

Multi-user Session Control in the Next Generation Wireless System

Eduardo Cerqueira
University of Coimbra
Polo II – Pinhal de Marrocos
3030029, Coimbra, Portugal
+351239790000

ecoelho@dei.uc.pt

Paulo Mendes
DoCoMo Euro-Labs
Landersbergerstr 312
80687, Munich, Germany
+4989568240

mendes@docomolab-euro.com

Edmundo Monteiro
University of Coimbra
Polo II – Pinhal de Marrocos
3030029, Coimbra, Portugal
+351239790000

edmundo@dei.uc.pt

ABSTRACT

Next generation IP wireless systems are envisioned to be heterogeneous and to provide ubiquitous services to mobile users with different quality of service requirements. Furthermore, in order to attract and keep customers, mobile operators are expanding their portfolio with the inclusion of publish-subscribe services, such as real-time multimedia sessions. This paper presents a signalling application layer based on the *Next Steps in Signalling* (NSIS) framework that aims to provide the control of sessions to multiple users across heterogeneous wireless systems, called *Multi-user Session Control* (MUSC). Supported by a session-aware signalling protocol, named MUSC-P, which is highlighted in this paper, MUSC controls the session quality of service mapping and adaptation, the ubiquitous access of mobile users to the available published sessions and the session connectivity between networks with different address realms. The protocol is detailed, its functionalities are presented, and an initial performance evaluation shows the protocol efficiency concerning the control of session setup to multiple users.

Keywords

Multi-user session, Quality of Service, Heterogeneous Wireless Systems

1. INTRODUCTION

Next generation IP wireless systems are envisioned to be heterogeneous, to handle new types of wireless group communication services with different Quality of Service (QoS) requirements, and destined to large audiences, such as publish-subscribe services [1]. Examples of these services are IP television, push media, remote learning, software distribution and many other applications such as entertainment and business presentations.

* Copyright Notice as required in ACM Word Template

In publish-subscribe services, a publisher is the entity responsible to make available a set of content that might be of interest to a group of mobile users. The interested users may access the available information by subscribing the announced services. After subscription, the published content is distributed to multiple users without the latter having to fetch it repeatedly. In this context, 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 multi-user session-flow at the same time, reducing the operational costs and enhancing efficiency in the usage of network resources. Figure 1 shows a generic definition of publish-subscriber services.

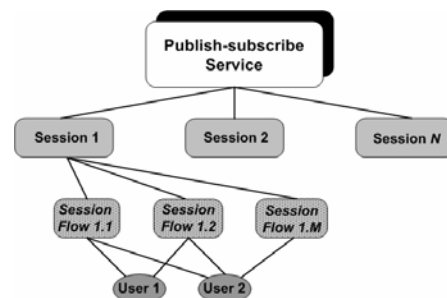


Figure 1 - Publish-subscribe services

The distribution of multi-user session content to multiple mobile subscribers with different devices and requirements may be done through very different access technologies (e.g. *Worldwide Interoperability for Microwave Access* (WiMAX)) and transport technologies (e.g. IP multicast) [2]. Furthermore, the ability to provide end-to-end QoS control as well as to allow ubiquitous services to mobile users are the keys to the success of the next generation IP wireless systems [3].

However, to allow the distribution of published multi-user sessions throughout heterogeneous wireless networks with end-to-end QoS control, it seems to be mandatory the control of session quality level mapping in receiver-driven and source-initiation approach, taking into consideration the QoS characteristics and network conditions of the path from the publisher to subscribers, allowing the control of QoS in environments with inter-domain/cluster asymmetric routing. Furthermore, in order to increase the user satisfaction, it is necessary the QoS support of

mobile users, allowing the user to keep acceptable quality level in on-going multi-user sessions, independently of its movement. However, when the QoS mapping control can not ensure the minimal QoS committed for the session, for instance due to the mobility of users to a congested access network, a QoS adaptation is required to adapt the session to the current network conditions, avoiding the waste of network resources.

In addition, due to the existence of clusters with different address realms (e.g. IP unicast or IP multicast, where in the field of IP multicast, may have different address realms associated with several multicast models such as *Any-Source Multicast* (ASM) or *Source-Specific Multicast* (SSM)) along the session path, it is needed a connectivity control mechanism to manage the creation of distribution trees cluster by cluster, avoiding any discontinuous point between clusters, and allowing the construction long of inter-cluster distribution trees with efficient shape [4]. Besides the connectivity control, it is required the ubiquitous access control for mobile users to the available published session.

This paper highlights the functionalities and operations of a session-aware signalling protocol supported by a signalling application layer, called *Multi-user Session Control* (MUSC), which is based on the *Next Steps in Signalling* (NSIS) framework. MUSC is a solution to provide multi-user session control in the next generation IP wireless systems. MUSC is supported by a QoS architecture named *QoS Architecture for Multi-user Mobile Multimedia* (Q3M), where it is responsible to control the ubiquitous access of groups of mobile users to the available services, and to support inter-cluster QoS mapping of multi-user sessions over heterogeneous systems. Furthermore, it also provides the control of session connectivity between clusters with different address realms. The MUSC signalling protocol, called MUSC-P, is used to exchange control information between network agents in order to give support for the above mentioned operations. MUSC-P operates in a receiver-driven and source-initiated model and takes a soft-state approach to maintain per-session and per-flow state, contributing to the system robustness. Furthermore, in order to analyse the impact of MUSC signalling associated with the control of session setup for multiple users, an initial performance evaluation is presented.

Even though the MUSC scheme is presented in the context of the Q3M architecture, it supports open interfaces, allowing mobile operators to use mobility, connectivity and resource control schemes of their choice. However, in order to show MUSC functionalities embedded in an architecture that provides QoS, seamless mobility in an heterogeneous wireless system, MUSC will be presented in the context of the Q3M architecture. The remainder of this paper is organized as follows. Section 2 presents relevant related work. The Q3M architecture and its main components are described in Section 3. A detailed description of MUSC is shown in Section 4. Section 5 illustrates example of operations. A preliminary evaluation involving MUSC is presented in Section 6. Conclusions and directions for ongoing and future work are summarized in Section 7.

2. RELATED WORK

This section highlights some related work used in this paper. The *Session Initiation Protocol* (SIP) [6] and the *Real-Time Streaming Protocol* (RTSP) [7] are *Internet Engineering Task Force* (IETF) standards signalling protocols used to control the streaming and

the access of users to published sessions. However, these end-to-end protocols are only suitable to control one-to-one or one-to-few sessions supported by unicast environments, lacking in terms of QoS control and in the support to one-to-many sessions, such as real-time multicast streaming.

In addition to SIP and RTSP signalling protocols, the NSIS IETF working group is working on a new signalling protocol suite [8]. NSIS is based on a two layer paradigm. The lower layer is called *NSIS Transport Layer Protocol* (NTLP) and is responsible for delivering signalling messages between a pair of agents. The upper layer is an application-specific layer called *NSIS Signalling Application Layer* (NSLP). QoS NSLP [9] and *Network Address Translation* (NAT)/Firewall NSLP [10] are examples of NSLPs.

QoS NSLP supports a signalling protocol to request resource reservation for single-user unicast sessions in heterogeneous environments. It also uses soft-state peer-to-peer refresh messages to increase system robustness. However, due to its single-user signalling nature (including session definition and the object used by the QoS NSLP messages to describe the QoS parameters desired by each unicast flow, called QSPEC [11]), the QoS NSLP does not provide a signalling solution that can be used to control publish-subscribe multi-user sessions. Therefore, QoS NSLP can not be used in situations where sessions may be composed of a set of flows and destined to groups of users, as happens with scalable multicast video streaming.

In addition, the NAT/Firewall NSLP is a connectivity control mechanism designed to request the dynamic configuration of NATs and/or firewalls along the session path, allowing hosts in a private network to communicate with destinations on an external network and vice versa. Even though the use of NATs solves the address translation problem between public and private networks, it does not give support for the control of multi-user sessions connectivity between clusters that support different address realms, such as the support for mapping of unicast flows to multicast trees or multicast-trees to multicast trees when clusters support different IP multicast models.

Furthermore, in order to support the quality level of sessions across heterogeneous clusters, a QoS mapping scheme is needed to assure the mapping of sessions into the most suitable network service class supported by each cluster or between clusters along the session path. This problem has been addressed by previous approaches [12, 13], which performed QoS mapping for single-user sessions. However, these approaches use a centralized control scheme that requires the installation of proprietary modules in the end-system which severely constrains scalability and increases complexity. In addition, Mammeri [14] proposed a QoS mapping scheme that provides end-to-end QoS guarantees for unicast flows across *Integrated Services* (IntServ) and *Differentiated Services* (DiffServ) models. Based on the analysis QoS requirements of each flow and the *Service Level Agreement* (SLA) supported by each domain, the mapping mechanism maps each flow of a DiffServ domain to an IntServ domain and vice versa. However, this proposal is dependent of the underlying QoS model, and only supports single-user sessions.

In situations where it is not possible to assure the minimal QoS committed for a session or for some flows of that session, for instance due to the unavailability of network resources, a QoS adaptive control mechanism should be used in order to adjust the session quality level to the current network conditions. However,

existing solutions require the implementation of proprietary modules on the end-hosts or need devices in specific places in the network, which are responsible for adapting the content coding to the network bandwidth [15, 16]. Furthermore, most solutions are focused in multicast-aware environments, with the scope limited to multimedia sessions [17]. Therefore, previous approaches are not suitable to adapt sessions to the current network conditions in heterogeneous networks, because they reduce the flexibility and increase the complexity of the environments.

The analysis of related work has shown that none of the approaches satisfy all requirements to control multi-user publish-subscribe sessions in the next generation wireless system. Therefore, to overcome the limitations identified, and to enable session QoS mapping and adaptation, connectivity and mobility support across heterogeneous clusters, the University of Coimbra is working with DoCoMo Euro-labs in a session-aware signalling protocol that is in accordance with the NSIS framework and requirements, called *Multi-user Session Control Protocol* (MUSC-P). MUSC-P is supported by the *Multi-user Session Control* (MUSC) component, which is part of the *QoS Architecture for Multi-user Mobile Multimedia* (Q3M) architecture, which will be shortly described in the next section.

3. THE Q3M ARCHITECTURE

The Q3M architecture aims to control publish-subscribe services to groups of mobile users across heterogeneous clusters, providing seamless mobility to users possessing different QoS requirements. The Q3M architecture provides the inter-cluster QoS mapping and adaptation control, manages the resources of multiple network services classes, the IP multicast connectivity inside a cluster and the mobility of users inside and between clusters without revealing to them any discontinuity point.

The Q3M architecture uses an edge-networking approach to control 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 Figure 2. A Q3MA is called access-Q3MA, ingress-Q3MA, or egress-Q3MA, when located at wireless access points, at the entry of a cluster or at the exit of the session in a cluster, respectively.




Figure 2- Q3M architecture

The referred components are the *Multi-User Session Control* (MUSC) that controls the ubiquitous access of multiple users to the announced multi-user sessions, inter-cluster QoS mapping and adaptation and connectivity between clusters. The *Multi-service*

Resource Allocation (MIRA) controls the resources of multiple network services/classes, and the IP multicast connectivity inside a cluster. The *Cache-based Seamless Mobility* (CASM) manages the mobility of users inside and between clusters.

The Q3M architecture defines internal interfaces for the interconnection of its components, and external interfaces to allow each one of them to interact with existing standards such as the *Session Description Protocol* (SDP) [18], SIP, *Internet Group Management Protocol* (IGMPv3) [19], *Multicast Listener Discovery Protocol* (MLDv2) [20] and the elements of a DiffServ network. The location of the Q3M components and interfaces in a TCP/IP protocol stack is described in Figure 3.

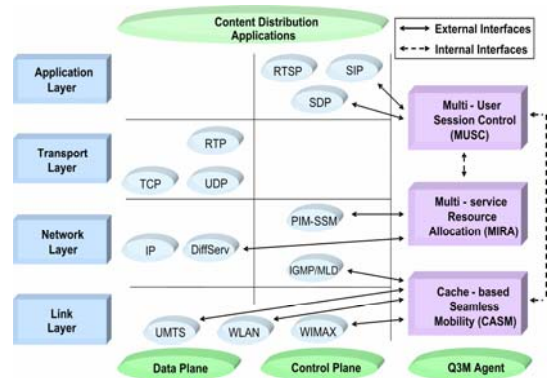


Figure 3 - Q3M components in a TCP/IP stack

3.1 Q3M Components

This section presents a brief description of CASM, MIRA and MUSC, where 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 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 mapping and 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).

MIRA controls the resources of multiple network services based on the DiffServ model inside a cluster and based on *Service Level Specification* (SLS) between clusters [22]. 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. When required, MIRA

uses a soft-state signalling protocol, which operates from ingress to egress to control bandwidth and multicast resources

MUSC controls the access of users to available multi-user sessions, performs the multi-session signalling control (e.g. signalling to setup or release a session), and aims to provide end-to-end QoS and connectivity between clusters. 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 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 and assuring QoS in environments with inter-cluster asymmetric routing.

4. MUSC OVERVIEW

MUSC is being designed to be proposed as new a signalling layer protocol aiming at the NSIS framework. As said before, MUSC controls the ubiquitous access of mobile users to the available multi-user sessions and aims to provide QoS mapping and adaptation, and connectivity between clusters, where the interoperation between network agents is done by MUSC-P. Following the NSIS proposal, MUSC assumes that each multi-user session has a global identifier defined by the publisher, independently from host location or IP address. Each multi-user session is described in a *Session Object* (SOBJ) identified by a unique session identifier, which can be composed by a set of flows, where the QoS parameters desired by each flow are described in a QSPEC object. Therefore, each receiver gets from the publisher the description of the available sessions, containing their SOBJs (QSPEC object of each flow, including its priority, bit rate, tolerance to loss, delay and jitter), and sends it, using SDP, to an agent in the access cluster to which the receiver is connected at the time of the session setup. MUSC is composed by four control mechanisms and one signalling protocol as follows:

- The QoS Mapping Control mechanism does session level QoS mapping from the session QoS requirements, described by the QSPEC object of each flow, into network service classes in heterogeneous environments, ensuring a clear separation between session level and network operations;
- The QoS Adaptation Control mechanism adapts multi-user sessions to the current network conditions, by dropping or adding the low priority flow(s) of sessions. Thus, receivers subscribe to the available sessions and this mechanism adjusts the number of flows of a session according to the conditions of the networks;
- The Connectivity Control mechanism provides the session connectivity between clusters that supports different address realms, by mapping for instance unicast flows to multicast trees or multicast trees to

multicast trees (e.g. when clusters support different multicast models). This avoids any discontinuous point and allows the construction of inter-cluster distribution trees with efficient shape;

- The MUSC Access Control mechanism is responsible to control the access of mobile users to the announced services, where it is done in access-Q3MAs, based on the interaction with the SIP protocol or the CASM;
- The MUSC-P is the signalling protocol used to exchange control information between Q3MAs. This protocol is essential to support the MUSC operations.

In addition, MUSC supports open interfaces for a tighter communication with external protocols or mechanism, such as other Q3M components and IETF standards.

4.1. MUSC Interfaces

MUSC is composed by four interfaces as follows:

- The *mobility interface* is used on the access-Q3MA to give support for mobility operations and to provide access to mobile users. In the Q3M architecture, when a new session is accepted by MUSC, it triggers CASM that creates a cache for the new session, or associates it with an existing one (the opposite operation is triggered by MUSC when a session ends). In a handover scenario CASM triggers MUSC in the predicted new access-Q3MAs, by transferring to it the session context (including the SOBJ) received from the old agent.
- The *resource allocation controller interface* is used to provide QoS and connectivity control. In the Q3M architecture, this interface is implemented in ingress and egress Q3MAs providing end-to-end connectivity and QoS over heterogeneous clusters. MUSC triggers MIRA by querying it about the QoS characteristic of network services and the available bandwidth capacity of each service class. After selecting a network service for the current session-flow(s), MUSC requests to MIRA the allocation of the resources required by each flow of the session. When there are available network resources for accommodating all flows, MIRA triggers MUSC informing that the request is accepted and furnishes information about the unicast flow or multicast tree to which the session-flow was associated in a cluster or between clusters, allowing MUSC to control the connectivity of the sessions. This interface is also used to control the session quality level, for instance, when a network class has not enough available bandwidth to assure, at least, the minimal packet loss rate of the session, MUSC is requested by MIRA to reallocate the session into other network service class or to adapt the session to the current network condition. When the network capability becomes available again, MUSC is triggered to re-establish the number of flows of a session in order to support the maximum session quality level. Furthermore, when there are no more interested users in a session or flows, MIRA is triggered by MUSC to release the resources associated with the flows of that session. Finally, when inter-cluster route

changes, MUSC is activated by MIRA to install state on the new path and to remove state on the old path of the affected sessions. The above procedures may require the update (request, modify or release) of state in Q3MAs (upstream and/or downstream) which is accomplished by the use of MUSC-P messages.

- The *access control interface* resorts on SIP to allow users to join/leave a session. Initially, applications compose a SIP request message containing SDP information to join an announced session or leave a session. For a join procedure, it includes the SOBJ and a QSPEC object to describe the QoS requirements desired by each flow. This message is sent to the SIP server in the cluster of the user that checks its validity, and forwards this message (based on previous registration and SIP location service) to the suitable access-Q3MA. The reception of this message activates MUSC access control mechanism. After the conclusion of MUSC inter-cluster operations, this interface is used to inform the application, by means of a SIP message via a SIP server, that its request was accepted or not.
- The *transport interface* is used to exchange MUSC-P messages along the data path, where MUSC uses the services offering by *General Internet Signalling Transport (GIST)* [23] to control the transport of its messages. Therefore, when a MUSC agent wants to initiate a communication, for instance setup a session, it constructs a correspondent MUSC-P message and triggers GIST to control the message delivery from the source to the destination. The opposite operation is performed in the destination agent, where GIST triggers MUSC-P furnishing the transported message.

4.2. MUSC Protocol

MUSC-P uses four message types to exchange control information between Q3MAs as follows:

- *SessRequest*: Unicast upstream message with the IP router alert option that can be used to setup or modify sessions. The former is sent by an access-Q3MA towards the Q3MA in which the requested session is activated (Q3MA closest to the source, or at the first Q3MA in which the requested session is found in the way towards the source). For the latter procedure, this message can be generated by any Q3MA in order to give support for QoS mapping and adaptation decisions (e.g. when used by an upstream Q3MA to adapt the session to the network condition, by dropping or adding flows);
- *SessResponse*: Unicast downstream message that can be used to setup or modify sessions. In order to setup a session, a *SessResponse* message is sent by the Q3MA in which the requested session is activated towards the access-Q3MA. This message is also used to modify sessions, where it can be stated by any Q3MA in order to give support for QoS mapping/adaptation or robustness operations (e.g. inter-cluster route changes). Furthermore, it gives support for connectivity control procedures by carrying information, for instance about the unicast flow or multicast channel identifier to which

each flow of a session was associated by MIRA, in each ingress-Q3MA point. Finally, it includes a confirmation about the request (acceptance or rejection);

- *SessAnnounce*: Unicast message sent periodically by an ingress-Q3MA to the other Q3MAs in the same cluster. This message is useful to announce information about the current sessions in an ingress agent, improving the MUSC-P efficiency in what concern session setup;
- *SessRefresh*: Unicast message that is sent periodically by downstream agents (egress and access-Q3MAs) to upstream Q3MAs in which are the ingress points of the sessions as a way to refresh their state.

Figure 4 shows the end-to-end MUSC-P message sequence to setup a session for a first receiver in an access-Q3MA (all posterior requests for the same session in the same access-Q3MA are processed locally and not end-to-end). Therefore, in order to setup a session, a *SessRequest* message is sent by an access-Q3MA towards the ingress-Q3MA in which the requested session is activated. Upon receiving a *SessRequest* message and performing local procedures, such as QoS mapping and connectivity control by interacting with MIRA, the ingress-Q3MA sends a *SessResponse* message to the next downstream Q3MA in the way towards the access-Q3MA (the IP address of the next downstream Q3MA is furnished by MIRA). Then, each Q3MA, in the downstream path, that receives this message will accomplish the same procedures as explained before. The *SessResponse* message is finished when it reaches the access-Q3MA. Moreover, the above mentioned messages can be generated by any Q3MA in order to give support for QoS mapping and adaptation operations. For instance, a QoS adaptation mechanism uses a *SessResponse* messages to signal downstream agents as a way to reduce the number of flows of a session.




Figure 4 - MUSC-P message sequence to session setup

Furthermore, starting from the setup of the first session in an ingress-Q3MA of a cluster, *SessAnnounce* messages are periodically sent by the ingress agent to the other Q3MAs in the same cluster in order to announce its ongoing sessions (it is assumed that the each Q3MA maintains a list with their neighbours). For instance, in order to setup a session in multi-homed clusters with inter-cluster asymmetric routing, the distribution of this message allows to an access-Q3MA to send a *SessRequest* message directly to an ingress-Q3MA of its cluster rather than towards the source. This control avoids that a session has more than one ingress point in the same cluster, reducing the waste of network resources and minimizing the session setup delay. Furthermore, when a session ends in an ingress-Q3MA, this stops its announcement and the state stored about that session in the Q3MA neighbours is removed by soft-state. Finally, in order to contribute to the system robustness, *SessRefresh* messages are periodically sent by downstream agents (egress and

access-Q3MAs) to upstream agents in which are the ingress points of the sessions as a way to refresh their state. If a session's state is not refreshed in a certain period of time by using SessRefresh message, its state is removed by soft-state.

In addition to the messages, MUSC supports two databases to control multi-user sessions, called Local Database (MUSC_LDB) and Cluster Database (MUSC_CDB). The former is responsible for maintaining state about the current sessions-flows in an agent (including its quality level) and for controlling the connectivity of sessions between clusters, where the state is refreshed by MUSC-P SessRefresh messages. The latter is responsible for maintaining state about the activated sessions in ingress-Q3MAs and it is updated by SessAnnounce messages. As explained before, MUSC uses a soft-state approach to maintain per-session and per-flow state in its databases, contributing to the system robustness. Therefore, two different timers are used to control the propagation of periodic messages and to remove state in the MUSC databases. The default values for sending SessAnnounce and SessRefresh messages and for the time that waits before deleting state are 30 seconds (named *refresh_timeout* timer) and 90 seconds (named *clean-up_timeout* timer) respectively. The values were chosen based upon *Border Gateway Protocol* (BGP) *keep_alive_time* variable and *BGP hold_time* variable [24]. However, the suitable values for these times will be studied in future work.

5. EXAMPLE OF MUSC OPERATIONS

This section presents some examples of the MUSC operations in the Q3M architecture, by interacting with MIRA and CASM in three different scenarios with one source publishing one session, which is required by several mobile subscribers. The examples are not exhaustive and do not attempt to cover all details of MUSC operations. Moreover, it is assumed that the clusters are NSIS-aware and the transport of MUSC-P messages is controlled by GIST. However, in order to simplify the explanation, the interaction between MUSC and GIST will not be described.

5.1. Session Setup

Figure 5 shows how MUSC-P operates to setup the same session for several receivers. Upon receiving a session announcement, via *Session Announcement Protocol* (SAP) or HTTP, the application on host *R1*, at *T1*, uses a SIP (SIP-INVITE) message to subscribe the multi-user session *S1*. It is assumed that this message is sent to the SIP server in user's cluster that checks its validity. If the user is not authorized, the server replies a correspondent message informing that the request is rejected. Otherwise, the SIP server forwards this message to the user's access-Q3MA.

After receiving the request, MUSC access control in Q3MA-B retrieves the SOBJ and based on the session identifier, it verifies in the MUSC_LDB if the requested session is activated locally. Then, it checks in the MUSC_CDB if the session is ongoing in the cluster. Since the session is not present locally nor in cluster *C1*, MUSC-P is requested to send a SessRequest message, containing the SOBJ (which includes the QSPEC object of the session-flow), towards the source *S1* with IP Alert enabled (since *C1* is a multi-homed cluster, this message can exit through Q3MA-C or Q3MA-D). However, the interior routers forward the message to Q3MA-C. The *IP Alert Option* in SessRequest messages allows MUSC agents to examine more closely the content of the message, stopping its progression if they have

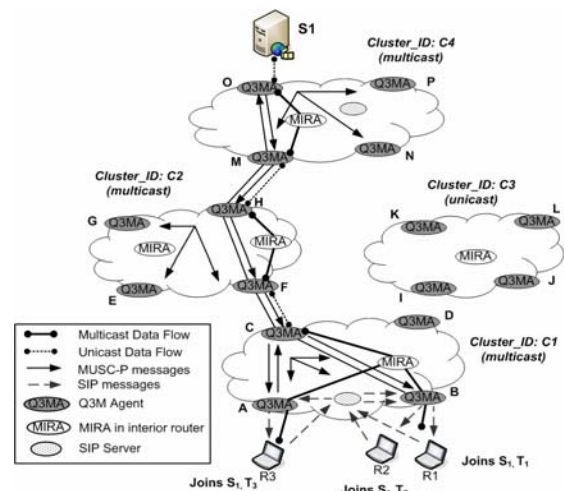


Figure 5 - Session setup

information about the session in their MUSC_LDB, which happens in Q3MA-O.

In Q3MA-O, MUSC receives the SessRequest message and verifies that it is activated locally in MUSC_LDB (it is assumed that based on previous negotiation, the source application informs MUSC about the published unicast session-flow(s) and MUSC keeps it in its MUSC_LDB). After that, MUSC interacts with MIRA to query about network services in the path towards Q3MA-B. Based on the response and the QoS parameters described in the QSPEC object, MUSC maps the requirements of *S1* into the most suitable network service class and requests MIRA to configure the required bandwidth and to allocate a distribution channel for the session-flow. In IP multicast clusters, MIRA allocates a 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 flow(s) of session will be associated by MUSC.

When MIRA finishes its job, MUSC updates its MUSC_LDB (based on the connectivity information furnished by MIRA, MUSC translates all incoming unicast flows addressed for *S1* to the correspondent multicast tree) and sends a SessResponse message to the next downstream Q3MA furnished by MIRA (in this case Q3MA-M) with the SOBJ and local address allocated by MIRA for the session. After that, since the Q3MA-O is the ingress point of the session in *C4*, it sends a SessAnnounce message to its Q3MA neighbours (M, N and P) in order to update their MUSC-P_CDB with information about its current sessions.

After all ingress operations, in Q3MA-M MUSC receives the MUSC-P SessResponse message and updates MUSC_LDB with information about the SSM channel used by this session-flow inside the cluster *C4*. Furthermore, MUSC interacts with MIRA in a similar way as described before. However, since unicast is always used between clusters, MIRA obtains the unicast IP address of the next Q3MA (Q3MA-H) and allocates a pair of source and destination ports for each flow of the session. Based on the connectivity information provided by MIRA, MUSC will translate all packets coming from the intra-cluster SSM channel to the source and destination IP addresses and ports to be used for this session-flow between *C4* and *C2*.

The interaction between MUSC and MIRA occurs in all Q3MAs and only in Q3MA-H and Q3MA-C the SessAnnounce message is triggered, because they are the ingress point of the session in the clusters *C2* and *C1* respectively. However, in Q3MA-B, since it is an access-Q3MA, MIRA configures the network and the multicast tree only on the outgoing interface. However, if unicast is used in the wireless interface, MUSC maps each session-flow to a list of interested users, and performs the required packet duplication, increasing the usage of resources. Furthermore, it is assumed the existence of a layer two technology, which provides service differentiation with definition of classes on the wireless link, such as 802.11e or 802.16 [25]. In addition, MUSC triggers CASM to activate the local cache for the session and furnishes information concerning session context (including the SOBJ), after which MUSC access control replies the SIP request (SIP-200 OK message). In multicast environments, MUSC access control includes in the SDP protocol information concerning the multicast channel to be used in the access cluster. Thus, the application can join the correct multicast channel by using the IGMPv3/MLDv2.

After the session setup for *R1*, the receiver *R2*, at *T2*, sends a SIP message to subscribe *S1*. Since the requested session was already activated for *R1* in the same access-Q3MA (it is verified based on query in MUSC_LDB), 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.

Finally, at *T3*, the host *R3* in the Q3MA-A wants to join *S1*, which is ongoing in the cluster. Receiving the SIP request, MUSC verifies that the requested session is not current locally, but it is presented in the cluster (based on query in MUSC_CDB, where the information was previously announced). After that, MUSC-P is triggered to accomplish the request and sends a SessRequest message addressed directly to the Q3MA-C rather than towards the source. The rest of the session setup process is done only in *C1* and in a similar way as explained above. The intra-cluster control performed by MUSC avoids that a same session has more than one ingress point for the same cluster, reducing the waste of network resources and minimizing the session setup delay.

5.2. Inter-cluster Mobility

Figure 6 illustrates an example of the configuration of a session on a new path, due to inter-cluster mobility, where the receiver is moving to a cluster (*C3*) without IP multicast capability. This example uses the same scenario described previously in Figure 5, but the cluster *C2* was removed to simplify the explanation. It is assumed that CASM detects that *R1* is moving away from the access point of Q3MA-B, and predicts in Q3MA-I, the next most probable access-Q3MA. After handover prediction, CASM signals the predicted Q3MA-I (procedure not illustrated in the figure) with session context information, including the SOBJ where the QoS desired by each flow is described in the QSPEC object. In the new predicted access-Q3MA, CASM triggers MUSC, which setups the session on the new path to Q3MA-I, including the pre-reservation of network resources by MIRA. CASM action in Q3MA-I ends up with its configuration, and a reply is sent to CASM in Q3MA-B. After that, in Q3MA-B CASM takes a handover decision to Q3MA-I and informs MUSC about the movement of the mobile device, allowing MUSC to

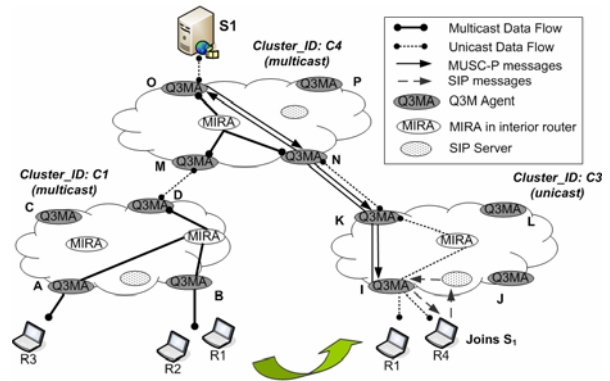


Figure 6 - Inter-cluster mobility and session multicasting

adjust the number of local receivers and to remove the state of the session in the MUSC_LDB if necessary. In this case, the session is not removed, since *R2* is still attached to Q3MA-B.

After the accomplishment of the handover process of *R1*, a new receiver *R4* in the Q3MA-I wants to subscribe session *S1*. After receiving the SIP request, MUSC identifies that the requested session is activated locally in the MUSC_LDB and increases the number of local receivers for that session-flow. After that, MIRA is requested to configure the network service on the outgoing interface, and MUSC updates its translation state with the IP address of *R4* and replies the correspondent SIP message to receiver confirming the operation. The packet duplication is only required in Q3MA-I, reducing the waste of network resources.

5.3. Inter-cluster Route Changes

In order to contribute to the system robustness, MUSC controls the reestablishment of session quality level when inter-cluster route changes (e.g. due to link failure). It is assumed that MIRA detects inter-cluster rerouting events and triggers MUSC to perform the session setup on the new path and tear down the reservation on the old path.

Figure 7 shows a simple example of inter-cluster route changes, due to a link failure between Q3MA-O and Q3MA-M. After detecting the failure, MIRA triggers MUSC in Q3MA-O to control the reestablishment of *S1* on a new path (MIRA furnishes information about the session-flow, including the IP address of the access-Q3MA and bit rate). MUSC performs its job in the same way as a normal downstream session setup, where it queries MIRA about the network services in the path towards access-Q3MA that replies it and informs that the Q3MA-N is the egress point for this session towards the Q3MA-B. Based on the response, MUSC updates its MUSC_LDB with information about the new downstream Q3MA for *S1* towards cluster *C1* and does not require MIRA to configure network resources, because *S1* is already current to Q3MA-N (in this scenario there is no service degradation due to route changes). Then, MUSC generates a SessResponse message to the Q3MA-N that receives this message, updates its MUSC_LDB, and requests MIRA the available class of services in the inter-cluster link. Based on the response and on the QSPEC object, MUSC selects the most suitable service class and triggers MIRA to allocate network resources and a distributed channel for *S1* in the link between Q3MA-N and the next Q3MA (in this case Q3MA-D). Upon receiving a response, MUSC updates its MUSC_LDB with

session connectivity information and forwards the SessResponse message to the Q3MA-D.

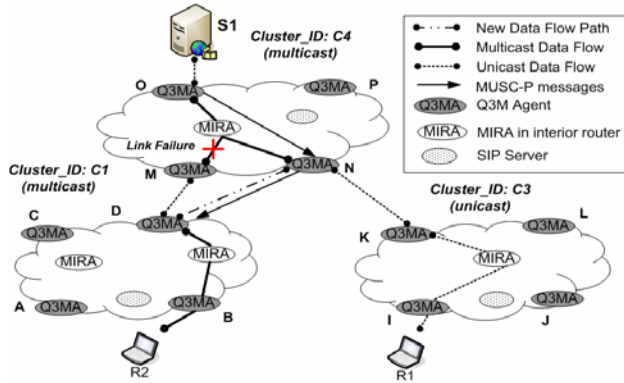


Figure 7 - Inter-cluster route changes

When Q3MA-D receives the message, verifies that *S1* is already activated towards the access Q3MA-B (based on query in its MUSC_LDB), updates the MUSC_LDB with connectivity control information (including the IP address of the new ingress point for *S1*) and finishes the message. Finally, the state stored by Q3MA-M will be removed by soft-state because no SessRefresh message addressed to *S1* will arrive before the expiration of *cleanup_timeout* interval (MUSC removes state in MUSC_LDB and interacts with MIRA to release its state about the deleted session).

However, if during the reestablishment of the session on the new path, the minimum session quality level required for the session can not be assured in the most suitable service class selected by MUSC (e.g. the minimal packet loss rate), a QoS adaptation is performed to adapt the session to the current network conditions. For instance, MUSC can be configured to reallocate the session into another class of service that assures the full bandwidth committed for each session-flow or to refuse the incoming session. The former avoids session blocking, while degrading the overall QoS assurance. With the latter procedure, MUSC-P is required to inform the upstream and/or downstream Q3MAs about the dropped session or flows of a session, avoiding waste of network resources, when a session can not be accommodated in a cluster or between clusters.

6. PERFORMANCE EVALUATION

This section presents an initial evaluation of MUSC, focusing on the signalling operations used to control session setup to multiple users. The analysis of the session setup time is important because if the signalling delay is large, it induces high delay especially in handover scenarios.

Three sets of simulations were done using the *Network Simulator 2.28 (NS)* in order to show the applicability of MUSC-P SessAnnounce messages, the comparison of the delay and signalling overhead caused by MUSC and SIP to session setup for multiple users, and the delay caused during the session setup when some QoS assurances are given to the session (MUSC is completed with MIRA, in order to allocate network resources to the existing session). To accomplish these evaluations, several experiments were done and different scenarios were generated by BRITE [26], where the intra and inter-cluster links have a bandwidth of 10 Mb/s and their propagation delay is attributed randomly by BRITE according distance between the edges of

each link. The size of messages used in the simulations is listed in Table 1 (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. In addition, it is assumed that the session was previously announced to several receivers and identified by a unique identifier. Furthermore, each cluster is composed by fourteen routers and one SIP server is placed in the cluster of the source and another one in the access cluster.

Table 1 - Message size for session setup

Message	Protocol	Length (bytes)
SIP INVITE	SIP with SDP	620
SIP 200 OK	SIP with SDP	450
SessRequest	MUSC-P	84
SessResponse	MUSC-P	88
SessAnnounce	MUSC-P	44
Reserve	MIRA protocol	70
Response	MIRA protocol	52

The first simulation analyses the efficiency of MUSC-P SessAnnounce messages in multi-homed clusters, based on the generic scenario illustrated in Figure 5. Six different scenarios of four clusters were generated and twelve experiments have been carried out (two for each scenario). Firstly, it is measured the session setup time to accomplish the requests generated by two receivers, *R1* and *R2*, respectively, when the SessAnnounce messages are not enabled in the ingress points of the clusters. Then, in another sets of tests, it is analyzed the session setup time for the same requests, when the SessAnnounce messages are activated in the ingress points of the clusters. Each receiver is connected to different access-routers in the same access-cluster. Figure 8 shows the sum of the session setup times in each scenario with and without the use of SessAnnounce messages.




Figure 8 - Sum of session setup times with and without MUSC-P SessAnnounce messages

The results show that the use of SessAnnounce messages decreases the delay to establish the session for *R2* in 66% of the tests. The session setup delay for *R1* is the same in each scenario, because MUSC exchanges end-to-end messages to setup the session for a first receiver. On average, the SessAnnounce message reduces in 20% the time to establish the session for *R2*, and in the best scenario (scenario 3), the receiver *R2* has to wait less than 9 ms for the session setup rather than 31 ms. In scenarios 4 and 6, the SessAnnounce messages did not contribute to reduce the session setup time, because the upstream SessRequest message was forwarded (by interior routers) to the Q3MA that is the ingress point for the session in the cluster. Besides minimizing the session setup time, the use of SessAnnounce messages reduces the signalling overhead caused by MUSC-P messages for the establishment of the same session for a first receiver placed in another access-router. Despite the overhead caused by

SessAnnounce messages, in this simulation, the use of these messages reduced in approximately 40% the signaling overhead used to install the session for *R2*, because MUSC-P messages are only exchanged inside the access-cluster. On average, it minimized in 32% the bandwidth consumption originated by MUSC-P messages to accomplish the request for *R2*. With and without *SessAnnounce* messages, the overhead consumed to session setup for *R1* is the same, because in each experiment it is the first receiver that is requesting the session.

In addition, a second simulation compares the session setup time of MUSC and SIP. We use the results of previous studies that have investigated traffic distribution on the Internet in order to make realistic assumptions [27]. These studies show that 26%, 40%, 26% and 9% of users 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 to compare MUSC with SIP, we performed 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 MUSC) and, then, SIP is used together with MUSC. Figure 9 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 9 - 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-P 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 the first receiver in an access-router, being all posterior requests processed locally in the access-agent. If the same session-flow is already present in the ingress-router of the access-cluster, which happens when receiver *R4* 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. The reduction of session setup time to a group of users achieved by MUSC is important specially in mobility scenarios, because if a mobile user is moving to an access-router in which its session is activated, MUSC aims to decrease the session setup time, giving support for achieve a seamless handover.

Figure 10 illustrates in more detail the results obtained when there are four clusters between the receivers and the source. 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-P 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-P 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 Figure 10.




Figure 10 - Variation of the number of visited routers

The last simulation addresses the session setup time when some QoS assurance is required. This evaluation, which results are illustrated in Figure 11, 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.




Figure 11 - Sum of session setup times for MUSC and MIRA

Figure 11 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.

7. CONCLUSION AND FUTURE WORK

This paper proposed a new NSIS NSLP called MUSC, which provides a solution to control multi-user sessions in the next

generation wireless system. MUSC is composed by four control mechanisms and one signalling protocol. Supported by its protocol, MUSC controls the session quality of service mapping and adaptation, the ubiquitous access of mobile users to the available services and the session connectivity between clusters with different address realms. Furthermore, even though MUSC is being developed within the Q3M architecture context, it supports open interfaces, allowing mobile operators to use mobility, connectivity and resource control schemes of their choice. The initial evaluation presented in this paper shows the setup time reduction brought by MUSC, when compared to a solution using only SIP. Furthermore, it shows the impact that the extra QoS control provided by MUSC and MIRA has on the session setup time. The implementation of the MUSC control mechanisms and protocol are currently under development.

Future plans include the standardization of MUSC in the NSIS working group, and a detailed analysis of the MUSC control mechanisms, including their relationship with existing protocols and Q3M components. Further simulations will be done to verify the functional and performance behaviour of MUSC, such as scalability, the convergence time when a fault happens in a link, the MUSC behaviour due to the mobility of users and the impact of the QoS mapping, adaptation and connectivity control on the mobile user and/or network perspectives.

8. ACKNOWLEDGMENTS

Other authors from University of Coimbra (UC) contributed for this paper, namely Augusto Neto, Luis Veloso and Marilia Curado. This project is ongoing at the Laboratory of Communications and Telematics (LCT) of the Faculty of Science and Technology of the UC (FCTUC). It is supported by NTT DoCoMo Euro-labs, by the Portuguese Ministry of Science, Technology and High Education (MCTES), and by European Union FEDER - program POSC (projects Q3M and SAPRA).

9. 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, October 2005, pp. 42-48.
- [2] V. Saranga and J. Chen. *Comparative Study of Protocols for Dynamic Service Negotiation in the Next-Generation Internet*. IEEE Commun Magazine, vol. 44, issue 3, March 2006, pp. 151-156.
- [3] R. Prior *et al.* *Heterogeneous Signaling Framework for End-to-end QoS support in Next Generation networks*. IEEE System Sciences - HICSS '05, January 2005, pp. 305-315.
- [4] R. Chalmers, and K. Almeroth. *On the topology of multicast trees*. Networking, IEEE/ACM Transactions on, vol. 11, issue: 1, February 2003, pp. 153-165.
- [5] R. Hancock *et al.*. *Next Step in Signalling (NSIS): Framework*. IETF, RFC 4080, June 2005.
- [6] J. Rosenberg *et al.* *SIP: Session Initiation Protocol*. IETF, RFC 3261, June 2002.
- [7] H. Schulzrinne *et al.* *Real Time Streaming Protocol 2.0 (RTSP)*. IETF, Internet Draft, March 2006.
- [8] X. Fu *et al.* *NSIS: A New Extensible IP Signalling Protocol Suite*. IEEE Communication Magazine, Internet Technology Series, October 2005, pp-133-141.
- [9] J. Manner *et al.*. *NSLP for Quality of Service Signalling*. IETF, Internet Draft, March 2006.
- [10] M. Stiemerling *et al.* *NAT/Firewall NSIS Signalling Layer Protocol*. IETF, Internet Draft, April 2006.
- [11] J. Ash, A. Bader and C. Kappler. *QoS-NSLP QSPEC Template*. IETF, Internet Draft, March 2006.
- [12] D. Zagar, D. Vilendecic and K. Mlinaric. *Network resource management for the services with variable QoS requirements*. The 6th International Conference on Telecommunications, ConTEL, vol. 2, Issue 3, June 2003.
- [13] M. El-Gendy *et al.* *Paving the first mile for QoS-dependent applications and appliances*. IWQOS 2004, Twelfth IEEE International Workshop on, June 2004.
- [14] Z. Mammeri. *Approach for End-to-End QoS Mapping and Handling*. Wireless and Optical Commun. Networks. Second IFIP International Conference on, March 2005, pp 265-269.
- [15] T. Shanableh and M. Ghanabari. *Multilayer Transcoding with Format Portability for Multicasting of Single-Layered Video*. IEEE Transact. on Multimedia, vol. 7, February 2005.
- [16] R. Hayder, S. Mihaela and K. Shirish. *Scalable Video TransScaling for the Wireless Internet*. Journal on Applied Signal Processing, number 2, February 2004, pp. 265-279.
- [17] H. Chiu and K. Yeung. *Fast-response Receiver-driven Layered Multicast with Multiple Servers*. Commun, 2005 Asia-Pacific Conference on, October 2005 pp. 259-263.
- [18] M. Handley, V. Jacobson and C. Perkins. *SDP: Session Description Protocol*. IETF, Internet Draft, January 2006.
- [19] B. Cain *et al.* *Internet Group Management Protocol, Version 3*. IETF, RFC 3376, October 2006.
- [20] R. Vida and L. Costa. *Multicast Listener Discovery Version 2 (MLDv2)*. IETF, RFC 3810, June 2004.
- [21] H. Holbrook and B. Cain. *Source Specific Multicast for IP*. IETF, Internet Draft, October 2005.
- [22] P. Mendes, J. Andrés-Colás and C. Pinho. *Information Model for the Specification of QoS Agreements among Ambient Networks*. IEEE International Symp. on Personal Indoor and Mobile Radio Communications, September 2005.
- [23] H. Schulzrinne and R. Hancock. *General Internet Signalling Transport*. EITF, Internet Draft, February 2006.
- [24] Y. Rekhter and T. Li. *A Border Gateway Protocol 4 (BGP-4)*. IETF, RFC 1771, March 1995.
- [25] H. Wang, W. Li and P. Agrawai. *Dynamic Admission Control and QoS for 802.16 wireless MAN*. IEEE Wireless Telecommunications Symposium, April 2005, pp. 60 - 66.
- [26] A. Medina *et al.* *BRITE: An approach to universal topology generation*. IEEE MASCOTS '01, August 2001, pp. 346-356.
- [27] U. Steve *et al.* *Implications of interdomain traffic characteristics on the traffic engineering*. SAC '04: ACM symposium on applied computing, March 2004, pp. 339-346.