Experience-Based Network Resource Usage on Mobile Hosts

Arjan Peddemors[♦]*, Henk Eertink[♦] and Ignas Niemegeers*

[♦]INCA Group, Telematica Instituut
The Netherlands
{Arjan.Peddemors, Henk.Eertink}@telin.nl

* WMC Group, EWI Faculty, TU Delft
The Netherlands
I.Niemegeers@ewi.tudelft.nl

ABSTRACT

The network resources available on a mobile device may be better employed when taking into account the history of mobility patterns of the user. We propose an experience-based mechanism for network resource usage by applications on mobile hosts. It incorporates a machine learning component that adapts the network resource usage given a high-level metric, typically consisting of system as well as application related sub-metrics.

1. INTRODUCTION

An increasing number of mobile devices have multiple network interfaces, allowing them to obtain network connectivity in a concurrent fashion. With the user moving around, the availability and characteristics of network resources on these devices – especially those based on short range wireless links – change over time. Some applications running on these devices have very straightforward preferences in terms of network connectivity and may easily select the most suitable available path to match these preferences. A web browser, for instance, may select an outgoing path primarily based on (monetary) cost of generated traffic and on expected speed to establish its TCP connections to the selected web servers. This choice is likely to remain valid during the existence of such a connection, as it is typically short-lived. Other types of applications have, however, requirements and preferences that are less obviously translated into choices on network resource usage. For example, a monitoring application may need to upload sensory readings to a healthcare center on a regular basis for further analysis, and may accept a degree of data staleness if that reduces the cost of transmission. In that case, connections must be flexible in the amount of data carried during a certain time frame, and the application must balance the different performance aspects as network resource availability and capacity fluctuate.

An indicator such as staleness illustrates that *planning* plays an important role in the optimization of network resource usage. If currently only a low-speed, high-cost path is available, and it is likely that on short notice a high-speed, low-cost network will be in range, it is better to postpone upload until that time. This kind of planning obviously highly depends on the mobility patterns of the user who carries the mobile device and is only possible when these patterns have been *experienced* and gathered in the past. These patterns consist of changes in parameters that describe network resources as seen from the perspective of the mobile host. This information is cross-layer: available networks including technology specific parameters such as identifiers and signal strengths of in-range access points and base stations, currently activated links, currently established IP settings for these links, available network paths (routes), etc.

Furthermore, it is likely that any observed correlation between the values of different network parameters can contribute to a more optimal usage of network resources. For example, power preservation is often an important consideration on a mobile host. The continuous scanning for available 802.11 networks may consume much power and could be temporarily halted when in range of GPRS cells with certain IDs, because *experience* in the past has learned that it is not possible to associate with any 802.11 network in that case – possibly because the user is driving on the highway from home to work.

We propose an experience-based mechanism for network resource usage by applications on mobile hosts. It incorporates a machine learning component that adapts the network resource usage given a high-level metric, typically consisting of system as well as application related sub-metrics. Our proposition is that a learning, experience-based approach results in a substantially better network resource employment when compared with straightforward logic that does not take historical data into consideration. We want to conduct experiments with this mechanism on an existing mobile device platform with real users to find out potential drawbacks and penalties for a number of machine learning alternatives, such as required time to learn and processing overhead. Our existing Network Resource Model (NRM) and Network Abstraction Layer (NAL) are used as a basis for the proposed mechanism (see [2] and [1]).

2. MECHANISM OVERVIEW

We have defined a number of criteria for our learning mechanism. It must be capable of dealing with on-line input which means that data is seen little at a time and not stored for future reference. On a mobile host with continuous updates of network resources, this input would quickly become too large over time. Furthermore, we require the mechanism to permanently learn, so that it may adapt

CoNEXT'06, December 4-7, 2006, Lisboa, Portugal.

if the user changes his mobility pattern over time. We require it to learn in an unsupervised manner.

The current view of the experience-based mechanism is inspired by the concepts introduced by reinforcement learning [4], because this machine learning approach matches well with our criteria. It defines a continuous process that receives input from the NAL and the application, and that takes actions to influence available network resources as well as application operations. It is altered by the metric feedback: it learns by taking into consideration the influence of an action on this feedback. This process is depicted in figure 1. The NAL provides information on cross-layer network resource entities and their inter-relationships. It allows for the activation of network interfaces, links, scanning operations, etc. The application provides information on, for instance, staleness of data and allows the learning process to initiate the upload.

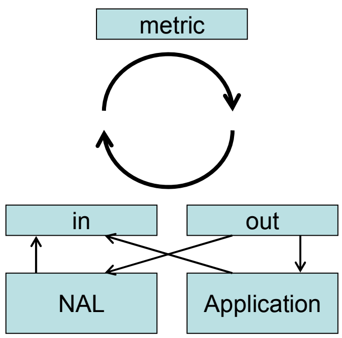

Figure 1: learning process receiving input from and taking action on the NAL and the application.

The metric may consist of a number of sub-metrics. For instance, it could be composed of a cost of network usage metric, a power consumption metric, and a metric for application sensory data staleness. This high level metric defines the target for optimization. Reinforcement learning is applied to various networking and telecommunications problems [3]. The process itself may be implemented using different algorithms, each with their own characteristics.

An important aspect is the balance between 'exploitation' and 'exploration': in order to learn new patterns, the process must allocate a part of the time and resources to explore new possibilities for optimization. This means that, at times, the algorithm must select actions that are – given the current experience – not the most optimal. Another crucial aspect is the bootstrap learning necessary to get to a point with reasonable performance. If this takes too long, the mechanism does not offer a practical solution. The experiments must provide an indication whether these penalties are acceptable.

3. EXPERIMENTS

We have planned two experiments with users carrying a mobile device. The first experiment focuses on gathering mobility patterns by simple running the NAL on a Windows CE device for a number of weeks with a group of approximately 10 office workers. This provides a detailed impression of observed 802.11 networks, cellular networks, Bluetooth devices, fixed USB links, IP configuration, etc. The objective is to have the users carry the device with them during their daily activities as much as possible – not only at the office but also during traveling, at home or elsewhere. It is likely that this generates different patterns for different users. These measurements will be used to experiment with candidate learning algorithms in an artificial environment. The best candidate is then used to execute the second experiment with a real application to validate the experience-based approach.

4. ACKNOWLEDGEMENTS

This research is supported by the Dutch Freeband Communication Research Programme (AWARENESS project) under contract BSIK 03025.

5. REFERENCES

[1] CoSphere NAL: Network Abstraction Layer Software for Windows CE, Retrieved on 29/9/2006 from http://cosphere.telin.nl/nal/

[2] A. Peddemors, I. Niemegeer, and H. Eertink, "An Extensible Network Resource Abstraction for Applications on Mobile Devices", Submitted to the *Second International Conference on COMmunication System softWAre and MiddlewaRE (COMSWARE'07)*, Jan. 2007

[3] L. Peshkin and V. Savova, "Reinforcement Learning for Adaptive Routing", In *Proceedings of the International Joint Conference on Neural Networks (IJCNN'02)*, 2002

[4] R. Sutton and A. Barto, "Reinforcement Learning", MIT Press, 1998