

CasCap: Cloud-assisted Context-aware Power Management for Mobile Devices

Yu Xiao
Aalto University, Finland
yu.xiao@aalto.fi

Pan Hui
Deutsche Telekom
Laboratories, Germany
pan.hui@telekom.de

Petri Savolainen
HIIT/University of Helsinki,
Finland
petri.savolainen@hiit.fi

ABSTRACT

Energy consumption of wireless data communications on mobile devices is growing due to the increase in mobile data traffic caused by the popularity of mobile Internet applications. To address this issue, we propose CasCap, a novel cloud-assisted context-aware power management framework, which takes advantage of the processing, storage and networking resources in the cloud to provide secure, low-cost and efficient power management for mobile devices. CasCap is featured by crowd-sourced context monitoring, function offloading to the cloud, and providing adaptations as services. We give three example scenarios where wireless data communications are adapted to network conditions, traffic characteristics and location information. Our preliminary results show the potential of using CasCap to improve energy efficiency in wireless data communications.

Categories and Subject Descriptors

C.2.1 [Computer Systems Organization]: Network Architecture and Design—*Wireless communication*

General Terms

Design

1. INTRODUCTION

The explosion in popularity of Internet applications, such as Facebook, YouTube, and Google maps, accompanied with the growing popularity of mobile Internet devices, is driving a boost in mobile data traffic. This boost carries with it a price in terms of energy demands in mobile devices. However, the battery technology has not been developing fast enough to satisfy the growing demands. Hence, improving the energy efficiency of wireless data communications on mobile devices has become more important than ever.

The energy efficiency of wireless data communications is highly dependent on the *context* of the mobile devices. For example, transmission cost per transmitted byte is relatively

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MCS'11, June 28, 2011, Bethesda, Maryland, USA.
Copyright 2011 ACM 978-1-4503-0738-3/11/06 ...\$10.00.

low if the signal strength of the wireless access network is high [11]. Another example of a context variable that has impact on the energy consumption is the traffic pattern. Less energy is consumed if the data is transmitted in the shape of bursts instead of transmitting the same packets at random intervals [3]. Therefore, it is important for power management to be context aware.

The context information that is relevant to the energy efficiency of wireless data communications can be collected from different sources, including mobile devices, wireless access networks, and Internet service providers. Previous works [1, 4, 3] focused on the results of power management using the context information that is specific to certain mobile applications or networks, whereas little attention has been paid to the actual collection, processing and distribution of the context information. These tasks, however, are far from trivial because of the diversity and the sheer number of different context information sources and consumers.

In this paper, we propose CasCap, a novel context-aware power management framework, taking advantage of the processing, storage and networking resources in the cloud to provide secure, low-cost and efficient power management for mobile devices. Three main features of CasCap are crowd-sourced context monitoring, function offloading, and adaptation as service.

Crowd-sourced context monitoring includes the collection of context information from distributed context sources and the processing of the large amount of context information into more meaningful information that can be used for power management for mobile devices. It enables mobile devices to schedule energy-efficient adaptations in a proactive manner based on collective intelligence that is formed during context sharing. For example, mobile devices that are connected to different access networks can periodically upload their signal strength traces to the cloud together with their own GPS co-ordinates to form a map of access networks with signal strengths. Now mobile devices can choose which access network to connect based on this signal strength map that they have downloaded from the cloud. Also using this map, devices that are on the move can save energy by pausing their data transmission before arriving to an area with a low signal strength.

Using the cloud as the basis for our crowd-sourced context monitoring framework allows us to deal with the highly variable, and potentially a large number of context information sources and consumers in an elastic way, as according to the computing-as-a-service principle [5], the processing power and

storage capacity of the system can be scaled up and down as needed.

Functionality offloading provides possibility for mobile devices to migrate context collection, context processing and/or adaptation scheduling to the cloud for energy savings. The savings include computational and communications cost. In practice, the scheduling can be offloaded to a third-party proxy or to a clone of the mobile device. By "clone" we refer to the whole or part of the augmented mobile OS in the cloud, as implemented in Cloudlet [10] or CloneCloud [2]. We argue that privacy and security concerns favor the clone-based approach over the traditional proxy-based one, since a clone dedicated to only one mobile device whereas a traditional proxy is shared by a number of mobile devices. Having your context information sent only to your private clone reduces the risk of compromising sensible context information.

Adaptation as service goes beyond the traditional proxy-based power management. Instead of using application-specific proxies, we propose offering the energy-saving functions that are traditionally provided by application-specific proxies as generic Internet services. For example, traffic shaping and transcoding can be offered as generic services that can be utilized in several mobile applications. In addition to avoiding the unnecessary work involved in implementing overlapping functionality in each application-specific proxy, certain energy optimizations such as traffic shaping will work much better if implemented with only one traffic shaper calculating the traffic burst schedule instead of a number of application-specific traffic-shaping proxies running at the same time.

2. BACKGROUND

Energy consumption of wireless data communications depends on the performance of data transmission. Generally, transferring at a higher data rate means higher utilization of wireless network interface(WNI), and is therefore more energy efficient [15]. The data rate is effected by the conditions of the access networks. For example, radio interference makes both the data rate and energy efficiency of the wireless network to decrease. A simple way of reducing the interference is to limit the number of connections to the access points that are using the same or neighboring channels. We apply this method in Scenario A of section 4 of this paper where the adaptation guides the mobile device to connect to the access point with a small number of users. Differently from [8], which estimates the interference based on the wireless traffic traces captured from access points and adapts the operations of the APs to the interference, our method focuses on the adaptations on mobile devices.

Signal strength of wireless access networks has been proved to have impact on the network throughput in both 3G [11] and WLAN[12]. With stronger signal strength, the network transmission can gain higher throughput and therefore consumes less energy. Signal strength can be associated with location. For example, Wifi-report [7] collected SNR from mobile devices and provided throughput statistics for a certain access point in each SNR range to mobile devices for access point selection. Similarly, in Scenario B, we use the collective knowledge gained from the performance information collected from mobile devices to guide the transmission adaptations. Compared to Wifi-report, we use real-time information in addition to historical information. In mobility

scenarios, previous work [12] has proposed to predict the SNR based on time-series statistical models. We extend the previous work by sharing the prediction models among mobile devices, so the mobile devices do not need to repeat trace collection and model training.

Proxy-based power management have been widely used. Examples of such services can be converting high-quality media stream into low-quality one for mobile devices at media servers [6], scheduling traffic at access points[9], or executing peer-to-peer sharing on a proxy [4]. Inspired by the fact that the traffic delivered in bursts is friendlier to the power saving mode (PSM) of wireless networks, proxies are introduced into the system to shape the traffic [1, 13, 3, 6] into bursts before forwarding it to mobile devices. In some cases like streaming, traffic shaping can save up to 65% [3, 1] more energy consumption compared to PSM. In the cloud traffic shaping can be provided as an adaptation service, so it can be discovered and subscribed by mobile devices and can be shared by Internet services.

3. SYSTEM OVERVIEW

As shown in Figure 1, CasCap encompasses three components: mobile devices, Internet services, and clones. Among them, utilizing the clone environment is optional, but recommended for improving energy savings and privacy.

On mobile devices CasCap consists of five functional components, namely, resource manager, context manager, scheduler, policy manager and communicator. Resource manager is a background service monitoring the resource consumption on the devices and collecting information from embedded sensors like accelerometer and GPS receivers. Context manager generates context information based on the information collected by resource manager and selects the context information to be uploaded to the cloud. It also downloads context information, which is missing but needed for adaptations, from the cloud.

The crowd-sourced context monitoring services running in the cloud, collect context information from mobile devices and other network elements, process the collected context information into a more meaningful form, and handle the context information queries from mobile devices. For example, a network mapping service provides the information of each WLAN access point, such as location, current number of connections and signal strength traces, based on the information provided by the mobile devices connected to each network access point.

The scheduler running on a mobile device adapts the mobile device to the changes in context basing its adaptation decisions on policies. The policy manager is in charge of the installation of new policies on the mobile device, updating the existing policies, and sharing the policies with other mobile devices. The sharing of policies is implemented in the same way as the sharing of context information. The policies that are managed by the policy manager on the mobile device define the actions that should be taken on the device itself and on Internet services whenever a specific change in context information occurs. The Internet services running in the cloud provide traffic shaping, transcoding, offloading and other adaptation functions.

Context sharing, policy sharing and using the Internet services in adaptations require wireless communications between mobile devices and Internet services. In CasCap the

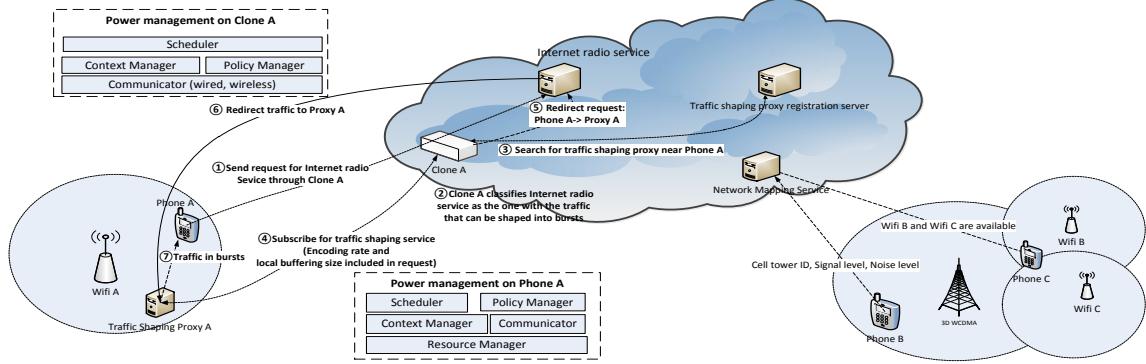


Figure 1: CasCap Architecture.

required wireless networking functionality on mobile phones is provided by the communicator component.

Some functions on mobile devices, such as the processing of context information, the security control in context sharing and policy sharing, and the scheduling of adaptations can be offloaded to clones for energy savings. If using the clone, the requests for Internet services from mobile devices are sent to the corresponding Internet services through the clone, and the clone will take care of power management of the wireless data communications involved in the requested services.

4. APPLICATION SCENARIOS

We take three moments from Alice's business trip as examples to show the scenarios where mobile devices would benefit from our framework.

Alice is a consultant. She is traveling to another city to meet her customers. During her trip, she uses her mobile phone among other things to access her employer's database, to exchange emails, and to browse the Internet.

Scenario A: Selecting the access point with the best performance

9.30am Alice is at a meeting. Her mobile phone sends a request to a network mapping service asking for the information about WLAN access points that are available in the meeting room. The network mapping service replies with a list of available access points, the number of clients currently connected to each access point, and the average data rate of each access point during the previous week. Based on this information, the mobile phone chooses an access point called 'guest', as its historical record shows a relatively high performance and there are not too many users connected to it at the moment.

Scenario B: Adapting network transmission to the signal strength of the wireless access network

11.45am Alice is going through a long U-style corridor. She has started downloading some analysis reports for the customer she is going to visit in the afternoon. The accelerometer on her mobile phone detects her movement. The movement triggers the signal strength prediction. The phone downloads the signal strength map of the wireless access point she is connected to from the network mapping service in the cloud. Based on the map and the estimation of Alice's movement, it is predicted that in 10 seconds Alice is going to enter a corner, where the signal strength is too low for file

downloading. In addition, it will take about 15 seconds for Alice to get through that corner. Hence, her mobile phone decides to pause the file download and switches off the network scanning for 15 seconds starting after 10 seconds.

Scenario C: Redirecting to the traffic shaping proxy that is closer to the mobile device

5:30 pm Alice starts to listen to Internet radio on her mobile phone while waiting for her flight at the airport. It is known based on historical information that the requested radio stream includes a lot of small packet intervals, which can be shaped into bursts for saving energy. The phone searches for a traffic shaping proxy and chooses the closest one. The phone then subscribes to the traffic shaping service giving the size of its buffer and the encoding rate of the Internet radio as parameters.

In the above three scenarios, the context collection from Internet services, the scheduling of adaptations, and the execution of adaptations are all conducted on mobile devices. As described in Section 3, these tasks can also be offloaded to clones for energy savings and privacy. In Figure 1 we show a clone-using workflow for Scenario C. Compared with the workflow without using clone, offloading to clones reduces the energy cost caused by application classification, service subscription, and traffic redirection. In addition, it eases the deployment of power management policies, as Internet service providers can deploy their power management policies on clones instead of mobile devices. The deployment through wired links on clones is more energy-efficient than the deployment through wireless links on mobile devices.

5. IMPLEMENTATION

We choose a Linux-based mobile OS, Maemo, as our experimental platform. We implement the resource manager as a daemon that monitors the WLAN interface and the GPS-based location information. The resource manager obtains the WLAN information, including BSSID, signal level and noise level, through ioctl calls with the macro SIOCGI-WSTATS as parameter. The liblocation library of Maemo location framework is used for gaining access to the location information.

Context manager uploads selected context information to the cloud. For example, when a connection between a mobile device and an access point is established or torn down, the context manager sends an anonymous notification with the BSSID of the access point in question to the network map-

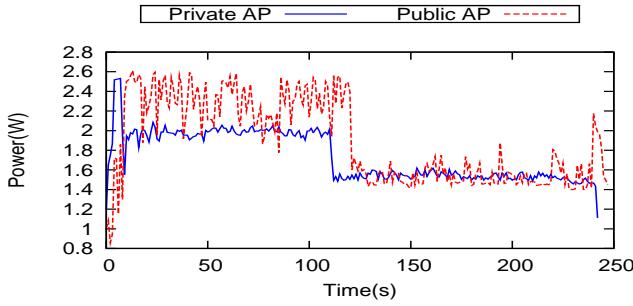


Figure 2: Power consumption of YouTube download through two WLAN access points with different number of users.

ping service. Based on these notifications, the network mapping service can keep an up-to-date count of the number of active connections to each access point. Context managers running on mobile devices can request the active connection count of access points from the network mapping service and use the information when choosing which access point to connect to.

The policy manager is a container for policies. We define policies as event-condition-action rules. Such a policy can be for example: "when the signal-to-noise ratio (SNR), is lower than a threshold, pause the ongoing transmission." In this scenario, SNR is the context to be monitored. The change in the context, the SNR becoming lower than the threshold, is defined as an event. Under the conditions of having ongoing transmission, the mobile device executes the action of pausing the ongoing transmission.

The scheduler component receives notifications about the changes in context from the context manager and Internet services. By mapping the changes defined in the notifications to the events described in policies, scheduler selects the corresponding policies from policy manager and invokes the corresponding actions.

Communicator implements the information exchange between mobile devices and their clones, and between mobile devices and Internet services. In the case of clone, there is a TCP connection between the mobile device and its clone. The messages exchanged between them follow a proprietary protocol. In the case of Internet services a HTTP request/response carries the messages. The messages can be formatted in SOAP or REST styles, depending on the Internet service in question.

The cloud-based Internet services that provide context information can be built and deployed using Google AppEngine. In our scenarios, a network mapping service is running in the cloud. It collects WLAN information including BSSID, location, connection notification, real-time signal level, real-time noise level, and traces of signal strength with location tags. The information is stored in a database and can be searched by BSSID and location. In addition, some existing location information services, such as Yahoo! Fire Eagle¹ and CTrack [14] can be integrated into our framework in the future.

6. PRELIMINARY RESULTS

To investigate the potential in energy savings by using the

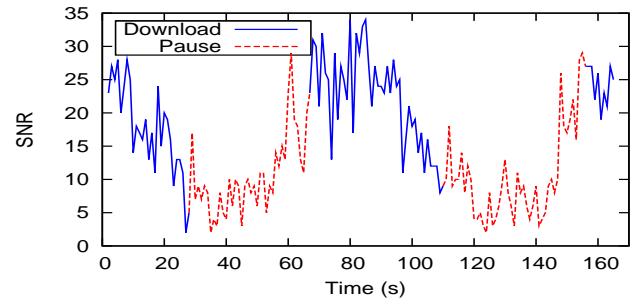


Figure 3: WLAN SNR trace collected from a U-style corridor.

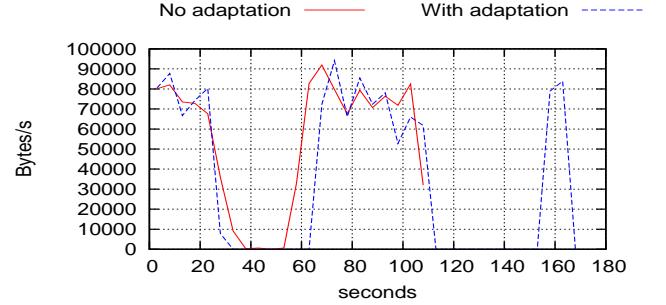


Figure 4: Comparison of throughput with/without adaptation.

techniques proposed in the three scenarios, we carried out power measurement for file downloading and streaming to mobile devices in different access points.

Scenario A: We downloaded a 9284KB video from YouTube through two different WLAN access points. The first access point was a private one with only one active user, and the other one was a public one with multiple users. The interference in the first one was ignorable, while that in the latter was causing packet loss and retransmissions. As shown in Figure 2, with less interference, the network performance in terms of download time is better, and the power consumption is relatively low. With a network mapping service running in the cloud, mobile devices can query the number of clients connected to each access point, and choose the one with less connections without wasting energy and time in trying all the candidates.

Scenario B: We collected the traces of signal and noise levels at 1Hz while walking through a U-style corridor at a fixed speed. Signal-to-noise-ratio (SNR) was calculated, as shown in Figure 3, and tagged with position. We downloaded a 5.3MB file from www.openss7.org using wget. We compared the performance of the downloads with and without adaptations. With the adaptation based on location information got from the signal strength traces shown in Figure 3, the mobile device pauses the downloading when it enters the areas where SNR is predicted to be lower than 15 and turns off the WNI during the pause. As shown in Figure 4, the downloading was paused for in total 85 seconds. The WNI spent 83.637 seconds in RECEIVE mode, which cost 98.775 Joules. Without adaptation, the WNI spent 85.07 seconds in RECEIVE mode and another 20.29 seconds in IDLE mode trying to repair the connection, which cost 118.404 Joules. Hence, the downloading without adap-

¹<http://fireeagle.yahoo.net/>

tation consumed around 20% more energy. This result can be further improved with more accurate location estimation, because, as shown in Figure 3, the error in location prediction causes the delay in resuming the download at the end of the pauses.

In this scenario, a network mapping service is expected to collect the traces of SNR and location and provide the traces to mobile devices through Internet. In some cases where location information is not available, it is possible to apply statistical model for SNR prediction. Using cloud mobile devices can share the models. Hence, even for the first connection to an access point, the mobile device can apply adaptation based on the model.

Scenario C: Hoque et al. [3] showed that the location of traffic shaping proxy has impact on the energy savings. They deployed proxies in different cities and compared the power consumption of mobile devices while using different proxies. Their results show that using the proxy which is the closest one to the mobile device can gain around 40% more energy savings than using the farthest one.

7. DISCUSSION

There are three challenges that need to be solved in the future. 1) Live session migration. Take Scenario C as an example, the Internet radio stream might need to be migrated from one traffic shaping proxy to another, when the mobile device moves to another access network, or when proxies need to balance their workload. Hence, mechanisms are needed to guarantee the performance of network transmission during migration. 2). Validity of context information. Some context information stored in the cloud might become invalid after some time. For example, the distribution of signal strength in a certain access network will change if new access points are deployed in the neighborhood. If the information is stored on mobile devices or clones in a distributed manner, evaluating the validity requires the support for content searching among these devices or clones. However, we have not seen feasible solution from the literature. An alternative is to store all the data in centralized points. In that case, the challenges come from the complexity in storing and processing the large amount of data and the issue of single failure point.

8. CONCLUSIONS

We proposed CasCap, a framework which uses the resources in the cloud to provide secure, low-cost and efficient context-aware power management for mobile devices. The main features of CasCap include crowd-sourced context monitoring, function offloading, and adaptation as service. Our design is motivated by the scenarios from the typical usage of mobile data communications. The preliminary results of the example scenarios show the potential in improving energy efficiency through CasCap. In addition, our framework provide a novel way for third-party service providers to develop and deploy power management services, which we believe will push the growth of energy-efficient Internet services further.

9. REFERENCES

- [1] S. Chandra. Wireless network interface energy consumption: implications for popular streaming formats. *Multimedia Syst.*, 9(2):185–201, 2003.
- [2] B.-G. Chun, S. Ihm, P. Maniatis, M. Naik, and A. Patti. Clonecloud: Elastic execution between mobile device and cloud. In *EuroSys 2011*, April 2011.
- [3] M. Hoque, M. Siekkinen, and J. K. Nurminen. On the energy efficiency of proxy-based traffic shaping for mobile audio streaming. In *CCNC 2011*. IEEE, 2011.
- [4] I. Kelenyi and J. Nurminen. Cloudtorrent-energy-efficient bittorrent content sharing for mobile devices via cloud services. In *CCNC 2010*, pages 1–2, January 2010.
- [5] A. Michael, F. Armando, G. Rean, A. D. Joseph, K. Randy, K. Andy, L. Gunho, P. David, R. Ariel, S. Ion, and Z. Matei. Above the clouds: A berkeley view of cloud computing. In *Technical Report EECS-2009-28, EECS Department, University of California, Berkeley.*, 2009.
- [6] S. Mohapatra, N. Dutt, A. Nicolau, and N. Venkatasubramanian. Dynamo: A cross-layer framework for end-to-end qos and energy optimization in mobile handheld devices. *Selected Areas in Communications, IEEE Journal on*, 25(4):722–737, may 2007.
- [7] J. Pang, B. Greenstein, D. McCoy, M. Kaminsky, and S. Seshan. Wifi-reports: Improving wireless network selection with collaboration. In *MobiSys 2009*, Jun 2009.
- [8] L. Ravindranath, C. Newport, H. Balakrishnan, and S. Madden. Improving wireless network performance using sensor hints. In *NSDI 2011*, pages 21–21, Berkeley, CA, USA, 2011. USENIX Association.
- [9] E. Rozner, V. Navda, R. Ramjee, and S. Rayanchu. Napman: network-assisted power management for wifi devices. In *MobiSys 2010*, pages 91–106, New York, NY, USA, 2010. ACM.
- [10] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies. The case for vm-based cloudlets in mobile computing. *IEEE Pervasive Computing*, 2009.
- [11] A. Schulman, V. Navda, R. Ramjee, N. Spring, P. Deshpande, C. Grunewald, K. Jain, and V. N. Padmanabhan. Bartendr: a practical approach to energy-aware cellular data scheduling. In *MobiCom 2010*, pages 85–96, New York, NY, USA, 2010. ACM.
- [12] R. Sri Kalyanaraman, Y. Xiao, and A. Ylä-Jääski. Network prediction for energy-aware transmission in mobile applications. *International Journal on Advances in Telecommunications*, 3:72–82, November 2010.
- [13] E. Tan, L. Guo, S. Chen, and X. Zhang. Psm-throttling: Minimizing energy consumption for bulk data communications in wlans. In *ICNP 2007*, pages 123–132, oct. 2007.
- [14] A. Thiagarajan, L. S. Ravindranath, H. Balakrishnan, S. Madden, and L. Girod. Accurate, low-energy trajectory mapping for mobile devices. In *NSDI 2011*, Boston, MA, March 2011.
- [15] Y. Xiao, P. Savolainen, A. Karpanen, M. Siekkinen, and A. Ylä-Jääski. Practical power modeling of data transmission over 802.11g for wireless applications. In *E-Energy 2010*, pages 75–84. ACM, april 2010.