A SCALABLE AND BUSINESS-ORIENTED FRAMEWORK FOR INTER-DOMAIN QUALITY-OF-SERVICE

Vítor Jesus, Susana Sargento and Rui L. Aguiar
Instituto de Telecomunicações, Universidade de Aveiro, Aveiro, Portugal

Keywords: Quality-of-Service, interdomain, 4G networks, federation.

Abstract: This paper presents and discusses a framework for end-to-end Quality-of-Service in a network operator environment. We focus on the inter-operator segment. Contrary to usual approaches, we consider that, in the current state-of-the-art, the interdomain QoS problem complexity resides on the business relations between administrative domains. Therefore, we attempt to decouple the technical problem of providing end-to-end traffic assurances from the business problem of setting up partnerships. We combine the two aspects to build a framework that allows a smooth integration of the inter-domain, the intra-domain and the access segments. We also provide an IntServ-over-DiffServ architecture that is able to cope with mobility scenarios.

1 INTRODUCTION

Quality-of-Service (QoS) remains an open problem despite all research so far. Although many techniques and architectures have been proposed, current commercial Internet lacks basic building blocks such as general support for traffic discrimination[1]. This can be explained over two dimensions: the need for QoS and the "all-or-nothing" implementation issue.

The actual need for QoS is still somewhat debatable, considering that networks today still didn't reach enough convergence to force the need for traffic differentiation. We still have two types of networks: legacy networks built around well-known services such as voice, and the general-purpose Internet that, in theory, can carry all types of traffic. While typically the Internet is used for applications that run over a "best-effort" network, applications that need QoS use legacy networks such as ISDN, ATM or GSM. It is, however, a weakening argument since the desired convergence is quickly approaching. A notable example is the maturing specifications of IP Multimedia System (IMS, from 3GPP) over which TISPAN (from ETSI) will build, considered the core of next-generation commercial networks.

On another perspective, "killer applications" that demand for QoS only recently have gained the attention of mass-market players. This is the case of VoIP, already a prevalent technology in the business environment. Others, such as Video-on-Demand or IPTV will be pushing the need for QoS.

The second set of arguments concern the "all-or-nothing" characteristic of QoS: QoS makes only sense if deployed e2e. The problem adds in deployability complexity if we note that a typical user will not tolerate domain discrimination. This means that users expect to have the same service for every possible destination. Therefore, following this rationale, this is a two-fold problem: a QoS-enable Internet must have the collaboration of all domains composing it and it must comprise all trenches between user equipments.

However, this all-or-nothing argument is currently not so dramatic for one main reason. On one hand, we still have a large installed capacity in core networks (hundreds of Gbps); on the other hand, access networks have limited capacity (on the order of the Mbps). The combination of these two factors limits the expectations of users so that one can still speak of generalized overprovisioning. However, either by caution or by forecast, it is expected that overprovisioning won't last. A quite strong example is the recent trend to carry TV over

[1] We distinguish the particular problem of traffic differentiation from the overall problem of traffic discrimination, that, in this context, includes all actions needed to setup and carry a flow that needs QoS assurances.

Jesus V., Sargento S. and L. Aguiar R. (2007).
A SCALABLE AND BUSINESS-ORIENTED FRAMEWORK FOR INTER-DOMAIN QUALITY-OF-SERVICE.
In *Proceedings of the Second International Conference on Wireless Information Networks and Systems*, pages 282-288
DOI: 10.5220/0002151502820288

the Internet. This will push broadband (tens of Mbps) everywhere and we will assist to real and effective traffic discrimination.

This work will discuss the scenario where overprovisioned domains won't be possible and e2e QoS will be required by end customers. The main problem we tackle is how to support interdomain QoS in a mesh of many different domains in a scalable and business-friendly way. In section 2 we show that the complexity of the problem has important roots in the business environment of the Internet. We also discuss related work. In section 3 we outline our architecture and in section 4 we detail the most important building blocks. In section 5 we give our concluding remarks and leave some open questions that will be subject of future work.

2 PROBLEM STATEMENT

In we show the two generic topologies of a WAN: one hierarchical and the other meshed.

Even though the Internet is architecturally a mesh of domains, current topology shows some hierarchy (Meireles, 2004). On the bottom layer, access networks, to which users are directly connected, are connected to core networks, most likely, under the same administration. The next level could be fitted to national backbones (for Portugal, the GigaPIX (fccn.pt) is a good example) and the higher layers are tier-2 or tier-1 domains that, typically, are only transit domains, aggregating traffic from many core networks and typically spanning over the globe. However, the exact domains that each small domain is connected to may vary so that a mesh of domains comes up. Perhaps the best way to capture the Internet topology is to classify each domain according to its nature (see Figure 2):

- access domains, connecting directly end users.
- core domains, aggregating end-user traffic from a set of access networks under the same administration;
- transit domains, connecting core domains.

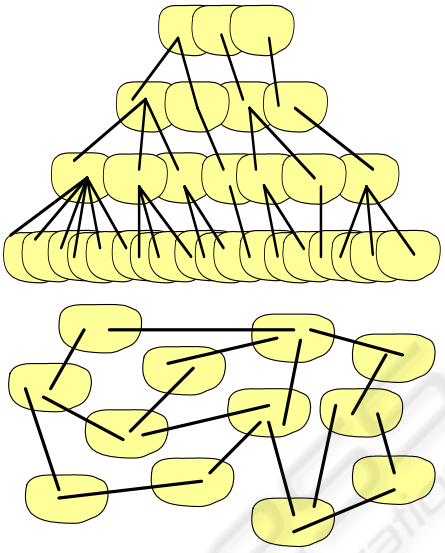

Figure 1: Generic topologies of the Internet. On the top, a hierarchical topology; on the bottom, a mesh of domains is shown.

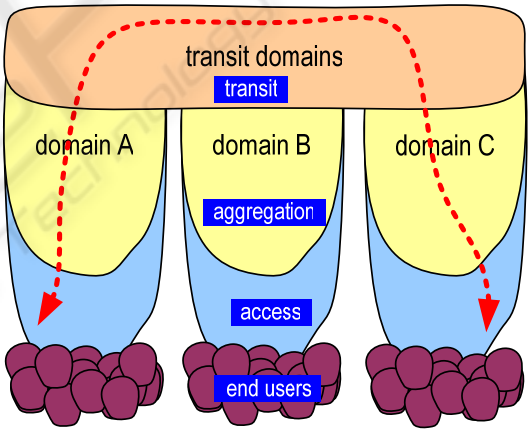

Figure 2: Segment decomposition of domains. The dotted line shows a connection between two end users.

2.1 Related Work

Technically, e2e QoS in a packet-switched network has two main components: per-node packet handling and session setup. Both components have several solutions that span from mimicking circuit-switched networks (e.g., ATM, IntServ/RSVP (Braden et al., 1994), MPLS (ietf.org)) to loose assurances (e.g., DiffServ (Blake et al., 1998), probing-based). Each solution has its own problems and is typically suitable in a certain context of scope. For example, IntServ/RSVP framework is well-known to have scalability problems but is flexible enough to handle QoS at a finer level; on the other hand, DiffServ is

far more scalable than IntServ but at the cost of not being able to handle well per-flow QoS.

A clear trend is to consider an IntServ-over-DiffServ framework (see, e.g., (Sargento et al., 2006)) which is also the approach taken by next-generation networks such as IMS or TISPAN. In such approach, the stored state incurred in each QoS reservation is progressively diluted as we enter the core network. At the access network (including the user equipment), per-flow management is realized by a combination of the way the access technology works (e.g., channels of WCDMA, time-slots in IEEE 802.11e) and per-flow state storage in the interface between the access network and the core network. On the core network, typically managed by a Bandwidth Broker, a flow can be mapped to a certain DiffServ code point (DSCP) and packets are handled in a per-aggregate basis. As packets move up in the hierarchy, aggregation of aggregates is supposed to happen (e.g., using MPLS paths) in order to avoid prohibitive scalability problems: a single link of a tier-1 transit domain may carry flows of hundreds of millions of users. Efficient techniques for aggregating flow bundles (including setup signalling) in the interdomain segment have been proposed (Pan et al., 2000). The main conclusion to be taken is that, from a technical point-of-view, interdomain QoS inherits from intra-domain QoS solutions and, from an architecture point-of-view, there is a clear vision of how a solution will be.

However, the problem is still open from the point-of-view of signalling. It is clear that the same solutions for intradomain could be applied (e.g, RSVP (Braden et al., 1997) or NSIS (ietf.org.../nsis-charter.html)) or similar such as the proposal of the QBone project that provides a specific signalling protocol (SIBBS (qbone.internet2.edu)) for interdomain QoS session setup. The IETF WG NSIS is already addressing such problem and some proposals already exist (Cordeiro et al., 2006). One key feature of NSIS, as of today (work in progress), is to allow independence for the QoS model the signalling traverses. As we'll see, we consider this aspect to be of utmost importance. However, we won't use real-time interdomain signalling.

2.2 A Model for Interdomain QoS

We argue that a key piece is still missing: a widely accepted model for interdomain QoS. One should note that an interdomain model for legacy networks exists for several decades: in PSTN, two network providers would agree on respecting ITU-T guidelines. There is, however, a major difference between the Internet and PSTN: PSTN has a single, well-known and well-regulated service (voice). This is in sharp contrast with the Internet where not only there may be several possible services, but also no well-known service is currently universally accepted (except, possibly, services similar to voice).

Yet another important difference is related to the exact notion of interdomain "services". While a "service" may be defined in terms of end-user (e.g., the 4 3GPP classes), "service" is a much different concept between domains, for the reasons expressed above: while a "voice call" may have a clear meaning in the access network, a bundle of aggregates of "voice calls" in a interdomain link represents a much different concept. There are reasons to believe that well-known interdomain services won't be defined in the short-term:

- the Internet is a worldwide mesh of domains whose technologies may be very different and the same service will be implemented in much different ways. Although the concept of service should be independent of the supporting technology, for transit domains, the technology set the available services.
- the set of services that can be provided by a general-purpose data network such as the Internet is expected to be extremely large.
- while end-user services may be changed frequently (following market trends), transit services will be much slower to create and discontinue.
- while transit domains will define its services in term of packet bundles, access domains will primarily define services in terms of end users. There must be, then, some mapping between one and another. This will cause a fast dynamic system that typically cannot be tracked by standardization bodies.

3 PROPOSED ARCHITECTURE

The set of previous considerations set the grounds over which our architecture is laid off. We distinguish three dimensions in a QoS-enabled network: the user, the business and the technicalities. These dimensions are unified by the concept of *Service-Level Agreement* (SLA) – see Figure 3. Each type of SLA has three main components:

- a legal component, consisting of a contract whereby a party states the conditions under which the service is provided

- a control component, consisting of 5 generic functions: the typical A4C functions (Authentication, Authorization, Accounting, Auditing and Charging) added and a fifth function providing means to check for service compliance (e.g., remote measurements for transit services).
- and the service level specification that depends on the type of SLA.

The envisioned types of service are the following:

- *transit-SLA* (tSLA): in QoS terms, this has the common meaning of an agreement between two domains concerning the transport services with the SLS also having the common understanding [14]
- *user-SLA* (uSLA): this is a formalization of the agreement between the user and service provider
- *peering-SLA* (pSLA): On one hand, this is the general SLA from where other SLAs derive since it specifies a peering relationship between two parties. It is here used as a peering relationship between two service providers.

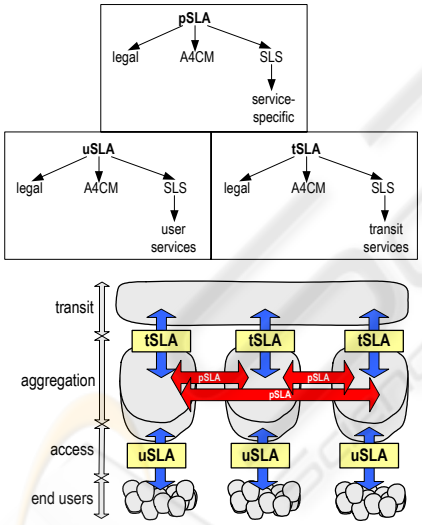

Figure 3: The SLA as the unifying concept for QoS.

An example is the way in which VoIP will be deployed. There will be, most likely, pair-wise agreements between all VoIP operators, worldwide. It is yet unclear how these agreements will be settled but probably there will be a mix between direct peering (pair-wise) and using intermediary such as clearing houses (Figure 4). One fact is that VoIP termination is clearly just one possible service out of many possible.

In this context, a Clearing House (CL) would settle agreements with other domains and represent them for other domains. For several reasons, a domain may chose to be represented by an external party. It would then delegate SLA handling to that 3rd party.

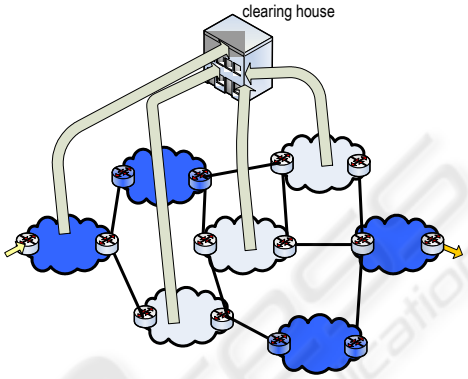

Figure 4: Two types of settlements between domains: direct, pair-wise (darker domains) and by means of a hierarchy of clearing houses (one shown).

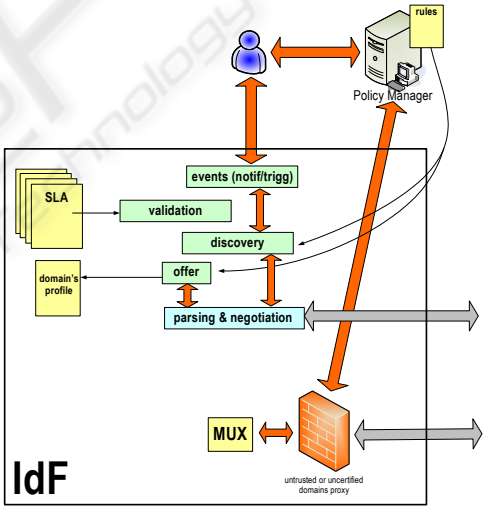

Figure 5: The Inter-domain Function responsible for automatic settlement of agreements.

3.1 The Interdomain Function

The mentioned possibility of on-demand peering plays an essential role in our architecture. In order to automate discovery and offer of transit and termination services, we define a special function: the Inter-Domain Function (IdF) – Figure 5.

There are three main functions of IdF:

- the *crawling function*[2] find domains with which there are no SLAs and that may be, in the future, transit or termination domains;
- the *negotiation function*, used to settle a new SLA or a new service;
- the *offer function* that advertises the domain's own profile to the outside.

These functions will be used to build e2e QoS services as we'll explain in the next section.

3.2 End-to-End QoS

Following our strategy, there are three main steps to setup a QoS session:

1. since QoS is a packet transportation service, establish peering agreements with all domains between source and destination
2. contract resources enough to carry the expected volume of services
3. setup in real-time QoS sessions

Hence, before two domains are able to setup a QoS session, they must acquaint with one another and negotiate a pSLA (or a tSLA if a transit domain is involved). The crawling function serves this purpose: it is to do the following generic actions (see Figure 6 and Figure 7):

- search the routing database (from the interdomain routing protocol, e.g., BGP) for destinations and/or Autonomous Systems. It will find transit-only domains, terminating-only domains, domains that are both.
- for relevant destinations, the seeking IdF (sIdF) will find the IP address to the offering IdF (o-IdF) (perhaps through a BGP extension or a DNS well-known name).
- upon contacting the domain's IdF (or its representative such as a mandated clearing house), o-IdF will trigger its offer function.
- s-IdF will download the offering domain's profile from o-IdF and eventually build services from it. The domain's IdF can be simply a server holding the domain's profile listening on a well-known port (the Extensible Provisioning Protocol [15] may be here a good solution).
- these services will be the subject of a SLA exchange between the two domains.

During negotiation, additional information may be needed such as they may exchange trust-related information (e.g., references). If references, they may be checked to verify the assertions of the o-IdF.

- s-IdF will then buy a certain amount of resources. This amount is related to the users it supports.

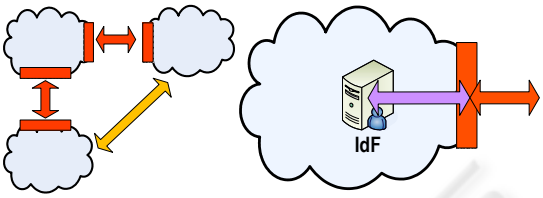

Figure 6: Use of IdF for peering. Shown three domains, each implementing an IdF that interfaces other domains.

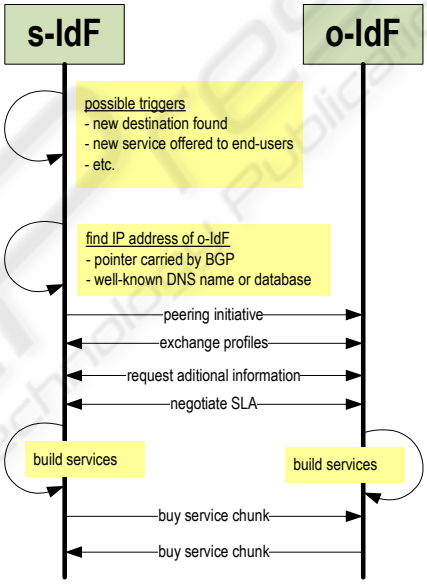

Figure 7: Peering setup.

3.3 Real-Time Operations

Our QoS architecture will follow an IntServ-over-DiffServ strategy for the following reasons:

- it is able to provide flow-granularity at UNI (user-network interface)
- DiffServ in the core network provides scalability while allowing for advanced services at the cost of expensive and complex traffic engineering operations (e.g., aggregating different PHBs)
- since DiffServ does not require specific signalling along the path (in theory, just packet marking) it is much easier to interface different technologies (e.g., SDH on one side and point-to-point ethernet on the other)

[2] This name was coined after the algorithms that scan websites for content indexing.

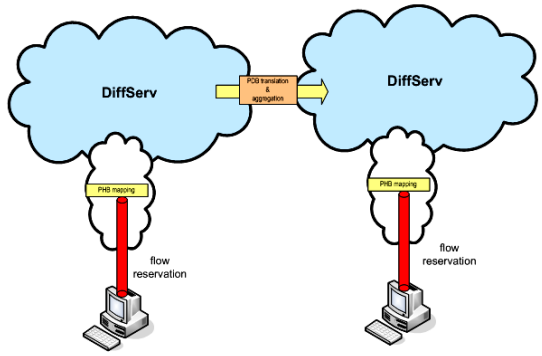

Figure 8: Top-level view of our QoS architecture.

Show in Figure 8 is a top-level view of our proposed architecture. A user requests QoS for a certain flow through specific signalling (an example will be provided). The network will map that flow on to a suitable PDB (packet marking). Upon leaving the domain, packet will be remarked at the interdomain interface in order for the next domain to know how to handle the packet. Packer remarking will continue until it reaches the destination. The previous discovery and negotiation process comes into play now and set the rules by which a specific packet gets a pre-specified treatment. Hence, no need for any kind of interdomain real-time signalling which is the main source of scalability problems. Figure 9 shows an example of packet remarking.

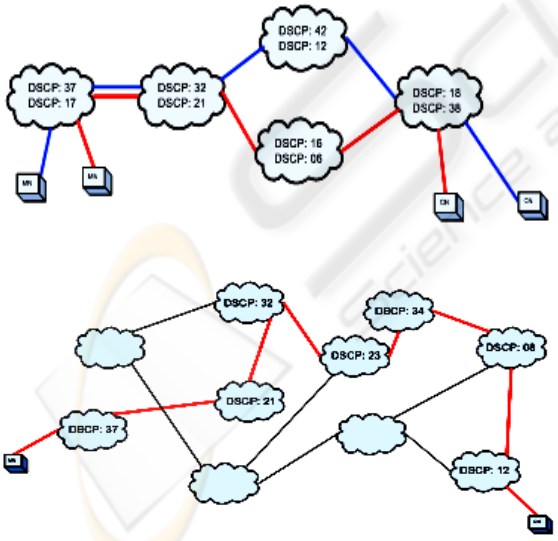

Figure 9: DCSP concatenation and re-marking. top: a single flow; bottom: two flows with different QoS requirements going through different transit domains.

Although not mentioned, DiffServ architectures are enhanced if resources are managed by a Bandwidth Broker (BB). A BB will be the main authoritative entity for flow admission.

Shown in **Error! Reference source not found.** is a possible signalling diagram for requesting QoS for a specific flow using NSIS. The user uses it User Equipment (UE) to send a request to the network. This message is intercepted by an Access Router (AR, in IMS, e.g., it would be a P-CSCF) that contacts a Bandwidth Broker (here called Core Network QoS Broker – CNQoSB) in case it detects the flow won't be terminated inside the domain. The CNQoSB, if authorizing it, will run an SLA routing algorithm.

This SLA-routing algorithm is responsible for determining whether it has any peering agreement with the terminating domain and – in case yes – the best path for the packets according to its SLA database. The best path for the flow is the one that will cross domains which the source domains has a tSLA. Out of the possibilities the best one will be chosen (criteria such as capacity utilization, price, resilience, etc.)

It will then deliver an authorization message to the requesting AR and will include a DSCP to mark the packets with. AR will then forward the NSIS request to the terminating domain (assuming that the pSLA with the terminating domain demands the use of NSIS). The destination user will then (upon authorization) proceed to reserve the needed resources and signal its serving domain that will provide the actions needed to reserve resources. An answer will then come back where resources will be committed.

It is important to notice that the diagram shows that no interdomain reservation is actually made. The reservation signalling is needed only to reserve resources on the source and destination access networks (per-flow reservations). In between segments use DiffServ and interdomain segments use pre-negotiated traffic pipes. It is up to the sourcing domain to control its available contracted resources even in remote domains that belong to the e2e path.

4 CONCLUSIONS AND FUTURE WORK

We have presented a framework for e2e QoS with special focus on interdomain QoS. As discussed, this is still an open problem since it involves a business component that engineering can't solve on its own.

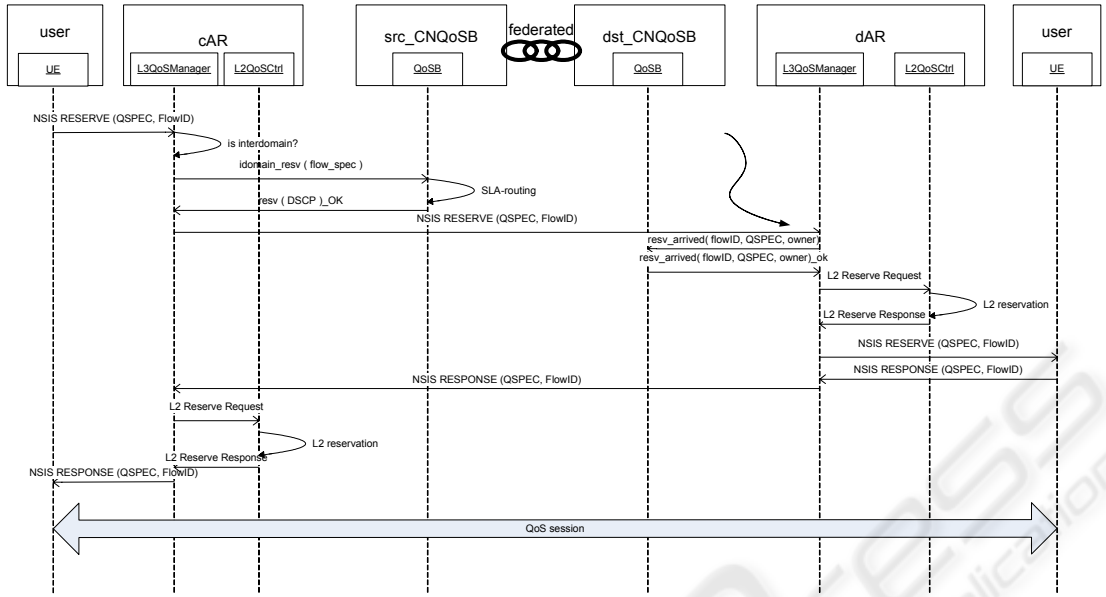

Figure 10: End-to-end QoS session setup signaling diagram.

We have attempted to decouple such business component from the technical component. Further, we showed an architecture that is scalable and fast while using common intradomain techniques and protocols.

There is, however, many open issues. The exact definition of QoS service was not discussed and this is tightly coupled with what is a domain profile. Besides the complex issue of negotiating a complex object such as a SLA, this will lead us to traffic engineering complexities such as guaranteeing that a transport service can be flexible enough to accommodate end user services and per-flow management. Another complex and open issue is billing and monitoring. On one hand, there must be means for a domain to check if an SLA is being fulfilled. On the other hand, a domain must be able to charge another domain for its traffic. This means that each packet will have to carry some kind of signature in order to prove its precedence. Overall, security issues will also be involved.

ACKNOWLEDGEMENTS

This work was partially funded by IST-FP6 Integrated Project Daidalos II. The authors wish to thank their partners of WP32 and WP33 for their collaborative work.

REFERENCES

Meireles, T.H.: "Interconnection Agreements Between Internet Service Providers", MSc dissertation, University of Aveiro, November 2004

http://www.fccn.pt

Braden, R., Clark, D. and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.

http://www.ietf.org/html.charters/mpls-charter.html

Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z. and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.

S. Sargento et al., “Mobility through Heterogeneous Networks in a 4G Environment”, Wireless World Research Forum Meeting 17, Nov 06.

P. Pan, E. Hahne, and H. Schulzrinne, “BGRP: Sink-Tree-Based Aggregation for Interdomain Reservations,” KICS 2000.

R. Bless, “Dynamic Aggregation of Reservations for Internet Services”, ICTSM 10, Oct. 2002

R. Sofia, R. Guerin, and P. Veiga. "SICAP, a Shared-segment Interdomain Control Aggregation Protocol. High Performance Switching and Routing", HPSR 2003, Turin, Italy, June 2003.

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

http://www.ietf.org/html.charters/nsis-charter.html

http://qbone.internet2.edu/

Cordeiro, L, et al, "Hybrid on-path off-path approach for end-to end signalling across NSIS domains (HyPath)", Internet-draft, February 2006.