Global Internet Roaming with ROAMIP

Zoltán R. Turányi[b] zoltan.turanyi@ericsson.com
Csanád Szabó[a] szabocs@ttt-atm.ttt.bme.hu
Eszter Kail[a] eszterk@indigo2.hszk.bme.hu
András G. Valkó[b] andras.valko@ericsson.com
[a]HSN Lab, Budapest University of Technology and Economics
[b]Traffic Lab, Ericsson Research

Reachability and session continuity represent two distinct services that global mobility protocols should provide. Reachability is the possibility for Internet hosts to initiate sessions to mobile users. Session continuity refers to mechanisms that ensure that active transport or application layer sessions are not broken due to mobility. We present ROAMIP, a global mobility architecture that uses application layer solutions for global reachability and reuses transparent Mobile IP tunnelling mechanisms to ensure session continuity. ROAMIP eliminates long triangular routes, yet it is compatible with mobility unaware correspondent hosts. It is applicable to IPv6 as well as IPv4 networks. The ROAMIP architecture naturally lends itself to easy, gradual deployment and can reuse existing Mobile IP or SIP message formats. We conclude the paper by showing example traces from an experimental implementation.

I. Introduction

Micro-mobility protocols, such as Cellular IP [6], Hawaii [7] or Mobile IP Regional Tunnel Management [10] adopt a hierarchical mobility management architecture that separates local from global mobility. Movements inside a micro-mobility domain are handled locally and are hidden from the rest of the Internet. Mobility between micro-mobility domains, in contrast, is handled by a standard global mobility protocol. In this paper we limit our attention to global mobility.

Most of the existing micro-mobility protocol proposals assume Mobile IP [1] as global mobility protocol. Mobile IP takes the standpoint that mobility should be handled transparently in the network layer. It uses address translation to ensure that mobile nodes are reachable via permanent IP addresses. This allows mobile nodes to move without breaking active transport or application layer sessions. In addition, it is a generic approach that can support the mobility of any Internet node including clients, servers and routers. A number of authors argue, however, that Mobile IP is less appropriate to support Internet user mobility [8, 9]. In today's Internet user identity is increasingly separated from host IP addresses. This lowers the importance of reachability through a permanent IP address. Legacy client-server applications with mobile clients, in fact, do not need reachability sevice. Emerging push applications, on the other hand, tend to rely on application layer user identification and do not need transparent network layer mobility support to locate users.

These considerations motivate another class of solutions that, in contrast to Mobile IP, provide user mobility support as an application layer service [12, 13]. These protocols assume that mobile users have permanent identifiers that are not IP addresses. In addition, mobile nodes dynamically obtain IP addresses in each foreign network they visit. The binding between a mobile user's permanent identifier and his/her computer's actual IP address is stored by a global location service that operates in the application layer. Examples of such name services include (dynamic) Domain Name System (DNS), Session Initiation Protocol (SIP) registrars and proprietary protocols used in, for example, instant messaging services. Since mobile nodes have no permanent IP addresses, session continuity during mobility is either not supported [12] or is assumed to be provided by correspondent hosts. In [13], for example, correspondent hosts need to be SIP enabled and mobility aware. These requirements raise backward application/transport compatibility concerns if mobile nodes wish to access fixed WWW, FTP or other servers.

In this paper we present ROAMIP, a global user mobility architecture that combines the strengths of the approaches described above. We argue that global user mobility protocols need to provide two distinct services, that is, session continuity and global reachability. These two services are largely different in nature and therefore they should be handled separately. Session continuity is equally important for existing and future applications, transport protocols or correspondent hosts. To ensure backward compatibility, session continuity should be a transparent IP layer service. Reachability, on the other hand, is typically required by emerging interactive applications, such as mobile Internet telephony or instant messaging. Most of these applications use an application layer query to initiate sessions even in the case of non-mobile destinations. By mapping user identifiers to IP addresses these query mechanisms inherently provide mobility support and can be reused in a mobile environment to provide reachability. This precludes backward compatibility problems and motivates the application layer approach to reachability.

ROAMIP relies on IP layer tunnelling mechanism to ensure session continuity and on application layer directory services to provide global reachability. Mobile nodes have no permanent IP addresses, rather, they obtain a new IP address in each subnet they visit. The locally obtained IP address is advertised through location services via which mobile users wish to be reachable. This allows correspondent hosts to initiate sessions to mobile users without knowing that they are mobile.

Though a single standard for user identification and directory service is technically attractive, current tendencies indicate that a wide variety of name spaces and independent location directories will coexist. Depending on the applications they need, users will select one or several directories through which they wish to be reachable. The same user that is reached by IP telephony calls using SIP servers may also receive instant messages using his/her ISP's proprietary location service. Therefore, instead of defining its own location directory service, ROAMIP incorporates existing location mechanisms and is open for future ones.

ROAMIP is fully compatible with existing mobility unaware hosts and can easily interwork with Mobile IP. It is applicable to IPv6 as well as IPv4 networks and can support mobility between IPv4 and IPv6 domains. It can reuse the planned AAA infrastructure similar to the way currently envisioned for Mobile IP [5]. Triangular routes are eliminated without need for mobility support in correspondent hosts. ROAMIP builds on recent contributions to the IETF Mobile IP working group, particularly Route Optimization and local Home Agents [5].

This paper is structured as follows. In Section II we present an overview of relevant macro mobility proposals. In Sections III and IV we describe the base ROAMIP protocol and some possible options. Next, we show how ROAMIP can be implemented using standard Mobile IP or SIP message formats in Section V. Lessons learned from an experimental implementation are presented in Section VI. Finally, in Section VII we present concluding remarks.

II. Related Work

In [12] the Session Initiation Protocol (SIP) [11] is used to provide global user reachability. Each user is identified by an address that takes the form user@host. The host part of the address identifies the SIP server which keeps track of the user's location. Users moving to a different location register the new IP address with their SIP server. To initiate a session to a mobile user correspondent hosts need to contact the SIP server indicated in the destination's identifier. The SIP server relays SIP messages to the user's actual location. Alternatively, instead of relaying the messages, SIP servers can inform correspondent hosts about the current IP address of the mobile user. Ongoing SIP sessions of moving users are maintained by re-inviting the correspondent nodes using the new IP address. TCP and other non-SIP sessions break at handoffs and need to be re-established by the application.

The Host Mobility Management Protocol (HMMP) [13] extends the solution described above with support for TCP based applications. To this end, both mobile and correspondent hosts must contain a mobility aware SIP agent that monitors ongoing TCP connections. When a mobile host changes IP address, its agent notifies all TCP peers about the movement using SIP messages. Corresponding hosts are required to interpret these messages and use encapsulation to transmit the packets of an ongoing TCP connection to the new IP address. This solution does not require modifications in existing TCP implementations, but it assumes that all correspondent hosts are mobility and SIP aware.

Rather than using SIP, [9] proposes a new TCP option to handle mobility of TCP sessions. If a mobile host changes IP address, its TCP connections are re-established from the new location using the same sequence number space used at the old location. To prevent hijacking, peers create a security association on the initial establishment of connections. This security association is used to authenticate re-establishments. For mobile host reachability [9] proposes the use of dynamic updates to the Domain Name System (DNS).

Recent extensions to the Mobile IP protocol also depart from the original concept of permanent route anchor point. The Mobile IP Network Access Identifier (NAI) extension [4] allows mobile users to dynamically obtain home addresses upon power-up using the NAI as identifier. In [5] the authors introduce the notion of *local Home Agents* that are allocated in the domain where the mobile host first registers. The home domain of the host authenticates the user based on NAI but the tunneling function of Home Agents is performed by the local Home Agent. This solution eliminates triangular routes but it removes the possibility of being reachable through a permanent identifier. The NAI is only used by mobile users to identify themselves to the AAA infrastructure and not by correspondent hosts to reach mobile users.

III. ROAMIP

In what follows, we describe three main components of ROAMIP, that is, registration, reachability and session continuity. When arriving to a visited network, mobile nodes use the registration process to update location information in their home domains and in various location directories. By querying these directories correspondent hosts can initiate sessions to mobile users at their current location.

When mobile nodes move to a new network the registration process is repeated. In addition, the previous visited network is also informed about the node's new IP address. For session continuity, packets arriving to a mobile node's previously used IP address are forwarded to its current address. This mechanism provides the same transparent mobility support toward mobility unaware correspondent hosts that is provided by Mobile IP. Unlike Mobile IP, however, it limits tunneling to sessions that are active during a movement, which is advantageous if internetwork mobility during active sessions is rare.

III.A. Registration

Figure 1 illustrates the ROAMIP registration process. When a mobile node MN first contacts the visited network VN, it needs to perform authentication, obtain an IP address and register this address with its home domain. A local IP address can be obtained using, for example, Dynamic Host Configuration Protocol (DHCP). Authentication and registration functions are performed using a Local Server (LS) which in this respect is identical to a Mobile IP Foreign Agent. The Local Server notifies the mobile node's Home Server (HS) about the new IP address.

A mobile node's Home Server is similar to Mobile IP Home Agents in the respect that it is responsible for processing messages received from Local Servers. It does not forward, however, data packets sent to a mobile node. Rather, it contains a profile for each of its mobile users, listing the application specific location directories through which a particular user wishes to be reachable. For example, subscribers of an instant messaging service include the messaging server's directory in their profiles. Other mobile nodes may have Dynamic DNS entries or bindings in SIP registrars. When a HS is notified about the change of a user's IP address, it updates the appropriate location directories, as dictated by the user's profile.

The purpose of the registration process is to inform location directories of a mobile user's new IP address. This allows correspondent hosts to communicate with mobile nodes directly. Home and Local Servers do not forward data packets on behalf of mobile nodes, which represents a major difference compared to Mobile IP.

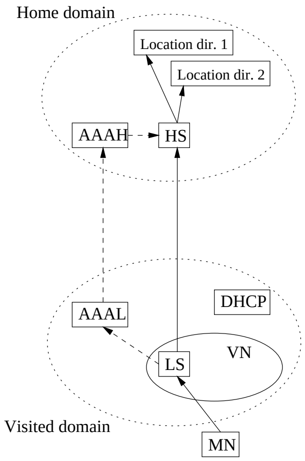

Figure 1: Registration messages in ROAMIP

When a mobile node first registers in a new domain, its registration and HS update may need to be authenticated for security and billing purposes. In Figure 1 this is performed using the AAA infrastructure, as proposed in [5] for Mobile IP. Upon receiving a mobile user's initial registration message, the Local Server contacts the local AAA server (AAAL) which, in turn, contacts the mobile user's home AAA server (AAAH). This message may also carry the HS notification message to avoid the need for two Internet traversals. In this case the HS is notified by AAAH instead of, as described earlier, directly by a LS. Subsequent registrations may use cached authentication information to avoid contacting the AAA servers.

III.B. Reachability

To reach client-server type services (e.g., WWW) the mobile node uses the locally obtained IP address. This results in direct routes between mobile nodes and correspondent hosts. In what follows, we describe the case where a session is initiated by a correspondent host CH toward a mobile node MN (see Figure 2).

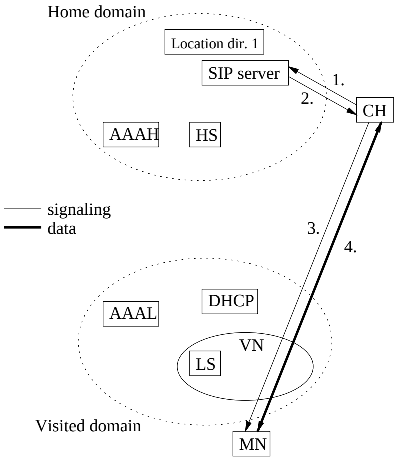

Figure 2: User data paths

Let us assume that the steps of authentication and registration have been completed as described previously. Without necessarily knowing that the destination is a mobile user, CH contacts the location directory as dictated by the application. To initiate an IP telephony call, for example, CH may send a SIP invite message (1) to the host indicated by the destination's SIP address. In our case this host is a SIP redirect server located in the mobile user's home domain. This server has previously been notified of the mobile user's current IP address using the mechanism described in Section III.A. This address is returned to CH in a SIP redirect (2) and is used by CH to contact the host directly (3). Data packets between the caller and the destination take the direct route between CH and MN (4). We note that this call setup procedure does not differ from calls directed to fixed destinations and does not require the caller to be mobility aware.

The example described above represents a typical case because the majority of Internet applications start from a textual identifier of the destination user, rather than from an IP address.

However, a few applications (e.g., a mobile WWW server) may want to be reachable through a permanent IP address. For these applications ROAMIP can incorporate a regular Mobile IP Home Agent, as illustrated in Figure 3. This Home Agent should be thought of as one of the location directories and is updated by the Home Server when the mobile node moves. ROAMIP mobile nodes that use a Home Agent are allocated a permanent IP address from the home network. Packets sent to this permanent IP address are captured and forwarded to the host's current address by the Home Agent. Figure 3 shows the resulting data path. As an implementation option, the Home Agent can be co-located or integrated with the Home Server.

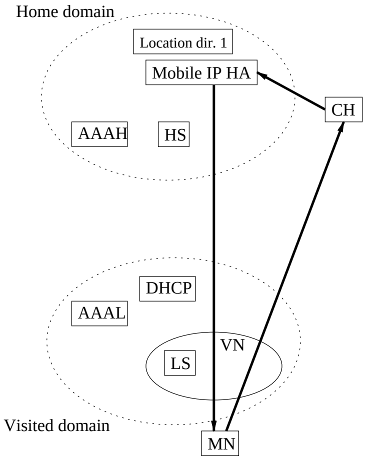

Figure 3: User data path when using Home Agent

III.C. Mobility and Session Continuity

In the example of Figure 4 a mobile node MN is engaged in an active data session with CH when it moves from one visited network VN1 to another network VN2 (step 1 in the figure). Upon arriving to VN2, the mobile node obtains an IP address and performs authentication and registration as described in Section III.A. Registration ensures that location directory entries corresponding to MN will be updated so that future sessions addressed to this user will reach MN directly. However, it is not a purpose of the registration process to notify correspondent hosts. In fact, correspondent hosts remain unaware of MN's move and the registration process is ignorant of correspondent hosts.

To ensure session continuity, MN includes the IP address of its previous LS (LS1) in the registration message sent to the Local Server of VN2 (LS2). LS2 sends an update message to LS1 and informs it about MN's new IP address (step 2). Upon receiving this message LS1 configures a local Forwarding Agent (FWA1) to capture future packets arriving to MN's old IP address and tunnel them to the new IP address.

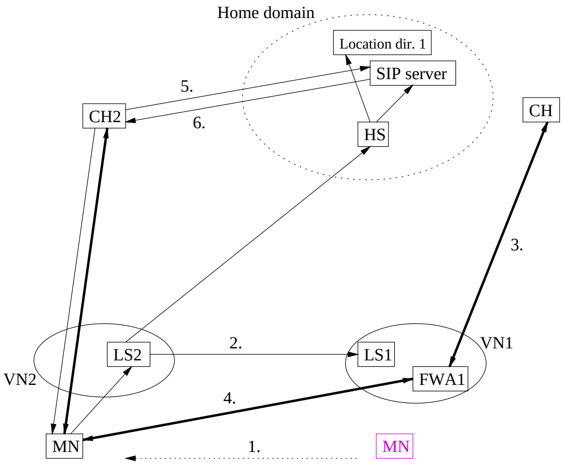

Figure 4: User data paths after moving to a new network

Figure 4 shows the data paths after MN moved. Since CH is not notified at mobility, it continues to send data packets to MN's previous IP address (3). These packets are captured by FWA1 and are tunnelled to MN's new IP address (4). The figure depicts a case with reverse tunnelling where uplink data packets belonging to this session are tunnelled to VN1 and from there on, routed to CH.

Such tunnelled routes are used only for those sessions that were initiated before MN moved to the new network. Sessions initiated after the move use MN's new IP address and take the direct path to/from MN. For comparison, Figure 4 also shows the path used by a session that CH2 initiated after MN moved to VN2. When CH2 contacted MN's location directory (5) to initiate this session, it obtained MN's new IP address (6) and therefore it communicates with MN without tunnelling.

This mechanism limits the impact of handoffs in both space and time. Mobility during active sessions implies that, in a geographical sense, the old and new networks are neighbors. In most cases this also means topological proximity, that is, tunnels from Forwarding Agents to Mobile Nodes will be short. In addition, tunnels are temporary and are used only by sessions started in the old network. Once all of the old sessions are over, the tunnel can be removed. In cases where neighboring networks are topologically distant or sessions started from the previous network last long, the tunnelling overhead introduced by ROAMIP will generally still not exceed the tunnelling overhead of Mobile IP with anchor Home Agents.

In the case of long sessions a mobile node may move multiple times during one active data session. In this case the mechanism described above is repeated at each move. The new Local Server notifies the previous LS but it does not have to know about networks that MN visited earlier. Data packets addressed to MN's earlier addresses will be forwarded along the chain of visited networks. In order to avoid multiple encapsulation of the same data packet, a FWA that receives a tunnelled packet from the previous FWA will replace the external addresses rather than encapsulate the packet once more. Long FWA chains are expected to be rare since mobile networks can cover large geographical areas and moving between networks is typically above data session time scale. One counterexample is telnet sessions that mobile users may wish to maintain for extended periods of time. Applications with typically long data sessions can incorporate the optimization method described in Section IV.A.

With the mechanism described above, ROAMIP provides the same transparent session continuity to mobility unaware correspondent hosts that Mobile IP does. ROAMIP separates the functions of Mobile IP Home Agents; the registration and reachability function is provided by Home Servers while packet forwarding functions are taken over by Forwarding Agents. In this respect Forwarding Agents can be considered as a distributed form of Home Agents. Anchor points are used only for those sessions where mobility really happened and tunnels only span distances that roaming mobile nodes really covered.

We note that in terms of operation ROAMIP Forwarding Agents also show similarity with Mobile IP Foreign Agents with binding cache [2]. A conceptual difference is that binding cache based smooth handoffs represent an optimization in case of distant Home Agents, while in ROAMIP Forwarding Agents provide the primary means of session continuity.

IV. Architecture Options

IV.A. Route Optimization

In a Mobile IP context route optimization [2] refers to the mechanism in which correspondent hosts are securely informed of a mobile node's actual location in order to shortcut triangular routes. This is of less importance in ROAMIP where correspondent hosts often have a direct route to/from mobile nodes using the base protocol. Suboptimal routes occur only when a ROAMIP mobile node moves during active data session. Applications with long-living sessions, however, can benefit of being notified when a mobile node moves.

Recently, a number of proposals have addressed mobility awareness in applications and transport protocols. In the case of SIP sessions, for example, when one of the endpoints move, it can notify its peer in a re-invite message [12, 13]. For TCP sessions [9] defines a mechanism to securely move one connection endpoint to a different IP address using TCP options.

All of these solutions are also applicable in a ROAMIP environment. A mobile node and its mobility aware correspondent are free to use any application specific mobility mechanism for route optimization. ROAMIP services for registration and session continuity are not affected by the use of such mechanisms.

IV.B. Placement of Local Servers

The role of Local Servers is to relay signalling messages and to support authentication in visited domains. These functions do not require that each subnet have a dedicated Local Server. Local Servers could be allocated for groups of subnets or a single Local Server could serve an administrative domain. If Local Servers are shared by several subnets then the address of the serving LS needs to be communicated to mobile nodes. This information can be inserted into DHCP messages or in an appropriate extension of router advertisements.

IV.C. IPv6

ROAMIP is applicable to IPv6 as well as IPv4 networks and hosts. Unlike in the case of Mobile IP, the entities and operation steps are identical in IPv4 and IPv6. The use of the AAA infrastructure, as described in Section III.A is also applicable in both cases. When used in an IPv6 environment, ROAMIP tunnels between Forwarding Agents and mobile nodes are replaced by source routing headers and IPv6 stateless autoconfiguration is used instead of DHCP. The similarities between ROAMIP for IPv4 and for IPv6 allow seamless mobility between IPv4 and IPv6 domains. ROAMIP can also easily integrate with Mobile IPv6.

V. Implementation Alternatives

We envision a mobile Internet that supports a diverse and growing set of applications. These applications may use a variety of independent location directory services. Mobile users should be able to be reachable through any of these directories without its visited networks being aware of the application. ROAMIP supports this flexibility because the registration process in a visited network is independent of the number and type of location directories used by a mobile user. The set of directories to update at registrations needs to be configured in the Home Server only.

To deploy ROAMIP, standardization needs to cover only those messages that affect visited networks. In what follows, we will show that existing standards can be largely reused. In particular, we show how ROAMIP can be implemented using standard Mobile IP messages with only slight modifications. We also show an alternative implementation of ROAMIP where we reuse messages from SIP.

V.A. Mobile IP Implementation

Figure 5 shows the mapping of the ROAMIP registration process to Mobile IPv4 messages. For ease of understanding we take an example where Local Servers are found in each subnet and are situated in the routers. Alternatively, Local Servers could be common to a domain comprising multiple subnets with only little modification to the process as described below. We also limit our attention to a case where the security model and movement detection follow the original Mobile IP model [1]. For example in the discussion we will assume that mobile nodes have manually established security association with their Home Servers.

Recent improvement proposals to Mobile IP, such as the use of a AAA infrastructure [5], could be directly incorporated but are not discussed here.

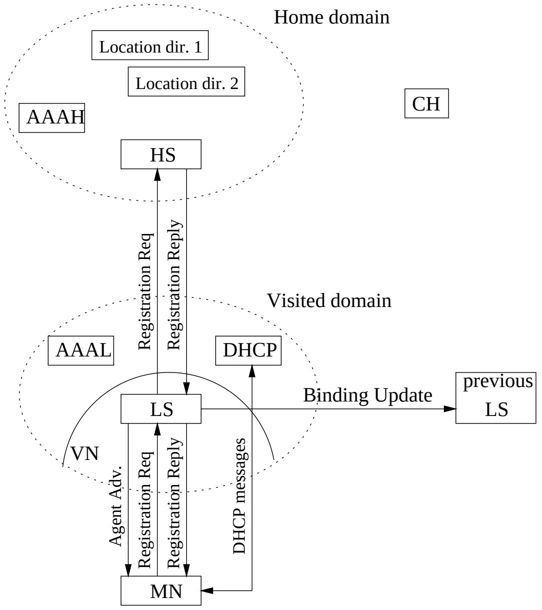

Figure 5: ROAMIP registration using Mobile IP messages

One can observe that in the registration process Local Servers act similar to Mobile IP Foreign Agents and Home Servers to Home Agents. We note again, however, that neither Local Servers nor Home Servers forward data packets on behalf of mobile nodes. In the registration process Mobile IP Registration Requests and Replies are used between mobile nodes, Local Servers and Home Servers, while Mobile IP Binding Update messages are used to notify a previously visited network for session continuity. In all of these messages the following general rules of mapping are used. Message fields containing a mobile node's home address are always set to zero. Fields originally used to carry care-of addresses contain a mobile node's IP address obtained locally. Home Agent address fields are used for Home Server addresses. In addition, each message must include a NAI extension. This is necessary because mobile nodes have no permanent IP addresses and therefore signalling structures (e.g., pending registrations or security associations) must be indexed by user identifiers. In what follows, we present a detailed mapping of ROAMIP registration on Mobile IP messages.

Local Servers emit periodic Mobile IP Agent Advertisements with the following parameters. The H and F bits are set to zero indicating that the message is not emitted by a Mobile IP Home or Foreign Agent. An additional bit needs to be defined to indicate ROAMIP Local Server. The message carries no care-of addresses.

Agent Advertisements are used by mobile nodes for layer 3 movement detection. To connect to a new network a mobile node first obtains an address using DHCP or stateless autoconfiguration in the case of IPv6. Next, it sends a Mobile IP Registration Request to the Local Server identified by the advertisement's source address. In this message the Home Agent address field carries the IP address of the Home Server and the home address field is set to zero. A NAI extension carries the NAI of the mobile user. In addition, a Previous Foreign Agent Notification Extension [2] is included to notify the previous Local Server. Other extensions, including authentication support, can be included as in Mobile IP.

The Local Server forwards the Registration Request to the Home Server. Upon receiving the request the Home Server updates location directories. Messages used in this step are not specific to ROAMIP and will vary for different directories. Next, the Home Server generates a Mobile IP Registration Reply to the Local Server with the Home Agent address field set to its own IP address. The Local Server relays this reply to the mobile node. Local and Home Servers may reject the registration similar to Mobile IP and this fact is communicated to the mobile node as in Mobile IP. ROAMIP registrations have a validity lifetime and must be renewed before they expire. Subsequent registrations are identical to the initial one.

Some applications may need to store additional information in a location directory besides the mobile user's current IP address. Such information elements may include terminal capabilities, application status or the current time zone of the user. To provide the necessary values, mobile nodes can include application specific extensions in registration request messages sent to Local Servers. These extensions are transparently passed by Local Servers to the Home Server which, in turn, uses them when updating the directories.

To achieve session continuity mobile nodes include a Previous Foreign Agent Notification Extension in the registration message when first registering at a Local Server. The previous Foreign Agent field of the extension contains the address of the previous Local Server, while the new care-of address field carries the IP address the mobile node obtained in the new network. Local Servers receiving such registration request messages send Mobile IP Binding Update messages to the specified previous Local Server on behalf of the mobile node similar to Foreign Agent Smooth Handoffs [2]. A NAI extension is added to the Binding Update messages to identify the mobile user. The home address field of the message is set to zero and the care-of address is set to the mobile node's new IP address. An additional bit needs to be defined in the header of Binding Update messages to distinguish ROAMIP specific binding updates. Upon receiving the message the previous LS returns a Binding Acknowledgment to the mobile node (not shown in Figure 5). When a Local Server or a mobile node wishes to delete a binding in a previous Local Server before its lifetime expires, it sends a Binding Update with the care-of address field set to zero.

V.B. SIP Implementation

Alternatively, ROAMIP can be implemented using SIP REGISTER messages as depicted in Figure 6. Local and Home Servers in this case are SIP registrars. After the mobile node MN moved to a new subnet, it uses DHCP to obtain an IP address and optionally the address of the Local Server. In case of IPv6 stateless autoconfiguration can also be used.

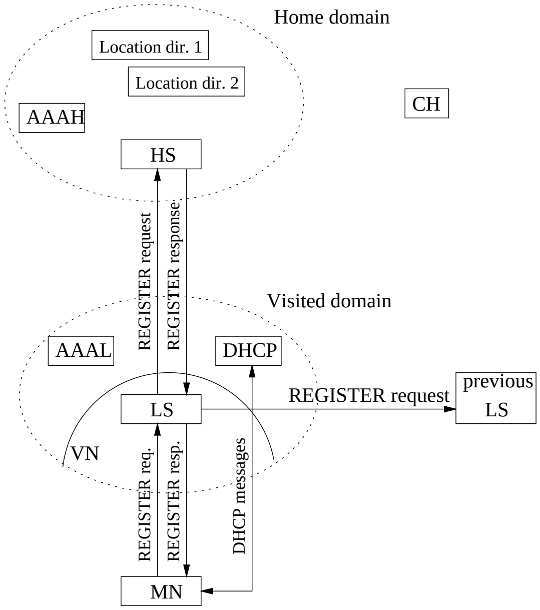

Figure 6: ROAMIP registration using SIP messages

To register its new IP address, MN sends a SIP REGISTER message to the Local Server. The mobile node can either directly send a unicast message if the address of the Local Server is available or it can address a (scoped) multicast message to the well known ”all SIP servers” multicast address. If SIP servers in the visited domain other than the Local Server receive the multicast registration message they should redirect the mobile node to the Local Server using a 3xx response with the Local Server in the Contact header. The Request-URI field of the REGISTER message contains the home domain of the mobile node, the To and From headers are both filled with the the SIP address (the NAI) of the mobile user, while the Contact header is used to store the new IP address of the mobile node. An expires parameter is appended to the Contact header to indicate the maximum lifetime of the registration.

When moving, mobile nodes must also de-register their previous IP address, as SIP registrars are able to maintain multiple registrations by default. This is performed by appending a second sip address to the Contact header with the old IP address and a zero expires parameter. REGISTER messages contain an Authorization header using, for example, the PGP authentication scheme with the private key of the mobile user.

Local Servers receiving SIP REGISTER messages from mobile nodes may perform authentication and authorization according to the local policies. A AAA infrastructure can also be utilized if available. Following any local processing, the Local Server forwards the REGISTER message to the Home Server identified by the Request-URI. If the AAA infrastructure is used for authentication, it can also be used to carry the REGISTER message to the Home Server similar to the Mobile IP case. Home Servers also authenticate the message, perform the registration and notify other location directories. Finally they respond with a SIP response message indicating success or failure as specified by SIP. In the response they specify a lifetime for the registration using the expires parameter of the Contact header. Local Servers forward this response to the mobile node and the registration is complete.

To inform the previous Local Server about the new IP address, mobile nodes append a new *Notify* header to the REGISTER message sent to Local Servers. In this header they specify the sip address of the previous Local Server. Upon receiving a successful response from the Home Server the new Local Server transmits a REGISTER message to the previous Local Server on behalf of the mobile node. The previous Local Server configures the Forwarding Agent of the previous network to forward packets arriving to the mobile node’s old IP address to the new IP address and responds directly to the mobile node. This way the mobile node is informed about the success of the registration itself and the configuration of the previous Forwarding Agent independently.

VI. Mobile Node Implementation Aspects

Mobile nodes represent the only component of the architecture that requires ROAMIP specific support in the operating system. After moving to a new network, packets belonging to a session opened in the previous network will arrive to mobile nodes using encapsulation. The operating system must decapsulate these packets before forwarding them to the applications. In the reverse direction packets need to be tunnelled to the previously visited network instead of being routed normally (reverse tunnelling). Unlike in the case of a Mobile IP mobile node with co-located care-of address, however, the tunnelling mechanism should affect only those sessions that were opened before moving to the new network. Data sessions opened in the new network should use the direct, non-tunnelled path to/from correspondent hosts. The need to route packets belonging to old and new data sessions differently requires changes in the operating system.

We have implemented a prototype of a ROAMIP mobile node in Linux 2.2.14. The implementation resides mainly in the kernel with a front-end in user space. The front-end is invoked after the mobile node has obtained an IP address in a new network. (Movement detection and DHCP are currently replaced by manual commands.) At this point the mobile node's network interface is reconfigured using the new IP address and a tunnel interface is created using the previous IP address. The kernel routing table is modified to contain the new network's default gateway. Using its newly obtained IP address the mobile node then sends a UDP packet to its previously visited network to configure the Forwarding Agent function. Finally, in order to distinguish existing data sessions from those that will be opened in the new network, all open sockets, except for listening sockets, are marked.

After a handoff is completed, packets transmitted via marked sockets are forced to use the IPIP tunnelling device as outgoing interface toward the previous network's Forwarding Agent. For this purpose we have modified the `ip_finish_output()` function to check for the ROAMIP mark in sockets before transmitting outgoing packets. Packets coming from new, that is, unmarked sockets are routed normally. The handling of incoming packets does not require kernel modification since in Linux packets arriving to the tunnel interface are decapsulated and are forwarded to the appropriate transport protocol by default.

The termination of tunnelled data sessions requires no additional action from a ROAMIP implementation. When sessions started before handoff are all closed, the corresponding marked sockets disappear and the ROAMIP tunnel will not be used further. At this point the tunnel interface can be removed and the mobile node can operate as if it had been switched on in the current network.

VI.A. Testbed Configuration

We have tested the ROAMIP session continuity mechanism in the testbed shown in Figure 7. The testbed contains two access networks, each of which consists of a single IEEE 802.11 wireless access point, co-located with the subnet routers (R1 and R2). The two access points are configured to use different radio channels in order to avoid overhearing in the air interface. The testbed contains a single mobile node with an IEEE 802.11 interface that is tuned dynamically to the radio channel of its current access point.

In this testbed all computers are Pentium PCs with Linux 2.2.14 operating system. The subnet routers, central router (R) and correspondent host use a standard, ROAMIP unaware kernel. The mobile node's TCP/IP stack has been modified as described earlier. Application software running in the mobile node and in the correspondent host are unaware of mobility.

VI.B. Experiments

Using the above configuration, we performed file downloads from a correspondent host to the mobile node using the `ttcp` utility. During the file transfer we manually triggered a handoff in the mobile node. Figure 8 shows packet traces of the file transfer captured in the correspondent host and in the mobile node. Black dots represent downlink data packets, while empty circles correspond to acknowledgements.

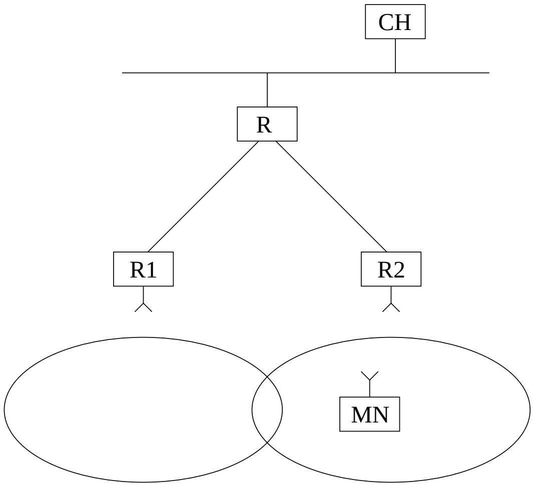

Figure 7: ROAMIP testbed

Handoff begins by the mobile node tuning its radio to the new access point and reconfiguring its interface using an IP address in the new subnet. From this point on, communication between the mobile node and the old access point becomes impossible over the air interface. In Figure 8 this appears as a stop of increasing packet and acknowledgment sequence numbers around 35 s. Communication with the mobile node using its old IP address is possible after the Forwarding Agent function has been configured in the old access point and open sockets have been marked for ROAMIP tunnelling in the mobile node. This point is indicated in the figure by a vertical line. The 200 ms delay up to this point can mostly be attributed to internal delays in our mobile node implementation. In this implementation most operations are performed from the user space front-end using scripts, which accounts for the large delay.

During this blackout period a number of TCP data packets and acknowledgments get lost. This results in a TCP timeout which accounts for the second part of the total silence period. Normal communication is resumed by the TCP source retransmitting a data packet around 35.5 s. We note that TCP timeouts could be avoided with a more sophisticated ROAMIP implementation.

The introduction of tunnelling may decrease path Maximum Transmission Unit (MTU), which may disturb ongoing TCP connections. This phenomenon is illustrated in Figure 9. This figure shows the same experiment as before with the exception that a larger initial TCP maximum segment size (mss) was used. In this case the first data packet retransmitted by the TCP source after handoff results in an ICMP error message (fragmentation needed) being sent back to the source. In response, the linux

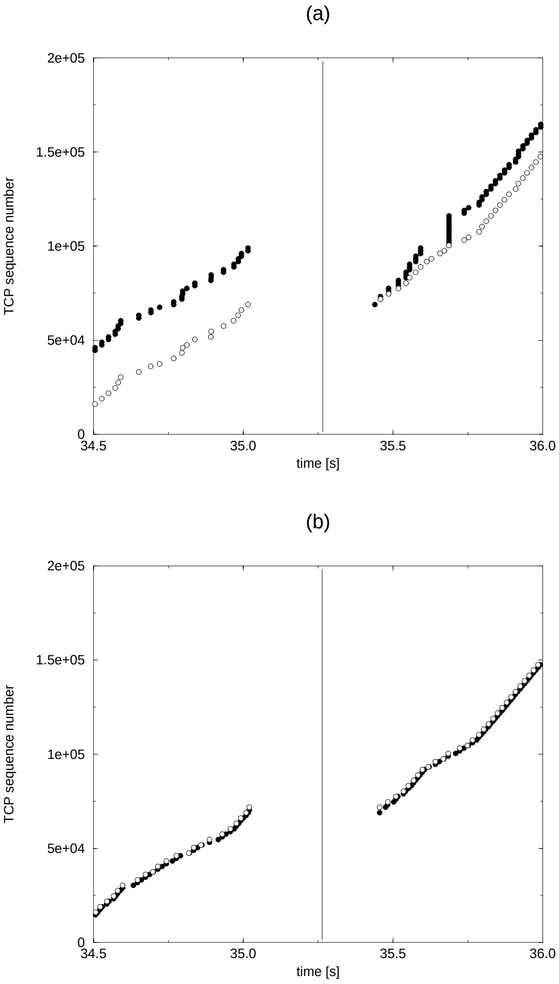

Figure 8: Handover during a file download captured (a) in the server, (b) in the mobile node

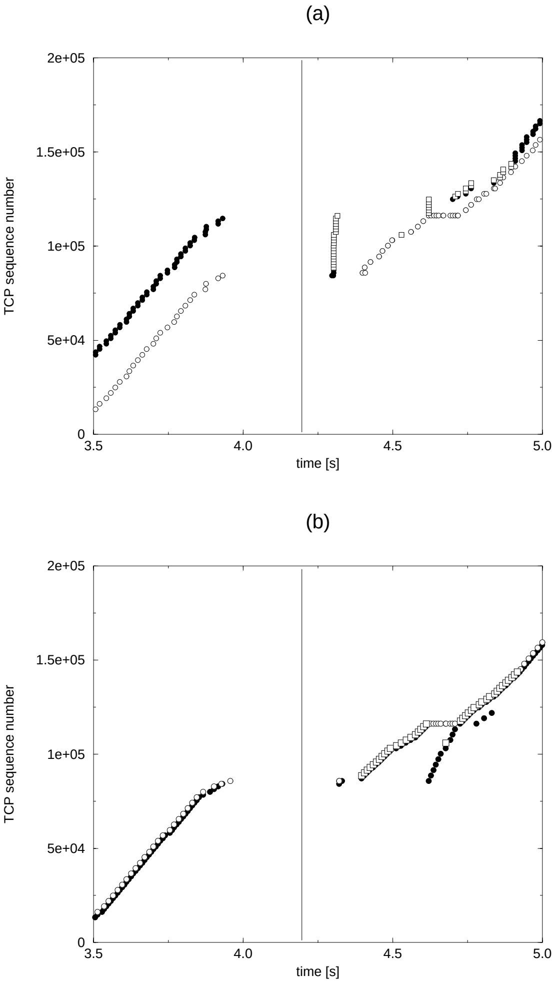

Figure 9: Handover with path MTU change during a file download captured (a) in the server, (b) in the mobile node

kernel retransmits unacknowledged packets in smaller TCP segments. In Figure 9 boxes correspond to these small segments. Normal operation is resumed around 4.9 s from where on the TCP source uses the new MTU. We note that this phenomenon can also be observed if the path MTU decreases due to reasons other than ROAMIP.

VII. Conclusions

We have presented ROAMIP, an inter-network mobility architecture that uses application layer protocols for user reachability and tunnelling for transparent session mobility. Tunnels, however, are used only for those data sessions during which mobility really happened and are limited to distances that roaming mobile nodes really covered. As a result, mobile nodes staying in a visited network can communicate with correspondent hosts without mobility overhead. This design is particularly advantageous if visited networks cover large geographical areas.

ROAMIP builds on the facts that user identity is increasingly separated from IP addresses and most applications need to map user identifiers to IP addresses even in the case of non-mobile destinations. Existing name-to-address mapping services are reused for mobile user reachability. This results in backwards compatibility with mobility unaware correspondent hosts.

The above considerations result in major conceptual differences compared to Mobile IP. ROAMIP mobile nodes have no permanent IP addresses. Consequently, Home Servers perform no packet forwarding on behalf of mobile nodes. Global reachability is provided via application level services rather than in the network layer. Mobile nodes normally communicate as if they were stationary and use tunnels only at movements.

Despite the differences ROAMIP can reuse existing Mobile

IP (or SIP) message formats and requires little standardisation effort.

References

[1] C. Perkins, editor, *IP Mobility Support*, Internet RFC 2002, October 1996.

[2] D. B. Johnson, C. Perkins, *Route Optimization in Mobile IP*, Internet Draft, `draft-ietf-mobileip-optim-09`, Work in Progress, February 2000.

[3] B. Aboba, M. Beadles, *The Network Access Identifier*, Internet RFC 2486, January 1999.

[4] P. Calhoun, C. Perkins, *Mobile IP Network Access Identifier Extension for IPv4*, Internet RFC 2290, March 2000.

[5] S. Glass, T. Hiller, S. Jacobs, C. Perkins, *Mobile IP Authentication, Authorization, and Accounting Requirements*, Internet Draft, `draft-ietf-mobileip-aaa-reqs-04`, Work in Progress, June 2000.

[6] A. Valkó, *Cellular IP: A New Approach to Internet Host Mobility*, ACM SIGCOMM Computer Communication Review, Vol. 29, No. 1, pp. 50-65., January 1999.

[7] R. Ramjee, T. La Porta, S. Thuel, K. Varadhan, S.Y. Wang, *HAWAII: A Domain-based Approach for Supporting Mobility in Wide-area Wireless Networks*, Proc. IEEE International Conference on Network Protocols, 1999.

[8] P. Karn, *Qualcomm white paper on mobility and IP addressing*, `http://people.qualcomm.com/karn/papers/mobility.html`, February 1997.

[9] A.C. Snoeren, H. Balakrishnan, *An End-to-End Approach to Host Mobility*, 6th ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '00), Boston, MA, August 2000.

[10] E. Gustafsson, A. Jonsson, C. Perkins, *Mobile IP Regional Registration*, Internet Draft, `draft-ietf-mobileip-reg-tunnel-02`, Work in Progress, March 2000.

[11] M. Handley, H. Schulzrinne, E. Schooler, J. Rosenberg, *SIP: Session Initiation Protocol*, Internet RFC 2543, March 1999.

[12] E. Wedlund, H. Schulzrinne, *Mobility Support using SIP*, Second ACM/IEEE International Conference on Wireless and Mobile Multimedia (WoWMoM '99), Seattle, Washington, August 1999.

[13] F. Vakil, A. Dutta, J-C. Chen, S. Baba, Y. Shobatake, *Host Mobility Management Protocol Extending SIP to 3G-IP Networks*, Internet Draft, `draft-itsumo-hmmp-00`, Work in Progress, October 1999.