

MEXT Working Group
Internet-Draft
Expires: January 11, 2009

K. Wong
MUST
A. Dutta (Ed.)
Telcordia
H. Schulzrinne
Columbia Univ.
July 10, 2008

Simultaneous Mobility Problem Statement
draft-wong-mext-simultaneous-ps-01

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 11, 2009.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

This document provides a problem statement for simultaneous mobility in an infrastructure-based mobility environment. It considers what the simultaneous mobility problem is and describes the conditions under which it may occur. It also illustrates the occurrence of the simultaneous mobility problem in the context of different mobility protocols such as Mobile IPv6. The simultaneous mobility problem defined here is strictly applied to scenarios when the intermediary nodes in the core are not mobile.

Table of Contents

1.	Introduction	3
1.1.	What is simultaneous mobility?	3
1.2.	Likelihood of Simultaneous Mobility	4
2.	Terminology	5
3.	Simultaneous Mobility Scenarios	7
4.	Applicability of simultaneous mobility	10
4.1.	Simultaneous Mobility in Mobile IPv4	10
4.2.	Simultaneous Mobility in Mobile IPv6	10
4.3.	Simultaneous mobility in SIP-based mobility	14
5.	Occurrence of Simultaneous Mobility Problems	17
6.	Security Considerations	18
7.	IANA Considerations	19
8.	Acknowledgments	20
9.	References	21
9.1.	Normative References	21
9.2.	Informative References	21
	Authors' Addresses	22
	Intellectual Property and Copyright Statements	23

1. Introduction

We consider mobility of end (or terminal) hosts in an infrastructure-based network like the global Internet. Simultaneous mobility is the special case when two end hosts are mobile and in communications with each other, and each moves away from one attachment point in the network to another at, or at about, the same time (we give a more precise definition in [Section 1.1](#)). Although it may not be typical (more typical would be the case when one of the two hosts moves and the other hosts within the network remain stationary during that time), such an event will occur in from time to time, and this must be handled properly by mobility protocols. The simultaneous mobility problem, when it occurs, results in the loss of either a binding update from a Mobile Host, or of some other control message prior to the sending of the binding update, because the binding update or prior control message is sent to a previous address of the other Mobile Host that is also moving at around the same time. Without the introduction of additional measures to rescue the communications session, this would be expected to result in the interruption of the communications session.

1.1. What is simultaneous mobility?

We begin by defining some terms that we will use to define the simultaneous mobility problem.

Two Mobile Hosts attached to two different points of attachment in the network are said to be in a communication session if they are actively exchanging data. A communication session may be in a normal state or an interrupted state. A session is in a normal state when data from one Mobile Host is arriving at the right location for the other Mobile Host, and vice versa. It is in an interrupted state otherwise. A handoff is a movement of a Mobile Host from a previous attachment point to a new attachment point, such that the old IP address does not route to the new attachment point, but a new IP address is needed to route there. The handoff time is the moment in time when it changes from being reachable at the previous attachment point to being not reachable at the previous attachment point. Given that time is continuous, it is assumed that only one handoff can occur at any given moment in time, i.e, that handoff times are unique. Two handoffs are consecutive if neither of the Mobile Hosts performs another handoff in between the two handoffs. Consecutive handoffs can be at the same mobile host or between two different mobile nodes.

We assume that binding updates do not contain information about future moves of the sending Mobile Host. While two Mobile Hosts are in a communication session, they get information on the location of

the other Mobile Host only from binding updates. In other words, they do not actively seek the location of the other Mobile Host, but only passively accept binding updates. Binding updates are sent directly to the most current known address (known by the sender) of the intended recipient Mobile Host. In general, the latency for binding updates to arrive is assumed to be much smaller than the average inter-handoff time, therefore, it is extremely unlikely that a binding update would be sent and the recipient Mobile Host would move twice before the binding update arrives at the previous network of the recipient. In some mobility protocols, there might be other control signaling prior to the sending of binding updates, e.g., for security purposes.

Then, the simultaneous mobility problem occurs when there are two Mobile Hosts in a communications session in normal state, and they both move such that one, or both, of the binding updates that they send to each other are lost through belated arrival, or if one or both of the binding updates do not even get sent because of the loss (through belated arrival) of some other control message prior to the sending of binding updates, and such that the return from interrupted to normal state is significantly delayed, or never happens.

For convenience, we also list some of these definitions in [Section 2](#).

[1.2](#). Likelihood of Simultaneous Mobility

Analysis in [[simultaneous-mob4](#)] shows that the likelihood of the simultaneous mobility problem occurring may be non-trivial, depending on the handoff latencies involved, and the frequency of handoffs.

2. Terminology

Communication session: Two Mobile Hosts attached to two different points of attachment in the network are said to be in a communication session if they are actively exchanging data. NB: We do not attempt to define here what "actively exchanging data" means. This may depend on the applications, and may be related to a timer set since the last received packet.

Normal state (of a communication session): A session is in a normal state when data from one Mobile Host is arriving at the right location for the other Mobile Host, and vice versa.

Interrupted state (of a communication session): A session is in interrupted state when it is not in normal state

Handoff: A handoff is a movement of a Mobile Host from a previous attachment point to a new attachment point, such that the old IP address does not route to the new attachment point, but a new IP address is needed to route there. NB: this is sometimes known as layer-3 handoff, in contrast to so-called layer-2 handoff

Handoff Time: The handoff time of a handoff is the moment in time when it changes from being reachable at the previous attachment point to not reachable at the previous attachment point.

Belated Arrival (of Binding Update or other control message): A binding update (or other control message) is considered to make a belated arrival at a network if the destination Mobile Host is no longer attached to that network.

Lost (Binding Update or other control message): A binding update or other control message is lost if it does not arrive at its intended destination. A binding update or other control message is "lost through belated arrival" if it makes a belated arrival and is consequently lost. NB: a binding update or other control message can be lost not necessarily through belated arrival, but through other possible causes of lost messages, e.g., network congestion, node failure, link failure. A binding update or other control message that has arrived late may also be retrieved by some other agent in the network without necessarily getting lost.

Consecutive Handoff:

Two handoffs are consecutive (with respect to a pair of mobile nodes) if neither of the Mobile Hosts performs another handoff in between the two handoffs.

3. Simultaneous Mobility Scenarios

In this section, we describe general simultaneous mobility scenarios in an infrastructure-based environment.

Figure 1 depicts the simultaneous mobility scenario of two mobile hosts in an infrastructure based network. The simultaneous mobility problem arises when Mobile Host 1 (MH 1) moves from one attachment point in domain 1 to another attachment point in domain 2 for example, while at the same (or nearly the same) time communicating Mobile Host 2 (MH 2) moves from one attachment point to another attachment point. The binding update information that must be passed between Mobile Host 1 and Mobile Host 2 will not reach the other prior to their movement.

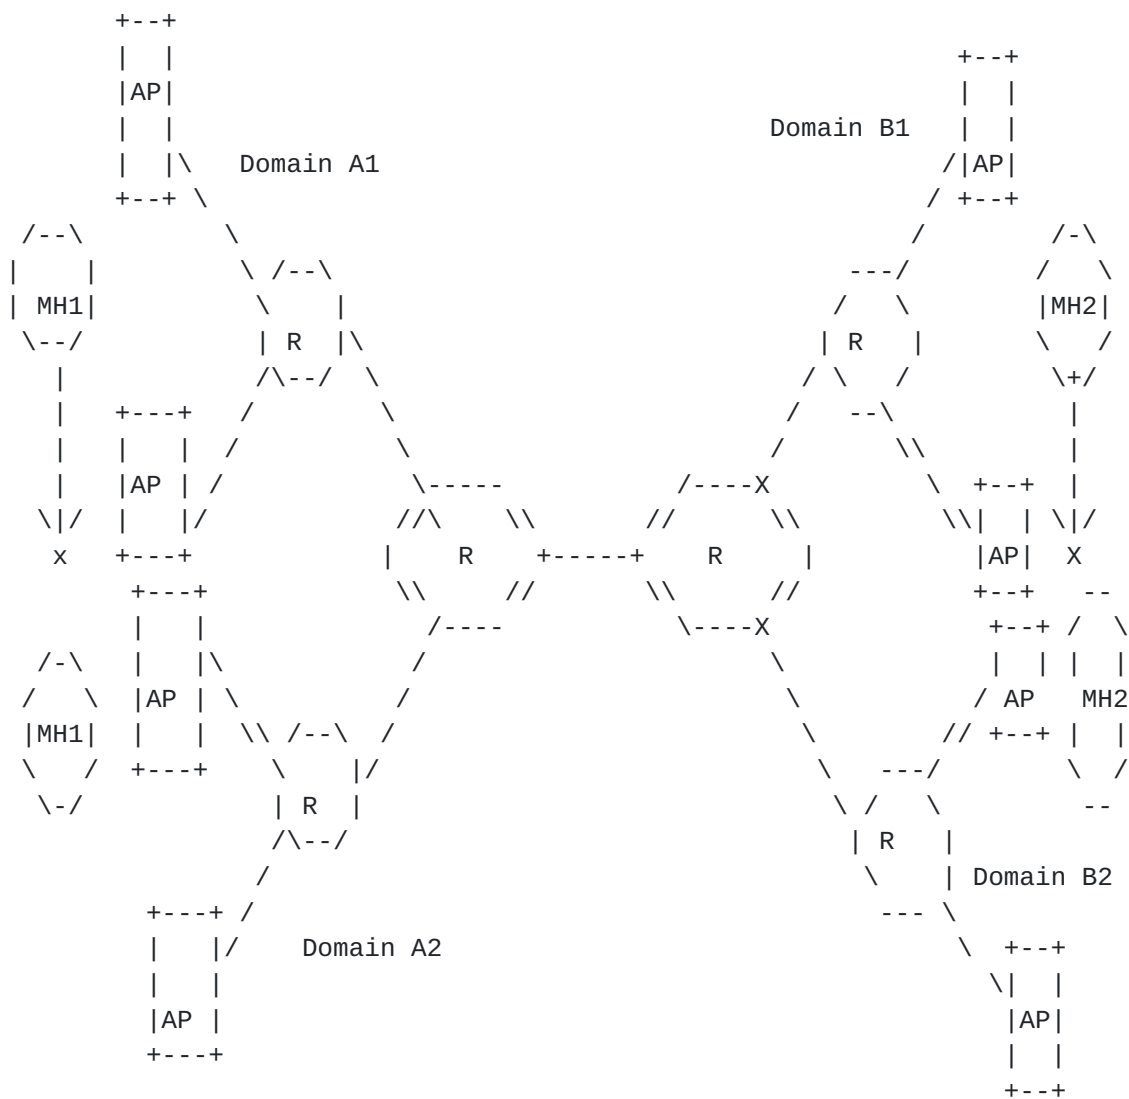


Figure 1: Simultaneous Mobility Scenario

Figure 2 shows the binding updates between two communicating hosts in an infrastructure-based environment. In the figure, Host A moves from domain A1 to domain A2 while host B moves from domain B1 to Domain B2 giving rise to the simultaneous mobility scenario. Under certain conditions, the mobiles may lose the binding updates from each other giving rise to problems encountered by simultaneous mobility.

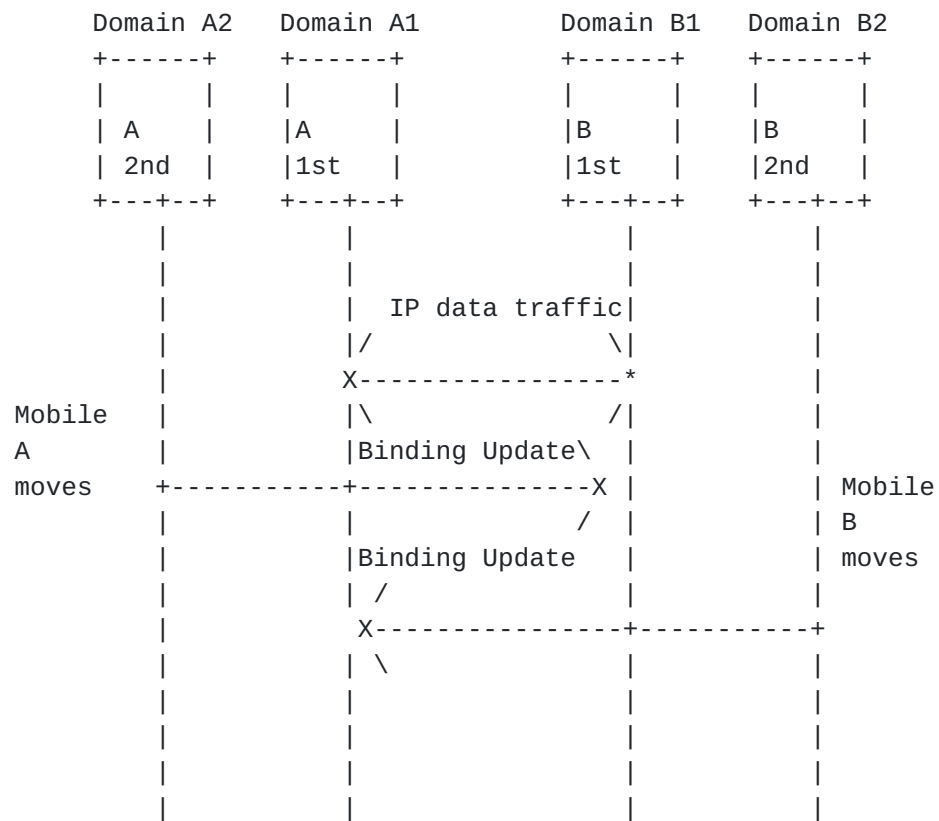


Figure 2: Simultaneous Mobility Flow: General Scenario

4. Applicability of simultaneous mobility

In this section, we describe how general simultaneous mobility arises for two different mobility protocols, such as MIPv6 [[RFC3775](#)] and SIP-based mobility [[SIPMM](#)] are used in an infrastructure-based environment. Both the protocols exhibit certain common features, such as the ability to send binding updates directly to the correspondent nodes.

4.1. Simultaneous Mobility in Mobile IPv4

Mobility extensions for Internet Protocol v4 (Mobile IPv4) does not have a problem with simultaneous mobility. By design, non-moving Correspondent Hosts are unaware of the mobility of Mobile Hosts. The home agent of the Mobile Host functions as an anchor point for the Mobile Host. No matter where the Mobile Host moves, packets for it always go first to its home network for interception and tunneling by its home agent. If it turns out that the Correspondent Host is also mobile, it will also have a home agent and packets from the Mobile Host will similarly be intercepted and tunneled to the appropriate network by its home agent. Since both home agents are stationary and can always be reached through IP routing simultaneous mobility does not present a problem to Mobile IPv4.

4.2. Simultaneous Mobility in Mobile IPv6

In this section, we describe how and under what conditions, simultaneous mobility affects the smooth operation of Mobile IPv6 protocol. In this current version of this draft, we consider only Mobile IPv6 as specified in [RFC 3775](#) [[RFC3775](#)]. We do not currently consider low latency handoff mechanisms, such as in [RFC 4068](#) [[RFC4068](#)].

Unlike Mobile IPv4, Mobile IPv6 has simultaneous mobility problems. The simultaneous mobility problem of Mobile IPv6, however, is somewhat subtle because of an in-built defense against the simultaneous mobility problem. As part of route optimization, binding updates are sent directly to Correspondent Hosts after a handoff. If a binding update is sent to the care-of address of a mobile Correspondent Host, then it is straightforward to see how this binding update might be lost through belated arrival, and how the simultaneous mobility problem would occur. In reality, as a Mobility Header message, the binding update MUST NOT be sent to any care-of address that the Mobile Host might have for a mobile Correspondent Host, but must be sent to the permanent home address of the mobile Correspondent Host. This binding update can then be forwarded by the mobile Correspondent Host's home agent to its current location. Similarly, Care-of-Test Init (CTI) and Home-Test Init (HTI), when

sent to a Correspondent Host that is also mobile, MUST NOT be sent to the care-of address for that Correspondent Host, but both go to the home address, and thus, are forwarded to the Correspondent Host by the home agent of the Correspondent Host. Thus, this serves as an in-built defense against the simultaneous mobility problem, since the home agent of the receiving host is involved. The defense, however, is only partial, because it needs the Home Agent registration of the receiving host to be successfully completed before the control message can be properly forwarded to the receiving host. Therefore, the interval of time in which the receiving Mobile Host is vulnerable to the simultaneous mobility problem is from when it transitions to interrupted state, until it can obtain a care-of address and complete successful home registration.

Figure 3 shows an example call flow for simultaneous mobility problems when applied to MIPv6. In this example, one side is unable to complete return routability because of simultaneous mobility. In particular, A is unable to initiate return routability with the HTI and CTI messages, because these arrive at B's home agent before B's home agent has received B's latest care-of address. Thus, these arrive at B's previous network shortly after B has moved to its latest network. Meanwhile, B performs its own return routability procedure, and this is successful, for the following reasons: (a) since A has attempted to initiate its own return routability procedure, A has already successfully registered its latest care-of address with its home agent; (b) B's CTI and HTI messages as usual are sent to A's home address and are correctly forwarded by A's home agent to A; (c) A sends the Home Test message to B's home address, and A sends the Care-of Test message to B's care-of address that is obtained from the source address of the CTI (as implied by [RFC 3775](#)).

In the worst case, A does not retransmit its lost CTI and HTI messages, and the communications session is broken. In the best case, the Mobile IPv6 implementation in A implements a retransmission scheme. [RFC 3775](#) suggests that such a retransmission might initially use a 1 second timeout (INITIAL_BINDACK_TIMEOUT). Longer timeouts might result in applications timing out and communication sessions thus breaking, even if both Mobile Hosts eventually receive the latest Binding Updates from each other.

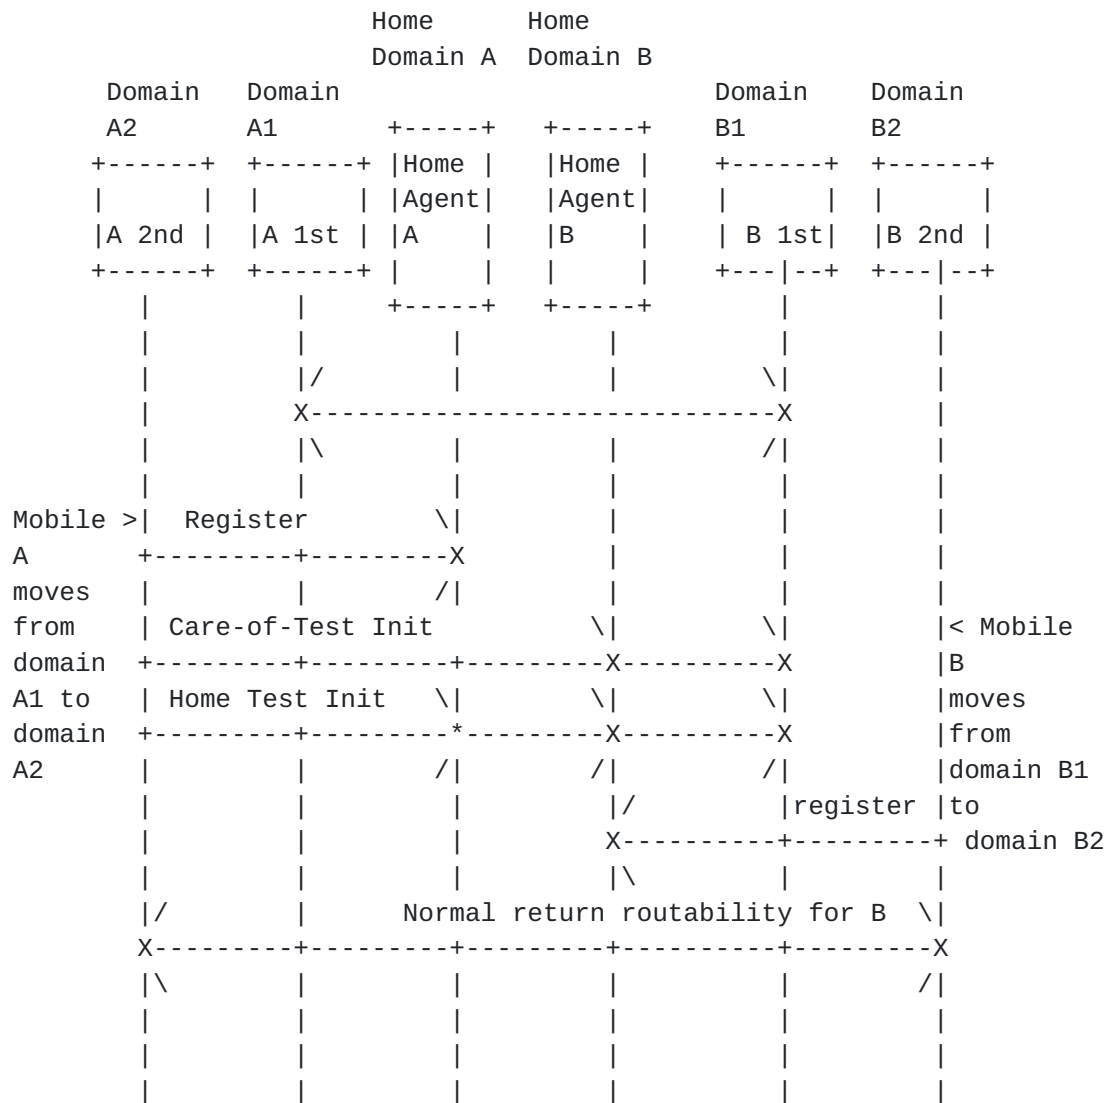


Figure 3: Simultaneous Mobility in MIPv6

In the same scenario, another possible manifestation of the simultaneous mobility problem may occur, where A's binding update is lost instead of its CTI and HTI messages. In this case, B moves a little later, so A's return routability procedure completes as normal, while B is still reachable at its previous network. Then, B moves to the new network, but before B's binding update to its home agent has reached the home agent, A's binding update (to B) arrives at B's home agent, and is forwarded to the previous network where B is no longer reachable, and thus it is lost.

4.3. Simultaneous mobility in SIP-based mobility

In this section we illustrate how simultaneous mobility problem also affects the SIP. Session Initiation Protocol (SIP) is a protocol that enables setting up voice calls, two-way multimedia sessions, teleconferencing, instant messaging and other types of communication using the Internet. SIP is a protocol in the Application Layer of the OSI Reference Model to enable the establishment, management and termination of such real-time communications sessions over an Internet Protocol (IP) network. SIP is defined by [RFC 3261](#) of the Internet Engineering Task Force (IETF). SIP negotiates the features and capabilities of a communication session at establishment of the session reducing time necessary for setting up communication sessions. Mobility of the end host is handled very naturally by SIP using existing SIP signaling mechanisms. For example, to initiate a session, the SIP INVITE message is sent by the initiating device to the other device. The extension of SIP to handle mid-session mobility specifies that when one of the two devices moves, it sends a re-INVITE message to the other device, informing it of its new location (e.g., its new IP address). In addition to the re-INVITE message sent directly to the device that moved (Mobile Host) to the other device (Correspondent Host), the Mobile Host also registers its presence in the new network with a SIP server in its home network. This allows other potential Correspondent Hosts to find the Mobile Host.

Figure 5 depicts the flow of data resulting from a mishandling of simultaneous mobility in SIP. In the basic SIP mobility scheme, when a Mobile Host moves to a new network, it sends a re-INVITE message to the Correspondent Host, as well as a REGISTER message to its home network SIP server. There are two options for the part of the re-INVITE, i.e., it could go through the inbound proxy of the Correspondent Host, or it could go directly to the Correspondent Host. This second option will not work properly if there is simultaneous mobility of the Mobile Host and the Correspondent Host. Clearly, if the Correspondent Host moves at the same time, the re-INVITE may be lost (and similarly, the re-INVITE from the Correspondent Host to the Mobile Host could also be lost). As the home SIP servers for both are stationary and always reachable through IP routing, the registrations are not affected by simultaneous mobility. It would appear the both hosts could now obtain the new location of the other party through the SIP servers, analogous to how the home agent provides such information in MIP. However, in MIP, the home agent quickly discovers that the Correspondent Host needs the updated binding information, because data packets from the Correspondent Host are routed to the home network and intercepted by the home agent.

In the SIP case, on the other hand, the Correspondent Host may not immediately contact the SIP sever. Instead the re-INVITE may time-out and may be tried again several more times to the wrong network by virtue of its built-in retransmission mechanism. The crucial difference is that the data path and signaling path are separate in the case of SIP, unlike that of MIP. For simplicity, automatic retransmissions of lost re-INVITE messages are not shown and neither are messages like ACK that acknowledge the SIP OK. Simultaneous mobility of two end-hosts is not usually an issue for pre-session mobility in SIP. However, during a transition before a SIP session set-up is complete, simultaneous mobility may present problems. As shown in FIG. 5, it may so happen that one of the signaling messages (INVITE, OK, ACK) does not reach the other party and gets lost, despite the SIP server keeping an up-to-date registration for both parties.



Figure 5: Simultaneous Mobility in SIP-based environment

5. Occurrence of Simultaneous Mobility Problems

(THIS SECTION IS TO BE REVISED SOON)

In this section, we consider when the simultaneous mobility problem may or may not occur, and we state some of the results of the analysis in [[simultaneous-mob2](#)]. In these results, the term "location proxy" shall be used (applied to a Mobile Host) in a broad sense, to mean a network function that is used to locate the Mobile Host. Different categories of location proxies are described in [[simultaneous-mob2](#)], but the categorization is not needed in this present draft.

In an infrastructure-based environment, given a pair of consecutive handoffs, one for each of two mobile nodes in a communications session in normal state (up till the first handoff), in the absence of location proxies for either mobile node (or there might be location proxies but they are not used/involved), any binding update sent by the earlier moving mobile node will be lost through belated arrival, if and only if the binding update does not arrive at the other mobile node before it moves.

Given a pair of consecutive handoffs, one for each of two mobile nodes in a communications session in normal state (up till the first handoff), in the absence of location proxies for either mobile node (or there might be location proxies but they are not used/involved), if the binding update from the node that moved first is lost through belated arrival, the binding update from the node that moved second will also be lost through belated arrival.

Given a pair of consecutive handoffs, one for each of two mobile nodes in a communications session in normal state (up till the first handoff), in the absence of location proxies for either mobile node (or there might be location proxies but they are not used/involved), the simultaneous mobility problem does not occur if and only if the binding update from the node that moved earlier reaches the other node before that node moves.

Simultaneous mobility problem does not occur in MIP if it does not use Route Optimization. No binding updates are sent directly to mobile nodes. So, there are none to be lost through belated arrivals. Instead, Home Agents serve as up-to-date intercepting location proxies that forward data to the right location as soon as Mobile IP registration occurs after each handoff.

6. Security Considerations

The mobility protocols MIPv6 or SIP can use the existing security mechanisms designed for each of the protocols for any related extension.

7. IANA Considerations

This document has no actions for IANA.

8. Acknowledgments

Authors would like to acknowledge anonymous reviewers of the simultaneous mobility papers, and Rute Sofia and Kilian Weniger for helpful comments on the first version of this draft.

9. References

9.1. Normative References

- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.
- [RFC4068] Koodli, R., "Fast Handovers for Mobile IPv6", [RFC 4068](#), July 2005.

9.2. Informative References

- [simultaneous-mob1]
Wong, K. and A. Dutta, "Simultaneous Mobility in MIPv6", Proceedings of EIT 2005 May 2005.
- [simultaneous-mob2]
Wong, K., Dutta, A., and H. Schulzrinne, "Simultaneous Mobility: Analytical Framework, Theorems and Solutions", Journal of Wireless Communications and Mobile Computing June 2007.
- [simultaneous-mob3]
Kravets, R., Carter, C., and L. Magalhaes, "A cooperative approach to user mobility", ACM Computer Communications Review October 2001.
- [simultaneous-mob4]
Wong, K. and W. Woon, "Analysis of Simultaneous Mobility under Asymmetric Conditions", Proceedings of Milcom 2007 October 2007.
- [SIPMM] Schulzrinne, H. and E. Wedlund, "Application Layer Mobility Using SIP", ACM MC2R.

Authors' Addresses

K. Daniel Wong
Malaysia University of Science and Technology
GL 33, Block C, Kelana Square, 17 SS 7/26, Kelana Jaya
Petaling Jaya, Selangor 47400
Malaysia

Phone: +603 7880 1777 x282

Email: dwong@must.edu.my

Ashutosh Dutta
Telcordia Technologies
1 Telcordia Drive
Piscataway, NJ 08854
USA

Phone: +1 732 699 3130

Email: adutta@research.telcordia.com

Henning Schulzrinne
Columbia University
Department of Computer Science
450 Computer Science Building
New York, NY 10027
USA

Phone: +1 212 939 7004

Email: hgs@cs.columbia.edu

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

