

Henrik Petander, Eranga Perera, Aruna Seneviratne, Yuri Ismailov, An experimental evaluation of Mobile Node based versus infrastructure based handoff schemes, in Proceedings of the 8th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 2007), Helsinki, Finland, 18-21 June 2007, pp. 1-4.

© 2007 IEEE

Reprinted with permission.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of Helsinki University of Technology's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org.

By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

An Experimental Evaluation of Mobile Node based versus Infrastructure based Handoff Schemes

Henrik Petander, Eranga Perera, Aruna Seneviratne and Yuri Ismailov

Abstract— The rate at which the Internet is becoming mobile is unprecedented. This has increased the demand for continuous connectivity even while moving from one network to another at very high speeds. Moving from one network to another gives rise to a handoff process which often incurs packet losses and severe end to end transport protocol performance degradations for the Mobile Node. Most research on IP mobility has focused on minimizing the delays of the handoff process with network infrastructure based approaches. A different way of minimizing the impact of the handoff is to enable the Mobile Node to connect to multiple access networks simultaneously, allowing it to perform Make-Before-Break handoffs. In this paper, we compare the performance of these two alternatives, focusing on the use of Fast Handovers for Mobile IPv6 framework on the infrastructure side and on the other hand Make-Before-Break handoffs using two network interfaces. Both of these schemes require proactive handoffs for optimal performance. The results show that the use of two interfaces for Make-Before-Break handoffs provides increased handoff performance over Fast Handovers for Mobile IPv6.

Index Terms—Mobile IPv6, Make-Before-Break Handoffs, Fast Handovers for Mobile IPv6

I. INTRODUCTION

UBIQUITOUS computing is emerging as an exciting new paradigm with a goal to provide services anytime anywhere. For ubiquitous computing to be a reality, mobile users need to be able to connect to the Internet seamlessly. Seamless connectivity requires that a Mobile Node (MN) can move between networks of the same or different types without impact on on-going communications. This handoff process is not instantaneous and the latency associated with the handoff may result in packet loss. Packet loss leads to application performance degradation and potentially to breaking or resetting of connections. Hence, one of the fundamental requirements for mobile computing is to minimize the impact of handoffs that a Mobile Node performs.

Manuscript received November 19, 2006.

H. Petander, E. Perera and A. Seneviratne are with the National ICT Australia, Bay 15, ATP Eveleigh, NSW 1430, Australia, (phone: +61-2-83745540; fax: +61-2-8374-; e-mail: henrik.petander@nicta.com.au, eranga.perera@nicta.com.au and aruna@nicta.com.au).

Y. Ismailov is with Ericsson research Networks and Systems, Torshamnsgatan 23 - 164 80 Kista, Sweden (email: yuri.ismailov@era.ericsson.se).

Over the years, many researchers have developed methods of mitigating the effects of handoffs in systems where the Mobile Node is connected to only one Access Point (AP) at a time. In such systems, the Mobile Node would have to break the connection to the current network before reattaching itself to the new network (Break-Before-Make handoff). This is due to most Mobile Nodes being incapable of listening to multiple Access Points simultaneously. A Mobile Node that could connect to its current and new access networks simultaneously could perform a Make-Before-Break handoff (MBB).

Use of Make-Before-Break handoffs has up to now been limited to mobile telephone networks, which employ them at the link level. There have been efforts to add this capability to the 802.11 link layer by utilizing two wireless interfaces [1], [2]. Make-Before-Break handoffs on the link layer effectively hide the link layer handoff latency from on going communications. However, when using a network layer mobility management protocol, such as Mobile IPv6 [3], there are other, more significant, components that contribute to the overall handoff latency, as shown in [4]. Further, mobility between different IP networks can not be handled on the link layer alone. Therefore, we have proposed to implement Make-Before-Break handoffs on the IP layer [4]. Employing Make-Before-Break handoffs on the IP layer has the additional advantage of being independent of the underlying access network technology.

The Fast Handovers for Mobile IPv6 (FMIPv6) [5] protocol provides an access technology independent way of emulating Make-Before-Break handoffs by buffering and using a localized forwarding scheme. However, the performance benefits from the emulated Make-Before-break handoffs are offset by the increased complexity required from the network infrastructure.

In this paper, we compare the performance of our previously proposed Make-Before-Break handoff protocol with the emulated Make-Before-Break handoffs of the Fast Handovers for Mobile IPv6 protocol in predictive mode through an empirical performance evaluation in our test bed environment. Based on the experimental results we discuss the applicability of the two schemes. This study augments the understanding of Make-Before-Break handovers at the IP layer.

The rest of the paper is organized as follows. In Section II, we present the background and related work. In Section III, we give an overview of our mobility test bed and the results of

the experimental comparison study. We conclude the paper in Section IV with discussion on the applicability of the schemes.

II. BACKGROUND AND RELATED WORK

A. Fast Mobile IPv6

Fast Mobile IPv6 allows a Mobile Node to perform a predictive handoff from its previous Access Router (pAR) to a new Access Router (nAR), if it can anticipate the handoff event before disconnecting from the pAR. The predictive handoff allows the pAR to start forwarding packets to the nAR at the start of the handoff. The nAR buffers the packets and delivers them to Mobile Node when it attaches to the nAR at the end of the handoff. This procedure allows the Mobile Node to avoid packet loss during a handoff.

In case the Mobile Node loses connectivity with the pAR before performing a predictive hand-off, the Mobile Node (MN) will perform a reactive handoff after connecting to the nAR. In both predictive and reactive mode, the Mobile Node establishes forwarding from its previous Care-of Address on the link of the pAR to its new Care-of Address on the link of the nAR. As a part of the handoff, the pAR and the nAR exchange state information for the Mobile Node, such as quality of service state, network access service state and security associations. This state transfer mitigates the need for the Mobile Node and the nAR to establish the state after the Mobile Node has connected to the nAR. The state transfer together with the localized forwarding scheme reduces the handoff latency even in the case of a reactive handoff.

B. Make-Before-Break handoffs using two interfaces

Advanced wireless network technologies, such as CDMA support Make-Before-Break handoffs on the link layer [6]. However, most wireless networks deployed today, such as IEEE 802.11 WLAN and GPRS do not support Make-Before-Break handoffs. There have been efforts to add this capability to the 802.11 link layer by utilizing two wireless interfaces [1],[2]. However, as discussed in Section I, employing Make-Before-Break handoffs at the IP layer provides better performance. Further, having a scheme which is independent of the access technology, allows it to be used both in vertical and horizontal handoffs. We proposed such a scheme in [4] to enable a Mobile Node to perform lossless Make-Before-Break handoffs. In our proposed scheme, a Mobile Node equipped with two interfaces uses one interface for active traffic and the other one for scanning for networks which can provide better connectivity. When such a network is found, the Mobile Node performs a link layer handoff to the new network using the scanning interface. After finishing the link layer handoff, it attaches to the new Access Router and configures itself with a new IP address while still receiving packets on the old active interface. After establishing IP layer connectivity, the Mobile Node would perform an IP layer handoff, i.e. move traffic over to the new IP address. Finally, the role of the previous active interface will be reverted to scanning for available

networks.

III. EMPIRICAL PERFORMANCE COMPARISON OF MAKE-BEFORE-BREAK AND FMIPv6 HANDOFFS IN WLANS

A. Mobility Test Bed and Implementation

In our test bed, the Mobile Node roams between two visited networks while communicating with a Correspondent Node (CN). The IEEE 802.11b wireless access networks are implemented using two Access Routers, which provide connectivity via two Access Points: a Cisco 1200 series Access Point and a Prism54 based PCMCIA card on the second Access Router in Access Point mode. The two Access Routers run FMIPv6 software [7]. The Access Routers, the Mobile IPv6 Home Agent (HA) and the CN are interconnected using an emulated Wide Area Network (WAN). We use NISTNet [8] to emulate a WAN (Internet) by introducing network latency between the nodes. The MN is equipped with two WLAN interfaces: a Prism 2.5 based PCMCIA card and an integrated Intel IPW2100 card. It runs a modified version of MIPL Mobile IPv6 software [9] which supports IP layer Make-Before-Break handoffs and the fmip6.org MN software [11] in predictive mode.

We use a latency of 40ms between the Home Agent (HA) and the CN and a latency of 40ms between the CN and the Access Routers to emulate cross ISP traffic. Additionally, we introduce a latency of 10ms between the HA and the Access Routers to emulate the case in which the HA and Access Routers are connected via the same ISP.

In all scenarios, the Correspondent Node acts as the source and the Mobile Node as the sink. This emulates a mobile user receiving data from a server. For UDP traffic, we use a Constant Bit Rate (CBR) stream of 100kbit/s consisting of packets with a payload of 100 bytes to measure the received data rate at the Mobile Node. For TCP we use variable bit rate traffic and the default window size under Linux (16kB) to measure the sequence number progression at the Mobile Node which accurately depicts the TCP throughput visible to an application.

B. Performance Comparison of IP layer Make-Before-Break Handoffs with Fast Handovers for Mobile IPv6

In the Fast Handovers for Mobile IPv6 handoffs, the Mobile Node is equipped only with the Intel card and the Prism 2.5 based card is deactivated. In the Make-Before-Break handoff, the Mobile Node initially uses the Intel card for active traffic and the Prism card as the scanning interface. After the handoff, all traffic is via the Prism card.

1) Performance Comparison with UDP

We first measure the UDP received data rate during a handoff, and the results are depicted in Figure 1. During the handoff, starting approximately at $t=2.7s$, the data rate for Fast Handovers for Mobile IPv6 drops down to zero. This is due to the Mobile Node disconnecting from the previous Access Router (pAR) and starting the link layer handoff. During the link layer handoff, the Mobile Node does not receive packets which are delivered to the new Access Router (nAR) and

buffered. After the Mobile Node attaches to the new Access Router, it receives the packets from the buffer at approximately $t=3.5s$. This is visible as a sharp increase in the received data rate. On the other hand, with the Make-Before-Break handoffs, the handoff does not have a visible impact on the received data rate, since the coverage of the previous and new network overlap sufficiently to allow for a loss less handoff. This is a requirement for the predictive mode of Fast Handovers for Mobile IPv6 as well as IP layer Make-Before-Break handoffs. In case of insufficient overlap of the access networks, the Mobile Node would need to revert to the reactive mode of Fast Handovers for Mobile IPv6 or a Break-Before-Make handoff.

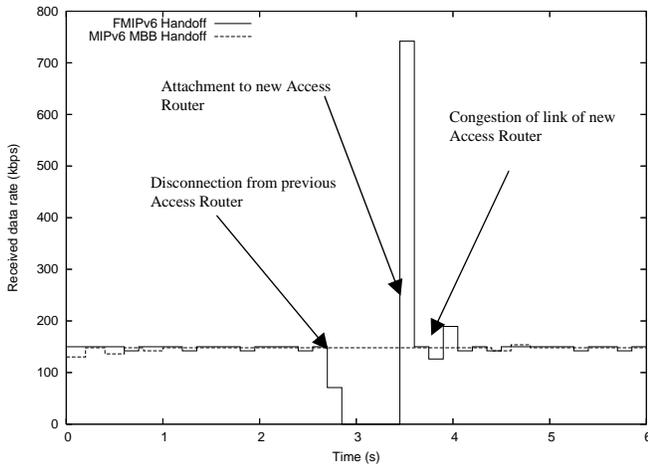


Fig 1. Received data rate comparison for UDP CBR traffic of 100kbps.

2) Performance Comparison with TCP

We compare the TCP performance for Fast Handovers for Mobile IPv6 and IP layer Make-Before-Break handoffs using two different scenarios. The first scenario emulates the case in which the wireless link does not act as a bottleneck for the connection. The bandwidth available to the TCP transfer is limited to 2Mbps between the ARs and the HA and the CN. In this scenario, the Fast Handovers for Mobile IPv6 handoff does not have a large impact on the TCP performance as depicted in Figure 2. At $t=4.7s$, the MN disconnects from the pAR and packets are buffered by the nAR. At $t=5.4s$, the MN connects to the nAR which delivers the buffered packets from the buffer. This results in the temporarily increased rate of TCP approximately between $t=5.4-5.6s$. The increased rate is enabled by the extra bandwidth available on the wireless link of the nAR.

In Figure 3, the Make-Before-Break handoff does not have a negative impact on the TCP performance. In fact, the handoff results in a temporary increase in the TCP performance due to the MN sending TCP acknowledgments via the nAR while receiving TCP data via the pAR. This effect is explained in more detail in [4].

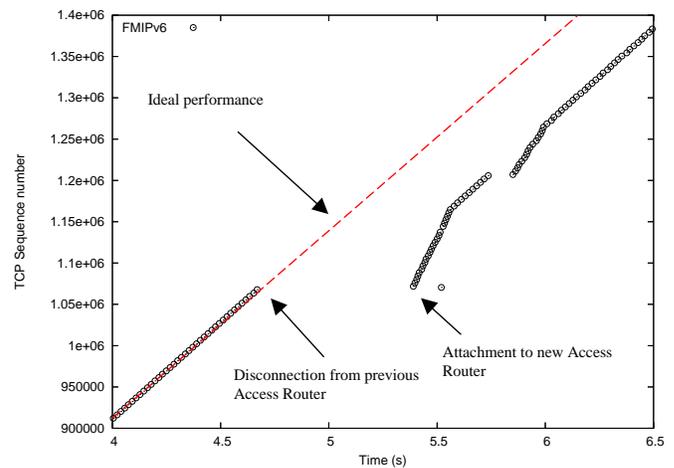


Fig. 2. Fast Handovers for Mobile IPv6 handoff performance for TCP when wireless link is not the bottleneck.

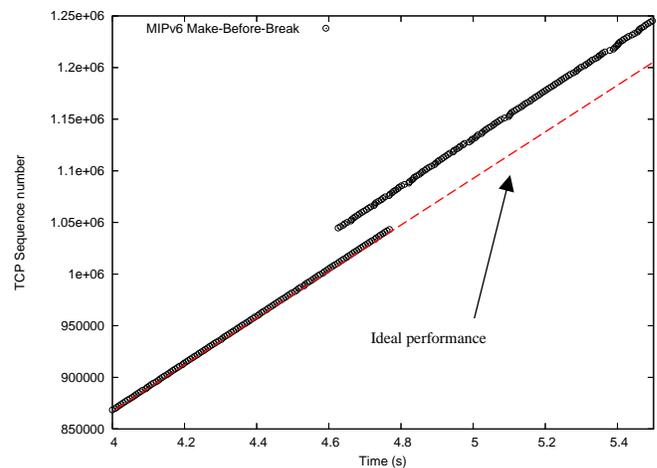


Fig. 3. IP layer Make-Before-Break handoff performance for TCP when wireless link is not the bottleneck.

We evaluate the performance of the two schemes in a second scenario in which the wireless link acts as the bottleneck for the TCP connection. This is depicted in Figure 4. This would typically be the case in many wireless network deployments. The wireless link of the nAR is slower (3Mbps) than that of the pAR (4Mbps) due to increased background traffic.

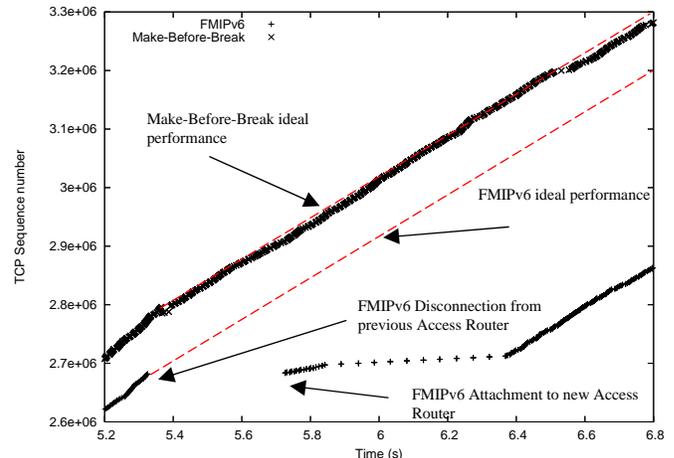


Fig 4. Fast Handovers for Mobile IPv6 and IP layer Make-Before-Break handoff performance for TCP when the wireless link acts as the bottleneck.

The Fast Handovers for Mobile IPv6 handoff results in a TCP slow start due to the congestion on the new link from the emptying of the buffer. The impact of the congestion is increased by the reduced bandwidth available at the new link. In the case of the IP layer Make-Before-Break handoff, the reduced bandwidth leads to 3 TCP segments being resent.

We compare the impact of the handoff on TCP progress in Fig. 5. In the first scenario, the IP layer Make-Before-Break handoff has a positive impact on the progress of the TCP transfer and increases the handoff performance by 129% when compared with Fast Handovers for Mobile IPv6. The TCP rate control algorithm amplifies the effects of the temporary disconnection of 200ms by a factor of 1.24 for Fast Handovers for Mobile IPv6 in this scenario.

In the second scenario, the IP layer Make-Before-Break handoff has a minor impact on the connection due to the congestion on the new link. This impact is still 98% smaller than the negative impact from Fast Handovers for Mobile IPv6. In the second scenario, the TCP rate control algorithm together with the link saturation amplifies the effects of the period of disconnection by a factor of 3.8 for Fast Handovers for Mobile IPv6.

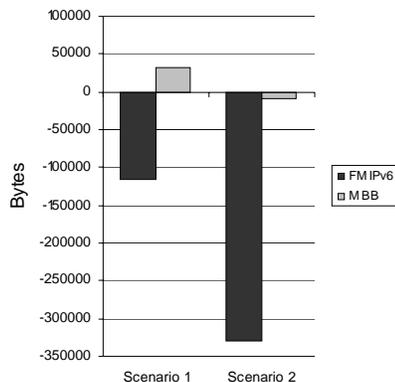


Fig 5. Impact of a handoff on TCP progress.

IV. CONCLUSIONS

In this paper, we compared the performance of Fast Handovers for Mobile IPv6 and IP layer Make-Before-Break handoffs. The empirical evaluation showed that with the IP layer Make-Before-Break handoffs the impact of the handoff is minimal. Fast Handovers for Mobile IPv6 can provide good performance in a network environment, in which there is sufficient link capacity for effective use of buffering. However, when this is not the case, the performance of the protocol degrades considerably when compared with IP layer Make-Before-Break handoffs.

In the case of intra-technology handoffs, a Mobile Node would require an additional interface of the same type for achieving Make-Before-Break handoffs. However, a Mobile Node roaming between multiple access network technologies would be equipped with a radio interface for each technology and could perform inter-technology Make-Before-Break

handoffs without additional hardware capabilities. In the former case of intra-technology handoffs, the additional capabilities would increase the cost and power requirements of the Mobile Node. However, the ability to perform seamless handoffs without infrastructure support may offset this cost. Further, in handoffs between different operators' networks and handoffs in legacy networks, the infrastructure support for handoffs would not exist and Mobile Node based solutions would be the only alternative for Mobile Nodes requiring seamless handoff performance.

A potential application of the IP layer Make-Before-Break handoffs technology would be Mobile Routers serving vehicular networks, which are powered by an external power source and thus are less limited by the increased cost and power requirements of multiple network interfaces. Further, one could argue that the additional cost of an extra radio receiver will decrease with the increasing integration level of radio chipsets.

In conclusion, both Fast Handovers for Mobile IPv6 and IP layer Make-Before-Break handoffs have their strengths and weaknesses. We analyzed the performance of the protocols to discover their applicability to different scenarios.

REFERENCES

- [1] K. Ramachandran et al, Make-before-Break MAC Layer Handoff in 802.11 Wireless Networks, IEEE International Conference on Communications (ICC), 2006.
- [2] V. Brik, A., Mishra, S. Banerjee, "Eliminating handoff latencies in 802.11 WLANs using Multiple Radios: Applications, Experience, and Evaluation", in proceedings of Internet Measurement Conference, CA, USA, Oct., 2005.
- [3] D. Johnson, C. Perkins, J. Arkko, "Mobility Support in IPv6", RFC3775, June 2004.
- [4] H. Petander, E. Perera, K. Lan, A. Seneviratne, "Measuring and Improving Performance of Network Mobility Management in IPv6 Networks", IEEE Journal on Selected Areas of Communications, Special Issue on Mobile Routers and Network Mobility, Vol. 24, No.9, Sep. 2006.
- [5] R. Koodli, "Fast Handovers for Mobile IPv6", RFC 4168, July 2005.
- [6] D. Kim, D. Sung, "Characterization of Soft Handoff in CDMA Systems", IEEE Transactions on Vehicular Technology, Vol. 48, No. 4, July 1999.
- [7] The fmipv6.org implementation of FMIPv6, <http://www.fmipv6.org>.
- [8] M. Carson and D. Santay, "Nist net: A linux-based network emulation tool", ACM Computer Communication Review, 2005.
- [9] MIPL Mobile IPv6 for Linux, <http://www.mobile-ipv6.org>.