21

Effects of User Mobility on a TCP Transmission

Anne Fladenmuller[1], *Ranil De Silva*[2]
[1] *School of Electrical Engineering,* [2] *School of Computing Science*
University of Technology, Sydney
P.O.Box 123, Broadway, NSW 2007, Australia
e-mail: {anne,ranil}@eng.uts.edu.au
Tel:+61 2 9514 2460 Fax:+61 2 9524 2435

Abstract

This paper studies Mobile IP performance during handoffs. We have conducted experiments with TCP and mobile IP to observe the effect of handoff on transmission reliability. This paper shows that although Mobile IP may be appropriate for current applications, its long handoff periods make it unsuitable for the future.

Keywords

Mobility, Handoff, Mobile IP, TCP

1 INTRODUCTION

With the increasing number of portable computers and the development of wireless networks, mobile computing has become more popular. Although there are still many technical problems to obtain seamless mobility, data can now be transparently transmitted to the host independently of the location given by its IP address. This is achieved through mobile routing protocols like Mobile IP.

With such mobile protocols, users can connect their laptop on diverse networks and receive data without changing their IP address or modifying their network configuration. They can also move between different networks without having to restart their applications. Therefore, when a mobile host switches networks, a handoff occurs to adjust the mobile routing functionality. Several protocols provide such services and all require fast handoffs.

Basing our work on the mobile IP standard, we have conducted some experiments, with varying network conditions, observing the effect of handoffs on higher layer protocols like TCP.

Performance of Information and Communication Systems U. Körner & A. Nilsson (Eds.)

2 MOBILE COMPUTING

Mobile computing introduces a number of problems to current Internet working standards. We identify three main problems dealing with mobility in the Internet - mobile routing, wireless protocol support and mobile application support.

1. The first problem deals with the IP address which has two uses in the Internet - to provide a unique identification for a network interface and secondly to provide routing information about this interface. When a computer becomes mobile, the IP address is still used to identify the network interface but it no longer indicates the location of the mobile computer. This causes normal Internet routing to fail. A number of different techniques have been proposed to solve this problem - one such solution is the Mobile IP standard.
2. The second problem involves protocol support for wireless networking technologies which have been integral in the development of mobile computing. Wireless network technologies have different network characteristics to fixed networks and traditional protocols result in poor performance when operating over wireless networks.
3. The third problem identified with mobile computing is support for mobile applications. Mobile computers currently are likely to be disconnected from the network for large periods. This could be the result of power saving measures on power-limited computers or due to a lack of network connectivity when moving. Mobile applications must be able to survive this and to maximise time when connected on the Internet. In addition, with mobility, applications will evolve and will have to support new services like location dependent behaviour. For this to occur there must be proper services present to aid application development.

This paper is focused on the first problem of mobile routing and in particular one facet of it - the handoff that occurs when a mobile computer moves between different networks.

3 MOBILE ROUTING PROTOCOLS

The aim of mobile routing protocols is to hide the movement of the mobile host to the upper layer protocols and applications. A number of solutions have been proposed all based on the principles of relaying packets from the home network to a foreign network before passing the packets to the mobile host. Routing in the opposite direction is assumed to take the normal routing path.

Mobile IP (IETF 1996) achieves the re-routing through encapsulation of IP headers. If the mobile host is at its home network, then packets can be routed to it using normal routing. If the mobile host moves to a foreign network, the mobile host registers with its home agent to forward any packets addressed to the mobile host via the foreign agent. The packets arriving are encapsulated in a new IP header and sent to the foreign agent. The packet is routed through the network using this new header. At the foreign agent, the new header is removed and the packet is sent to the mobile host. When the mobile host returns to its original network, it deregisters with the home agent and packets are again routed to its normal location.

Many of the alternate solutions also attempt to solve wireless networking problems at the same time. For example I-TCP (Bakre *et al.* 1995) uses a non standard Mobile IP implementation developed at Columbia University to solve mobile routing and at the same time uses two separate TCP links to provide different services for varying wireless and fixed environments.

A different solution for handling mobile routing is Snoop (Balakrishnan *et al.* 1995) which attempts to use multicast addresses to hide the location of mobile computer.

4 NETWORK SWITCHINGS

When the mobile computer moves into a new network, mobile routing services will have to change to reflect this. These changes generally require an exchange of packets called a handoff and during this period, normal transmissions to the mobile host are disrupted. We are interested in the effects of mobile routing handoffs on transport protocols.

Handoffs take place at two levels. The first is the lowlevel handoff that involves the mobile host moving to a new network. In terms of a fixed network, this may consist of plugging the mobile computer onto the network while in an wireless network environment it may simply consist of moving into a new cell. The second level involves the mobile routing handoff, that detects the mobile host has moved into a new network and handles changes to redirect traffic to the mobile host.

It has been shown in (Caceres *et al.* 1994) that handoffs* have a negative effect on TCP performance over wireless networks. Packets get lost during the switching of networks and that triggers the TCP congestion control algortihm. As a result the transmission throughput is decreased and the performances drops.

In order to determine whether this degradation of performance is due to the wireless link or more generally to a special network configuration, we compare the effects of mobile IP handoffs under different conditions. In the next section we describe the results of our experiments to show the effect of handoffs over wireless and fixed networks. Thus we intend to show that it is not possible to improve the performances without modifying TCP, mobile IP or both.

---

*the non-standard Columbia implementation of Mobile IP was used

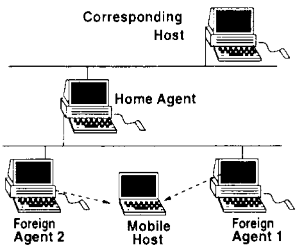

Figure 1 Testbed configuration

5 EXPERIMENTS

5.1 Network configuration

Our experimental testbed consists of a mobile host, two foreign agents and a home agent deployed in a normal office environment as shown in Figure 1. The PCs (486s and pentiums) used for these tests were running Linux (version 2.0.30) and the Mobile IP v1.0 developed at the State University of New York, Binghampton (Dixit *et al.* 1996). We have chosen this implementation as it complies with the IETF Mobile IP draft (revision 16). Tcpdump is used to observe the data transmission during the Mobile IP handoff.

The home and foreign agents and the corresponding host are all connected through fixed LANs (Ethernet 10Mbps). The mobile host can be connected to the foreign agent using wireless links (WaveLan 2Mbps) or fixed networks (Ethernet 10Mbps).

5.2 Wireless vs Fixed Network Handoffs

The experiment done with wireless links between the mobile and the foreign agents is shown in figure 2. This graph presents the packets sent and the acknowledgments exchanged during the transmission. The dotted line corresponds to the registration phase occurring between the foreign and the home agent. Similarly the figure 3 represents the experiment done in a fully wired environments.

In both cases, it takes about 3 seconds for the transmission to be normally reactivated although the registration phase between the home and the foreign agent takes only 0.5 seconds. This difference can easily be explained as the handoff period consists of several operations of which one is the registration. The low-level handoff during which the mobile host is disconnected from both networks lasts nearly 1 second for both cases. The discovery period for the mobile host to detect that it has moved into a new network can take up to a second as this is how often advertisements are sent by foreign agents.

Adding the disconnection, the discovery and the registration times, we ob-

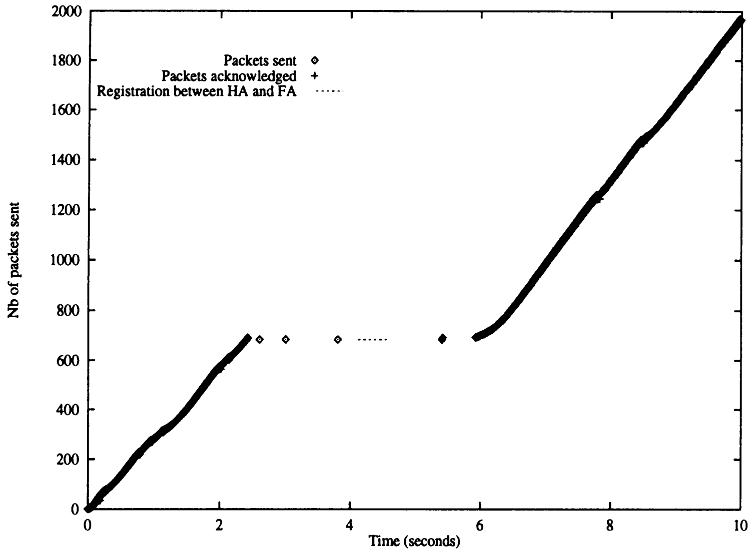

Figure 2 Handoff in wireless environment

tain a handoff period of 2 to 3 seconds for both experiments. But once the handoff is finished we can notice the transmission does not immediatly recover. This delay, of nearly one second, is the result of TCP's congestion control mechanisms : the exponential backoff and slow start algorithms.

In TCP (Stevens 1994), each packet has to be acknowledged to guarantee the reliability. If after a certain time* the acknowledgment has not been received, the packet is retransmitted. To prevent network congestion, the timeout value is doubled for each unsuccessful retransmission. This behaviour is called the exponential backoff and can be observed on both figures just after the disconnection at T=2.5s. After the disconnection, it is necessary for a data packet to be correctly received to resuscitate the connection. Due to successive timeouts occurring during the handoff period, the exponential backoff algorithm results in long delays before retransmitting a data packet. Therefore after the registration, there can be a period of no activity until a retransmission occurs.

Furthermore the slow start algorithm is designed to prevent TCP from transmitting its full window size when the underlying network is congested. It is based on the assumption that if a packet is lost during transmission, it is due to congestion and as a result TCP immediates reduces its current window size. It can be observed from the curves on both graphs that rate of transmission is slower after the handoff.

As a conclusion of these experiments, we have shown that the use of a wireless link does not increase the handoff time. However, if one of the mobile IP

*RTO: Retransmission TimeOut value is approximately 3 x roundtrip time

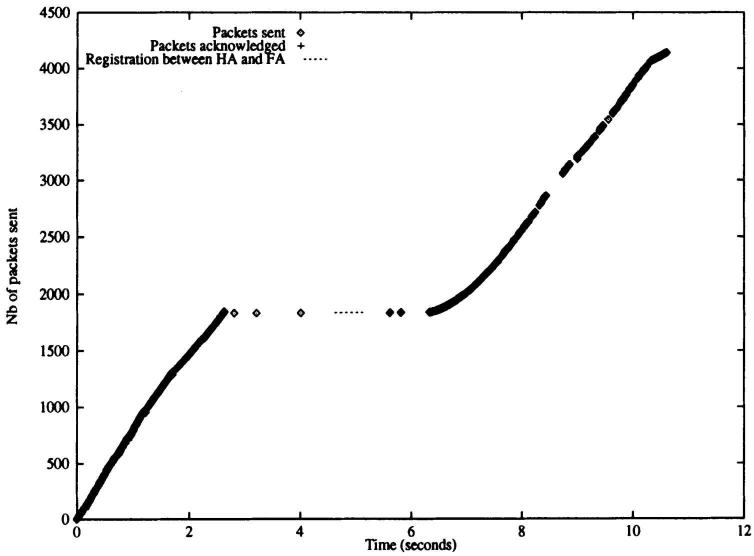

Figure 3 Handoff in wired environment

registration message is lost due to the poor link quality* then the handoff for a wireless link might be longer. The only difference between the experimental results shown is the throughput of the transmission.

An important issue to raise is the poor performances of TCP. One third of the handoff time is due to the unsuitable congestion control algorithm of TCP. In order to avoid slow start, it is necessary that the sender receives an acknowledgment before the timeout occurs. In local area networks, due to the short roundtrip delays and given the handoff period, it is unlikely to prevent timeouts from occurring. It is then important to check if the roundtrip time can in some circumstances be greater than the handoff time.

5.3 Increasing the roundtrip time

We have conducted similar experiments as in section 5.2, but we have chosen the corresponding host in order to increase the roundtrip time and hence the retransmission timeout value to help avoid TCP's slow start. For our experiments, the corresponding host was in France while the rest of the computers remained in Australia.

The results obtained in figure 4 show that the throughput is 100 times lower than the one obtained in figure 3. In this configuration, we had a timeout time of roughly 1.5 seconds, where as the handoff period remained at 3 seconds. Since the timeout time was less the handoff period it was not possible to avoid

*This occured only once during our experiments.

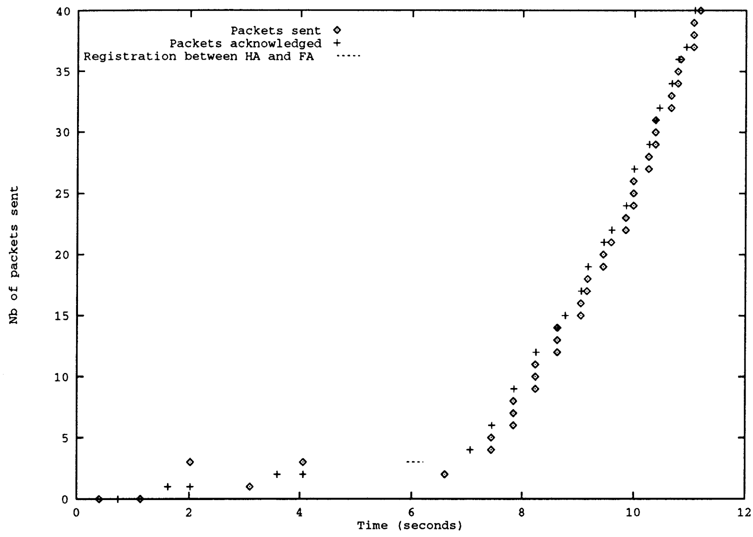

Figure 4 andoff when transmitting in a wide area network

the slow start. We believe that our choice of network configuration is about the worst that can be achieved currently on the Internet. This suggests that it may not possible to avoid slow start during a Mobile IP handoff in any realistic network.

6 CONCLUSION

The work presented in (Caceres *et al.* 1994) shows that Mobile IP handoffs degrade transmission over wireless links. We have extended this conclusion to wired links and we have shown the congestion control algorithm would in any case be triggered. Buffering packets at the foreign or home agent may appear as a good solution to reduce loss during handoff but it can not prevent timeouts occurring thus triggering the slow start algorithm.

Although current applications may not be adversely affected by Mobile IP handoffs, the problem will become more significant in the future. As users become more mobile, the frequency of handoffs will increase. In a pico-cell environment, if handoff takes too long, user may reach the next cell before its completion. Therefore it will be necessary in the future to improve handoff performances which will require the modification of both Mobile IP and TCP.

REFERENCES

Bakre, A. and Badrinath, B. (1995) Indirect TCP for mobile hosts. In *Proc. of 15th Internation Conf. on Distributed Computing Systems*, May.

Balakrishnan H. Seshan S. Amib E. and Kratz R. (1995) Improving TCP/IP performance over wireless networks. In *Proc. of 1st Intl. ACM Conf. on Mobile Computing and Networking (MOBICOM)*, November.

Caceres R. and Iftode L. (1994) Improving the performance of reliable transport protocols in mobile computing environments. In *JSAC, special issue on Mobile Computing Networks*, 1994.

Dixit A. and Gupta V. Mobile IP for Linux (version 1.00). Technical Report, Dept. of Computer Science, State University of New York, Binghampton.

Internet Engineering Task Force IP Mobility Support. Technical Report, April.

Stevens W. R. TCP/IP Illustrated Volume 1, Addison-Wesley, 1994.

7 BIOGRAPHY

Anne Fladenmuller received her Ph.D. from University Pierre et Marie Curie in 1996. She then completed a post doctorate at the School of Electrical Engineering, University of Technology, Sydney where she has since taken up the position of Lecturer. Her research interests include mobile communications, network protocols and QoS Management.

Ranil De Silva received his B.CompSci(Hons) from Bond University in 1993 and has just submitted his Ph.D. thesis at the School of Electrical Engineering, University of Technology, Sydney (UTS). He is currently employed by the School of Computing Science, UTS as a Research Assistant. His main research interests are in adaptive protocols and mobile computing environments.