Efficient End-to-End Mobility Support in IPv6

Christian Vogt, Mark Doll
Institute of Telematics, Universität Karlsruhe (TH), Germany
Email: {chvogt | doll}@tm.uka.de

***Abstract*— The efficiency of Mobile IPv6 Route Optimization in terms of propagation latencies and packet overhead is contrasted by significant handoff delays. Much analytic effort has recently been spent on reducing these delays, but little practical experience has yet been gathered. This paper compares the efficiency of the combination of two proposals, Early Binding Updates and Credit-Based Authorization, with that of standard Route Optimization. This is based on measurements for RTP/UDP voice traffic and TCP file transfers, which were taken in an experimental testbed.**

I. Introduction

Along with a growing appreciation of the Internet as a communications platform for business, academia, and entertainment services rises the desire for anytime, anywhere network access. This development is driven by new, real-time and/or multi-media applications, such as audio and video streaming, IP telephony, or video conferencing, which are of particular convenience when ubiquitously accessible.

Given these developments and new requirements, the Internet Engineering Task Force (IETF) added mobility support [1] to the IPv6 protocol suite in mid 2004. The conceptually new approach taken therein compared to older IPv4 mobility mechanisms is that Mobile IPv6 performs Route Optimization to minimize propagation latencies and packet overhead. Route Optimization allows mobile peers to communicate via a direct path; in Mobile IPv4, all packets are diverted through a stationary home agent. While this constitutes the primary strength of Route Optimization, a problem is that arbitrary peers do not generally pre-share credentials based on which mobility management could be secured. So how can a node's connection be protected from getting hijacked by an impersonator? And how to keep an ostensibly mobile node from flooding a victim by claiming the victim's IP address?

As will be explained in a later section, the solution adopted in Route Optimization is based on testing a mobile node's reachability at the addresses that it uses. Address tests reduce nearly all threats of Route Optimization to those which already exist in non-mobile IPv6 or IPv4 [2]. On the other hand, the tests are unfavorable with respect to handoff latency. In fact, Route Optimization introduces handoff delays, at IP layer, of up to four round-trip times, depending on the implementation.

Different approaches have been proposed to reduce the handoff latency of Route Optimization. Some mechanisms involve local enhancements within the mobile nodes' access networks. Others function purely end to end. In all cases, there is a trade-off between the performance benefit an optimization yields and the additional costs it has in terms of infrastructure or pre-configuration requirements [3]. One of the first proposed enhancements *without* such prerequisites was a combination of Early Binding Updates and Credit-Based Authorization. This has since been discussed and reviewed both in the academic community [4][5] as well as within the IETF [6][7]. However, as with most related work, little real-life experience has yet been made and publicly documented. This paper is the first that supports the estimated benefits of Early Binding Updates and Credit-Based Authorization by performance results obtained in an experimental testbed.

Following this introduction, section II describes Mobile IPv6 with a focus on Route Optimization. Section III explains Early Binding Updates and Credit-Based Authorization. An evaluation of experimentation results follows in section IV. Section V attends to related research conducted elsewhere. The paper concludes in section VI.

II. Mobility Support in IPv6

IP addresses traditionally serve both end-point identification in transport protocols or applications as well as locating interfaces for routing. This is problematic when nodes are mobile: Changes in IP connectivity necessitate IP-address reconfiguration and so cause ongoing transport connections or application sessions to break. Mobile IPv6, specified in RFC 3775 [1], solves the ambiguity problem of IP addresses by using two of them per mobile node. Packets are routed based on a temporary *care-of address*, which the mobile node replaces when it moves to a new access network. A static *home address*, with a prefix from the mobile node's *home network*, serves as an identifier at upper layers.

The mobile node can directly receive packets at its home address while it stays at home, so no care-of address is needed during that time. When it moves to a different network, the mobile node requests a dedicated router in the home network, its *home agent*, for proxy service. This *home registration* establishes a bidirectional tunnel between the home address and current care-of address so that the mobile node can continue to communicate through the home address from remote. A pair of home address and care-of address is called a *binding*. It is the mobile node's responsibility to update a binding whenever the care-of address changes.

As bidirectional tunneling causes encapsulation overhead and increases packet-propagation times, Mobile IPv6 provides a mechanism for *Route Optimization*: When the mobile node receives the first encapsulated packet, it initiates a *correspondent registration* with the sender of this packet. A route-optimized packet carries the care-of address within the IPv6

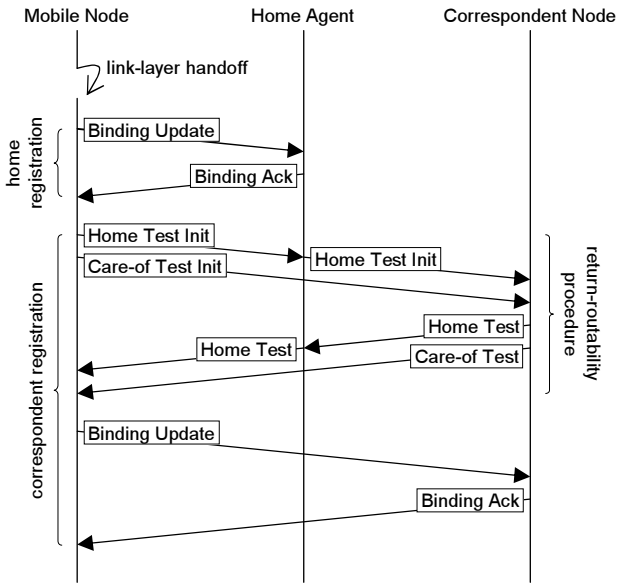

Fig. 1. Home and correspondent registration in Mobile IPv6

header during transit. The home address is located within an IPv6 Destination Options or Routing extension header, depending on whether the mobile node or the correspondent node sent the packet. Both end nodes swap the addresses when the packet traverses the IP layer so that transport protocols and applications can access the home address as usual.

Preventing misuse of registrations is typically straightforward in the case of home registrations: Mobile nodes and their home agents are administered by the same authority and pre-share an IPsec security association or credentials to bootstrap one. On the other hand, there is generally no such relationship between a mobile node and an arbitrary correspondent node [2]. So how can a mobile node authenticate itself during a correspondent registration? And how can the correspondent node verify whether the mobile node is indeed reachable at a new care-of address? Failure to properly authenticate the registrant introduces a vulnerability to connection hijacking and man-in-the-middle attacks. And unless the new care-of address undergoes a reachability check, it could be spoofed in order to redirect traffic towards a flooding target.

Mobile IPv6 uses the *return-routability procedure* to protect correspondent registrations despite the lack of pre-shared credentials. This is based on the following two observations: First, authentication in the context of Mobile IPv6 essentially boils down to verifying a node's ability to receive packets at the home address. A reachability check at the home address hence authenticates a registering mobile node. Second, a reachability check at the care-of address prevents redirection-based flooding attacks and so authorizes a mobile node to claim that care-of address.

Figure 1 shows the messages exchanged during home and correspondent registrations. The home registration consists of a Binding Update message, notifying the home agent of the new care-of address, and a Binding Acknowledgment message to indicate registration success or failure. Both messages are authenticated and usually encrypted through IPsec ESP.

During the return-routability procedure, the mobile node reverse-tunnels a Home Test Init message to the home agent, which forwards the message to the correspondent node, and sends a Care-of Test Init message directly to the correspondent node. Each of the messages causes the correspondent node to generate a secret token: The *home keygen token* is sent to the home address within a Home Test message and forwarded by the home agent to the care-of address. The *care-of keygen token* is sent directly to the care-of address. By knowledge of the home and care-of keygen tokens, the mobile nodes proves its ability to receive packets at the home address and care-of address, respectively. Specifically, it authenticates the Binding Update message that it subsequently sends to the correspondent node with a key derived from both tokens. This allows the correspondent node to bind the two addresses. The final Binding Acknowledgment message affirms or rejects the correspondent registration.

III. Optimized End-to-End Mobility Support

The benefit of Route Optimization in using a direct route between communicating peers is very much in opposition to a high latency during handoff. According to figure 1, registrations may take up to four round trips if mobile nodes wait for the home registration to complete before they initiate the return-routability procedure: two round trips between the mobile node and the home agent, one between the home agent and the correspondent node, and another one between the mobile node and the correspondent node. The two dominant open-source Mobile IPv6 implementations [8][9] operate this way in order to guarantee that the home agent is aware of the current care-of address when it processes Home Test Init and Home Test messages. Such *conservative* behavior is henceforth contrasted with the more *optimistic* approach of executing the home registration and the return-routability procedure in parallel [10]. Conservative Route Optimization avoids a useless return-routability procedure in case the home registration fails. This comes at the cost of an additional round trip when the home registration is successful. Optimistic Route Optimization requires one round-trip time of signaling time less, but may run a return-routability procedure in vain should the corresponding home registration fail.

Round trips through the home agent are particularly undesirable when one or both of the peers roam far from home, e.g., at a conference venue. One refers to this as the *"Two-Japanese-in-America" problem*. The problem is worst with conservative Route Optimization; optimistic Route Optimization mitigates it to some limited extent. More substantial improvement can be obtained by a combination of Early Binding Updates [6][4] and Credit-Based Authorization [7][5].

Early Binding Updates eliminate the handoff delay caused by the return-routability procedure in the way shown in figure 2: The mobile node initiates a *proactive home-address test* prior to handoff. It may do so periodically, whenever the most recently obtained home keygen token is about to expire, or on

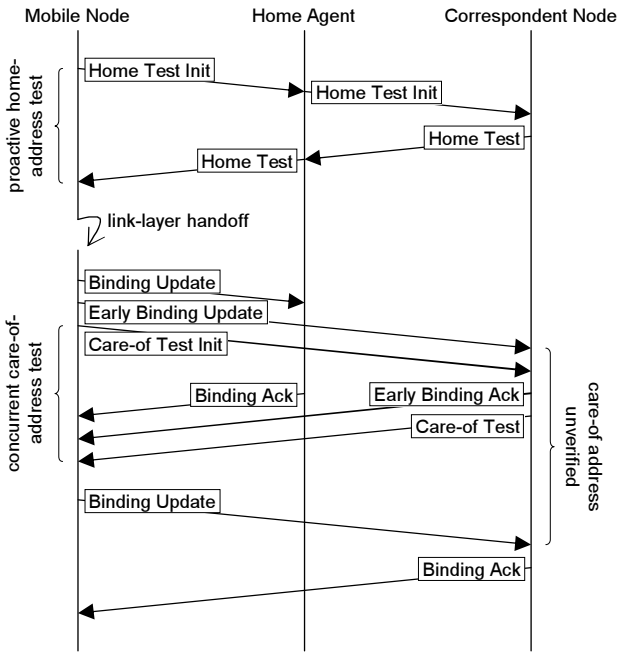

Fig. 2. Home and correspondent registration with Early Binding Updates and Credit-Based Authorization

the basis of link-layer triggers indicating imminent handoff. After the handoff, the mobile node first registers the new care-of address with an Early Binding Update message, prompting an Early Binding Acknowledgment message. A *concurrent care-of-address test* is then performed while bidirectional communications are already being resumed.

Authenticity of the Early Binding Update message is protected based on a key derived from a home keygen token only. This lacks a reachability proof for the new care-of address, so the correspondent node labels the address *unverified* at first. The status of the address changes to *verified* when the correspondent node receives the Binding Update message, which is signed by a key generated from both the home and the care-of keygen token, after completion of the concurrent care-of-address test.

Credit-Based Authorization helps a correspondent node to prevent misuse of unverified care-of addresses for redirection-based flooding attacks. The correspondent node maintains a byte counter per mobile node. The counter increases by the size of each packet received from the mobile node or, in an advanced mode, sent to the mobile node while the care-of address is verified. The counter decreases by the size of each packet sent to the mobile node while the care-of address is unverified unless this would cause it to take a negative value. The packet is typically dropped in this latter case, but it may also be directed to the home address or buffered until the care-of address becomes verified. Credit-Based Authorization ensures that a correspondent node does not send more data to an unverified care-of address than the node assumed to be present at that address has either previously sent or provably received, depending on the mode of operation. This outrules

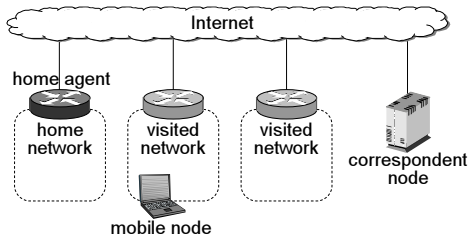

Fig. 3. Testbed topology

amplified flooding attacks against unverified care-of addresses and hence permits secure concurrent care-of-address tests [7].

The combination of Early Binding Updates and Credit-Based Authorization reduces the handoff latency introduced at IP layer to a single round-trip time since both the home- and the care-of-address test are moved to a non-critical period during which they do not delay communications.

IV. Performance Analysis

The performance of Mobile IPv6 Route Optimization with Early Binding Updates and Credit-Based Authorization relative to that of conservative and optimistic Route Optimization has been verified in an experimental testbed. This section summarizes and analyzes the results obtained from experiments with RTP/UDP voice traffic and TCP file transfers.

A. Experimentation Environment

The experimental testbed consists of five FreeBSD nodes playing the roles of the mobile node, home agent, correspondent node, and two routers for access to visited networks. Figure 3 illustrates the testbed topology. The mobile node may attach to its home agent or to either of the two access routers in the visited domains. The exterior interfaces of all three routers and the correspondent node connect to the ”Internet”. The properties of global Internet routes are reproduced by FreeBSD’s DummyNet facility. This limits bandwidth to 1024 kbps and imitates end-to-end round-trip times of between 40 and 200 milliseconds, depending on the experiment.

Mobility is handled by Kame-Shisa [8], a two-part Mobile IPv6 implementation including a kernel patch for performance-critical packet processing as well as userland daemons for control and signaling. Kame-Shisa implements conservative Route Optimization. For the purpose of the experiments described herein, the software was modified to also support optimistic Route Optimization as well as Route Optimization with Early Binding Updates and Credit-Based Authorization. The three protocols are labeled ”conserv”, ”optim”, and ”ebu/cba”, respectively, in the figures below.

Routers multicast Router Advertisement messages within the home and visited networks in intervals of between 30 and 70 milliseconds [1]. Hence, when the mobile node changes IP connectivity, it receives the first Router Advertisement message from the new router after an expected 25 milliseconds. Movement detection is based on this advertisement in conjunction with IPv6 Neighbor Unreachability Detection indicating absence of the old router. During Neighbor Unreachability

Detection, the mobile node solicits the old router three times, interspaced by a configurable time during which the mobile node listens for a response. A 10-milliseconds pause is used in these experiments. If no response appears, the mobile node selects a new care-of address and updates its bindings at the home agent and correspondent node. Movement detection hence takes 30 milliseconds in the best and 100 milliseconds in the worst case. The expected duration is 55 milliseconds. The mobile node is assumed to use Optimistic Duplicate Address Detection [11] to avoid address-configuration delays.

Much contemporary, experimentative work on IP mobility focuses on a particular data-link and medium-access technology, frequently adopting the IEEE 802.11 standard. Results from such experiments have the convenient property that they shed light on performance achievable in a certain real-life environment. On the other hand, it is generally infeasible to convey the results to different technologies. And although IEEE 802.11 prevails today, it is questionable whether this standard can accommodate the rigid requirements of delay-sensitive applications [12][13]. Sufficient performance may be achievable only through further technological improvements [14]. Research also shows that link-layer characteristics may vary strongly even for a *specific* technology, given different cell loads or user application and mobility patterns [15]. While focusing experiments on a certain technology is a must for research on cross-layer interactions, it might unintendedly narrow down the results' representativity otherwise. In this study, a deliberate decision was therefore made to abstract from link-layer specifics. This is realized through static, wired connections between the mobile node and the routers, where communications can be selectively en- and disabled through FreeBSD's IPFW2 firewall and MAC filter.

B. Voice Traffic with RTP over UDP

Voice traffic is modeled as a bidirectional 64-kbps constant-bit-rate data stream, split into chunks of 10 milliseconds length. Each chunk is prepended by IPv6, UDP, and RTP headers to form a packet of 164 bytes length, including the IPv6 Destination Options and Routing extension headers required for Route Optimization.

Figure 4 juxtaposes the handoff latency observed by a voice-over-IP (VoIP) application for conservative and optimistic Route Optimization as well as Route Optimization with Early Binding Updates and Credit-Based Authorization. Mobile nodes measure the handoff latency as the period between reception of the last packet at the old care-of address and the time when the first packet is delivered to the new care-of address. The diagram is based on 500 handoffs for each of the three mobility protocols. End-to-end round-trip times are 200 milliseconds in this case. The measurements clearly reflect the handoff latency claimed by Mobile IPv6, which dominates the overall handoff latency observed by the application: Conservative Route Optimization requires 3.5 end-to-end round-trip times to renew the mobile node's care-of address at the correspondent node. An additional one-way time elapses until the correspondent node's first data packet reaches

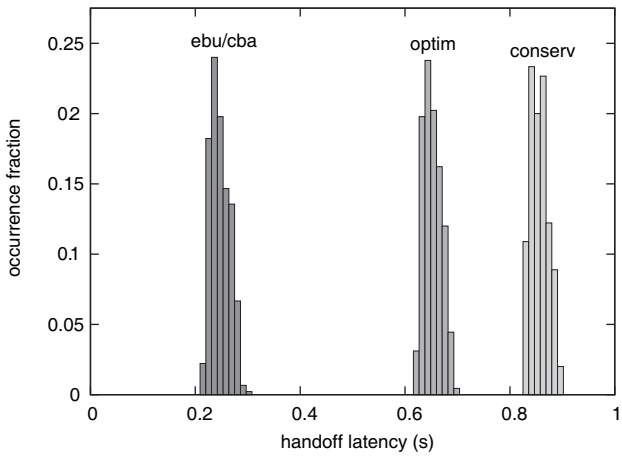

Fig. 4. Handoff latency for voice traffic and 200ms round-trip times

the mobile node at the new care-of address. Including the expected latency of movement detection yields a total handoff latency of roughly 855 milliseconds.

Optimistic Route Optimization spares one end-to-end round-trip time and so produces an overall handoff latency of about 680 milliseconds. Route Optimization with Early Binding Updates and Credit-Based Authorization updates a binding within a single one-way time. Adding the propagation time of the first packet delivered to the new care-of address plus the time required for movement detection yields a total handoff latency of approximately 280 milliseconds. Further experiments with different round-trip times corroborate these relationships between the three protocols, but are for brevity purposes not included in this paper.

C. File Transfers with TCP

Figure 5 compares the three protocols under evaluation with respect to the data delivered during a 60-seconds file transfer and five handoffs. It shows averages and 95% confidences from 20 experiments per protocol and round-trip time. Data flows from the correspondent node to the mobile node, with TCP Reno providing transportation. While the measurements verify an expected dependency on the round-trip times, they also show that the dependency is lowest for Route Optimization with Early Binding Updates and Credit-Based Authorization. What is striking is that this can lead to a performance gain of more than 70 percent compared to conservative or optimistic Route Optimization. Likewise noteworthy, optimistic Route Optimization evidently fails to provide any noticeable improvement. What makes Early Binding Updates and Credit-Based Authorization perform so much better?

The answer to this question is found in TCP's loss-recovery mechanisms. For reliability, TCP uses a retransmission timer to estimate when previously dispatched segments have been lost and ought to be resent. Specifically, when a TCP sender does not get positive feedback from the receiver for the duration of a retransmission timeout, all data sent, but not yet acknowledged, will be retransmitted. This strategy has

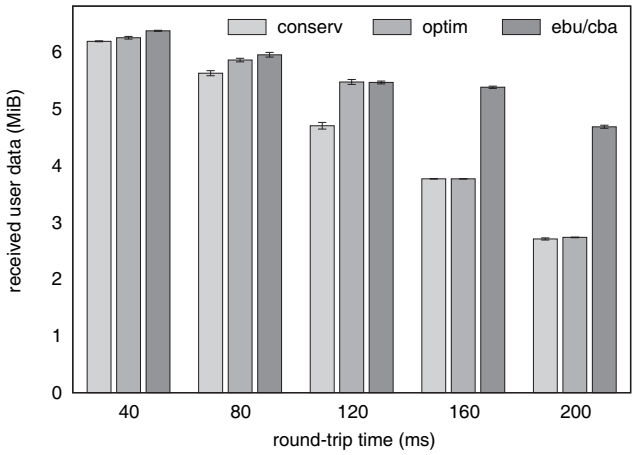

Fig. 5. User data received after 60 seconds and 5 handoffs

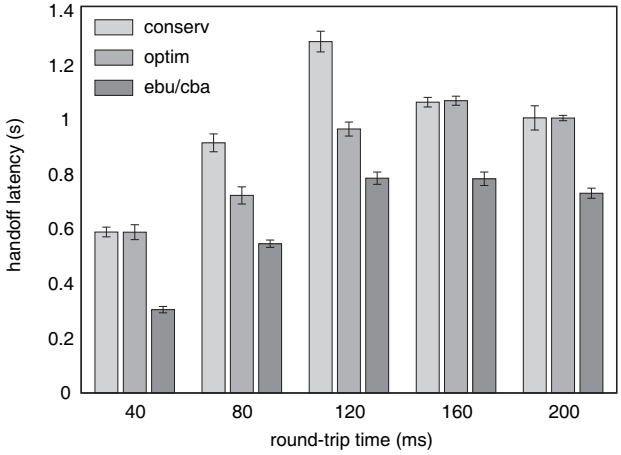

Fig. 7. Handoff latency for TCP connections

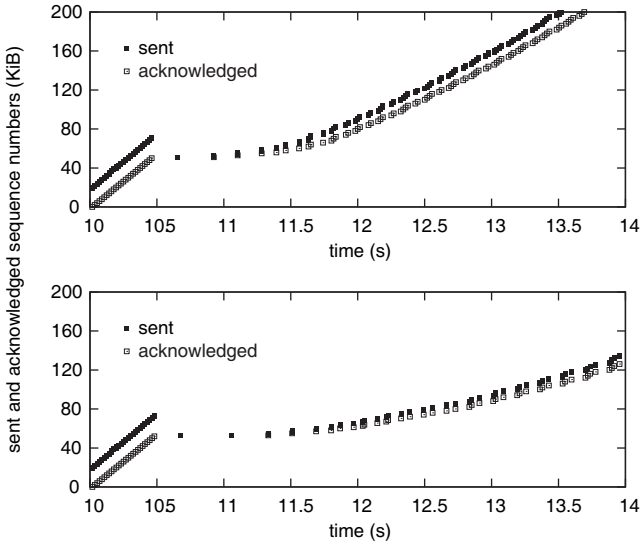

Fig. 6. Handoffs for 160ms round-trip times cause a single TCP retransmission timeout with Early Binding Updates and Credit-Based Authorization (top) and two successive ones with conservative Route Optimization (bottom).

undesired side effects if TCP operates over Route Optimization. When a mobile TCP receiver changes IP connectivity, all packets currently in flight to the old care-of address are lost. An additional one-way worth of data is lost while the receiver updates its binding at the TCP sender. The total loss roughly corresponds to the maximum data volume that TCP can transmit without receiving an acknowledgment. The sender consequently stalls and runs into a retransmission timeout. The timeout period is a function of the measured round-trip time and the variation in the samples.

If the binding update completes prior to expiration of the retransmission timer, the TCP sender uses the new care-of address when it resends the lost segments. Otherwise, TCP directs the lost segments to the old care-of address and times out yet again. The timeout period doubles for each successive retransmission, up to a certain limit. To make things worse, TCP does not exponentially ramp up its transmission rate in slow-start mode after the second retransmission timeout as it usually does after a single timeout. Instead, TCP operates in congestion-avoidance mode, accelerating transmission by only one segment per round-trip time. This happens because the point of transitioning from slow start to congestion avoidance is defined as the time when the amount of outstanding, unacknowledged data equals half of what this amount was at the time of the last retransmission timeout. As the sender probes network conditions with only one segment after the first timeout, the threshold is set to a minimum of two segments when the timer expires a second time. Figure 6 illustrates this effect with two exemplifying TCP traces from scenarios with 160-milliseconds round-trip times.

The advantage of Route Optimization with Early Binding Updates and Credit-Based Authorization compared to conservative and optimistic Route Optimization is a much higher probability to complete a binding update prior to the first retransmission timeout. The mean handoff latencies and 95% confidences shown in figure 7 substantiate this on the basis of 100 handoffs for each of the three protocols and five round-trip times. Here, the handoff latency is defined as the period between the transmission of the first lost packet sent to a stale care-of address and the first acknowledged retransmission of that segment.

The advanced retransmission strategies of TCP NewReno and Selective Acknowledgments are of little help in mobile environments. Though these mechanisms allow for efficient recovery from accumulated packet loss, they still require that later segments are successfully delivered and trigger acknowledgments. This does not happen when an entire window worth of data is lost during handoff.

V. Related Work

Contributors to handoff latency are multifold, and so are the approaches to mitigate them. Fast Handovers for Mobile IPv6 (F-MIPv6) [16] amend access routers so as to enable a mobile node to discover new points of IP attachment and configure a new care-of address prior to handoff. Local rerouting allows

the mobile node to temporarily continue communications through the old care-of address subsequent to handoff until remote bindings have been updated. Hierarchical Mobile IPv6 (H-MIPv6) [17] eliminates end-to-end signaling from handoffs within a certain domain. Local mobility anchor points handle movements of visiting mobile nodes transparently to exterior home agents and correspondent nodes.

Research shows that both F-MIPv6 and H-MIPv6 can substantially reduce handoff latencies and packet loss, although higher signaling or encapsulation overhead may defeat these benefits when contention on the access link is high [18]. Local rerouting generally renders F-MIPv6 superior to H-MIPv6 in terms of packet loss [19], but may at the same time cause reordering and thereby disturb TCP connections [20]. The relative handoff latency of the two protocols depends on the topological locations of the mobility anchor point and geographically adjacent access networks.

The high performance benefits achievable through enhancements within the access networks come at the cost of required infrastructure upgrades, however. Also, *inter*-domain handoffs may not benefit from the infrastructure even if it exists due to lack of roaming agreements between network-access providers. End-to-end optimizations do not have these constraints [3]. At the cost of some performance, they provide an independence which can be of great advantage in many roaming scenarios. One such protocol [21], currently under discussion within the IETF, uses cryptographically generated home addresses to avoid home-address tests. Like the mechanisms analyzed in this paper, it applies Credit-Based Authorization for early use of new care-of addresses during reachability verification. [22] replaces the return-routability procedure by cryptographic authentication based on secrets shared between mobile and correspondent nodes. Handoff latencies can so be reduced where the required credentials exist. The technique provides no protection against spoofed care-of addresses other than policy recommendations, however, so a care-of-address test may still be required and hence limit performance benefits.

VI. Conclusions

This paper evaluates the efficiency of a combination of Early Binding Updates and Credit-Based Authorization, two earlier proposed enhancements to Mobile IPv6 Route Optimization, based on measurements obtained from an experimental testbed. The performance results are related to those of conservative Route Optimization, which appears to prevail amongst today's Mobile IPv6 implementations, and a slightly more efficient optimistic variant. Experiments were conducted for RTP/UDP voice traffic as well as TCP file transfers.

The results indicate that Early Binding Updates and Credit-Based Authorization effect a significant reduction in handoff latencies for both UDP and TCP traffic. In the case of TCP, lower latencies lead to fewer retransmission timeouts, faster adaptation to available network resources, and thus increased overall throughput.

Acknowledgment

Special thanks are due to Ralf Beck, Roland Bless, Daniel Jungbluth, Max Laier, and Constantin Schimmel, in alphabetical order, for their vigorous support regarding this project.

References

[1] D. Johnson, C. E. Perkins, and J. Arkko, "Mobility Support in IPv6," IETF Request for Comments 3775, June 2004.

[2] P. Nikander et al., "Mobile IP Version 6 Route Optimization Security Design Background," IETF Request for Comments 4225, Dec. 2005.

[3] C. Vogt and J. Arkko, "Taxonomy and Analysis of Enhancements to Mobile IPv6 Route Optimization," IETF Internet Draft draft-irtf-mobopts-ro-enhancements-04.txt (work in progress), Oct. 2005.

[4] C. Vogt et al., "Early Binding Updates for Mobile IPv6," in *Proceedings of the IEEE Wireless Communications and Networking Conference*, vol. 3. IEEE, Mar. 2005, pp. 1440–1445.

[5] C. Vogt, "Credit-Based Authorization for Concurrent IP-Address Tests," in *Proceedings of the IST Mobile and Wireless Communications Summit*, June 2005.

[6] ——, "Early Binding Updates for Mobile IPv6," IETF Internet Draft draft-vogt-mobopts-early-binding-updates-00.txt (work in progress), Feb. 2005.

[7] C. Vogt and J. Arkko, "Credit-Based Authorization for Mobile IPv6 Early Binding Updates," IETF Internet Draft draft-vogt-mobopts-credit-based-authorization-00.txt (work in progress), Feb. 2005.

[8] "Kame-Shisa," Mobile IPv6 for FreeBSD 5.4. [Online]. Available: http://www.kame.net/newsletter/20041211/shisa.html

[9] V. Nuorvala, H. Petander, and A. Tuominen, "Mobile IPv6 for Linux (MIPL)." [Online]. Available: http://www.mobile-ipv6.org/

[10] C. Vogt, "Samitas Review of draft-irtf-mobopts-ro-enhancements-00," IETF MIP6 mailing list, http://www.atm.tut.fi/list-archive/MIPv6-2005/msg00677.html, June 2005.

[11] N. S. Moore, "Optimistic Duplicate Address Detection for IPv6," IETF Internet Draft draft-ietf-ipv6-optimistic-dad-07.txt (work in progress), Dec. 2005.

[12] A. Mishra, M. Shin, and W. Arbaugh, "An Empirical Analysis of the IEEE 802.11 MAC Layer Handoff Process," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 2, pp. 93–102, Apr. 2003.

[13] J.-O. Vatn, "An Experimental Study of IEEE 802.11b Handover Performance and Its Effect on Voice Traffic," Telecommunication Systems Laboratory, Department of Microelectronics and Information Technology, KTH, Royal Institute of Technology, Stockholm, Sweden, Technical Report TRITA-IMIT-TSLAB R 03:01, July 2003.

[14] J.-P. Ebert et al., "Paving the Way for Gigabit Networking," *IEEE Communications Magazine*, vol. 43, no. 4, pp. 27–30, Apr. 2005.

[15] N. Montavont and T. Noël, "Analysis and Evaluation of Mobile IPv6 Handovers over Wireless LAN," *Mobile Networks and Applications*, vol. 8, no. 6, pp. 643–653, Dec. 2003.

[16] E. Rajeev Koodli, "Fast Handovers for Mobile IPv6," IETF Request for Comments 4068, July 2005.

[17] H. Soliman et al., "Hierarchical Mobile IPv6 Mobility Management (HMIPv6)," IETF Request for Comments 4140, Aug. 2005.

[18] X. Pérez-Costa, M. Torrent-Moreno, and H. Hartenstein, "A Performance Comparison of Mobile IPv6, Hierarchical Mobile IPv6, Fast Handovers for Mobile IPv6 and Their Combination," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 7, no. 4, pp. 5–19, Oct. 2003.

[19] Y. Gwon, J. Kempf, and A. Yegin, "Scalability and Robustness Analysis of Mobile IPv6, Fast Mobile IPv6, Hierarchical Mobile IPv6, and Hybrid IPv6 Mobility Protocols Using a Large-Scale Simulation," in *Proceedings of the IEEE International Conference on Communications*. IEEE, June 2004.

[20] R. Hsieh et al., "Performance Analysis on Hierarchical Mobile IPv6 with Fast-Handoff over End-to-End TCP," in *Proceedings of the IEEE Global Telecommunications Conference*. IEEE, Nov. 2002.

[21] J. Arkko, C. Vogt, and W. Haddad, "Applying Cryptographically Generated Addresses and Credit-Based Authorization to Mobile IPv6," IETF Internet Draft draft-arkko-mipshop-cga-cba-02.txt (work in progress), Oct. 2005.

[22] C. E. Perkins, "Securing Mobile IPv6 Route Optimization Using a Static Shared Key," IETF Internet Draft draft-ietf-mip6-precfgkbm-04.txt (work in progress), Oct. 2005.