

A Client-based Handoff Mechanism for Mobile IPv6 Wireless Networks

Leo Patanapongpibul
Laboratory for Communication Engineering
University of Cambridge
Cambridge, United Kingdom
lbp22@eng.cam.ac.uk

Glenford Mapp
Formerly of AT&T Laboratories - Cambridge Ltd
now at Middlesex University
United Kingdom
gsmapp@ntlworld.com

Abstract

Mobile IP enables mobile computers to roam transparently in any network. However, the current proposed protocol specification does not support a suitable handoff mechanism to allow a mobile computer to change its point of attachment from one network to another. This paper describes a handoff mechanism for a mobile host which makes use of Internet Protocol version 6 (IPv6) and Mobile IP without the need to introduce a new mobility management protocol or make changes to the network infrastructure.

1. Introduction

The proliferation of mobile computers has created a need for transparent mobility. Internet Protocol version 4 (IPv4) is widely used in all networks but is a relatively old protocol originally designed for wired networks. With the advent of wireless computing, new problems have emerged which challenge the capabilities of IPv4.

Over the years, the research community has introduced new methods to overcome these problems and support mobile networking. Perkins [12] introduced Mobile IP for IPv4 to support mobile hosts roaming away from their home network domain, thereby allowing them to retain active network sessions without having to restart their network services.

In the first Mobile IPv4 proposals, there were problems with triangular routing, security and other wireless networking issues, including the need to add new components to the IPv4 network infrastructure. The IETF Mobile IP working group was created to solve these problems and refine the protocol. Mobile IP for IPv4 is now an Internet standard (RFC3344) whereas Mobile IP for IPv6 is on course to becoming a standard.

2. Motivation

Since IPv6 was designed to replace IPv4, considerations for introducing new functionalities and improving on IPv4 were taken into account. IPv6 routers have built-in functions eliminating the need for a Foreign Agent. Triangular routing and tunneling required for Mobile IPv4 can now be avoided through IPv6 routing headers. Security problems are intrinsically solved with improved addressing architecture and scalability issues are overcome with its 128-bit address space.

Despite all the benefits from IPv6, Mobile IP still needs some minor refinements. One such refinement tackled in this paper is the handoff mechanism required when a mobile host wishes to roam. The IEEE 802.11b wireless LAN standard is used as a case study to demonstrate a new client-based handoff mechanism for Mobile IPv6.

Currently it is not possible for mobile computers to transparently roam away from its home administrative domain with today's IEEE 802.11b wireless LAN products. There are two solutions to overcome this problem, one method is to incorporate a protocol in the wireless access points or routers to assist mobile hosts with seamless handoffs as they move from one point of attachment to another, i.e. at the subnet and network domain scale. The second solution is to allow the mobile host to decide when it should move over to a new point of attachment. The later solution avoids the problem of interoperability between different vendor's products. The handoff mechanism we proposed is based on this solution.

This paper introduces a handoff policy and a Mobile IP registration enhancement to efficiently handle mobile host handoffs with minimal disruption to user traffic. The new handoff mechanism is evaluated experimentally and compared to traditional Mobile IP handoffs.

3. Related Work

Mobile IP enables mobile computers to roam seamlessly in different administrative domains. However, the protocol has its disadvantages when mobile hosts rate of handoff increases within the foreign administrative domain and from one foreign domain to another. The protocol broadly describes movement detection based on Router Advertisements (RA) as an indication of when to perform a handoff and also suggests that the mobile host can implement its own policy, i.e. a link-layer “roaming” protocol [7], to help with the handoff decision process. A number of approaches have been taken in order to fulfill the need for such a process and reduce the handoff latency.

One approach taken was the introduction of micro-mobility protocols which were broadly aimed at improving the transparent roaming of mobile hosts at the subnet level of a network domain. A number of these solutions have been proposed since the introduction of Mobile IP. Campbell [2] has written a survey of micro-mobility protocols. Further to this, an IETF working group, called *Seamoby*, was formed to resolve the complex interaction of parameters and protocols needed for seamless handoffs. The two main issues being dealt by this working group are the dormant mode host alerting problem (i.e. paging) [8], and context transfers between nodes in an IP access network (i.e. handoff) [9]. The work proposed in this paper can complement micro-mobility protocols and Seamoby efforts, introducing minimal additions to the mobile host for less signaling overheads.

The IETF MobileIP working group has an Internet Draft proposing a protocol for supporting fast handoffs: FMIPv6 [10]. The protocol aims to reduce the handoff latency caused by the movement detection and the Mobile IP registration process. When a handoff is imminent, the later problem is solved by keeping the mobile host ongoing traffic alive with the current access router while the Mobile IP registration process is carried out with the new access router. Like the MIPv6 draft, FMIPv6 also broadly suggests a trigger for the “Handoff Initiation” which may derive from specific link layer (L2) events or policy rules. These triggers are not specified in the draft and thereby would benefit the work discussed in this paper.

A number of simple methods have been proposed, some modifying network entities, which directly tackle the handoff latency in Mobile IP. These are based on hierarchical or multicast handoff mechanisms. Càceres [1] obtained experimental results for the performance of a minimal hierarchical handoff scheme. This had the advantages of not having the complexity of extending routes or anticipating handoffs to improve the handoff latency. A simple multicast technique proposed by Helmy [5] suggested minor modifications to Mobile IP, and validated the method by simula-

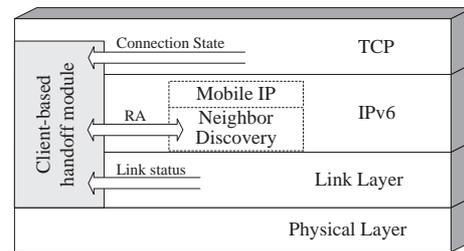


Figure 1. Relative position of the client-based handoff mechanism in the TCP/IP protocol stack.

tion. Both the hierarchical and multicast techniques require modifications to the network entities whereas our work is a simple extension made to the mobile host.

There has been no implementation that avoids the modification of Mobile IP or the standard network entities in the network domain for mobile hosts to roam transparently with a low handoff latency. Our implementation intends to address this problem.

4. Mechanism Overview

There are two ways to provide a suitable handoff mechanism for mobile hosts. The first is to make modifications or add extensions to the entities in the network infrastructure. Routers or base stations can be changed so that they will only send RAs to the mobile host when a handoff is necessary as opposed to periodically sending RAs. However, this means the approximate location and signal strength of the mobile host need to be cached in nearby routers or base stations. Additional signaling may be required in order to enable such a system to operate correctly. This method has the advantage of offering a complete mobility management protocol for the network infrastructure, but has the disadvantage of introducing greater complexity.

The second way is to make modifications or add extensions on the client-side, i.e. the mobile host. In this case, it is the client that decides when a handoff is appropriate. This necessarily implies loss of control to some extent on the network domain’s side. The advantage of this, however, is its apparent simplicity and scalability, which are the reasons why our own mechanism is based on this approach.

We will now describe our mechanism in detail and how it tackles the following issues:

1. Controlling and forcing handoffs
2. Determining the best link
3. Handing off at the appropriate time

The client-based handoff mechanism is illustrated in Fig. 1 as a module in the TCP/IP protocol stack.

4.1. Controlling and Forcing Handoffs

The mobile host initiates a handoff every time it receives a Router Advertisement (RA) from any base station. Our handoff module provides the mobile host with the capability of filtering RAs to avoid the default processing of handoffs. Thus, the handoff module only forces handoff when required.

4.2. Determining the best link

To enable the mobile host to select the best point of attachment, we have introduced a *RA cache* in the handoff module. This provides the mobile host with the capability of choosing the best link from the cache. A policy, based on prioritizing RAs, was devised to assist with the best link decision. The two most important criteria used to determine the priority of the RAs stored in the cache are:

- the link signal strength, i.e. signal quality & SNR level
- the time since the RA entry was last updated

The two less important criteria are:

- the number of hops to the access router
- whether or not the access router is link-local

4.3. Handing off at the appropriate time

Although the application of the aforementioned criteria will generally yield a higher data throughput by handing off to the best link, there are cases when it is advantageous to trade off a potential increase in signal strength against maintaining an active data connection. In order to adopt our mechanism accordingly, the handoff module takes into account the state of TCP connections. Hence, when a handoff is necessary, an open TCP socket will cause the threshold value of the signal strength criterion to be lowered and the handoff to be delayed. In this way, disruptions to TCP connections can be avoided if the difference between the current link quality and the threshold level is minimal. Once the signal strength drops below the lower threshold value or there are no open TCP sockets, the RA Cache entry flagged with the highest priority is passed to the IP packet handler for processing.

The handoff module depends on a link status handler which monitors the link connectivity. This avoids the need to decrease the RA interval in the access routers in order to improve the detection speed of a link disconnection as suggested in previous Mobile IP Internet Drafts [7].

5. The Handoff Process

There are two situations where handoff can be initiated:

Scenario 1: The current mobile host's point of attachment (base station) has a failure or becomes out of range, preventing any data transmission or reception causing the mobile host to perform a hard handoff. This potentially causes some packets to or from the mobile host to be dropped during the process. Handoff is initiated by a handler, which monitors the link status, upon the detection of a link disconnection. The next available RA from the RA Cache is then immediately processed.

Scenario 2: The signal strength of the wireless link between the mobile host and current base station reaches a predefined threshold. A soft handoff is initiated by the link status handler. The handoff module checks for the TCP connection status before deciding to perform a handoff. If there is an active connection, the signal strength threshold is set to a lower value. If there is no active TCP connection or the signal strength is below the lower threshold value, the next available RA stored in the RA cache is immediately processed. There is virtually no packet loss in a soft handoff.

In both scenarios, provided that there is at least one alternative RA in the cache, the RA with the highest priority in the RA cache is forwarded to the IP packet handler for processing. If there is no RA available, handoff is delayed and the IPv6 Neighbor Unreachability Detection is invoked to probe for a point of attachment in the network. This is done by forcing the sending of a periodic Router Solicitation to request for RAs.

6. Implementation

In order to evaluate the system, a testbed was implemented. This was based on a number of desktops and handheld personal computers (mobile hosts) running Linux with IPv6 support and Mobile IPv6 extensions (MIPL) [11]. The decision to use MIPL was based on its completeness and its open source nature as compared to other Mobile IP implementations [3]. The topology of the testbed is shown in Fig. 2. The base stations in the network domain uses the IEEE 802.11b standard, and the wireless interface installed on the Pocket PC is set to promiscuous receive mode. The network in the diagram is a native IPv6 network. The two network domains are indirectly connected through 6BONE [6] facilitating testing of the Mobile IP capability of the network.

In Fig. 2, each base station (A, B and C) belongs to a different subnet. With the existing IPv6 implementation, the mobile host can move between different base stations (not illustrated) in subnet A without any disruption to its network connection. However, when the mobile host moves from, say, base station A to B, the terminal needs to re-establish

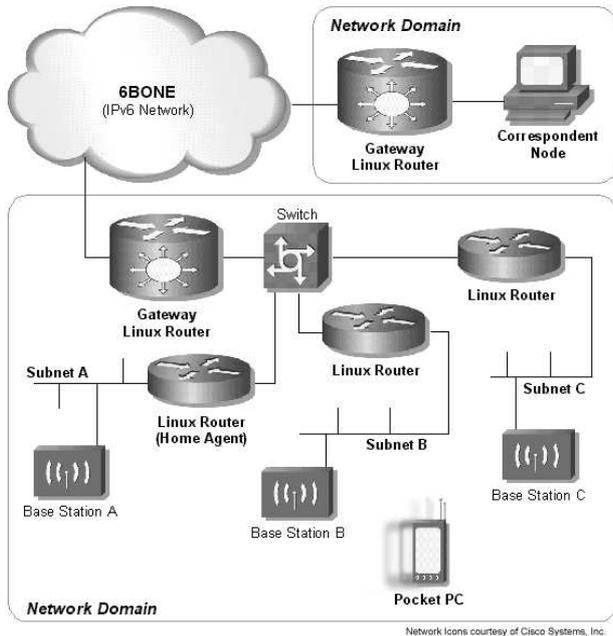


Figure 2. Mobile IPv6 Testbed

network connections – in effect, terminating any active data session such as File Transfer Protocol (FTP) file transfers.

The proposed handoff mechanism was developed as a dynamically loadable Linux kernel driver module for the mobile hosts. The module is tied to the IPv6 Neighbor Discovery protocol to monitor for IPv6 RAs from nearby Home Agents received via the wireless interface, and to monitor the IPv6 packet handler for any active TCP sessions via the same interface. As described previously, these factors are used by the mobile host to make a handoff decision. The handoff extension also restricts the processing of RAs by only allowing the Neighbor Discovery protocol to process RAs when a handoff is required.

6.1. The Experiments

Two experiments were carried out based on both scenarios described above but without FMIPv6.

1. The first experiment was to analyze the effect on the average packet loss when the UDP datagram size and the handoff frequency were varied. The characteristic of the link and the bottleneck at the mobile host was studied to help select a UDP datagram size for the next experiment. Handoff was forced on the mobile host between base stations B and C to achieve a defined handoff frequency in the range of zero to ten handoffs per minute. To achieve a defined handoff frequency, experiments on UDP data streams have shown that hav-

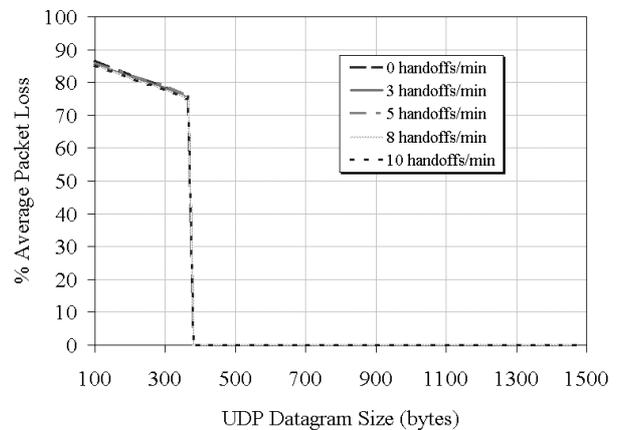


Figure 3. Behaviour of a UDP data stream from the correspondent node to the mobile host.

ing an uniform or a variable handoff interval does not effect the final result.

2. The second test was to observe the effect of the average throughput for a number of handoff frequencies with and without the handoff extension for Mobile IP. The UDP datagram size (excluding UDP, IP and link layer headers) was chosen to be 1424 bytes, based on the results from the first experiment, for optimum link utilization – the IPv6 specification states that the link MTU (maximum transmission unit) must not be greater than 1500 bytes. Similarly to the first experiment, handoff was forced between base stations B and C. Both hard (Scenario 1) and soft (Scenario 2) handoffs were studied for this experiment.

In both experiments a video data stream was set up between the correspondent node (see Fig. 2) and Pocket PC (a mobile host) to monitor the behavior of User Datagram Protocol (UDP) packets in a data stream with a variable handoff frequency. The correspondent node was a UDP source and the mobile host was a UDP sink. The average round trip time between the mobile host and correspondent node was 2.2ms regardless of the mobile host's current point of attachment. The mobile host was tested under the condition to roam outside its home subnet A and forced to handoff only between two foreign subnets B and C at random time intervals.

The access routers were configured to send RAs every 3 to 4 seconds. These were the minimum possible values that can be set according to the Neighbor Discovery for IPv6 RFC document.

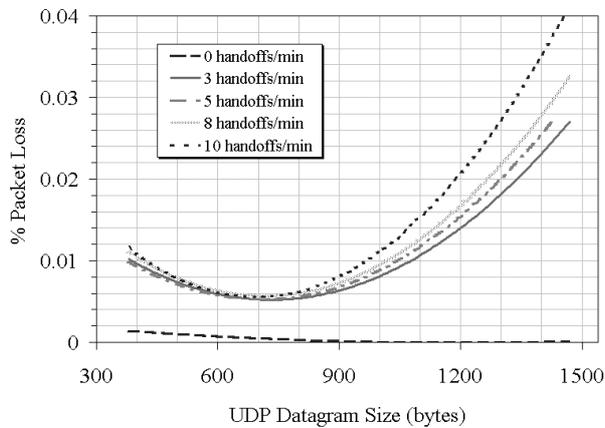


Figure 4. A detailed graph of Fig. 3

The range of the RA interval was also varied above the minimum values but has proven to have no significant effect on the average throughput and packet loss because of the RA cache in the mobile host. There has been previous work [4], with experimental results, which investigated the effect of the RA period on the data throughput. This work has shown that the higher the RA frequency, the less likely the handoff latency will effect a UDP data stream. However, there is a tradeoff between reducing the RA interval and the throughput of the data stream.

Because the experiments were conducted on a testbed, the results were affected by traffic from 6BONE, and by the limited processing power of the mobile hosts needed to process incoming packets.

6.2. Results and Discussion

The percentage packet loss for the corresponding datagram size from the first experiment were plotted and shown in Fig. 3 and Fig. 4. The data points for each handoff frequency were joined by a line to create the graph in Fig. 3. The graph illustrates a high percentage packet loss for UDP datagrams less than 380 bytes in size. This is caused by the limiting processing capability of the mobile host. Packets were generated from the correspondent node at a rate which the mobile host could not maintain.

The percentage packet loss of UDP datagrams between 380 to 1470 bytes is illustrated in a separate graph, Fig. 4. This graph shows second order trend lines drawn through a plot of data points. For a UDP payload larger than 1424 bytes, fragmentation of the packets causes an increase in the packet loss and the reduction of the throughput with higher handoff frequencies. At zero handoffs per minute, there appears to be nearly no packet loss compared to tests with handoffs. Although the handoff frequency was increased,

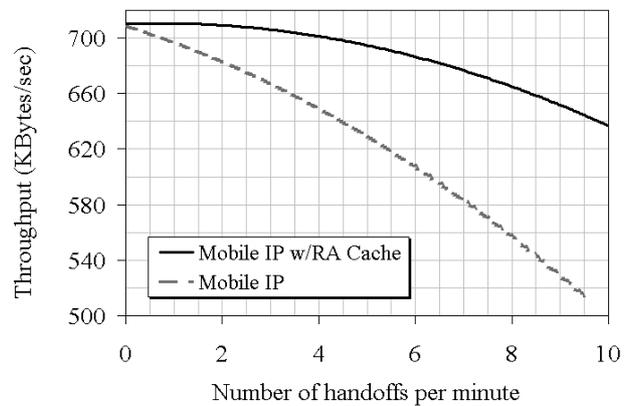


Figure 5. The hard handoff behavior of Mobile IP and Mobile IP with the handoff module for 1424-byte UDP datagrams.

the result did not show any significant packet loss and reduction in throughput. This experiment was carried out to illustrate the mobile host limitations, and a suitable datagram size for UDP data streams.

The results of the second experiment are plotted in Fig. 5. The graph clearly shows an improvement in the handoff latency when the RA Cache in the handoff mechanism is present for hard handoffs. This is because the mobile host perceives the RA interval to be zero as compared to the minimum possible RA interval without the handoff module. The delay in the mobile host receiving the RA with Mobile IP for hard handoffs can be seen from the reduced throughput with increasing handoff frequencies shown by the graph.

Experimental results for soft handoff show that the throughput of Mobile IP with and without the RA Cache are similar. They were averaging at 710 Kbytes/second for any handoff interval. This is because handoff was forced even though the signal strength of base station B and C (see Fig. 2) were above the predefined threshold level.

The Mobile IP with RA Cache graph shows a drop in throughput at higher handoff frequencies. This is caused by the Mobile IP registration time. The latency of this period was measured to be approximately 1.5 seconds. If there was no latency in this process, the throughput should be constant at 710 Kbytes/sec.

The utilization of the link can also affect the experimental results. The speed of the link, determined by the link capacity used, can greatly improve the throughput of the data connection. To verify that the mobile host was using the full link capacity, particularly in the second experiment, since the datagram size was chosen to be 1424 bytes, equa-

tion 1 was used to calculate the link utilization. Because a is inversely proportional to the datagram size L , a larger L will result in a higher U . The distance d of the mobile host from the base stations was within one meter and the data rate R of the link was set to the maximum value of 11 Mbits/sec to ensure a high U . Data transmission was through air, hence V is the speed of light (3×10^8 m/s). The resulting link utilization U , calculated from equation 1, was 100%.

$$U = \frac{1}{1+a} \quad (1)$$

where

$$a = \frac{\text{Length of data link in bits}}{\text{Length of frame in bits}} = \frac{d/V}{L/R}$$

In summary, the graphs show a clear indication of the improvement of Mobile IP with the handoff module. The integration of the FMIPv6 Internet Draft could also mean a reduction in the handoff latency due to the Mobile IP registration process, hence further decreasing the packet losses.

7. Future Work

The handoff protocol described in this paper is part of an effort to allow hosts to roam while away from their home domain without experiencing significant disruption to their network connections. The handoff module and the Mobile IP Fast Handoff (FMIPv6) Internet Draft are important components to support stateless mobile computing and multimedia mobile terminals.

The next step in this research will be to perform an exhaustive analysis of the handoff mechanism and compare the results with FMIPv6. Preliminary results have shown our mechanism reduces the handoff latency of TCP connections by 50%. Post analysis results may introduce additional enhancements which could be made to our mechanism.

8. Conclusion

The client-based handoff mechanism presented in this paper is a simple solution to provide a controlled handoff technique and a reduction in the handoff latency for IPv6 networks with Mobile IP support. The solution offers a decision making mechanism, known as “triggers,” for handoffs and a method to reduce the mobile host dependability on the router advertisement period and router solicitation. The concept of a RA cache has been proven to reduce the handoff latency in our testbed.

The IETF Mobile IP working group proposed additional techniques for faster handoffs [10] to the base Mobile IP protocol by reducing the latency in the Mobile IP registration process. The work described in this paper can complement these extra techniques providing an even faster handoff.

Scalability issues may be a problem if the mobile host can send or receive data to any base stations in the network, but some base station vendors have support for intelligent handoffs of mobile hosts within the same subnet. In cases when the base stations do not have support for intelligent handoffs, the work in this paper can be used to solve such a problem in wireless networks.

9. Acknowledgments

The authors would like to thank James Scott for helpful discussions and comments. Further gratitude goes to Andy Hopper for his guidance. This project is funded by the UK Engineering and Physical Sciences Research Council and AT&T Laboratories.

References

- [1] R. Cáceres and V. Padmanabhan. Fast and scalable wireless handoffs in supports of mobile Internet audio. *Mobile Networks and Applications*, 3(4):351–363, 1999.
- [2] A. Campbell and J. Gomez-Castellanos. IP Micro-Mobility Protocols. *ACM SIGMOBILE Mobile Computing and Communications Review*, 4(4):45–53, October 2000.
- [3] M. Dunmore and C. Edwards. Survey and Evaluation of MIPv6 Implementations. Deliverable 32603/ULANC/DS/4.1.1/A1, 6NET, May 2002.
- [4] N. Fikouras, A. Konsgen, and C. Gorg. Accelerating Mobile IP hand-offs through Extended Link-layer Information. In *Proceedings of the International Multiconference on Measurement, Modelling and Evaluation of Computer-communication Systems*. VDE Verlag, 2001.
- [5] A. Helmy. A multicast-based protocol for ip mobility support. In *Proceedings of NGC 2000 on Networked Group Communication*, pages 49–58. ACM Press, 2000.
- [6] R. Hinden, R. Fink, and J. Postel. IPv6 Testing Address Allocation. IETF Request for comments, RFC 2471, December 1998.
- [7] D. Johnson, C. Perkins, and J. Arkko. Mobility support in IPv6. IETF Internet-Draft, October 2002.
- [8] J. Kempf. Dormant Mode Host Alerting (IP Paging) Problem Statement. IETF Request for comments, RFC 3132, June 2001.
- [9] J. Kempf. Problem Description: Reasons For Doing Context Transfers Between Nodes in an IP Access Network. IETF Request for comments, RFC 3374, September 2002.
- [10] R. Koodli. Fast Handovers for Mobile IPv6. IETF Internet-Draft, September 2002.
- [11] Mobile IP for Linux (MIPL) Implementation by HUT Telecommunications and Multimedia Lab. <http://www.mipl.mediapoli.com/>.
- [12] C. Perkins. IP Mobility Support. IETF Request for comments, RFC 2002, October 1996.