP downlink throughput and standard deviation in Figure 3. As we can see from the figure, we can improve TCP throughput by at least 200% when connecting to the local server instead of the remote server, because the capacity of a wireless link between smartphones and WiFi APs is usually higher than the available capacity between WiFi APs and the Internet [7]. Although similar observations have been made by recent studies of vehicular Internet access [6], *another benefit of separating the wireless links from wired ones is that we may reduce data transmission time, and thus save energy consumption on smartphones.*




Fig. 3. The mean downlink TCP throughput of a smartphone moving at walking speed.




Fig. 4. The mean TCP throughput for a smartphone and a netbook at driving speed.

The differences in throughput when walking towards different directions may be caused by the location of home AP and the antenna direction of smartphone, as it was held in hands during the experiments. Thus, the duration of WiFi connectivity is ~ 75 seconds when walking towards the east and ~ 25 seconds longer (100 seconds) to the west. As we conducted all the experiments in the wild, there are several other co-channel APs in the experimental area and other users may use them for extensive data transfer during the experiments.

3) *Moving Smartphones at Driving Speed:* Although moving speeds may not affect 3G throughput too much, they do impact the performance of WiFi-based vehicular Internet access. Slow vehicle speed usually leads to high WiFi throughput [5]. We compare the performance of smartphones and netbooks for vehicular Internet access. The experimental setup shown in Figure 2 is the same as the setting of local server with static IP address for the walking scenario. We did not use netbooks for the walking experiments because we do not consider it as a typical usage scenario of netbooks. As shown in Figure 4, netbooks perform much better than smartphones, achieving higher throughput (1.8 vs. 1.0 Mbps).

Vehicular Internet access and the mobility problems of WiFi have been widely investigated [2], [4], [8]. We summarize the devices and antennas in use in our technical report [10]. With no exception, external antennas were used for all of PCs. It is well-known that antenna plays an important role in WiFi-based Internet access. For example, Deshpande et al. [5] reported that a 12 dBi antenna provides better connectivity than 5 and 7 dBi antennas. However, due to the limited size of smartphones, they may not be able to use external antennas. Thus, mobile data offloading for smartphones on vehicles is much more challenging than that for PCs and that for those with pedestrians.

4) *Energy Concern for Smartphones:* To investigate the energy consumption, we further measured the energy consumption of ~ 20 MB data transfer from a server to Nokia N900 and Samsung Nexus S smartphones through 3G and WiFi networks, using iperf TCP. This experiment is conducted in Finland through another 3G network with higher throughput. We increased the distance between the smartphones and WiFi AP to achieve low throughput WiFi data transfer. As we can see from the results in Table II, when WiFi throughput is lower than 3G, the data transfer through WiFi networks consumes more energy than that of 3G. *These results verify*

TABLE II
MEASURED ENERGY CONSUMPTION OF 20 MB DATA TRANSFER.

	N900		Nexus S	
	Energy Joule	Throughput Mbps	Energy Joule	Throughput Mbps
3G	109.4	1.89	65.40	1.99
WiFi	116.0	0.422	191.7	0.302

that the solutions that utilize every offloading opportunity without considering the energy consumption may reduce the battery life of smartphones.

C. Insights

Our measurements offer the following insights:

(1) For WiFi-based mobile data offloading, the number of open-accessible APs is very low, verifying the trend that less and less WiFi APs are open. Therefore, we propose *collaboration among cellular operators, WiFi service providers and end-users to increase the offloading opportunities*.

(2) The performance of WiFi-based Internet access for moving smartphones is worse when connecting to remote servers than connecting to a local server, and is also worse than that of PCs (e.g., netbooks). To address this challenge, we propose to *prefetch predictable data traffic at WiFi APs*, which can increase the amount of offloaded mobile traffic and save energy consumption on smartphones.

For the low throughput presented at driving speed, we choose to evaluate our proposal for smartphones only at walking speed (in Section IV). We note that although similar studies have been conducted earlier [2], [6], [15], the focus is different. Our findings provide valuable insights in this domain and lay the foundation for the proposed collaborative MADNet architecture.

III. MADNET ARCHITECTURE

A. Overview

We advocate that mobile data offloading is win-win-win for cellular operators, WiFi service providers, and end-users, and that they should cooperate to make it feasible and effective. With offloading solutions, cellular operators may be able to meet the rising traffic demand without significantly increasing their CAPEX (capital expenditure) and OPEX (operating expenses). Moreover, they can provide mobile data offloading as a value-added service, and thus expand their user base. The offloading schemes may also bring third-party WiFi service providers more customers without service contracts, and thus extra revenue. Finally, end-users can benefit from mobile data offloading through higher data rate and longer battery life for their smartphones.

We aim to design an energy efficient approach to offload real-time streaming traffic for smartphones that are *in motion*. When people walk, although they usually do not read news or watch video, they often listen to music. Also, more and more people watch video using smartphones when riding on metro buses. According to Cisco Visual Network Index³, most

³http://www.cisco.com/en/US/netsol/ns827/networking_solutions_sub_solution.html (verified in July 2012)




Fig. 5. Communications between the three MADNet components.

of the mobile data traffic in 2015 will be either music or video streaming. Although streaming traffic cannot tolerate much delay, it is feasible to predict the delivery of *non-live streaming* traffic. When mobile users listen to music online, they usually make a playlist to include all the songs they like. With the playlist, we can know the expected streaming traffic for the next several minutes. The same is true for video streaming. Once a mobile user starts to watch video online, the future streaming traffic is predictable. Therefore in MADNet, smartphones may start data streaming immediately over 3G networks and meanwhile try to find WiFi APs that can prefetch data and deliver data with lower energy cost.

B. Collaborative Mobile Data Offloading

MADNet has the following three key design goals:

- Reducing the *energy consumption* of smartphones, which is realized by the energy-aware offloading decisions;
- Offloading the mobile data traffic for smartphones, which is enabled by the *collaborative* architecture;
- Making the offloading process *transparent* to end-users.

We show in Figure 5 the three major software components of MADNet: a client module on smartphones (Smart-Client), a proxy on WiFi APs (Proxy-WiFi) and another proxy in the cellular access networks (Proxy-Cellular). To coordinate the split of data traffic between these two networks, we set up two control channels, as shown in Figure 5.

MADNet makes offloading decisions based on the contextual information of smartphones. First, Proxy-Cellular knows the locations of the neighboring WiFi APs of a given smartphone, because these locations are fixed (and known through services such as wgle.net) and some of them are deployed by cellular operators. Second, as people are creatures of habit and usually take similar paths everyday, it is feasible to predict their mobility patterns using history information [18].

When a mobile user wants to request content from the Internet, Smart-Client will send the geographical location and moving speed and direction to Proxy-Cellular. Smart-Client can easily retrieve this information from the wireless interface and various sensors (e.g., WiFi, GPS, and accelerometer). To follow the common practice of “Wake on Wireless” [24] and thus reduce energy consumption of smartphones, Smart-Client turns on the WiFi interface and the sensors only when users issue content requests.

Using the proposed offloading decision algorithm, Proxy-Cellular determines whether the offloading of

Algorithm 1 Energy-Aware Offloading Decision Algorithm

Require: The power of data reception P_{3G} for 3G and P_W for WiFi.

Require: The head and tail energy E_T of 3G and E_{oo} for the offloading related overhead.

- 1: Predict the throughput B_{3G} for 3G network and estimate the offloading capacity C_W and throughput B_W of a WiFi AP.
- 2: Predict the prefetching capability F of this WiFi AP.
- 3: Calculate the WiFi offloading duration C_W/B_W and the time to receive the same amount of data through 3G network C_W/B_{3G} .
- 4: **if** $F \geq C_W$ and the following inequality holds

$$E_T + P_{3G} \cdot C_W/B_{3G} - P_W \cdot C_W/B_W > k \cdot E_{oo} \quad (1)$$

then

- 5: Offload mobile data traffic to this WiFi AP.
 - 6: **else**
 - 7: Repeat the above for other available APs.
 - 8: **end if**
-

cellular traffic to a given WiFi AP can potentially save energy on smartphones and notifies `Proxy-WiFi` to prefetch the content if possible. With the context information from `Smart-Client` as input, the heavy computational tasks, including positioning and mobility prediction are hence offloaded to `Proxy-Cellular`. After receiving information from `Proxy-Cellular` about the WiFi AP (e.g., MAC address, ESSID and time to associate) for offloading, `Smart-Client` can initialize the WiFi association and download prefetched data accordingly.

In this paper, we focus on the case that each `Proxy-WiFi` serves one `Smart-Client` at a time to investigate the effectiveness of our proposal. We note that we can easily extend the MADNet architecture to support multiple devices simultaneously by integrating efficient scheduling algorithms.

C. Energy-Aware Offloading Decision

We propose an energy-aware mobile data offloading decision algorithm for smartphones in Algorithm 1. Although WiFi is generally more energy efficient than 3G [3], offloading mobile data traffic to WiFi networks may cause extra energy consumption to get location information and to associate with the WiFi APs that are predicted to be available.

MADNet performs WiFi-based mobile data offloading only when the receiving of prefetched data from WiFi APs (instead of streaming data through 3G networks) saves more energy than the extra energy consumption overhead discussed above. We describe this requirement rigorously in inequality (1) of Algorithm 1, where k is a parameter to accommodate measurement errors. For small value of k , the estimation errors may cause more energy consumption on smartphones due to offloading. On the other hand, we may lose some offloading opportunities if k is too large. We set k to be 1.1 tentatively for the experiment.

The energy-aware offloading decision is affected by the throughput of 3G and WiFi networks. For the measurement study of Wiffler [2] in Amherst, the downlink median TCP throughput is 600 Kbps for 3G and 280 Kbps for WiFi. In this case, although offloading 3G traffic to WiFi networks can reduce 3G usage, it may cause more energy consumption on smartphones, as the duration of WiFi data transmission doubles that of 3G. In another measurement study by Deshpande et al. [5], WiFi offers substantially higher median throughput than 3G, ~ 2000 Kbps vs. ~ 500 Kbps,⁴ respectively. For this scenario, WiFi-based mobile data offloading may potentially reduce the energy consumption of smartphones.

We need to know the predicted throughput of 3G networks and the WiFi offloading capacity (i.e., the number of bits we can offload to WiFi networks) when calculating the above energy saving. Through an eight-month measurement study, Yao et al. [26] show strong correlation between cellular throughput and location for 3G HSDPA networks, and thus it is feasible to predict 3G throughput using history and location information. As pointed out in Wiffler [2], we can also estimate the offloading capacity of WiFi networks using existing work like BreadCrumbs [18]. To accommodate the estimation errors, we use a lower value (e.g., the 30th percentile) for WiFi throughput and a higher value (e.g., the 70th percentile) for 3G throughput, instead of the median. The maximum prefetching capability F is defined as the product of the Internet backhaul throughput for the WiFi AP and the prefetching duration (from the notification time of prefetching to the expected dissociation time of `Smart-Client` from the WiFi AP). To better utilize the capacity of a WiFi link, the result of the offloading decision will be negative if this prefetching capability is smaller than the estimated offloading capacity.

To calculate the saved energy, we assume that 3G and WiFi interfaces keep at the same data reception power level during the data transfer (C_W/B_{3G} and C_W/B_W), which is conservative and gives the lower bound of actual energy saving. This assumption is reasonable for 3G interfaces. Due to the radio resource control of cellular networks, after transmitting or receiving a packet the 3G radio stays at high power state and drops to low power only when the interface has been inactive for several seconds [3]. This state transition also introduces significant head (from low power to high at the beginning) and tail (from high power to low at the end) energy, which is considered as E_T in (1) of Algorithm 1.

In a nutshell, MADNet avoids the offloading of mobile data traffic to low throughput WiFi networks and enables smartphones to select the most energy efficient WiFi AP if possible. It adopts data prefetching at WiFi APs which can further improve the utilization of WiFi channels and reduce energy consumption on smartphones. Note that the performance of streaming applications is not affected by these offloading decisions, because when offloading is not possible or feasible (e.g., caused by incorrect mobility prediction) MADNet relies

⁴The contradictory results reported in the above studies may be caused by the fact that Wiffler [2] used an 802.11b radio for all the experiments, which limits the maximum PHY bit rate to be 11 Mbps.

completely on 3G networks for data transfer.

IV. IMPLEMENTATION AND EVALUATION

We implement a prototype to explore the gains of mobile data offloading for smartphones, and evaluate its performance in a live environment with smartphones *in motion*.

A. MADNet Implementation

1) *Proxy-WiFi*: We implement *Proxy-WiFi* as a service daemon on Linux platform. Once *Proxy-WiFi* starts it will connect to *Proxy-Cellular* and wait for instructions of data prefetching for the downlink traffic. After *Proxy-WiFi* has prefetched data contents on behalf of mobile users, it will save them at a temporal buffer and feed them to mobile users. The amount of data to prefetch is determined by the estimated WiFi offloading capacity. As we focus on real-time streaming applications, we implement *Proxy-WiFi* for downlink traffic in the current prototype.

2) *Proxy-Cellular*: *Proxy-Cellular*'s major functions are to estimate the location and predict the mobility of end-users, make offloading decisions, forward content requests to corresponding APs and coordinate the authentication between smartphones and APs when necessary. We implement a WiFi fingerprint based localization scheme, similar to the standard RADAR approach [1]. We also implement a modified version of the KNT mobility prediction algorithm [23] by considering the regularity of human mobility. The output of the mobility prediction algorithm is a group of WiFi APs that are predicted to be visited by *Smart-Client*. We run the offloading decision algorithm described in Algorithm 1 to select the WiFi APs that can achieve the maximum energy reduction. If the predicted WiFi-AP set is empty or none of the WiFi APs can satisfy the requirements specified in Algorithm 1, *Proxy-Cellular* will notify *Smart-Client* to keep using the 3G network. We implement *Proxy-Cellular* on several MADNet servers, since currently it is hard to access and run our code on the base stations of operational cellular networks.

3) *Smart-Client*: We implement a music streaming application with *Smart-Client*, called *MStreamer*. Once a playlist is given by a mobile user, *MStreamer* will start to stream the first music over 3G networks. At the same time, *Smart-Client* sends the context information to *Proxy-Cellular* to make the offloading decision. By design, both WiFi and GPS can be used to obtain positioning context. To save energy *Smart-Client* prioritizes WiFi over GPS and will use GPS only if WiFi positioning fails. If *Proxy-Cellular* chooses to offload streaming traffic to WiFi networks, it will notify *Smart-Client* to submit the whole playlist and then forward the prefetching instructions to *Proxy-WiFi* running on the WiFi AP that the smartphone will visit soon. Once the smartphone enters the coverage area of this WiFi AP (determined by a timer set by *Proxy-Cellular* based on mobility prediction), it starts to download the prefetched songs and save them into a temporal

buffer. The *MStreamer* can then play the next music directly if it is in the buffer and avoid the steaming over 3G networks.

Although the cooperative architecture has several requirements on WiFi APs (e.g., storage space to prefetch data), we note that it is feasible on modern APs which are programmable and have USB ports to attach storage devices.⁵ Moreover, smartphones usually have several GB storage space (32 GB for Nokia N900), which is enough to buffer prefetched data.

B. Performance Evaluation

Our evaluation consists of two parts: 1) we first evaluate the impact of mobility prediction and data prefetching on the performance of MADNet; 2) we then estimate the energy savings of mobile data offloading for music streaming applications using the measurement data obtained from static devices, since it is hard to measure directly the energy consumption for moving smartphones. We did the experiments along the Limingantie street in Helsinki, Finland. The experimental setup is similar to that shown in Figure 2 and we use a roadside AP to evaluate the outdoor scenario.

We emphasize that MADNet can avoid handoffs between different WiFi APs, thanks to the coordination using the control channel via the 3G network. Thus, performance evaluation with a single AP should be enough for our purpose. Moreover, benefiting from the cooperative architecture of MADNet, there is no fundamental difference between our WiFi AP under test and others, such as commercial APs and encrypted home APs.

1) *Difference Between Prefetching and Downloading*:

To evaluate the potential energy savings, we measure the amount of prefetched data (determined by the estimated WiFi offloading capacity) and the actual amount of data that *Smart-Client* can download during the association with the WiFi AP (i.e., the actual offloading capacity). Since we calculate the estimated offloading capacity by considering the location of users, their future mobility, and history information about WiFi throughput, our evaluation takes into account the effectiveness of WiFi-fingerprinting localization, mobility prediction and prefetching. We run the experiments 10 times for each direction, to the east and the west, and plot the results in Figure 6 and 7. If *Smart-Client* finishes the downloading of prefetched data when it is still connected with the WiFi AP, it downloads some dummy files to measure the actual offloading capacity. As we can see from these figures, for around 85% of the runs, we can fully utilize the WiFi link (i.e., the amount of downloaded data is smaller than that of prefetched data with average gap around 4.8 MB).

2) *Locations to Start Association*: We also measure the location where *Smart-Client* starts the WiFi association and plot the results of 10 runs for each direction in Figure 8. Our field study shows that the coverage range of the deployed WiFi AP is about 30 meters. On average, when walking to the east, *Smart-Client* starts to associate with the AP when it was around 33.3 meters away from the AP. The average

⁵For example, (NaDa) [25] has been proposed to leverage the computing and storage capabilities on ISP-controlled home gateways to reduce energy consumption of Internet-based video streaming services.

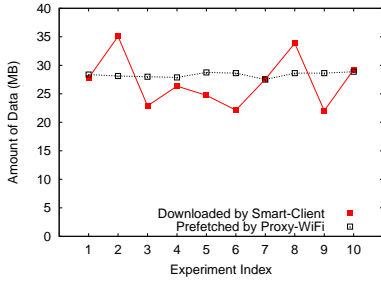


Fig. 6. The amount of prefetched and downloaded data when walking to the east.

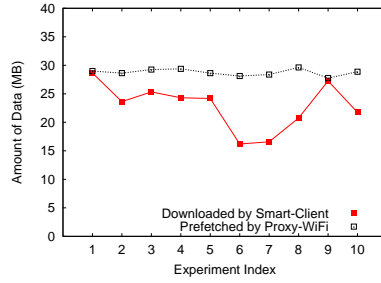


Fig. 7. The amount of prefetched and downloaded data when walking to the west.

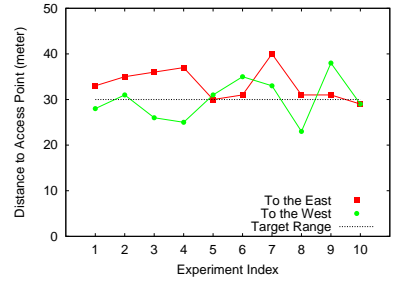


Fig. 8. Distance between smartphone and WiFi AP when Smart-Client starts WiFi association.

distance was 29.9 meters when walking to the west. The above results demonstrate that WiFi fingerprinting localization and existing mobility prediction schemes are good enough to support WiFi-based mobile data offloading.

3) *Energy Savings*: As we focus on the offloading of downlink cellular traffic in this paper, we summarize our measurements for power of 3G and WiFi downlink data transmission with TCP, P_{3G} and P_W , and extra energy consumption of WiFi offloading in Table III, using Monsoon Power Monitor (<http://www.msoon.com/LabEquipment/PowerMonitor/>) with 5 KHz sampling rate. We repeated each measurement 10 times and report the average results and standard deviations. The N900 data is used to estimate the energy saving of our proposal.

The offloading related overhead E_{oo} in Algorithm 1 includes E_{W_on} and E_{W_off} to turn on and off WiFi radio, E_{scan} for WiFi localization based on radio beacons (e.g., RADAR [1]), and E_{asso} of association with the WiFi AP specified by Proxy-Cellular. As we use static IP address configuration, instead of DHCP, we can save the energy consumption of acquiring IP address, which is ~ 4.8 Joules for Nokia N900. We emphasize that another benefit of using 3G as a control channel is that smartphones can scan only the channel of the specified WiFi AP and avoid complete scanning of all possible channels. For N900, the head and tail energy E_T of 3G is ~ 5.4 Joules, which is used in our evaluation.

E_{GPS} shows the energy consumption of AGNSS (Assisted Global Navigation Satellite System) location method with assistance data from external location servers. It takes 14 seconds on average for a cold-start GPS to get the first accurate fix using AGNSS. For the location method using only internal GPS (i.e., GNSS), the cold-start duration is around 20 seconds with ~ 6.30 Joules energy consumption. Thus, we prefer the localization schemes that leverage WiFi fingerprint to GPS.

We also present the results for Nokia E71 (measured using its Energy Profiler application) and Samsung Nexus S (measured using Monsoon Power Monitor) in Table III. We used AGNSS for E71. The duration to get the first accurate fix ranges from 6.7 to 43 seconds and thus the average and standard deviation of E_{GPS} are high for E71. E_{W_on} of E71 includes the energy consumption of turning on its WiFi interface, association with an AP and the automatic

configuration of IP address through DHCP, as it is hard to separate these operations on E71. Since AGNSS does not work on our Nexus smartphone, E_{GPS} of Nexus is for the GNSS location method and thus is higher than other two (using AGNSS). These results verify that WiFi fingerprint localization is more energy efficient than GPS. These results also show that the receiving power of WiFi interfaces is lower than that of 3G for these three smartphones. For the real deployment of MADNet, we can get the power values using either existing online estimation tools (e.g., PowerTutor [28]) or offline profiles built for different types of smartphones.

We estimate the energy saving of offloading mobile data traffic to the deployed outdoor AP using the results in Table III combined with the above experimental results. We present the energy savings for the walking experiments in Table IV, where the first row of the results is for walking to the east and the second row is for the west. E_{3G} and E_W are the energy consumption of data transfer through 3G and WiFi networks. We turn on the WiFi radio twice during the offloading, first for WiFi localization and second for WiFi association.

During WiFi fingerprinting based localization, Smart-Client scans neighboring APs three times to collect WiFi beacons. n is the number of associations to connect to the WiFi AP. Thus the value of E_{oo} is set as

$$E_{oo} = 2 \times E_{W_on} + 2 \times E_{W_off} + 3 \times E_{scan} + n \times E_{asso}.$$

The results presented in Table IV show that MADNet can achieve more than 80% energy saving on smartphones when offloading mobile traffic to WiFi networks, benefiting from the collaborative MADNet architecture.

We also note that we did not consider the energy consumption of control information exchange between Smart-Client and Proxy-Cellular due to the small amount of data exchanged (usually less than 1KB). As these information is transferred in parallel with the music streaming, the head and tail energy is thus avoided. Moreover, we use static IP address configuration enabled by the control channel of MADNet and thus avoid the energy consumption for getting an IP address through DHCP.

To gain more insights, we also present the estimated energy savings in Table V for scenarios where 3G throughput is set

TABLE III
AVERAGE MEASURED POWER (WATT) AND ENERGY CONSUMPTION (JOULE) RELATED TO MOBILE DATA OFFLOADING.

Device	OS	P_{3G}	P_W	E_{W_on}	E_{asso}	E_{W_off}	E_{scan}	E_{GPS}
N900	Maemo (Nokia)	1.10±0.017	0.645±0.023	0.18±0.025	0.28±0.13	0.13±0.021	0.53±0.077	4.0±1.3
E71	Symbian (Nokia)	1.33±0.023	1.28±0.032	6.4±0.19	n/a	0.13±0.025	0.11±0.036	9.0±6.5
Nexus S	Android (Google)	0.891±0.022	0.658±0.16	0.27±0.019	0.25±0.049	0.29±0.016	0.27±0.017	10±1.3

TABLE IV
MEASURED OFFLOADING CAPACITY, AVERAGE THROUGHPUT OF 3G AND WiFi NETWORKS AND ESTIMATED ENERGY SAVING.

	C_W	B_{3G}	B_W	E_{oo}	E_{3G}	E_W	Saving
	MB	Mbps	Mbps	Joule	Joule	Joule	%
East	27.2	0.8	3.5	3.4	304.6	39.3	85.98
West	22.9	0.8	3.2	3.5	257.3	38.9	83.52

TABLE V
ESTIMATED ENERGY SAVING FOR DIFFERENT 3G THROUGHPUT.

	0.5 Mbps	1.0 Mbps	1.5 Mbps	2.0 Mbps	3.0 Mbps	5.0 Mbps
East	90.7 %	81.7 %	72.9 %	64.2 %	47.5 %	16.3 %
West	89.1 %	78.5 %	68.2 %	58.2 %	38.9 %	3.4 %

virtually to be 0.5, 1.0, 1.5, 2.0, 3.0, and 5.0 Mbps. The key observation is that when 3G throughput increases, the gain of energy consumption by offloading mobile data to WiFi networks decreases due to the shorter 3G transmission duration at higher throughput.

V. DISCUSSION

A. Distributed Content Caching

Usually, WiFi service providers prefer to over-subscribe the backhaul connections from the APs to Internet and thus these backhaul links become the communication bottleneck. We aim to explore the limited storage space (e.g., several GB through mounted USB drives) on WiFi APs to cache data locally. The challenge here is to understand the content access patterns from mobile users, without which it is hard to decide the right locations to cache the right content. Similar to data prefetching, caching data on local APs can also fully utilize the bandwidth of WiFi links.

B. Incentives

Social participation is becoming an enabling factor for more and more mobile applications. How to integrate an effective incentive scheme into MADNet to encourage active participation is another challenging problem. Cellular operators may offer reduced subscription fees to users who are willing to share their home APs with others and thus increase WiFi AP availability. There may also be issues related to the pricing and billing models for traffic offloading between different cellular operators and WiFi providers. For example, in countries like Canada where usage-based billing has been introduced for Internet access, people cannot share their home APs with others for free and mobile users may need to pay for the traffic offloaded to WiFi networks (probably cheaper than 3G traffic).

C. Security and Privacy

For security and privacy, when MADNet redirects mobile data traffic to WiFi APs, cellular networks can pass the

identity information of end-users to WiFi networks and thus enable various security models to prevent illegal uses of these APs. Meanwhile, we may require mobile users to tunnel their packets to one of their trusted points and handoff the access control responsibility to that endpoint [22]. Thus, all the offloaded traffic will go through that trusted point and can be identified if the mobile users illegally download music or video. For example, SWISH [14] has utilized a similar technology to secure the shared WiFi networks and to protect the privacy of mobile users. Note that mobile data offloading is transparent to content/cloud service providers and MADNet will not reveal the end-user's location to them.

VI. RELATED WORK

A. Cellular Traffic Offloading

Among several existing schemes for cellular traffic offloading, Femtocells as an extension of the macrocells of cellular networks were originally proposed to offer better indoor services. When indoor users switch from macrocells to femtocells, femtocells can potentially offload cellular data traffic. However, the femtocell signal may interfere with nearby macrocell transmissions, since they work on the same spectrum as macrocells [27]. Recently, how to offload cellular traffic through mobile-to-mobile opportunistic communications is also investigated as a complementary offloading solution [9].

WiFi is another attractive technology for cellular traffic offloading. For example, Korhonen et al. [12] analyze the latest trend of network controlled offloading and compare industrial standardization solutions. Hou et al. [11] propose a transport layer protocol to offload 3G data traffic to WiFi hotspots for vehicular access networks. Ristanovic et al. [20] propose an algorithm, called HotZones, to offload delay-tolerant cellular content to WiFi APs and evaluate the performance through trace-driven simulations.

Wiffler [2] aims at maximizing the amount of 3G traffic offloaded to WiFi networks for PCs on vehicular networks. In contrast to Wiffler, MADNet targets at smartphones which are much more challenging than PCs (as shown in Section II). MADNet also advances the state-of-the-art by taking the energy consumption of smartphones into account when making offloading decisions. In addition, Wiffler uses only open WiFi APs while MADNet also considers the support from cellular operators, WiFi providers, and end-users to increase the offloading opportunities.

Using a trace-driven simulation, Lee et al. [13] show that WiFi networks can offload around 65% of mobile data traffic for the traces collected from about 100 iPhone users. The simulator assumes that data traffic can be offloaded to WiFi networks *whenever a user connects* a smartphone to a WiFi AP

in the traces, thus the offloading decision is actually made by these users. Compared to their work, our proposed offloading solution is transparent to users. Furthermore, we implement a prototype on Nokia N900 smartphones and evaluate its performance in live environment.

A recent measurement study examining the relationship of WiFi versus 3G usage also explored the antenna limitation of smartphones [15]. Although the target is different from ours, their findings confirm that a dedicated offloading scheme for smartphones is necessary.

B. Leveraging Multiple Radio Interfaces on Mobile Devices

Nowadays smartphones are typically equipped with multiple radio interfaces, including Bluetooth, WiFi and 3G cellular, and there are several existing systems leveraging these interfaces for energy efficiency and better throughput performance. For example, CoolSpots [19] explores policies that enable a mobile device to automatically switch between WiFi and Bluetooth interfaces by considering their different transmission ranges, and thus reduces the energy consumption of wireless communication. MAR [21] is a multi-homed mobile access router that exploits multiple channel access technologies provided by different service providers. The MAR implementation demonstrates the benefits of aggregating link capacity from three cellular operators for different applications.

Comparing to previous work, MADNet uses 3G networks to provide primary data channels and to facilitate offloading procedures. At the same time, it employs WiFi interfaces to prefetch predictable data and thus improves energy efficiency of data transfer.

VII. CONCLUSIONS AND FUTURE WORK

In this paper we propose the MADNet mobile data offloading architecture that leverages WiFi networks to migrate mobile traffic from cellular networks. The design choices are motivated by experimental results from war-driving and war-walking and measurement study of WiFi-based Internet access for moving smartphones. By considering the battery constraints of smartphones, we design the first energy-aware offloading decision algorithm with the goal of improving energy efficiency for smartphones. We show that the potential energy saving of offloading depends on the throughput of 3G and WiFi networks, and also the amount of data we can offload. We confirm the feasibility of our proposal through prototype implementation on Nokia N900 smartphones and evaluate the performance in the wild, which verifies the effectiveness of MADNet.

ACKNOWLEDGEMENT

This work is supported in part by the Internet of Things (IoT) program of Tivit funded by Tekes, the Academy of Finland Grant No. 139144, and by the US NSF Award CNS-1010789. The PhD research of Aaron Yi Ding is supported by the Academy of Finland and the Nokia Foundation. The

authors would like to thank Yonghao Li and Peng Liu for their support with the experimentation, and thank the SECON 2013 referees for their helpful comments.

REFERENCES

- [1] P. Bahl and V. N. Padmanabhan. RADAR: An In-Building RF-Based User Location and Tracking System. In *Proc. of INFOCOM 2000*.
- [2] A. Balasubramanian, et al. Augmenting Mobile 3G Using WiFi. In *Proc. of MobiSys 2010*.
- [3] N. Balasubramanian, et al. Energy Consumption in Mobile Phones: A Measurement Study and Implications for Network Applications. In *Proc. of IMC 2009*.
- [4] V. Bychkovsky, et al. A Measurement Study of Vehicular Internet Access Using In Situ WiFi Networks. In *Proc. of MobiCom 2006*.
- [5] P. Deshpande, et al. Performance Comparison of 3G and Metro-Scale WiFi for Vehicular Network Access. In *Proc. of IMC 2010*.
- [6] P. Deshpande, et al. Predictive Methods for Improved Vehicular WiFi Access. In *Proc. of MobiSys 2009*.
- [7] F. R. Dogar, et al. Catnap: Exploiting High Bandwidth Wireless Interfaces to Save Energy for Mobile Devices. In *Proc. of MobiSys 2010*.
- [8] D. Hadaller, et al. Vehicular Opportunistic Communication Under the Microscope. In *Proc. of MobiSys 2007*.
- [9] B. Han, et al. Mobile Data Offloading through Opportunistic Communications and Social Participation. *IEEE Transactions on Mobile Computing*, 11(5):821–834, May 2012.
- [10] B. Han, et al. Energy-Aware Collaborative Mobile Data Offloading for Smartphones. Technical report, Department of Computer Science, University of Maryland, July 2011. Available at <http://www.cs.umd.edu/~bohan/MADNet.pdf>.
- [11] X. Hou, et al. Moving Bits from 3G to Metro-Scale WiFi for Vehicular Network Access: An Integrated Transport Layer Solution. In *Proc. of ICNP 2011*.
- [12] J. Korhonen, et al. Toward Network Controlled IP Traffic Offloading. *IEEE Communications Magazine*, 51(3):96–102, March 2013.
- [13] K. Lee, et al. Mobile Data Offloading: How Much Can WiFi Deliver? In *Proc. of CoNEXT 2010*.
- [14] D. Leroy, et al. SWISH: Secure WiFi sharing. *Computer Networks*, 55(7):1614–1630, May 2011.
- [15] S. Liu and A. Striegel. Casting Doubts on the Viability of WiFi Offloading. In *Proc. of CellNet Workshop 2012*.
- [16] G. Maier, et al. A First Look at Mobile Hand-held Device Traffic. In *Proc. of PAM 2010*.
- [17] A. J. Nicholson, et al. Improved Access Point Selection. In *Proc. of MobiSys 2006*.
- [18] A. J. Nicholson and B. D. Noble. BreadCrumbs: Forecasting Mobile Connectivity. In *Proc. of MobiCom 2008*.
- [19] T. Pering, et al. CoolSpots: Reducing the Power Consumption of Wireless Mobile Devices with Multiple Radio Interfaces. In *Proc. of MobiSys 2006*.
- [20] N. Ristanovic, et al. Energy Efficient Offloading of 3G Networks. In *Proc. of MASS 2011*.
- [21] P. Rodriguez, et al. MAR: A Commuter Router Infrastructure for the Mobile Internet. In *Proc. of MobiSys 2004*.
- [22] N. Sastry, et al. Architecting Citywide Ubiquitous Wi-Fi Access. In *Proc. of HotNets-VI 2007*.
- [23] U. Shevade, et al. Enabling Enabling High-Bandwidth Vehicular Content Distribution. In *Proc. of CoNEXT 2010*.
- [24] E. Shih, et al. Wake on Wireless: An Event Driven Energy Saving Strategy for Battery Operated Devices. In *Proc. of MobiCom 2002*.
- [25] V. Valancius, et al. Greening the Internet with Nano Data Centers. In *Proc. of CoNEXT 2009*.
- [26] J. Yao, et al. An Empirical Study of Bandwidth Predictability in Mobile Computing. In *Proc. of WiNTECH 2008*.
- [27] J.-H. Yun and K. G. Shin. CTRL: A Self-Organizing Femtocell Management Architecture for Co-Channel Deployment. In *Proc. of MobiCom 2010*.
- [28] L. Zhang, et al. Accurate Online Power Estimation and Automatic Battery Behavior Based Power Model Generation for Smartphones. In *Proc. of CODES/ISSS 2010*.