A Resource Reservation Protocol Supporting QoS-aware Multicast Trees for Next Generation Networks

A. Neto, E. Cerqueira, A. Rissato, E. Monteiro

University of Coimbra
{augusto, ecoelho, anderson, edmundo}@dei.uc.pt

P. Mendes

DoCoMo Euro-Labs
{mendes@docomolab-euro.com}

Abstract

It is expected that next generation of networks will handle new types of services which are destined to large audiences and have different QoS requirements. From the transport point of view, multicast is the most suitable technology for the required group communication services, since it avoids packets duplication and saves network resources. However, QoS–aware multicast content deliver raises several problems, such as the control of QoS trees in environments with asymmetric routing. This paper presents Multi-servIce Resource Allocation (MIRA), a multicast-aware resource reservation protocol for class-based networks that consider routing asymmetries. MIRA controls the resources of network classes associated to multicast sessions considering the QoS characteristics of the latter and network conditions of the available classes in the source-to-receivers path. A detailed description of the functionalities of MIRA and a conceptual analysis against RSVP and RMD-QoSM are presented. Finally, the session setup time, and the signalling and state overheard in comparison with RSVP are analyzed based on simulations.

1. Introduction

It is expected that the next generation of networks will support real-time Quality of Service (QoS) aware group communication services, such as IPTV, radio, file distribution and push media, where the content can be distributed simultaneously to groups of users. From the transport point of view, multicast is the most suitable technology for group communications since it allows several users to get the content of the subscribed sessions at the same time without packet duplication, saving thus the network resources. From the QoS point of view, the session path must be computed taking into account the QoS characteristics of the multicast session and the current conditions of the path from the sender to receivers. From the routing point of view, the control of QoS-aware multicast trees should be done taking into account the common routing asymmetries [1]. In what concerns the coexistence of QoS and multicast over IP networks, the Differentiated Services (DiffServ) model [2] provides a scalable traffic differentiation and QoS assurances, while the Protocol Independent Multicast for Source-Specific Multicast (PIM-SSM) model [3] controls multicast trees without the address allocation problems and complexity of Any-Source Multicast (ASM) model [4].

However, the traffic differentiation provided by DiffServ is not enough to ensure QoS resources for the incoming multicast packets, mainly due to the dynamic bandwidth behaviour. Static resource reservation for DiffServ classes can either lead to waste of bandwidth or to the blocking of application requests. Furthermore, while DiffServ controls the Classes of Service (CoSs) resources in a source-initiated approach, PIM-SSM creates shortest-path trees in a receiver-initiated approach, based on the receiver-to-sender path. Hence, asymmetric routes may compromise the creation of QoS-aware multicast trees, since the path from the sender to the receiver may be different from the opposite direction. Any solution to combine the control of DiffServ and multicast must have open interfaces for allowing interaction with controllers of inter-network Service Level Specifications (SLS), as well as to support resource allocation for wireless technologies.

This paper introduces the requirements to control multicast-aware DiffServ resources in networks with asymmetric routing and presents a new solution called Multi-servIce Resource Allocation (MIRA). MIRA is integrated in the QoS Architecture for Multi-user Mobile Multimedia (Q3M) [5] which provide QoS and seamless mobility in a heterogeneous wireless system. In the Q3M context, MIRA controls the resources of multiple CoSs and IP multicast trees, being the resources of inter-network links controlled based on

SLSs. MIRA extends the NSIS [6] Signalling Layer Protocol (NSLP) for QoS signalling [7] with multicast support. Control information is exchanged among MIRA agents by using MIRA protocol. Moreover, open interfaces are supported that allows mobile operators to use mobility and connectivity control schemes of their choice.

The lower complexity of MIRA in relation to the Resource ReSerVation Protocol (RSVP) [8] and the Resource Management in DiffServ-QoS Model (RMD-QoSM) [9] is conceptually demonstrated. Moreover, MIRA and RSVP signalling and state overhead are analyzed by simulations. The remainder of the paper is organized as follows. The next section analyses related work. The Q3M architecture is described in Section 3. Section 4 describes MIRA in details, while an overview of the MIRA protocol is illustrated in Section 5. A preliminary conceptual and simulation-based evaluation is presented in Section 6. Finally, Section 7 summarizes our conclusions and provides indications about future work.

2. Related Work

Our analysis of related work started with RSVP, which is a standard solution for resource reservation with multicast support. The main drawback of RSVP is the poor scalability due to its per-flow resource management (although aggregate signalling enhancements). Moreover, RSVP's multicast support is not functional because besides increasing its complexity, the performance is compromised since the RSVP sources first probe the network about the data path toward the receivers, which them initiate the operations for reserving the required resources and merging the multicast trees in the reverse path.

On the other hand, existing proposals on QoS-aware Multicast Routing Protocols [10-12] are focused on the computation of the best QoS paths from source to the receivers, taking into account QoS constraints (e.g., available bandwidth). However, these proposals do not provide resource reservation support. To overcome the above limitation, the approach QMRP: A QoS-Aware Multicast Routing Protocol [13] relies on RSVP for the accomplishment of the reservation process on the probed path, which increases the state and signalling overhead.

The NSIS IETF working group aims to provide a general signalling protocol suite based on a two layer paradigm. The lower layer, called NSIS Transport Layer Protocol (NTLP), uses existing transport and security protocols under a common messaging layer, the General Internet Signalling Transport (GIST) [14].

The upper layer is an application-specific layer called NSIS Signalling Layer Protocol (NSLP). The QoS-NSLP and the Network Address Translation (NAT)/Firewall NSLP are examples. The QoS-NSLP signals QoS information to reserve resources for unicast sessions in heterogeneous environments. Multicast support was not included to reduce its complexity in relation to RSVP. Based on QoS-NSLP, the Resource Management in DiffServ (RMD) [15] approach was presented to the NSIS working group as a QoS model called RMD-QoSM. It adds admission control to DiffServ networks and allows external triggers for dynamic resource control. However, being based in QoS-NSLP, RMD-QoSM does not support multicast, and moreover it is focused on the control of Expedited Forwarding (EF) type of DiffServ classes, denoting its lack of support to multiple DiffServ CoSs.

In order to overcome the identified limitations and to allow the dissemination of group communication sessions with QoS-aware multicast support in heterogeneous environments with asymmetric routing, the University of Coimbra is working with DoCoMo Euro-Labs in a multicast-aware per-class resource reservation solution. This solution, called Multi-servIce Resource Reservation Protocol (MIRA), was built as part of the QoS Architecture for Multi-user Mobile Multimedia (Q3M) architecture in accordance with the NSIS framework.

3. Q3M Architecture Overview

The Q3M architecture controls group communication sessions across heterogeneous clusters, providing seamless mobility to users with different QoS requirements, by using an edge-networking approach. A Q3M-aware edge node implements one or more of three components in an element called *Q3M Agent* (Q3MA), denoted as access-Q3MA, ingress-Q3MA or egress-Q3MA. Access-Q3MAs are located at wireless access-routers, while the ingress and egress agents are located in the edges of inter-network links.
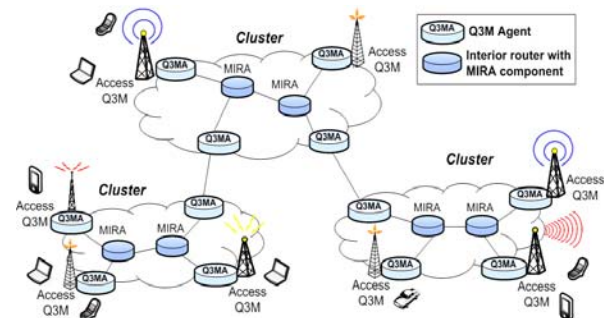

Figure 1 – Q3M Architecture

Besides MIRA, the Q3M architecture encompasses two other functionalities: the control of multi-user sessions and the seamless mobility control. The *Multi-User Session Control* (MUSC) [16] component controls mobility of multi-user sessions across heterogeneous networks based on the basic multicast remote subscription method, augmenting it with the control of QoS (mapping and adaptation) and connectivity translation between networks. The end-to-end mobility management provided by MUSC is complemented by the *Seamless Mobility of Users for Media Distribution Services* (SEMUD) [17] component. SEMUD provides seamless mobility of users inside and between access-networks through the use of buffers in terminals and caches in access-Q3MA. SEMUD predicts the most probable set of access-routers, triggering MUSC after transferring the relevant session context to the collocated Q3MAs. MUSC sets up the session on the new possible paths including triggering MIRA in Q3MAs. After the handover, the resources reserved on the old path and on the possible new paths that were not selected are indicated for releasing.

Due to the existence of different transport connections and to avoid re-construction of the end-to-end session path caused by the mobility of hosts or failure in a network agent, the Q3M architecture assumes that each session has a global identifier, independently of the host location or IP address. Furthermore, sessions can be composed by a set of flows with well defined priorities, rates and identified by a local identifier. Therefore, the separation between session and flow identifier allows the control of connectivity among network with different locator realms, independently of the movement of the users or underlying transport technologies.

4. MIRA Overview

MIRA aims to control the resources of CoSs and multicast trees in a synchronized manner taking into consideration the existence of asymmetric routes. MIRA assures the quality level requested for each flow of a multi-user session (e.g. by MUSC), by adjusting the resources of the selected CoS. Moreover, MIRA supports the construction of QoS-aware multicast trees for each session-flow by manipulating the *Multicast Routing Information Base* (MRIB). The CoS bandwidth and the MRIB are updated in a single-pass operation from the ingress to the egress located in the direction of the access-router of the receiver that asked for the session. In the ingress-router and in each interior router, the configuration of the selected CoS is

done in the outgoing interface indicated by the unicast RIB. After configuring the CoS, MIRA updates the MRIB with the IP address of the previous visited router. The configuration of the MRIB is not done in ingress-routers, since the Q3M architecture uses unicast between networks. In an egress-router, MIRA configures the selected CoS taking into account the SLS established with the neighbour network, and triggers PIM-SSM, after configuring the MRIB.

Edge-agents keep the definition of session-flows, information about available CoSs (e.g., tolerance to loss, delay and jitter), information about the edge-to-edge per-class reservations, and the list of interior-routers involved in the reservation-paths. While the edge-to-edge per-class information is used for fast admission control (accept or deny a request without signalling the network) and to support QoS mapping and adaptation functionalities (e.g., by MUSC), the list of interior-routers is used for resilience. Updated information about the edge-to-edge per-class state is acquired periodically by querying network resources associated with each class inside a network or by checking SLSs between networks. Although edge-agents are statefull, interior-agents store only per-class reservation state for a lightweight control and packet forward processing optimization.

The signalling between pair of edge-agents is done using the MIRA *Protocol* (MIRA-P), which extends QoS-NSLP with CoS, multicast and path-related information. MIRA-P operates by soft-state, where periodic messages are sent for state maintenance, to collect network resource capability information for admission control and to detect re-routing events. In addition, MIRA supports open interfaces for exposing it to external applications or mechanisms via a tighter communication, such as other Q3M components and IETF standards.

4.1. MIRA Interfaces

The *Application Interface* exposes MIRA to external applications or mechanisms, allowing it to receive requests for both, information about available CoSs and per-class resource control. For instance, for resource allocation MIRA needs information about the session (identifier and IP address of the access-Q3MA to which the session is destined to) and each of its flows (identifier, selected CoS and desired bandwidth).

MIRA uses the *Resource Control Interface* to interact with DiffServ control plane mechanisms. This interface can be used both to collect information about the available CoS (including its bandwidth capability), and to configure the packet scheduler (for per-class

bandwidth reservations) and the traffic conditioner or marker.

The *Multicast Control Interface* is used to allocate multicast channels to session-flows (from the pool of SSM channels of each ingress-agent), to update the MRIB, and to trigger PIM-SSM for creating the QoS-aware multicast trees.

The *Transport Control Interface* enables the propagation of signalling messages by interacting with NLTP/GIST for passing or receiving messages. Moreover, MIRA uses this interface to interact with the RIB to retrieve the outgoing interface for processing and detect re-routing events.

4.2 MIRA Functionalities

For resource allocation, MIRA requires the flow identification, desired bandwidth, selected class and the IP address of the destined access-Q3MA. Next, MIRA allocates an SSM channel to the session-flow in case of networks with PIM-SSM and accomplishes the bandwidth reservation by configuring the packet scheduler in the outgoing interface provided by the unicast RIB. After succeeding, MIRA signals the network towards the egress-Q3MA in the direction of the access-Q3MA. Based on the signalled information, all interior routers accomplish the reservation and update the MRIB with the IP address of the previous visited router. Besides updating the MRIB, MIRA may not perform the reservation in the egress-Q3MA during the edge-to-edge signalling since the selection of the inter-network route should be done in accordance with an inter-network mechanism (e.g. MUSC or the inter-domainQoSM [18]). For instance, in the Q3M architecture MIRA is triggered again by MUSC in the egress-agent after the selection of the next network. However, in the Inter-domainQoSM, MIRA may configure the egress-router during the edge-to-edge signalling. After the operations in the egress-agent, a confirmation message is sent direct to the requester ingress-Q3MA, which then reports the requester application.

MIRA controls the state storage and system robustness through the Refresh functionality. When the refresh timer (30 seconds by default) is reached at each ingress-Q3MA, MIRA locally retrieves the session-flows which share the same DSCP, reserved-path and egress-Q3MA state. For each matching, MIRA aggregates the resources into a single message and signals the egress-Q3MA. Each interior-Q3MA checks whether the next hop is in conformance with the reserved-path carried in the message. When the next hop provided by the local RIB is different from the one filled in the reserved-path, the current interior-agent transforms the refresh message into a reservation message and sends it to the next hop indicated by the RIB. Otherwise, MIRA decrements the local variable *Tear BW* of the CoS indicated in the message by the signalled value. Each local *Tear BW* variable is set with the bandwidth allocated to a CoS when the clean-up timer (90 seconds by default) is re-initiated, and its value is used to decrement the bandwidth of the CoS when the clean-up timer reaches zero.

To explicit release resources, MIRA must be triggered at an egress-router, where the network is signalled based in the MRIB. Upon signalled, the agent first checks whether it is a branch point of the indicated session-flow by checking the number of outgoing interfaces configured with the provided SSM channel. When a branch is detected, the release message is forwarded upstream with a flag notifying all Q3MAs to ignore the release operation. When the message reaches the ingress-Q3MA, the state is updated and a confirmation message is sent to the requester egress-Q3MA.

4.3 MIRA Protocol

MIRA-P extends QoS-NSLP with new message-specific flags and objects. From the objects defined in QoS-NSLP and QSPEC [19], MIRA re-uses the *Request Identification Information* (RII), INFO_SPEC and QSPEC objects. The latter is only composed by *QoS Desired* and *QoS Available*. As QoS-NSLP does not consider multicast, we conceived the *Multicast Specification* (MSPEC) object for carrying the multicast communication channels. Moreover, the *Reserved-path* (RSVPATH) object carries the IP addresses of the interior-routers.

All functionalities of MIRA are supported by only two messages: RESERVE and RESPONSE, being transported as payload of GIST messages. Since no signalling association between MIRA agents is required, GIST operates in the *Datagram Mode* (D-Mode). Figure 2 shows the sequence of MIRA-P messages.
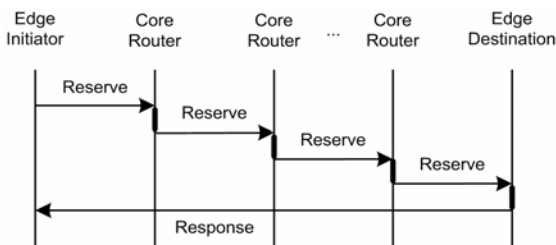

Figure 2. MIRA-P downstream/upstream message sequence

The edge initiator signals the session path with a RESERVE message, and the correspondent result is feed backed by the edge destination through a RESPONSE message. As soon as a Q3MA detects a failure, a RESPONSE message is sent in the reverse-path of the RESERVE message.

5. MIRA Illustration

To illustrate MIRA operating in the ambit of the Q3M architecture, consider an environment composed by network clusters *C1* and *C2,* where source *N1* supplies session *S1*, which is composed by flow *F1* and subscribed by the mobile user *R1*. The illustration is not exhaustive and does not attempt to cover all details of MIRA operations (e.g., interactions MIRA/GIST).

5.1 Session Setup

The session setup is shown in Figure 3, initiating when MUSC in Q3MA-A, receives a SIP message from *R1* to subscribe *S1*. MUSC notices that *S1* is neither presented in Q3MA-A nor in *C2*, and propagates a message to *N1*. At Q3MA-R, MUSC verifies that *S1* is locally active and queries MIRA about the available CoSs. After selected the most suitable CoS to *F1*, MUSC triggers MIRA to allocate the required resources. MIRA firstly allocates *G1* as the multicast channel for *F1*, and the SSM tuple is identified in *C1* by <Q3MA-*R, G1*>. After confirmed that *S1* is not activated in Q3MA-R and the selected CoS has enough bandwidth to accommodate *F1,* MIRA reserves the resources for the selected CoS, adds the state of *F1* in the *S1* entry and sends a RESERVE (step *1* in Figure 3) towards Q3MA-A.
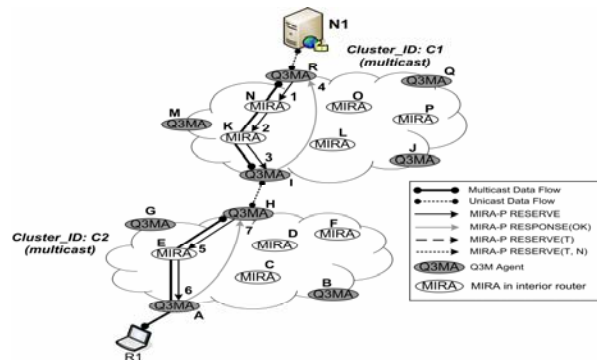

Figure 3. MIRA-P operations in a reservation scenario

Upon succeeded the admission control verifications and accomplished the reservation in MIRA-N, the IP

address of Q3MA-R is added to the local MRIB and the RSVPATH object in the RESERVE message is updated with the IP address of Q3MA-N. The RESERVE message (step *2* in Figure 3) is propagated to MIRA-K, which performs the same operations as Q3MA-N. After admitting *F1*, Q3MA-I update its MRIB and triggers PIM-SSM to build the multicast tree <Q3MA-R, G1>. Finally, Q3MA-I sends a RESPONSE messages (step *4* in Figure 3) directly to Q3MA-R informing the success of the operation. Based on this message, Q3MA-R updates the state of *F1* and confirms MUSC informing the SSM channel.

At Q3MA-I, MUSC triggers MIRA for configuring the required resources in the inter-cluster link. MIRA works based on the inter-cluster SLSs, instead of performing any signalling. After confirmed, MUSC updates its state allowing packet mapping between *C1* and *C2* and signals Q3MA-H. At this agent, MIRA is triggered to perform in *C2* the same set of operations as in *C1*, after which it notifies MUSC about the usage of the SSM tuple <Q3MA-H, *G1*> in *C2*. This allows MUSC to perform the connectivity translation between the set of unicast flows used in the inter-cluster link (IP address of Q3MA-I, plus IP address of Q3MA-H, plus a port for each flow) and the SSM channel allocated in *C2*. At Q3MA-A, MUSC updates its state, informs the host about the success of the session setup and notifies it about the SSM channel that should be used to join *S1* in *C2*. Based on this information, the application can use IGMPv3 or MLDv2 to join the multicast group.

5.1 Intra-cluster Mobility and System Robustness

Consider SEMUD detecting *R1* moving away from Q3MA-A and predicted Q3MA-B as the next access-Q3MA, the Q3MA-B is signalled with the session context information and MUSC is triggered, which then triggers MIRA in Q3MA-H, as illustrated in Figure 4a. MIRA does not reserve resources in Q3MA-H since *F1* is active and thus, only sends the RESERVE message (step *1* in Figure 4a) to Q3MA-B. MIRA reserves resources in the outgoing interface of MIRA-E to MIRA-C, since the latter is indicated as the next hop and *F1* is not activated. The same operations are performed in MIRA-C and Q3MA-B, after which Q3MA-H is signalled with the RESPONSE message (step *4* in Figure 4a), being informed about the success of the operation. Upon confirmed by MIRA, MUSC updates its state and triggers SEMUD in Q3MA-B to finish its configuration, being a reply sent to SEMUD in Q3MA-A. The performed pre-reservation of the

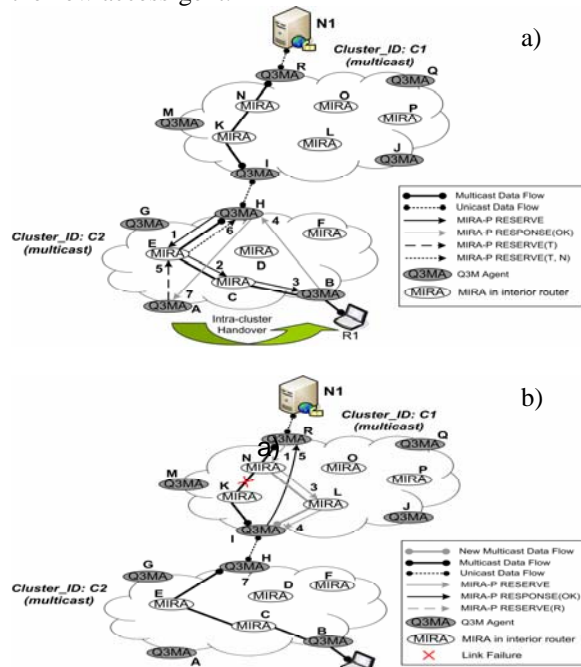resources decreases the latency when attaching *R1* in the new access-gent.

Figure 4. MIRA-P operations in a mobility scenario (a) and re-routing (b)

After the handover, SEMUD informs MUSC in Q3MA-A about the movement of *R1* and verifies that there are no interested receivers for *S1*. MUSC triggers MIRA in Q3MA-A to release resources related to *S1*, which then checks if the local node is a branch point. Since it is not, the resources of *F1* are released and a RESERVE message with a Tear-down flag (step *5* in Figure 4a) is sent to Q3MA-H. When detected that MIRA-E is a branch point, the resources of *F1* are released in the outgoing interface to Q3MA-A and the RESERVE message (step *6* in Figure 4a) is forwarded upstream with an extra flag indicating that all subsequent Q3MA must ignore this request. At Q3MA-H, MIRA removes the state of *F1* and Q3MA-A, and sends a RESPONSE messages (step *7* in Figure 4a) informing Q3MA-A about the success. Finally, PIM-SSM is triggered at Q3MA-A leave <Q3MA-H, *G1*>.

When the refresh time expires in Q3MA-R, MIRA composes and sends a RESERVE message with a refresh flag (step *1* in Figure 4b) to Q3MA-I. Upon signalled, MIRA-N detects the re-routing event since MIRA-K is indicated by the RSVPATH as the next hop and MIRA-L by the RIB. Thus, MIRA processes a normal reservation in the outgoing interface of MIRA-N to MIRA-L and propagates the RESERVE message with refresh flag turned *off* (step *2* in Figure 4b). Upon

succeeding, Q3MA-I triggers PIM-SSM and signals Q3MA-R with a RESPONSE message (step *4* in Figure 4b) informing the success.

6. Evaluation

Section 6.1 presents a conceptual evaluation, comparing the complexity of RSVP, RMD-QoSM and MIRA signalling events. Complementing this conceptual analysis, section 6.2 shows simulation results of MIRA against RSVP in terms of signalling and state overhead, and of the impact that MIRA operations have in the data and control planes.

6.1 Intra-cluster Mobility and System Robustness

For the analysis of MIRA, RSVP and RMD-QoSM, the complexity property of the reservations was determined by measuring the number of signalling events generated by each solution. The protocols were analyzed in two scenarios, with and without a successful reservation attempt, being illustrated in Figure 5.

Figure 5.a clearly shows that RSVP generates more events than RMD-QoSM and MIRA. Figure 5.b shows that, while RSVP and MIRA feed back a notification as soon as detected a reservation failure, RMD-QoSM signals the egress-router, which then notifies the ingress-router for releasing the resources. Although MIRA and RSVP have similar behaviour in the presence of reservation failures, MIRA is faster by initiating the resource releasing in the same router which detected the failure. RSVP needs to signal back the egress for releasing the resources while the ingress node gets this acknowledgement by time-out.
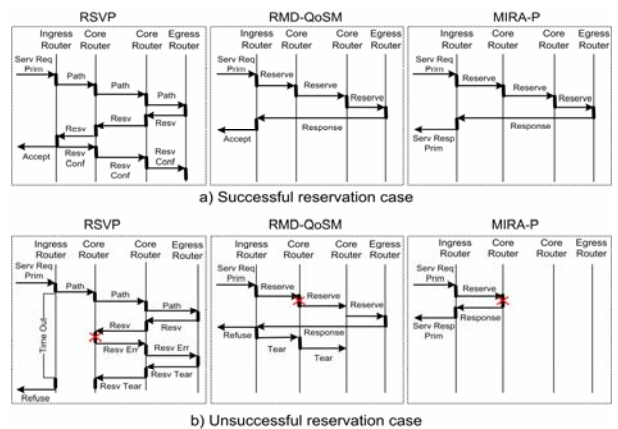

Figure 5. Analysis of RSVP, RMD-QoSM and MIRA-P in terms of signalling events

In summary, this conceptual evaluation shows that in successful reservations MIRA and RMD-QoSM are less complex than RSVP. However, the analysis of an unsuccessful reservation reveals the lower complexity of MIRA in relation to RMD-QoSM and RSVP.

6.2 Simulation analysis of MIRA and RSVP

Based on the conceptual analysis results, it would be interesting to perform a thorough study of MIRA against RMD-QoSM and RSVP. However, we did not consider RMD-QoSM in the simulations since it does not reach a mature state in the standardization track, and is not as broadly implemented as RSVP.

The simulations are performed over a topology composed by three sources (*N1*, *N2* and *N3*), two receivers (*R1* and *R2*), five edge routers and ten core routers with links of 10Mb/s and variable propagation delay randomly placed. Four DiffServ classes for *Best Effort* (BE), *Signalling* (SIG), *Assured Forwarding* (AF) and *Expedited Forwarding* (EF) packets are configured, and the queues served by the *Weighted Fair Queuing* (WFQ) discipline. Each class queue is initialized without resources, except the signalling which is set with an initial 1 Mb/s. The AF class is configured to allocate the out-profile packets in the BE queue.

Five sessions, composed by one flow, are available and enter by the same ingress. *N1* and *N2* supply the multicast sessions *S1* and *S2* with respective AF- and EF-like packets, and *N3* supplies three unicast sessions with BE-like packets. *R1* is linked to egress 1 and interested in *S1* and *R2* is linked to egress 2 and interested in *S2* and in the BE sessions. The joint of *R1* to *S1*, *R2* to *S2* and *R1* moving away from egress 1 to egress 2 with a handover time of 300ms, are the simulation occurrences. *S1* has 1.5 Mb/s of constant bit rate, being the committed rate 1Mb/s and *S2* has 2 Mb/s of constant peak bit rate. Each BE session has 3 Mb/s of bit rate and constant, exponential and poison traffic patterns.

6.2.1. MIRA and RSVP Overhead

Figure 6 shows the signalling and state overheads in function of the type of the node. The packet size of the signallings is traced by NS-2 and the amount of state is provided by each protocol implementation. As the topology has only one domain, MIRA does not reserve resources in the egress routers since there is no inter-domain link. Hence, we only measured the ingress and core routers for fairness.
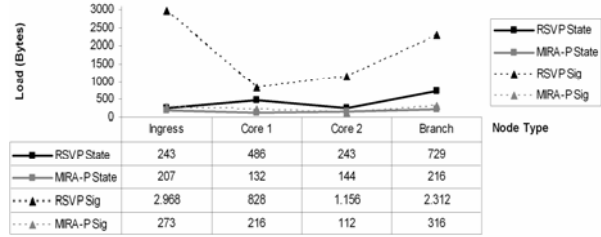

	Ingress	Core 1	Core 2	Branch
RSVP State	243	486	243	729
MIRA-P State	207	132	144	216
RSVP Sig	2.968	828	1.156	2.312
MIRA-P Sig	273	216	112	316

Figure 6. Overall number of signalled and stored bytes for RSVP and MIRA

The signalling overhead of RSVP in the network has an average of 8.5 times higher than MIRA, due to the operations for probing and merging. In addition, MIRA stores 14.8% less state than RSVP in the ingress node for this experiment. Moreover, RSVP stores 3.6 times more state than MIRA in core 1, 1.6 in core 2 and 3.3 in the branch node, once MIRA only requires path state at the ingress node, and RSVP in all nodes along the data paths.

6.2.2. MIRA Impact in the Data and Control Planes

Figure 7 shows the impact that MIRA has on the control and data planes in function of the simulation time. The control plane is analyzed in terms of the overall edge-to-edge reservations (upper graph), the signalling performed by MIRA in each crossed router (middle graph), and the data plane in relation to the traffic load of the sessions.
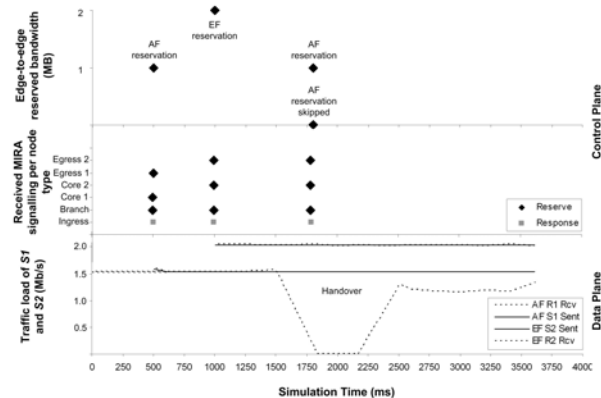

Figure 7. MIRA Impact on the Control and Data Planes

By analyzing the occurrences, we can see that MIRA succeeded the reservations for *S1* and *S2*. After the handover, MIRA skips the reservation in the ingress node since *S1* is already activated and accomplishes the reservation in the new path accordingly. Once the axle *x* has a short event occurrence, the middle graph gives the wrong idea that the events are accomplished at the same time. The

analysis of *S1* show a controlled reduction of the throughput seen by *R1* since MIRA reserved, after the handover, the committed rate (1Mb/s). Hence, the overload of the new path and the out-profile packets end up sharing networks resources with BE-like packets, which means that some of them are lost. The analysis of *S2* shows that *R2* received all its content because MIRA reserves the peak rate.

7. Conclusion and Future Work

This paper presents the requirements to control the resources for multi-user sessions in environments with asymmetric routing and proposes a new solution called *Multi-servIce Resource Allocation Protocol* (MIRA). MIRA accomplishes per-class resource reservations and configures the MRIB in order to force PIM-SSM to build QoS-aware multicast trees. Our conceptual and experimental evaluation demonstrated that MIRA has lower complexity than RSVP and RMD-QoSM. Moreover, it also demonstrated that MIRA allows a dynamic control of differentiated network resources while reducing the impact that occasional network congestion has in the throughput received by mobile terminals.

Future plans include the analysis of an over-reservation mechanism allowing MIRA to support reservations based on inter-network agreements and without per-flow signallings. A thorough study of MIRA in scenarios with distributed requests over multiple ingress-routers will also be performed. Last but not least, the specification of MIRA protocol will be presented to the NSIS working group as an extension of QoS-NSLP.

References

[1] H. Yihua et al. "*On routing asymmetry in the Internet*". In proceedings of IEEE *GLOBECOM´05,* St Louis, MO, US, Dec 2005.

[2] S. Blake *et al*. "*An Architecture for Differentiated Services*". *IETF*, RFC2475, Dec 1998.

[3] S. Bhattacharyya. "*An Overview of Source-Specific Multicast (SSM)*". *IETF*, RFC3569, Jul 2003.

[4] D. Thaler, B. Fenner and B. Quinn. "*Socket Interface Extensions for Multicast Source Filters*". *IETF*, RFC3678, Jan 2004.

[5] E. Cerqueira *et al*. "*A Unifying Architecture for Publish-Subscribe Services in the Next Generation IP Networks*". Published on *IEEE GLOBECOM MMCOM´06*, San Francisco, CA, US, Nov/Dec 2006.

[6] R. Hancock *et al*. "*Next Steps in Signaling (NSIS): Framework*". *IETF*, RFC4080, Jun 2005.

[7] J. Manner, G. Karagiannis and A. Mcdonald. "NSLP for Quality of Service Signaling", *IETF Internet Draft*, Mar 2007.

[8] R. Braden *et al*. "*Resource Reservation Protocol (RSVP)*". *IETF*, RFC2205, Sep 1997.

[9] A. Bader *et al*. "*RMD-QoSM - The Resource Management in DiffServ QoS Model*". *IETF*, Internet Draft, Mar 2007.

[10] M. Faloutsos, A. Banerjea and R. Pankaj. "*QoSMIC: quality of service sensitive multicast Internet protocol*". In proceedings of *ACM SIGCOMM'98*, Vancouver, British Columbia, Canada, Sep 1998.

[11] G. Bianchi *et al*. "*QUASIMODO: Quality of Service-Aware Multicasting over DiffServ and Overlay Networks*". *IEEE Network*, Jan/Fev 2003.

[12] C. Yung-Mu *et al*. "*QMNF: QoS-aware Multicasting Routing Protocol Using N-hop Dominating Flooding Approach on Active Network*". In proceedings of *ICACT 2005*, Phoenix Park, Korea, Feb 2005.

[13] C. Shiang, K. Nahrstedt and Y. Shavitt. "*QMRP: A QoS-Aware Multicast Routing Protocol*". *IEEE Journal On Selected Areas In Communications*, vol. 3, 2000.

[14] H. Schulzrinne and R. Hancock. "*GIST: General Internet Signaling Transport*". *IETF*, Internet Draft, Feb 2007.

[15] G. Karagiannis *et al*. "*RMD - A Lightweight Application of NSIS*". In proceedings of *The Networks 2004*, Vienna, Austria, Jun 2004.

[16] E. Cerqueira *et al*. "*Multi-user Session Control in the Next Generation Wireless System*". In proceedings of *The 4-th ACM International Workshop on mobility Management and Wireless Access*, Malaga, Spain, Oct 2006.

[17] L. Veloso *et al*. "*Seamless Mobility of Users for Media Distribution Services*", submitted at IEEE International Conference on Communications (ICC 2007), Glasgow, Scotland, Jun 2007.

[18] J. Zhang *et al*. "*InterDomain-QOSM: The NSIS QoS Model to Fulfill the E2E QoS Control in the ITU-T RACF Functional Architecture*". IETF, Internet Draft, October 2006, work-in-progress.

[19] J. Ash, A. Bader and C. Kappler. "*QoS NSLP QSPEC Template*". *IETF*, Internet Draft, Feb 2007.