Optimized Fast-Handoff Schemes for Application Layer Mobility Management

Ashutosh Dutta*[*], Sunil Madhani, Wai Chen
Telcordia Technologies Inc., 445 South Street, Morristown, NJ 07960

Onur Altintas, Toyota InfoTechnology Center, 4009 Miranda Avenue, Palo Alto, CA 94304

Henning Schulzrinne
Computer Science Department, Columbia University, New York, NY 10027

In order to ensure proper quality of service for real-time communication in a mobile wireless Internet environment, it is essential to minimize the transient packet loss when the mobile host (MH) is moving between different cells (subnets) within a domain. This paper introduces application layer techniques to achieve fast handoff for real-time (RTP/UDP) multimedia traffic in a SIP-based signaling environment. These techniques are based on standard SIP components such as user agents and proxies and provide a network-independent solution suitable for application service providers.

I. Introduction

In order to provide seamless mobility support to clients in a mobile wireless Internet environment several variations of mobile IP [1] have been proposed. Reference [2], [3] discuss several solutions for intra-domain mobility. Internet telephony uses SIP [4] to establish and tear down multimedia sessions. These multimedia sessions are mostly based on RTP and thus have different delay and error characteristics than standard TCP-based applications. Application-layer mobility management [5] provides an alternative mobility solution using SIP for such real-time traffic. This scheme does not depend on home agents (HAs) or foreign agents (FAs) in the home or visited network and can be easily deployed by any third-party application provider without depending on the cooperation of the ISP providing network-layer connectivity. SIP mid-call mobility uses SIP INVITE messages to inform the correspondent host (CH) about the new network attachment point.

When a mobile node moves between cells (subnets) within a domain, data in transit may be lost during the time that it takes to complete the mobile IP registration or SIP re-INVITE. In order to reduce the data loss when the communicating hosts are far apart, it is necessary to limit the movement indication to within the domain. There are several intra-domain mobility management solutions at the network layer [3, 6, 7, 8]. Most of these are variations of hierarchical mobility agents installed within a domain. That way, the binding update request does not travel all the way to the CH or HA, but rather is kept within the domain. Similarly, we want to extend SIP's mobility management to provide a similar intra-domain solution. Below, we propose several approaches that rely on intercepting the data traffic for a limited period during the MH's intra-domain mobility within the domain itself.

[*]*First author is a Part-time Ph.D student at Columbia University

II. SIP fast-handoff techniques

Each visited domain may consist of several subnets. Every move to a new subnet causes the MH to send a re-INVITE to the CH containing its new care-of address. If the re-INVITE request gets delayed due to path length or congestion, media packets will continue to be directed to the old address. We assume that the visited network has an outbound proxy. We enhance this proxy with the ability to temporarily register visitors . The visitor obtains a temporary, random identity from the visited network and uses it as its address-of-record to register with the registrar in the visited network. The MH informs the home registrar of this temporary address. It then only updates that registration with its current local IP address. This speeds up registrations, but does not address the "delayed binding update" issue. In this section, we describe several ways to achieve fast handoff using SIP, namely, using a SIP registrar and RTP translator or NAT, using the outbound proxy and B2BUA as a mobility agent. In-transit packets can be redirected to a unicast or multicast address based on the movement pattern of the mobile hosts and usage scenario. By limiting the signaling within the domain it helps reduce the transient data loss.

II.A. SIP registrar and RTP translator or NAT

Each subnet within a domain is equipped with an RTP translator [9] that provides application-layer forwarding of RTP packets for a given addres and UDP port to a given network destination. (RTP applications generally do not care about the source IP address of RTP packets, using just the synchronization source identifier (SSRC) to identify the source.) Figure 1 shows a sequence of operations when a mobile host moves from one network to another within a domain. SIP server here acts like a registrar. The visited-network registrar described earlier receives the registration update from the MH that has just moved, and immediately sends a request to the RTP translator in the network that the MH just left. The request causes the RTP translator to bind to the old IP address used by the MH and forward any

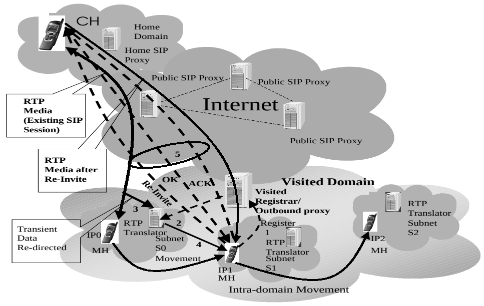

Figure 1: Intra-domain fast handoff with RTP translator

incoming packets to the new address of the MH. After a set interval or after no media packets have been received by the RTP translator, the RTP translator relinquishes this old address and removes the forwarding table entry, assuming that the re-INVITE has reached the CH. While elaborate mechanisms to explicitly turn off the translation could be devised, the spacing between RTP packets is sufficiently small so that RTP translator can quickly detect that no more packets are likely to arrive. The DHCP server should be configured to assign addresses on a least-recently-used basis, to avoid accidental re-assignment of the address to the new visitor. The SIP registration does not contain the media port number, so that this registration-based scheme only works if the media is directed to a well-known port. To avoid the problem, one can use a network address translator (NAT) instead of the RTP translator to rewrite the destination IP address from the old to the new subnet. We have implemented the mechanism using SIP-CGI (SIP- Common Gateway Interface) scripts in our registrar without causing any changes in the end systems.

II.B. SIP outbound proxy

SIP requests typically traverse a SIP proxy in the visited network, namely the *outbound proxy*. This outbound proxy can also support fast handoff, by using the data in the MH-to-CH re-INVITE to configure the RTP translator or NAT. The advantage of this approach is that the outbound proxy usually has access to the Session Description Protocol (SDP) information containing the MH media address and port, thus simplifying the configuration of the translator or NAT. On the other hand, this outbound proxy has to remember the INVITE information for an unbounded amount of time and become call stateful, since it needs the old information when a new re-INVITE is issued by the MH.

II.C. SIP back-to-back UA and Multicast Agent

Another way of providing fast-handoff is to use a back-to-back SIP user agent (B2BUA). A B2BUA consists of two SIP user agents where one user agent receives a SIP request, possibly transforms it and then has the other part of the B2BUA re-issue the request. A B2BUA in each domain needs to be addressed by the MH in the visited domain. The B2BUA issues a new request to the CH containing its own address as the media destination and then forwards the packets, via RTP translation or NAT, to the MH. This approach has the disadvantage that it requires some cooperation from the MH. As noted, the INVITE request needs to be addressed explicitly to the B2BUA, as otherwise end-to-end encryption of the body may prevent the B2BUA from inspecting it.Scoped multicast may help to avoid packet losses if the MH can predict that it will be moving to a new subnet shortly.

III. Implementation

We have implemented the register-based intercept in our mobile multimedia testbed. We have used modified version of Columbia University's SIP user agent, Linux 2.4.7 based router with iptables NAT and rtptrans functionality for packet re-direction, VIC and RAT tools for measuring video and audio packet loss and NIST delay simulator to emulate the distance and network congestion. We observed that the packet delay of the iptables-based NAT was less than 1ms, while the RTP translator added 4ms of delay.

References

[1] C. Perkins, "IP mobility support," Request for Comments 3220, IETF, Jan 2002

[2] A.T. Campbell, Javier Gomez, Sanghyo Kim et al, "Comparison of IP Micromobility Protocols," IEEE Wireless Communications, Vol 9 , Feb 2002

[3] S.Das, A. Misra, A.Dutta, A. Mcauley, et al., "An Intra-Domain Mobility Management for Next Generation Wireless Networks," IEEE PCS Magazine June 2002

[4] J.Rosenberg, H. Schulzrinne, et al, "SIP: Session Initiation Protocol, "Request for Comments 3261, Internet Engineering Task Force

[5] Henning Schulzrinne, Elin Wedlund, "Application Layer Mobility using SIP" ACM Mobile Computing and Communications Review, vol. 4, no.3, pp. 47-57, July 2000

[6] E. Gustafesson, A. Jonsson, and C. Perkins " Mobile IP Regional Registration" Internet Draft, draft-ietf-mobileip-reg-tunnel-06.txt, IETF work in progress

[7] K. El Malki et al., "Low Latency Handoffs in Mobile IPv4", Internet Draft, draft-ietf-mobileip-lowlatency-handoffs-v4-02.txt, IETF work in progress

[8] P. Calhoun et al., "Foreign Agent Assisted Hand-off," draft-calhoun-mobileip-proactive-fa-03.txt, IETF work in progress

[9] H. Schulzrinne, S. Casner, R. Fredrick, V. Jacobson, "RTP: A transport Protocol for Real-Time Applications" Request For Comments 1889, IETF