

Fine-Grain Adaptation Using Context Information

Iqbal Mohamed

University of Toronto, Toronto, ON, Canada
iq@cs.toronto.edu

ABSTRACT

In this research, I explore how distributed applications can achieve fine-grain adaptation in a mobile environment. I put forth the thesis that this can be achieved by exploiting context information, and support the idea with two applications: browsing the web on a mobile device, and continuous remote health monitoring of individuals.

KEYWORDS

Content Adaptation, Mobile Devices, Context

1 INTRODUCTION

Writing distributed applications for mobile environments is a challenging problem because of severe resource limitations on devices [1]. These include constraints on battery capacity, processing capability, screen real estate, I/O facilities and storage. The problem is further exacerbated by the considerable heterogeneity among different classes of devices (PDAs, cell phones, pagers, etc.). Further, as users move about naturally in the course of their activities, the mobile computers they carry with them experience significant variability in wireless connectivity. These issues severely limit the usefulness of mobile devices for applications such as browsing the World Wide Web.

The research community has expended significant effort into addressing these problems [2-6]. My work focuses on a specific approach, called automatic adaptation. At a high-level, the goal of adaptation is to customize content to better suit the prevailing resource environment on a mobile device. Of course, the fundamental question is how should content be adapted? Automatic adaptation is typically achieved by applying a set of static or dynamic (accounting for prevailing resource constraints, for instance) adaptation rules to content. Unfortunately, the amount of effort required to achieve fine-grain adaptation with rules is comparable to manually adapting content.

In my research, I have been exploring how fine-grain adaptation can be achieved without expending significant manual effort. I put forth the following thesis:

Applications running on a mobile device can use context information to determine the relevance of data in a given situation, and exploit this knowledge to perform fine-grain adaptation.

To support this thesis, I have primarily focused on the problem of adapting web content for mobile devices. A second application that I have considered is adapting the

fidelity of data generated by body-worn sensors, as it is being relayed to remote servers on the Internet via a mobile device. Two key differences between these distributed applications are the role of the user (web browsing involves the user while sensor data relay does not) and the direction of data flow (web browsing involves downloading of content onto the mobile device, while sensor data relay primarily involves upload).

2 BROWSING THE WEB ON MOBILE DEVICES

Fine-grain adaptation of web content requires taking into account how the user will utilize the content (intentional context, which we call usage semantics), as well as situational context, such as the device being used, location, network connectivity, etc. Though it is impractical to scale, one could create tailored adaptation rules for a particular object that account for situational context. However, such rule-based techniques are unable to take user intention into consideration. In this section, I first describe our work that takes usage semantics into account in the adaptation process. Next, I describe how our system can determine what situational context is most relevant for adapting content. Taken together, this work allows us to provide fine-grain adaptation to mobile users browsing the web.

2.1 Tailoring Content to User Intentions

We developed a new adaptation technique called Usage Aware Interactive Content Adaptation (URICA) that takes into account how relevant an object is expected to be for users [8]. URICA works as follows: when serving content for the first time, the adaptation system makes a default decision on how to customize the content. At any point, users who are unsatisfied with the system's adaptation decision can take control of the adaptation process and make changes until the content is suitably adapted for their purposes. The successful adaptations made by users are recorded and become part of the history that is used in making future adaptation decisions for other users. The collected history can then be used to determine the relevance of a particular object.

We applied the URICA technique to the problem of browsing the web on a mobile device over heterogeneous wireless links, such as GPRS, EVDO and Wi-Fi. As users move about, they have access to different types of wireless links, where links have distinct costs and throughput. We interpose a proxy in the network path between the mobile client and the web server, which can vary the fidelity of large web objects such as images. Our system can then allow users to minimize their expected browsing time,

which can be defined as the sum of object download times and the user's response times for interactively adapting content. All things being equal, the slower the link or the faster the user's interaction time, the more aggressive the fidelity reduction on objects. This is because in the aforementioned situations, serving an object at higher fidelity than that is required results in significant download time, while doing the opposite incurs a relatively smaller time cost of user interaction. Our system can also allow users to balance the effort of providing interaction with the expected monetary savings in bandwidth costs. Users can specify the minimum monetary savings that they should achieve for each interaction with the adaptation system. Users who find the interactive adaptation process more cumbersome than others would require higher monetary savings to be achieved for every interaction. Thus, our adaptation system can take into account the effort required to interact, and provide a tailored experience to the user based on the current monetary cost of the network link.

2.2 Tailoring Content to the User's Context

There is a lot of situational context that can potentially be relevant for adapting a specific content object. Examples include display size and the type of wireless link currently available. However, the context that is most relevant to adapt an object may vary (consider the relevance of location information in a map application, for instance). These present significant challenges in building a context-aware adaptation system. We addressed these by developing a technique called Feedback-driven Context Selection that uses the feedback that is naturally received in the course of operation in a URICA system, and automatically determines what context is most relevant for a particular type of adaptation. We added context-awareness to URICA (CA-URICA) so that it only uses the restricted history of the "community" of previously encountered users with the same relevant context [9].

To determine the efficacy of our techniques, we performed several controlled user studies with two different types of adaptation systems: image fidelity (for adapting to different network links) and page layout (to adapt for different display sizes). We observed that the influence of context on adaptation requirements can vary based on the content being adapted, as well as the type of adaptation being performed. Our experiments with CA-URICA showed that grouping users into communities based on context can lead to significant improvements in overall performance of the adaptation system.

3 RELAYING DATA FROM ON-BODY SENSORS VIA A MOBILE DEVICE

The second application that I am exploring is the adaptation of data coming from body-worn health sensors. The prevailing architecture for this type of data collection is a three-tier model, wherein sensor data is first relayed via short-range radio (e.g. Bluetooth) to a mobile device, and then transmitted to remote servers on the Internet via the more powerful radio on the mobile device [10].

We observed that in many cases, the readings from sensors follow a "normal" pattern in a given contextual environment. We developed the notion of context-aware

filtering of sensor data streams in order to reduce transmissions in cases where the observed data corresponds to the norm expected by the system in a given context. We implemented these ideas in a prototype middleware system called HARMONI, and used it to carry out a user study in which we collected heart rate data from four users in two different contexts: in the gym, and while doing light office work. Our evaluation showed that context-aware filtering can provide significant reductions in the uplink bandwidth requirements of the system.

4 STATUS

In my current research, I am trying to reduce the user intervention required in an Interactive Content Adaptation system. My approach is to use machine-learning techniques to find correlations between the adaptation requirements of different objects. When such correlations exist, the user can interactively adapt a single object and achieve fine-grain customization for all the other related objects.

I am also conducting a large scale, real-world deployment of an image fidelity adaptation system. The goal of this endeavor is to learn about the behavior of users performing interactive adaptation on web content outside a lab environment, and over an extended period of time.

ACKNOWLEDGEMENTS

The work described in this document has been done in collaboration with a number of individuals, who I would like to acknowledge and thank. The work on web adaptation is being done at the University of Toronto, in collaboration with Professor Eyal de Lara, Jin Zhang, Jim Cai, Sina Chavoshi and Alvin Chin. For the HARMONI project, my primary collaborators are Dr. Maria Ebling, William Jerome, and Dr Archan Misra of the IBM TJ Watson Lab.

REFERENCES

- [1] M. Satyanarayanan. Pervasive Computing: Vision and Challenges. IEEE Personal Communications, August 2001
- [2] B.D. Noble and M. Satyanarayanan. Experience with adaptive mobile applications in Odyssey. Mobile Networks and Applications, Vol 4 (1999)
- [3] A. Fox, S. D. Gribble, Y. Chawathe, and E. A. Brewer. Adapting to network and client variation using infrastructural proxies: Lessons and perspectives. IEEE Personal Communications, 5(4):10–19, Aug. 1998.
- [4] R. Chakravorty, A. Clark, and I. Pratt. GPRSWeb: Optimizing the Web for GPRS Links. MobiSys 2003
- [5] E. de Lara, D. S. Wallach and W. Zwaenepoel. Puppeteer: Component-based Adaptation for Mobile Computing. 3rd USENIX Symposium on Internet Technologies and Systems (USITS). San Francisco, California. March 2001.
- [6] Y. Dotsenko, E. de Lara, D. S. Wallach, and W. Zwaenepoel. Extensible adaptation via constraint solving. In Proceedings of the 4th IEEE Workshop on Mobile Computing Systems Applications, Callicoon, New York, June 2002.
- [7] I. Mohomed, A. Chin, J. Cai, and E. de Lara. Community-Driven Adaptation: Automatic Content Adaptation in Pervasive Environments. WMCSA 2004
- [8] I. Mohomed, J. Cai, and E de Lara. URICA: Usage-awaRe Interactive Content Adaptation for Mobile Devices. EuroSys 2006
- [9] I. Mohomed, J. Cai, S. Chavoshi, and E de Lara. Context-Aware Interactive Content Adaptation. MobiSys 2006
- [10] D. Husemann; C. Narayanaswami and M. Nidd, Personal Mobile Hub, 8th International Symposium on Wearable Computers (ISWC 2004), 31 October - 3 November 2004, Arlington, VA, USA