RSVP Support for Mobile IP Version 6 in Wireless Environments

George Fankhauser, ETH Zürich
Stathes Hadjiefthymiades, University of Athens
Neda Nikaein, Eurecom
Lorraine Stacey, Lucent Technologies

November 17, 1998

Abstract

This draft describes a specific problem encountered when using RSVP (Resource Reservation Protocol) over optimised routes in MIPv6 (Mobile IP Version 6). The address translation in the MIP's binding cache creates a mismatch between the flow-id of the packets sent from correspondent node to mobile node and the flow-id signalled by RSVP.

We discuss several solutions to this problem: (1) By modifying RSVP at mobile and correspondent nodes to become aware of MIPv6 addressing, we provide a simple repair that allows RSVP flows to be established between the fixed network and mobiles; (2a) By adding optional objects to RSVP messages, a performance enhancement is proposed to make handovers smooth and seamless; (2b) A different technique with the same goal is called flow extension and it provides flows with fixed flow-ids from the correspondent node into the wireless access network at the expense of forwarding traffic inside the access network, whenever the mobile node moves.

We conclude that the minimal solution (1) is a requirement in order to make MIPv6 and RSVP interoperable; our favored approach requires only minor changes to the correspondent and mobile node's RSVP and MIP specification. However, for well performing and uninterrupted operation we strongly recommend one of the solutions (2a or 2b) that support fast re-establishment or preservation of resource reservations when mobile nodes move.

1 Introduction

Wireless access to the Internet is becoming more and more popular due to the rapid spread of wireless LANs and other wireless access technologies with various speeds and ranges (IRDA-LANs, IEEE 802.11, Wireless ATM, CDPD, etc.). Devices used include mobile phones, PDAs, and notebook computers. Mobile IP could become a common platform for mobile access (regardless of the underlying access technology), providing solutions to the many security, routing, and addressing problems. For a good introduction to Mobile IP see [12].

More traditional services like telephony, radio and TV broadcasting, and other new Internet applications that require more than best-effort service are increasingly being carried on the Internet. Mobile users also want to enjoy multimedia and other realtime services. Therefore it is important to design Mobile IP to interoperate seamlessly with protocols that provide real-time services in the Internet.

This draft focuses on issues of MIPv6 and RSVP interoperability, more precisely on the operation phase where optimised routing between fixed and mobile hosts is used. We apply our considerations to MIPv6 because it uses route optimisation by default. However MIPv4 may also use this technique [13].

Although RSVP is standardised [7] it currently looks like it is going to be employed for access networks but not necessarily for end-to-end connections through the backbone networks. Solutions for bridging RSVP across the backbone are under development [6] and are being discussed in the IETF intserv and diffserv working groups. We assume here the usage of RSVP for wireless and fixed access networks and its transparent operation over the backbone.

1.1 Terms

This draft uses the terms as defined in section "Mobile IPv6 Terms" in the Internet Draft "Mobility Support in IPv6" [9].

1.2 Problem statement

1.2.1 The Unicast Case

An integral part of MIPv6 is the use of optimised paths between a mobile node and correspondent nodes. Routing of packets via the Home Agent is typically only a transient phase in MIPv6. If the mobile node has one or more packet streams (flows) controlled via RSVP, the PATH and RESV messages will also be transmitted via the direct route between the mobile node (MN) and correspondent node (CN). When a MN moves, the optimised route between the MN and CN will change. Furthermore the MN's care-of address will also change. Both the route and care-of address change impact the operation of RSVP.

In RSVP, PATH messages are sent end-to-end and RESV messages are sent hop-by-hop. When the *CN is a sender* it will transmit a PATH message where the address details are based on the MN's home address. However the outer header details will be modified by the MIPv6 binding cache to contain the MN's care-of address as the destination address. Even more importantly, to ensure PATH and RESV messages follow the same route, the PATH messages contain a RSVP_HOP object which collects the address of each outgoing interface the message traverses. The correspondent node and intermediate routers will each determine the outgoing interface based on the MN's home address. However the packet will actually be routed based on the MN's care-of address. Thus when the PATH message reaches the MN the routing information stored in the RSVP_HOP object will not correctly reflect the route followed by the PATH messages. Hence RESV messages can not be routed back to the CN and therefore RSVP state will not be created between the MN and the CN. Additionally, RESV messages sent back from the MN reference the home address of the

MN as the flow's destination in the SESSION object. Again, this is not the real flow-id of the flow's data packets, it should rather reference the MN's care-of address.

For the unicast case, the problem of a flow-mismatch occurs when the CN is the sender. When the MN is the sender, the PATH message will contain correct routing information when it reaches the CN because the MN directly addressed it to the CN. Hence the RESV message can be forwarded correctly along the reverse path to the MN. Also, the RESV message's SESSION object sent by the CN contains correctly its own (fixed!) address as the flow's destination. However, the PATH messages also contain the sender's address (MN) in the SENDER_TEMPLATE object. This address would be normally the home address of the MN and has to be changed to its care-of address.

Solutions to these 'change of route' and 'flow-mismatch' problems are described in this Internet Draft. Another important issue is that when the MN moves to a new subnetwork its care-of address changes. In standard RSVP state information is based on the MN's home address. However the data packets this state needs to apply to will be addressed with the MN's care-of address. Hence there is an issue of matching RSVP flow state with the data packets belonging to that flow at intermediate routers.

Yet another difficulty with MIPv6 is its applicability to wireless environments where the MN moves to new subnetworks in real time. In wireless environments the speed of re-routing is of major concern. The objectives are to minimise the loss of packets, and restore the end-to-end RSVP state as quickly as possible.

1.2.2 The Multicast Case

For the multicast case, there are two options for MNs to receive data. The first option is via its Home Agent. In this case the HA is also a multicast router. The HA tunnels the multicast packets to the MN. The second option is for the MN to join its groups directly in its foreign network. This option allows the MN to receive the multicast packets directly without passing by its HA. Since this draft considers the problem of using RSVP over optimised routes in MIPv6, only the second option is studied here.

Let's consider the case where the CN sends multicast data and at least one of the receiver's is a MN. An RSVP session is identified by the triple <destination address, protocol id and optionally destination port>. The destination address of a multicast group remains fixed regardless of the mobility of its receivers. In other words, the MN's exact address and its mobility have no impact on the multicast group address. Therefore, an MN making a handover to another network can be considered as a new node joining the group. That is, each time the MN moves to a new subnetwork it must leave and re-join the multicast group. The process of leaving and re-joining the multicast groups can be done as part of the handover procedure.

Now, we consider the case where the MN sends multicast data. IP unicast routing protocols depend only on the IP destination address. Some multicast routing protocols such as DVMRP [18] and MOSPF [10] build a multicast tree per source. These routing protocols construct the multicast tree depending on the source address as well as the destination address. Other algorithms like CBT [5] build a single shared tree per group. These protocols use only the IP destination address for packet forwarding.

Using the home address as the IP source address of the datagram leads the routers executing DVMRP or MOSPF to expect the datagram from the link used to reach the MN's home address. Therefore, sending multicast traffic in a foreign network using the MN's home address as the IP source address means that these routers receive the traffic via unexpected links. DVMRP drops such datagrams because it expects the packet to arrive to the router on the shortest path from the router to the MN's home address. MOSPF forwards the traffic based on an incorrect information. In both cases, some destinations may not be reached. In order to overcome such routing problems, we must use the MN's care-of-address in the IP source address.

RSVP is not aware of mobility. When building a PATH message, the IP source address of the filter specification is filled by the the MN's home address. However the IP source address of the outer IP header must be modified to the MN's care-of-address because of the routing problems described above. Therefore, there is a mismatch between the information in the RSVP message and in the outer IP header. The PATH message can be routed correctly thanks to the care-of-address in the outer IP header. The problem is RSVP_HOP object's last-hop-address field which has to be changed at the MN's RSVP daemon because it has been calculated based on the MN's home address and the group destination address.

1.2.3 Summary

Summarising, the following issues have to be addressed:

- Most importantly, we need a working solution, i.e., MIPv6 and RSVP need to be integrated. MIP provides transport layer transparency but does not work with protocols like RSVP. Since this requires changes to either MIP or RSVP, we would like to keep them as minimal as possible. We also favor a modular solution that allows for incremental deployment (e.g., hosts first, routers later).
- Multicast aspects of all possible solutions and extensions (even if optional) have to be considered.
- Performance aspects for wireless operation and their respective applications (e.g. mobile IP phones) are an important factor when evaluating the pros and cons of a mobile integrated services solution.

2 Possible Solutions

This section describes solutions and enhancements that resolve the flow-id mismatch between RSVP messages and the flow's data packets. The first solution provides a basic fix and some optional enhancements to restore reservations in a quick and efficient manner. The second proposal builds partially on top of the first and provides advantages to fast moving terminals inside wireless access networks by extending flows to new routers while keeping the RSVP session through the backbone unchanged. This approach makes use of RSVP tunnels.

2.1 Mobility Enhanced RSVP

2.1.1 Changing End-Nodes

A possible solution to the problem of correctly routing RSVP messages between CNs and MNs is to modify the RSVP daemon at the CN and MN to operate on the MN's care-of address instead of its home address. The RSVP daemons could learn the care-of address by consulting its MIPv6 binding cache. This means that RSVP state would be created based on the care-of address, and thus the path the traffic actually follows.

There are two ways that the care-of address can be obtained. One option is to modify MIP to *provide an interface* [1] that allows the RSVP module to look up the care-of address of a mobile node. An alternative is to modify MIP at the CN and MN only. In this case the MIP module needs to become RSVP aware and swap the home address in the PATH and RESV messages SESSION objects with the mobile nodes care-of address (among other fields, see Section A for details).

[1] It is current practice of MIPv4 implementors to use the routing table to store the MN's location at home agents [12]. If MIPv6 would use the same technique *consistently*, a simple routing socket lookup (PF_ROUTE) would reveal the care-of address when fed with the home address of the MN.

We recommend here the implementation of a clean interface in MIPv6 which must be used by the RSVP daemon. As we will see in the next subsection, this interface may also be extended for triggering PATH messages when mobiles change their location.

2.1.2 Changing Intermediate Routers (Alternative Approach)

An alternative, but similar approach means RSVP implementations at routers MUST be changed but changes are not needed at CNs[2] and MNs. In this solution outer header address information is passed up to the RSVP module at each router (outer header means the IP header transporting the PATH or RESV message, i.e., the whole packet and not only the payload is forwarded to the RSVP daemon). This allows the RSVP daemon in each intermediate router to learn the mapping between the mobile node's home address and current care-of address. The RSVP daemon should then base its calculation of the RSVP-HOP and filters on the care-of address. Given the router RSVP daemon maintains a mapping between the home and care-of address when the care-of address changes the router will still recognise the RSVP messages and traffic as belonging to the same reservation.

2.1.3 Evaluation of the two Alternatives

Our preferred solution, however, is the first approach. This is because it requires less changes. It requires to make a small change at CNs and MNs only, to enable the RSVP daemon to learn the MN's care-of address. This solution can be optimised by adding new RSVP objects. However these are not essential to the operation of RSVP with MIPv6. In contrast the second, alternative solution requires all routers to change their RSVP implementations.

2.1.4 Performance Enhancements

One minor disadvantage of this approach (using either implementation alternative) is that when the mobile node moves and thus obtains a new care-of address, all of the intermediate routers will assume this is a new RSVP reservation flow. Hence there may be situations where the new reservation is denied because the old reservation is still active and blocks resources. This problem could be overcome by introducing a new RSVP object to the RSVP messages the MN sends (that is, if the MN is a receiver, it will place the RSVP object in the RESV message). This would allow intermediate routers to recognise the reservation is the same even though the care-of address has changed. A key point to note is that if some or even all intermediate routers do not recognise this RSVP object this solution will still work. At those routers that don't understand the RSVP object the RSVP state with the new care-of address will be treated as a new independent flow and the previously reserved flow expires later.

Note if the home address was kept as the destination address, and the care-of address was stored in the new RSVP object this solution would require all intermediate routers to understand the new RSVP object instead of just changing the CNs and MNs and treat the new RSVP object as an option.

Another concern is the time required to modify the RSVP state so that traffic flows along the optimal path to the MN's new care-of address. In standard RSVP operation PATH and RESV messages are transmitted periodically. Hence there can be a significant delay between the MN's care of address changing and the next PATH message being sent. This extra delay can be avoided by having the arrival of a binding update message at the CN, trigger the transmission of a PATH message. Again, an interface between MIP and RSVP daemon should be used for this triggering.

[2] If a CN as a flow's sender exercises traffic control, it has, of course, to apply the same changes as intermediate routers.

2.1.5 Use of IPv6 Flow Label

One could argue that the mechanism described above is not required, since *IPv6 flow labels* (in conjunction with the source IP address) uniquely identify the traffic flow. If non-zero flow labels are employed it is still possible to identify the traffic flow. However, this won't solve the routing problem of PATH messages. Moreover IPv6 flow labels are optional and hence they can often be zero. If this is the case the flow-id mismatch problem still exists. Furthermore, if the flow-id is created by hashing on the destination address, it may also change when the care-of address changes.

2.1.6 Required or Optional Changes

To summarise the key components of this solution are:

- At correspondent nodes
  - RSVP daemon can obtain the mobile node care-of address from the binding cache
  - RSVP daemon places the mobile node care-of address as the PATH message's SESSION object's destination address
- At the MIP information aware intermediate router
  - On receipt of a PATH message store the mobile node care-of address (standard RSVP operation) and home address (optional for mobility support) in the RSVP daemons flow table.
  - Create a classifier entry based on the mobile node's care-of address (standard RSVP operation).
  - When a PATH message arrives with the same home address but a different care-of address update the flow state and filter information with the new care-of address (optional for mobility support).
- At standard intermediate routers
  - On receipt of a PATH message store the mobile node's care-of address in the flow state information.
  - Create a classifier entry based on the mobile node's care-of address
  - On receipt of a PATH message with a different care-of address for the mobile node create new flow state information and filters.
  - Remove the old flow state information on receipt of a PATHTear or when it times out.
- At mobile nodes
  - Since the MN is reachable under its care-of address, PATH messages that arrive with the care-of address as the SESSION object's destination address should not disturb regular RSVP operation
  - When an MN sends RESV messages the SESSION object must also contain the care-of address in order to correctly identify the flow on the optimised route.
  - RSVP daemon places the mobile node's home address in a new RESV MIP RSVP object (optional) for efficient recycling of resources.

The solution described above means RSVP implementations at CNs and MNs MUST be changed and RSVP implementations at routers SHOULD be changed.

2.1.7 Multicast Support

The same approach can be used to resolve the problem of multicast reservation when an MN is a sender. In this case, we can also use the care-of-address in the IP source address field of the SENDER_TEMPLATE. Therefore, the RSVP_HOP object of the PATH message is calculated correctly. However, this can cause the intermediate routers to consider the MN as a different source. In order to overcome this problem, we use the same approach as described above. We add a new RSVP object to the PATH message which contains the MN's home address. To summarise, we propose the following solution:

- At mobile nodes
  - When an MN sends PATH messages the SENDER_TEMPLATE object must contain the care-of address.
  - RSVP daemon places the mobile node's home address in a new PATH MIP RSVP object.
- At the MIP information aware intermediate router
  - On receipt of a PATH message store the mobile node care-of address (standard RSVP operation) and home address (optional for mobility support) in the RSVP daemons flow table.
  - Create a filter spec entry based on the mobile node's care-of address (standard RSVP operation).
  - When a PATH message arrives with the same home address but a different care-of address update the flow state and filter information with the new care-of address (optional for mobility support).
- At standard intermediate routers
  - On receipt of a PATH message store the mobile node's care-of address in the flow state information.
  - Create a filter spec entry based on the mobile node's care-of address
  - On receipt of a PATH message with a different care-of address for the mobile node create new flow state information and filters.
  - Remove the old filter spec information when it times out.
- At correspondent nodes
  - The CN is aware of the MN's care-of address via the binding cache. RSVP operates regularly by sending the RESV message.

2.2 Flow Extension

This section describes an alternative approach for the efficient operation of RSVP on top of MIPv6. This approach proposes an alternative to the use of new RSVP objects for the fast update of end-to-end RSVP connections. It mainly refers to the time periods following the occurrence of active terminal relocations in the network (i.e., handovers). It is also assumed that, prior to terminal relocations, RSVP flows have been established between the CN and the MN using the basic approach discussed in Section 2.1 (the MN is assumed to be in a foreign subnetwork).

When an MN attaches to a new subnetwork and acquires a new IP care-of address, the MIP-capable router in this subnetwork intercepts and suppresses the MIPv6 binding update message sent to the CN (we use the term MIP-router here for a router serving wireless access networks and performing the flow extension tasks). This causes the CN not to update its Binding Cache. This strategy is not applied to the Binding Update sent to the former MIP-router. This information can be used in rerouting best effort traffic destined to the old location of the MN, to its current location (through the use of an IP tunnel between former and current MIP router). The MN is then capable of receiving datagrams destined to its current IP address as well as the previous IP address. For best effort traffic destined to the MN, the previous IP address should be used. This packet forwarding technique is an enhancement to support loss-less handovers [8].

Prior to or during the relocation of the terminal, the old MIP-router also extends the existing RSVP flows to the new MIP-router. This task is performed by a mobility management (MM) entity operating within the router. The extension of downlink (CN $\rightarrow$ MN) flows is performed by the old MIP-router while the uplink flows (MN $\rightarrow$ CN) are handled by the new MIP router. For this last step, the new MIP-router needs to receive the characteristics of existing flows from the old router. For this task, specialised signaling between MIP-routers (MM entities in particular) is introduced (it is assumed that such signaling is treated as best effort traffic in the access network). The elongation of flows by the old MIP router to the current *avoids the invalidation of existing flows* caused by changes in the IP addresses of connection endpoints. [3]

The proposed elongation of the CN-MN path causes the route of the communication to be sub-optimal and, consequently, imposes additional, but relatively small, time delays. It was shown that consecutive elongations (leading to increased sub-optimality) will be needed rarely, as the number of inter-subnet MN relocations (i.e., handovers to subnetworks controlled by different MIP-routers; and hence changes in addresses) is very small during the lifetime of IP connections in a CPN environment [17]. Most relocations will result in attachments of the MN to base stations (radio transceivers) in the same subnetwork (also termed intra-subnet handovers). Such mobility events can be handled at the local addressing-level without affecting the communication path beyond the router, i.e. the care-of address of the MN remains the same. This elongation of data paths has been adopted in Wireless ATM technology for the handling of connections (VCs) during the occurrence of handovers. Relevant information can be found in [2], [3], [8].

2.2.1 Required or Optional Changes

Due to the flow extension approach modifications are confined to the end-nodes and mobile/wireless access network routers. Intermediate routers along the transmission path do not need to be changed.

- At correspondent nodes
  - The same basic modifications as described in Section 2.1 are needed.
- At standard intermediate routers
  - No changes needed here (of course, one could think of a combination of the two performance enhancment approaches, i.e., flow extension is used for high-frequent movements and path triggers and RSVP objects are used, e.g. when roaming between provider networks).
- At wireless access network routers (MIP-routers)
  - Communicating the IP address of the new MIP-router to the old MIP-router.
  - Transmitting the existing RSVP flow characteristics (flowspec) to the new MIP-router from the old MIP-router. To elongate the RSVP state from the old MIP-router to the new MIP-router (i.e., extend the flow) an RSVP tunnel could be used.

[3] If Guaranteed Service is employed as a service model the end-to-end delay calculation may have to be performed across the entire path between the CN and MN to take into account extra nodes and links that have been added to the communication path.

- Control and suppression of binding update transmission from MNs. If the design alternative mentioned below is used, this point becomes obsolete.

- At mobile nodes
  - The same basic modifications as described in Section 2.1 are needed.
  - However, since the flow will retain its flow-id over a long time period, binding updates do not need to trigger RESV messages, nor do we need to insert special RSVP objects.
  - Design alternative: It is also possible to suppress the binding update messages at the MN without considerable modifications to its MIPv6 module. Such approach improves bandwidth efficiency at the radio interface and reduces the complexity of MIP routers. One additional advantage of this alternative is that it does not cause disruptions in IP communication if the Binding Update message is included (as Destination Option header) in IP packets having, besides headers, standard payload data such as TCP/UDP. Disabling the transmission of Binding Update messages at the MN is also adopted in the MRSVP approach discussed in Section 3.2.

2.2.2 Using RSVP Tunnels for Flow Extension

Lastly, we are considering how the extension of RSVP flows could be accomplished in light of existing protocol specifications, (i.e., RSVP, MIPv6, IP encapsulation). RSVP operation over IP tunnels [16] provides a good basis for the implementation of the proposed scheme (see also Section 3.3). The old MIP-router uses regular IP tunnels for forwarding best effort traffic and RSVP tunnels for handling the extension of RSVP flows. The old MIP-router serves as the RSVP tunnel entry point in the downlink direction while the new MIP-router plays the role of the tunnel exit point (roles are inverted in uplink communication). The tunnel session is a separate RSVP session between the involved routers. Its characteristics are dictated by the characteristics of the flows that need to be extended. The original session (CN $\rightarrow$ MN / MIP-router) views the tunnel as a single communication link. The PATH and RESV messages of the end-to-end session are encapsulated at one tunnel end-point and decapsulated at the other. The end-to-end session and the tunnel session are associated at the entry/exit points of the tunnel. The tunnel may encompass one or more RSVP-capable nodes.

The overall scheme is based on the assumption that the new MIP-router is aware of the existence of RSVP flows and thus, suppresses only the binding update messages for active RSVP flows When the entire set of RSVP flows is terminated, the new MIP-router allows the propagation of the binding update signal to the fixed network. This restores the optimal communication between the MN and the CN regarding best effort traffic. (We might also propose to do this whenever flow extension has become too long or too slow).

2.2.3 Multicast Support

As we said before, if the MN, which is the receiver of the multicast data, changes its subnetwork, it must rejoin its groups in the new subnet. Now, if the new subnetwork already has members of the MN's groups with the same reservations, the MN can receive the data without any delay. If this is not the case, the MN can receive data from its old MIP-router by using an RSVP tunnel. The new MIP-router knows about the MN's group list and also about the presence of groups and their reservation style in its local network. Therefore, instead of trying to graft a path to the multicast tree, the new MIP-router asks the old MIP-router to forward the traffic destinated to this group via an RSVP tunnel. The same thing happens, if the new MIP-router has a member of the group but not with the specified reservation style.

Now if the MN is the sender of the multicast traffic (uplink direction), we always pass by the RSVP tunnel to reach the old MIP-router.

3 Previous Work

This section summarises selected recent research on providing QoS support in mobile and wireless networks and serves as a starting point for further reading.

3.1 Mobile IP Version 6

The description MIPv6 [9] encompasses a streamlined version of Mobile IP that makes use of new IPv6 features that help to improve mobility support. Most notable, foreign agents could be replaced by stateless address autoconfiguration and neighbor discovery. As stated earlier, optimised routes between CN and MN are the default operation in MIPv6. They can be implemented properly by using IPv6 routing headers.

3.2 Mobile RSVP (MRSVP)

MRSVP has been proposed in [15], as an extension to conventional RSVP, for the provision of QoS guarantees to MNs independently of their movement throughout the access network. MRSVP suggests making the required resource reservations in all the locations expected to be visited by the MN (mobility specification). MRSVP supports two service classes. The first one, called Mobility Independent, provides the agreed service in every location visited by the MN. The second, called Mobility Dependent, provides a high probability, but no guarantee, to receive the agreed service in the locations the MN may visit.

MRSVP supports two different reservation styles. The MN makes active reservations from its present location towards all the CNs it is communicating with. It also triggers the establishment of passive reservations from all the locations it may visit. Such reservations are made by proxy agents (remote), operating on the behalf of the MN which is not present at their subnetwork. As passive reservations are not used by the MN (data is not flowing through them) they may be used temporarily by other connections of the Mobility Dependent class. When the MN roams, passive reservations are switched to active and vice versa.

The MRSVP scheme provides mobility support for both unicast and multicast traffic.

3.3 Supporting RSVP/MIP Integration with RSVP-Tunnels

RSVP tunnels [16] provide signalling of RSVP messages in tunnels. Normally, RSVP messages in tunnels are not visible to routers due to encapsulation. This problem is similar to the MIP flow-mismatch problem, but with the difference, that in MIP the RSVP daemon is provided with the wrong information and in tunnels, the encapsulated RSVP messages feature the wrong protocol id (next header field) which makes them invisible.

The proposed solution for RSVP tunnels is to establish nested RSVP sessions between the tunnel's entry and exit points, i.e. new PATH and RESV messages (without the encapsulation headers!) are being sent between entry and exit point. This solution also focuses on modifications at the entry/exit points and tries to avoid changes to the intermediate routers along the tunnel path.

RSVP tunnels were already proposed to support RSVP connections in MIPv4 as a replacement for the regular IP tunnel between HA and MN. They could serve also as a basis for the various forwarding connections needed in our flow extension approach (Section 2.2).

RSVP tunnels could also be employed in our RSVP object approach (Section 2.1) for the transient phase where the MN has moved to a new care-of address and the old router is forwarding traffic to the new router, but the end-to-end state hasn't been adjusted yet. Thus commonality with MIPv4 - RSVP integration in terms of the use of RSVP tunneling is true for both approaches.

3.4 IPSOFACTO: IP Multicast over Wireless ATM

[1] describes how IP multicast can be natively supported in wireless networks over ipso-facto. To differentiate different types of traffic VCs are classified as unicast, multicast and broadcast VCs. Small extensions to the IGMP protocol are proposed. In addition, a cell-level error recovery scheme at the data link layer is described to enhance the packet-level throughput of the transport layer.

3.5 Designing a Wireless Broadband IP System with QoS Guarantees

[4] describes a wireless broadband IP system where a solution such as those described in this Internet Draft is required. This system, specified within the ACTS Magic WAND project is a native IP wireless access system providing 20Mbits/s, 5GHz radio access to an IP backbone. The objective of this system is to allow full exploitation of real-time IP applications in a mobile environment. This system comprises a mobile node, access point (implementing all radio dependent control functionality) and a mobility enhanced IPv6 router.

This work is also described in [14]. In this report, in addition to the overall design of the wireless broadband IP system, some of the MIP/RSVP integration problems are described in greater detail, including diagrams and message sequence charts.

3.6 Mobility Management in an IP-based Wireless ATM Network

In [8], Hadjiefthymiades et al. studied mobility management issues in a WATM architecture exclusively intended and optimised for IP traffic support. The considered system is a variation of the Magic WAND architecture. It is capable of supporting IP applications with specific QoS requirements in light of terminal mobility. The proposed architecture is based on specifications/standards like RSVP, MIPv6 and IFMP.

The architecture assumes two types of terminal relocations (handovers). Handovers confined to the same MIP router are treated at L2 (ATM in the case of a WATM system). Handovers involving more than one MIP routers (intersubnet handovers), necessitate the use of resource and mobility management schemes. The paper identifies the problems encountered in MIP-RSVP interaction and presents some initial thoughts for their resolution.

3.7 Wireless Multicasting in an IP Environment

[11] describes a framework for multicast communication in a wireless IP environment. It presents a group membership management mechanism optimised for wireless networks. It also studies the effect of terminal mobility on the multicast communications in a mobile IPv6 environment.

4 Conclusion

As stated in Section 2.1 the *minimal solution* to the problem of MIP/RSVP integration requires the modification and the interfacing of the RSVP daemon and MIP's binding cache at CNs and MNs. This solution requires less changes when compared to an approach that tries to fix flow-ids at intermediate routers. It is important to note that interfacing MIP and RSVP implementations requires changes to both standards. For proper multicast operation the same changes as for unicast are sufficient.

For advanced solutions, where performance and smooth handovers in wireless environments are considered, we have proposed two solutions:

1. *Triggers/Objects:* PATH messages are triggered on binding updates and home address objects in RESV messages enable intermediate routers to recognise connections and to reuse resources even when the care-of address (destination address of flow) changes.
2. *Flow Extension:* This approach keeps the reservation unchanged until it reaches the wireless access network.

A qualitative comparison of the two approaches shows the following result:

	Triggers/Objects	Flow extension
Changes to CN	yes (needed for minimal solution)	yes (needed for minimal solution)
Changes to intermediate routers	yes (RSVP MIP object extension and reuse of flow's resources)	no
Changes to MIP-router	no (forwarding of late packets is also an option here)	yes (binding update interception, flow forwarding)
Changes to MN	yes	yes
Changes to HA	no	no
Supports multicast delivery	yes	yes
Bandwidth efficient	yes	yes (we assume overprovisioning in) the access network)
End-to-end delay	always shortest path (but re-establishment of resources requires a round-trip)	slightly incr. delay
Lossless HO	yes (with forwarding of late packets)	yes
HO delay	roundtrip	faster
Impl. complexity	moderate	higher

A quantitative evaluation of these advanced solutions depends on traffic characteristics, network topologies, etc. and is subject to future investigations.

Although a minimalist solution would enable the operation of RSVP over MIP, we strongly recommend one of the solutions (2) that support fast re-establishment or preservation of resource reservations when mobile nodes move. Only such enhancements can guarantee well performing and uninterrupted operation.

Without quantitative evaluation we can just observe that Triggers/Objects is a quick and low complex solution that might be able to provide sufficiently good service, e.g. for mobile IP phones. The flow extension approach is a little more complex but has the advantage of faster deployment. In multi-provider environments, where we cannot control the whole path end-to-end, a solution that modifies only CNs, access network routers and MNs has a big advantage.

We should also mention, that a combination of the two enhancements is possible and useful for large wireless networks, roaming services etc.

Furthermore, the two performance enhancement approaches described in this draft could also be applied to the MIPv4 route optimisation approach since it is based on the same principles as used in MIPv6. However, we do not provide a detailed description of this topic and more study of the applicability of this draft to MIPv4 is required.

Acknowledgements

We would like to thank the following people for commenting and discussing early versions of this draft: Juha Ala-Laurila, Sarantis Paskalis, and Jukka Seppälä. Part of this work has been performed in the framework of the project ACTS AC085 The Magic WAND, which is partly funded by the European Community and the Swiss BBW (Bundesamt für Bildung und Wissenschaft). The authors would like to acknowledge the contributions of their colleagues from Ascom Systec AG, Compagnie IBM France, IBM Zürich Research Laboratory, Institut Eurecom, France, INTRACOM Hellenic Telecommunications, Lucent Technologies WCND, Nokia Mobile Phones and Nokia Research Center, Robert BOSCH GmbH, Swiss Federal Institute of Technology (ETH) Zürich, Tampere University of Technology, University of Athens, University of Lancaster, University of Ulm, and VTT Technical Research Center Finland.

References

[1] A. Acharya, R. Dighe, and F. Ansari. IP Switching over Fast ATM Cell Transport (IPSOFACTO): IP Multicast over Wireless ATM. In *Proceedings of IEEE ICUPC '98*, October 1998.

[2] A. Acharya, J. Lin, B. Rajagopalan, and D. Raychaydhuri. Mobility Management in Wireless ATM Networks. *IEEE Communications Magazine*, 35(11):100–109, November 1997.

[3] P. Agrawal, E. Hyden, P. Krzyzanowski, P. Misthra, M. Srivastava, and J. Trotter. SWAN: A Mobile Multimedia Wireless Network. *IEEE Personal Communications*, April 1996.

[4] Juha Ala-Laurila, Lorraine Stacey, Neda Nikaein, and Jukka Seppala. Designing a Wireless Broadband IP System with QoS Guarantees. In *Proceedings of ACTS Mobile Summit '98*, June 1998.

[5] A. Ballardie. Core Based Trees (CBT version 2) Multicast Routing - Protocol Specification. RFC 2189, September 1997.

[6] S. Berson and S. Vincent. Aggregation of Internet Integrated Services State. In *Proceedings of IWQoS 98*, Napa Valley, USA, May 1998.

[7] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin. Resource Reservation Protocol (RSVP) - Version 1 Functional Specification. RFC 2205, September 1997.

[8] Stathes Hadjiefthymiades, Sarantis Paskalis, George Fankhauser, and Lazaros Merakos. Mobility Management in an IP-based Wireless ATM Network . In *Proceedings of ACTS Mobile Summit '98*, June 1998.

[9] David B. Johnson and Charles Perkins. Mobility Support in IPv6. Internet Draft, draft-ietf-mobileip-ipv6-06.txt, work in progress, August 1998.

[10] J. Moy. Multicast Extensions to OSPF. RFC 1584, March 1994.

[11] N. Nikaein and C. Bonnet. Wireless Multicasting in an IP Environment. In *Proceedings 5th Intl. Workshop on Mobile Multimedia Communication MoMuc '98*, October 1998.

[12] C. Perkins. Mobile Networking Through Mobile IP. *IEEE Internet Computing*, January 1998.

[13] Charles Perkins and David B. Johnson. Route Optimization in Mobile IP. Internet Draft, draft-ietf-mobileip-optim-07.txt, work in progress, August 1998.

[14] Lorraine Stacey, Juha Ala-Laurila, Jouni Mikkonen, Jukka Seppälä, Stathes Hadjiefthymiades, Neda Nikaein, George Fankhauser, and Sarantis Paskalis. ACTS Project AC085 "Magic WAND", Deliverable 1D6 "IP Over Wireless ATM", http://www.tik.ee.ethz.ch/~gfa/papers/AC085_1D6.pdf, August 1998.

[15] A. Talukdar, B. Badrinath, and A. Acharya. MRSVP: A Resource Reservation Protocol for an Integrated Services Packet Network with Mobile Hosts. Technical report DCS-TR-337, Rutgers University, 1997.

[16] A. Terzis, J. Krawczyk, J. Wroclawski, and L. Zhang. RSVP Operation Over IP Tunnels. Internet Draft, draft-ietf-rsvp-tunnel-03.txt, work in progress, August 1998.

[17] M. Veeraraghavan and G. Dommety. Mobile Location Management in ATM Networks. *IEEE JSAC*, October 1997.

[18] D. Waitzman, C. Patridge, and S. Deering. Distance Vector Multicast Routing Protocol. RFC 1075, November 1988.

A List of Modified RSVP Objects

A.1 Minimal Changes to RSVP Objects for Correct MIP/RSVP Operation

- SESSION objects: All of these (i.e. in PATH and RESV messages need to make reference to the MN's care-of address (flow direction CN $\rightarrow$ MN). Flows to CNs keep CNs' address as a flow destination.
- RSVP_HOP objects: In PATH messages sent to MNs must contain the care-of address of the MN.
- FILTER_SPEC objects: Sender addresses in filter specs originated from MN's in RESV messages must contain care-of addresses.
- SENDER_TEMPLATE objects: In PATH messages sent to the CN must contain the sender's (MN) care-of address.

A.2 Changes to RSVP Objects with Triggers/Objects

- MIP_HOME_ADDR objects: Needed in RESV messages sent to CNs. They could be optionally replaced by constant IPv6 flow labels for flow recognition.

Authors' addresses

George Fankhauser
Computer Engineering and Networks Laboratory
ETH Zürich
ETH Zentrum
CH-8092 Zürich
Switzerland
email: gfa@acm.org

Stathes Hadjiefthymiades
Communication Networks Laboratory
Department of Informatics
University of Athens
TYPA Building, Panepistimioupolis,
Ilisia, Athens 15784
Greece
email: shadj@di.uoa.gr

Neda Nikaein
Mobile Communication Department
Institut Eurecom
2229, Route des Crête
B.P. 193 F-06904 Sophia Antipolis
France
email: nikaein@eurecom.fr

Lorraine Stacey
Lucent Technologies
Sigma, Windmill Hill Business Park,
Swindon, Wiltshire SN5 7EP
United Kingdom
email: lstacey@lucent.com