

A Unifying Architecture for Publish-Subscribe Services in the Next Generation IP Networks

Eduardo Cerqueira, Luis Veloso, Augusto Neto, Marilia Curado, Edmundo Monteiro
University of Coimbra
{ecoelho, lmveloso, augusto, marilia, edmundo}@dei.uc.pt

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

Abstract - Next generation IP networks are envisioned to be heterogeneous, to provide a wide variety of services, and to support mobility of users with distinct requirements. Moreover, mobile communications are expected to expand from telephony to publish-subscribe services, such as real-time multimedia. With this goal, the University of Coimbra is working with DoCoMo Euro-Labs on a unifying control architecture for multimedia publish-subscribe services over an IP-based mobile system, where our proposal aims to control the quality of service and mobility across heterogeneous networks with no perceived service degradation to the users. This article highlights the requirements of a control architecture for publish-subscribe services, and presents our proposal, called *QoS Architecture for Multi-user Mobile Multimedia* (Q3M). The main components and functionality of the Q3M architecture are analysed from the user and network perspectives, and some simulation results concerning the analysis of the session setup times are presented.

Index Terms: QoS, Multimedia, Mobility, Heterogeneity

I. INTRODUCTION

Mobile multimedia communications are contributing to enhance our life experience by a constant integration of an increasing diversity of services, such as multimedia publish-subscribe [1]. Examples of these services are entertainment, software distribution and business presentations.

In publish-subscribe services, a publisher is the entity responsible of making available a set of multimedia that might be of interest to a group of heterogeneous and mobile users. The interested users may access the available multimedia content by subscribing the announced services. After subscription, the published content is distributed simultaneously to multiple mobile users.

The distribution of multimedia to multiple mobile subscribers may be done through very different access technologies, such as cellular networks, the *Worldwide Interoperability for Microwave Access* (WiMAX) or the *Wireless Local Area Networks* (WLAN) [2]. Therefore, a common transport infrastructure is required to reduce operational costs and enhance efficiency in the usage of network resources. Due to the low cost and efficient multiplexing systems provided by the IP, the evolution is towards an all-IP mobile system that will support different wireless access technologies.

In addition to the heterogeneity brought by the different wireless access technologies, the next generation networks are expected to be heterogeneous also at the network layer, in consequence of the diversity of connectivity control mechanisms and different address realms. For instance, in the control of unicast mobility, the use of *Hierarchical Mobile IP* (HMIP) allows a mobile device to use a sequence of two global IP addresses, one inside and other outside each access network. An extreme example is the *Network Address Translator* (NAT), in which devices use addresses in a private address realm, and the NAT box deals with the discontinuity point. In the field of IP multicast, there are several different address realms associated with models such as the *Any-Source Multicast* (ASM) and *Source-Specific Multicast* (SSM).

To allow the distribution of published multimedia to multiple mobile users over heterogeneous environments, it seems to be mandatory the clustering of homogeneous network devices into domains/clusters, the creation of open interfaces between such clusters, to provide QoS and mobility control with no perceived service degradation for the users.

Therefore, a unifying architecture should be built as modular and decentralized as possible. The modularisation facilitates the inclusion of emerging technologies, while the decentralization permits a higher scalability. Both of these characteristics might be sustained by an edge-networking approach in which the functionality of each cluster is controlled by a group of organized edge devices, which support interfaces that allow an easy interaction of the architecture with existing standards.

To support multimedia publish-subscribe services over heterogeneous clusters, a unifying architecture should handle transport connections cluster by cluster. Hence, the first requirement is to stop the duality use of IP addresses as routing locators and identifiers for transport connections. In this way, each session has a global identifier, independent from hosts and IP addresses. In the case of multimedia publish-subscribe services, each published service can be defined by a set of multi-user sessions, where each session can be composed by set of flows that may require similar or different quality levels, importance and rates (this generic definition accommodates common multimedia encoders, such as H.264, MPEG-2 and

MPEG-4). Thus, multiple mobile users utilize the same session-flow at the same time.

From a functional point of view, a unifying architecture should allow publishers to announce available multimedia services and users to access them, and must provide an end-to-end quality and connectivity control over heterogeneous clusters, taking into account established inter-cluster *Service Level Specifications* (SLS) [3]. Inside each cluster, the architecture must control the allocation of network resources and the operation of multicast trees (in clusters that have IP multicast technology). Moreover, it is mandatory to have a seamless control of the mobility of users and publishers.

Summing up, from an architectural point of view, a unifying architecture for publish-subscribe services should control generic multi-user sessions over heterogeneous clusters and support the separation between session identifiers and IP locators. From a functional point of view, it should control multi-user sessions, the allocation of network resources, and the seamless mobility of publishers and users.

This paper aims to analyse the requirements of a unifying control architecture for multimedia publish-subscribe services over an IP-based mobile system, and presents our proposal denominated *QoS Architecture for Multi-user Mobile Multimedia* (Q3M). In addition, preliminary results concerning the Q3M session setup mechanism are presented. The remainder of this paper is organized as follows. Section II presents some related work. The Q3M architecture is described in Section III. Section IV shows the expected benefits and the preliminary evaluation of the Q3M architecture. Conclusions of the paper and directions for ongoing and future work are summarized in Section V.

II. RELATED WORK

Currently, there are IETF standards that address some of above mentioned needs, such as the *Differentiated Services* (DiffServ) model for a scalable traffic differentiation and QoS assurance, the *Real-Time Streaming Protocol* (RTSP) and the *Session Initiation Protocol* (SIP) to control the streaming and access to published services, the *Protocol Independent Multicast for the SSM* (PIM-SSM) to handle packet distribution to groups of users, and HMIP to control the local mobility in unicast networks. However, each of these solutions does not cover by itself all aspects necessary for the development of a unifying architecture to control connectivity, QoS, and seamless mobility over heterogeneous networks.

In addition to the IETF solutions above, other approaches of architectures for QoS, mobility and heterogeneous systems have been proposed [4-6]. Nevertheless, these approaches are focused on single-user sessions. Moreover, the mobility and QoS are controlled by a tunnel-based approach, which restricts network operations [4, 5], and the use of QoS brokers to control resources per-flow reduces scalability [6]. To overcome the limitations identified and to enable the deployment of end-to-end multimedia publish-subscribe services for mobile multi-users, the University of Coimbra is working with

DoCoMo Euro-labs in a new architecture called Q3M, that will be described in the next section.

III. THE Q3M ARCHITECTURE

The Q3M architecture uses an edge-networking approach to control multimedia multi-user sessions across heterogeneous clusters. An edge of the Q3M architecture implements one or more of three components, in an element called Q3M Agent (Q3MA), as shown in Fig. 1. A Q3MA is called access-Q3MA, ingress-Q3MA, or egress-Q3MA, when located at wireless access points, at the entry or at the exit of the session in a cluster, respectively.




Figure 1 – Q3M architecture

The referred components are the *Multi-User Session Control* (MUSC), the *Multi-service Resource Allocation* (MIRA) and the *Cache-based Seamless Mobility* (CASM). The Q3M architecture also defines open interfaces for the interconnection of its components, and interfaces to allow each one of them to interact with existing standards such as, SIP, PIM-SSM, *Internet Group Management Protocol* (IGMPv3), *Multicast Listener Discovery protocol* (MLDv2), and the elements of a DiffServ network.

The MUSC component controls the access of users to available multimedia multi-user sessions, and the inter-cluster QoS and connectivity (linkage of different address realms). The access control is done in access-Q3MAs, based on an interface with SIP. Besides, an interface to CASM is used to control the access of moving users. Moreover, in ingress and egress Q3MAs, the QoS mapping is based on the association between the session quality requirements of each session-flow and the available network services provided by MIRA. If such mapping is not optimal, a QoS adaptation to the current network conditions, by dropping low priority flow(s) of the session, is performed. Furthermore, if the current ingress or egress Q3MA is in the frontier between unicast and multicast address realms, MUSC maps unicast flows to multicast trees or vice-versa. A receiver-driven and source-initiated signalling protocol, called MUSC-P, is used to exchange control information between Q3MAs using a soft-state approach to maintain per-session and per-flow state, contributing to the architecture robustness. It is receiver-driven because it is triggered at access-Q3MAs, and it is source-initiated since MUSC starts its operations at the Q3MA nearest to the source,

or at the first Q3MA in which the requested session is found in the way towards the source. The latter functionality reduces the session setup time as well as signalling and state overhead for establishing multi-user sessions to groups of users.

MIRA controls the resources of multiple network services based on the DiffServ model inside a cluster and based on SLSs between clusters. MIRA also controls the shape of intra-cluster multicast trees. This is done taking into account route asymmetries by updating the *Multicast Routing Information Base* (MRIB) on each router and triggering PIM-SSM at egress-Q3MAs or to directly configure the SSM multicast channel on routers outgoing interfaces. Following the SSM model, MIRA reserves, in ingress-Q3MAs, multicast channels, which MUSC associates into sessions. Moreover, MIRA uses a signalling protocol, called MIRA-P, which operates from ingress to egress to control the bandwidth and multicast resources. Finally, CASM controls the mobility of users between heterogeneous clusters and inside them. Seamless mobility is achieved through the use of buffers in terminals and caches in access-Q3MA. After predicting the most probable set of access points, CASM uses a signalling protocol denominated CASM-P to exchange control information between old and new access-Q3MAs. In the predicted agents, CASM triggers MUSC, by supplying it with session related context. This allows MUSC to setup the session in the new possible paths including the control of QoS (by triggering MIRA), and performing QoS adaptation when required. At the old access-Q3MA, the information provided by MUSC on the new possible access router, combined with knowledge regarding the access technologies, gives support to the handover decision. When the handover decision is taken, the interaction between CASM and MUSC allows the release of the resources reserved on the old path, and on the new paths that the mobile device is not going to use (if no other users are subscribed to the session at those locations).

A. Example of Q3M Overall Functionality

Fig. 2 shows the interaction of the Q3M components in a scenario with one source publishing one multimedia session with one flow, which is received by mobile subscribers. The

establishment of a new session (on the left side), is started when the application on host *R1* uses a SIP message with SDP to subscribe the session (it is assumed the existence of a SIP server that forwards the message to the user's access-Q3MA, based on previous registration). Afterwards, the MUSC access control receives the message and triggers MUSC-P in Q3MA-B to send a request signalling message towards the source with IP Router Alert Option enabled. If any Q3MA has information on the required session-flow carried in the message, it stops its progression, which happens in Q3MA-K.

From Q3MA-K, and towards the access-Q3MA, MUSC interacts with MIRA to query about network services. Based on the response, MUSC selects the most suitable network service and requests MIRA to configure the required bandwidth and to allocate a distribution channel for the session-flow. In IP multicast clusters, MIRA allocates an SSM channel to the session-flow, while in clusters without PIM-SSM, MIRA selects a pair of ports to define a unicast flow from ingress to egress router, to which the session-flow will be associated by MUSC. When MIRA finishes its job, MUSC updates its state and sends a MUSC-P response message to the next Q3MA informed by MIRA (the egress router of the same cluster, or the ingress router of the next cluster in the direction of the access-router). Based on the connectivity information furnished by MIRA, MUSC translates all incoming unicast flows to multicast trees or vice-versa.

When the MUSC-P response message reaches Q3MA-B (access-Q3MA), MIRA is required to configure the network service and multicast tree only on the outgoing interface. If unicast is used in the wireless interface, MUSC maps the session-flow to a list of interested users, and performs the required packet duplication. The configuration of MIRA in the access-Q3MAs with information about the wireless interface technology, such as service differentiation capabilities (802.11e or 802.16 [7]), is being considered. Based on this information, MIRA can configure the service classes to be used on the wireless link. Afterwards, CASM is triggered by MUSC to activate/update the local cache for the session-flow and receives information concerning session context.




Figure 2 - Example of Q3M functionality

Finally, MUSC sends a correspondent SIP response message to the host informing it about the success of the session setup. In a multicast access-network, MUSC will include information about the SSM channel to be joined by the host in the SIP response message.

After the session setup for $R1$, the receiver $R2$ sends a SIP message to subscribe the same session as $R1$. Since the requested session was already activated for $R1$ in the access-Q3MA, MUSC adjusts its state about the number of local receivers for that session-flow and immediately replies with a SIP message to $R2$. This procedure reduces the session setup time and avoids end-to-end signalling to establish the same session for every other receiver in the same access-router.

The situation on the right side of Fig. 3 represents the inter-cluster mobility, in which $R1$ moves away from the access point of Q3MA-B. Based on information about the location of the base-stations, collected from inter-cluster agreements, and based on the condition of the mobile node (e.g. location, moving direction and velocity), CASM predicts in Q3MA-B that Q3MA-E is the most probable new access-Q3MA. After handover prediction, CASM transfers the session context to the predicted Q3MA, in which MUSC is triggered, leading to the pre-reservation of resources in the new path to Q3MA-E. CASM action in Q3MA-E ends up with the configuration of a cache to avoid packet loss, after which a reply is sent to CASM in Q3MA-B.

In Q3MA-B, CASM analyses the resources viability in the predicted Q3MA (Q3MA-E) and the signal-to-noise ratio to take a handover decision. After the handover decision, CASM informs MUSC about the movement of the mobile device, allowing the adjustment of the number of local receivers. During handover the application consumes packets from the $R1$ buffer while packets continue to arrive to the cache in Q3MA-E. When $R1$ attaches to the new access point, it synchronizes the packet reception with the cache in Q3MA-E.

IV. EXPECTED BENEFITS AND Q3M EVALUATION

The Q3M architecture allows users to access multimedia broadcast-like services via an IP network, in a similar way as they already access data and VoIP services. With Q3M, users are unaware of the network technology diversity, and experience seamless mobility. By implementing most of the network computation at the edges, a cost reduction is expected. Operators will also be able to provide services with high availability times due to the high robustness of the decentralized architecture. The concept of cluster grants operators more flexibility in the control of the structure of their networks, and the use of open interfaces allows them to use connectivity and resource control methods of their choice.

In order to analyse the initial impact of the Q3M components on the expectation of users, we present an analysis of the time required to setup a session, based on two sets of simulations done using the *Network Simulator 2.28* (NS). The first simulation compares the session setup time of MUSC and SIP (RFC 3261) when the session, identified by a unique identifier,

was already offered to several receivers. Furthermore, there are two SIP servers, one in sender's cluster and one in the receivers's cluster. The second simulation aims to further analyse the impact of the Q3M architecture on the session setup time, when QoS assurances are given to the session. In this case, MUSC is complemented with MIRA, in order to allocate network resources to the existing session.

We use the results of previous studies that have investigated traffic distribution on the Internet in order to make realistic assumptions [8]. These studies show that 26%, 40%, 26% and 9% of receivers are placed in a cluster at a distance of two, three, four and five clusters from their source, respectively. Thus, we use a scenario with five clusters, and a random topology generated by BRITE [9]. Each cluster is composed by fourteen routers. The intra and inter-cluster links have a bandwidth of 10 Mb/s and their propagation delay is attributed by BRITE according to the distance between the edges of each link. The sizes of the messages used in the simulations are listed in Table I (these messages are IPv4 address and are transported on top of TCP). We use only SIP *INVITE* and *200 OK* messages, because they are mandatory in the establishment of sessions.

Table I - Message Size for Session Establishment

Message	Protocol	Length (bytes)
SIP INVITE	SIP with SDP	620
SIP 200 OK	SIP with SDP	450
Request	MUSC-P	84
Response	MUSC-P	88
Reserve	MIRA-P	70
Response	MIRA-P	52

To compare MUSC with SIP, we perform four experiments, in which the receivers $R1$, $R2$ and $R3$ are connected to an access-router A (AR-A) and the receivers $R4$, $R5$ and $R6$ are linked to an access-router B (AR-B). Both access-routers belong to the same access-cluster and the receivers request access to the same session sequentially from $R1$ to $R6$. In each simulation the distance between the receivers and the source was varied in terms of the number of cluster (from 2 to 5). First SIP is used to allow the receivers to request access to the session (i.e. without the Q3M architecture) and, then, SIP is used together with MUSC. Fig. 3 shows the sum of the session setup times in each access-router, when the number of clusters between the sender and the receivers varies from 2 to 5.




Figure 3 – Sum of Session setup times for SIP and MUSC

The results reveal that the utilization of MUSC reduces the time to establish the same session in approximately 69% and 83% for receivers in AR-A and AR-B, respectively. These small setup times are possible because not only the size of the

MUSC messages is smaller than the size of SIP messages, but also because MUSC only exchanges end-to-end MUSC-P messages to setup the session for a first receiver in an access-router, being all posterior requests processed locally in the access-agent. Moreover, if the same session-flow is already present in the ingress-router of the access-cluster, which happens when receiver $R4$ of AR-B requests access to the session, the MUSC-P message sent due to the request of $R4$ is transported only inside the access-cluster, and not end-to-end, as occurs with SIP to perform each and any request.

Fig. 5 illustrates in more detail the results obtained when the distance between the receivers and the source is of four clusters. The results show that in the worst case, a MUSC-P message visits the same amount of routers than a SIP message. This situation occurs for the first receiver in an access-cluster. However, for the subsequent receivers in the same access-cluster, MUSC messages only travel locally inside the access-cluster, while SIP messages continue to have an end-to-end scope. This is illustrated for $R4$, for which the request generated MUSC message visited 75% less routers than when only SIP is used. Thus, MUSC contributes to decrease the setup time as well as signalling and state overhead when setting-up the same session to several receivers in the same cluster. The difference between using MUSC or only SIP, in what concerns the number of visited routers, is even higher when subsequent receivers join a session in an access-router in which another receiver of the same session is already attached. In this case there is no need to generate messages if MUSC is used, as shown by the examples of $R2, R3, R5$ and $R6$ in Fig 4.




Figure 4 - Variation of the number of visited routers

The second simulation addresses the session setup time when some QoS assurance is required. This evaluation, which results are illustrated in Fig.5, adds MIRA to the operation of MUSC. In this configuration, MUSC selects a generic class of service and requires MIRA to configure network resources to this class. Moreover, MIRA allocates, in each cluster, one SSM channel to each session-flow indicated by MUSC.

Fig. 5 shows that the total session setup times for receivers in AR-A is higher than in AR-B. This is caused by the interaction of MUSC with MIRA in all edge routers along the data path until the access-router is reached, in order to setup the session for $R1$. However, when the session is ongoing in the access-cluster and the first receiver of AR-B requests access to the session, this interaction occurs only inside the access-cluster. In addition, it shows that the inter-cluster operations have higher impact in the setup time than the overall intra-cluster ones. This is justified, because inter-cluster links

have higher propagation delay and MUSC needs time to configure the connectivity mapping between the SSM channels used in different clusters. Moreover, the setup time associated with MUSC includes the SIP communication delay between the terminal and the SIP server, and from the later one to the terminal access-router.




Figure 5 – Sum of Session setup times for MUSC and MIRA

V. CONCLUSION AND FUTURE WORK

The Q3M architecture provides a solution to support seamless multimedia publish-subscribe services in the next generation of IP networks. The proposed architecture controls multimedia multi-user sessions across heterogeneous networks, by using an edge-networking approach with open interfaces. The seamless characteristic of the Q3M solution is achieved by a modular integration of the control of multimedia communication sessions, network resources, and device mobility. The initial evaluation shows the setup time reduction brought by MUSC, when compared to a solution only using SIP, as well as the impact that the extra QoS control provided by Q3M has on the session setup time. Future plans include interaction of the proposed architecture with inter-cluster SLS control mechanisms and with the MBMS. Moreover, extra evaluation including QoS mapping and adaptation and CASM will be performed.

REFERENCES

- [1] M. Carugi, B. Hirschman, and A. Narita, "Introduction to the ITU-T NGN Focus Group Release 1: Target Environments, Services and Capabilities", *IEEE Commun. Magazine*, vol. 43, issue: 10, pp. 42-48, October 2005.
- [2] R. Prior, S. Sargento, D. Gomes, and L. Aguiar, "Heterogeneous Signaling Framework for End-to-end QoS support in Next Generation networks", *IEEE System Sciences - HICSS '05*, pp. 305-315, January 2005.
- [3] Paulo Mendes, Jorge Andres-Colas and Carlos Pinho, "Information Model for the Specification of QoS Agreements among Ambient Networks", *IEEE International Symposium on Personal Indoor and Mobile Radio Commun.*, Berlin, Germany, September 2005.
- [4] F. Greve *et al*, "FAMOUS: A Network Architecture for Delivering Multimedia Services to FAST MOving USers", *Wireless Personal Commun., Kluwer Academic*, vol. 33, issue: 3-4, pp. 281- 304, June 2005.
- [5] L. Shou-Chih *et al.*, "Architecture for mobility and QoS support in all-IP wireless networks", *Selected Areas in Commun., IEEE Journal on*, vol. 22, issue: 4, pp. 691- 705, May 2004.
- [6] C. Carneiro *et al*, "The DAIDALOS Architecture for QoS over Heterogeneous Wireless Networks", *IST Mobile & Wireless Commun. Summit*, June 2005.
- [7] H. Wang, W. Li and P. Agrawai, "Dynamic Admission Control and QoS for 802.16 wireless MAN", *IEEE Wireless Telecommunications Symposium*, pp. 60 – 66, April 2005.
- [8] A. Medina *et al*, "BRITE: An approach to universal topology generation", *In IEEE MASCOTS '01*, pp. 346-356, August 2001.
- [9] U. Steve *et al*, "Implications of interdomain traffic characteristics on the traffic engineering", *SAC '04: Proceedings of the 2004 ACM symposium on applied computing*, pp. 339-346, March 2004.