Providing Consistent Service Levels in IP Networks

Solange Rito Lima, Pedro Sousa, and Paulo Carvalho

University of Minho, Department of Informatics, 4710-057 Braga, Portugal
{solange, pns, pmc}@di.uminho.pt

Abstract. The use of Internet as an ubiquitous communication platform puts a strong demand on service providers regarding the assurance of multiple service levels consistently. Designing flexible and simple service-oriented management strategies is crucial to support multiconstained applications conveniently and to obtain a deployable and sustainable service quality offer in multiservice IP networks. In this context, this paper proposes the use of a self-adaptive QoS and SLS management strategy sustained by a service-oriented traffic admission control scheme to ensure the negotiated quality levels. A proof-of-concept of the proposed strategy is provided, illustrating its ability to self-adapt and control efficiently distinct QoS requirements in multiservice IP networks.

Key words: Multiservice Networks, Admission Control, Service Quality

1 Introduction

Providing service integration in IP networks assuring, at same time, consistent levels of service quality tend to require the adoption of specific service models and traffic control mechanisms to handle traffic with multiple requirements. The challenge is increased when considering end-to-end Quality of Service (QoS) delivery, involving multiple heterogeneous domains with negotiated Service Level Specifications (SLSs) between them to be fulfilled. In fact, the end-to-end QoS panacea will not be based on a single network service model attending to the diversity of business goals and technologies available. This means that the network control tasks, when in place, should be flexible enough to accommodate heterogeneity and service integration efficiently. At same time, simplicity is a major design goal and a key aspect for their deployment in real networks.

This paper addresses the problematic of efficient and versatile QoS/SLS management, proposing a service-oriented and self-adaptive management architecture to improve QoS guarantees and enforce SLSs fulfillment in multiservice networks. In this architecture, we identify and structure the main issues and tasks subjacent to the definition, building and control of network services both intra and interdomain. Attending to the key role of Admission Control (AC) in preventing QoS instability and service degradation, we specify a service dependent AC criteria adjusted both to explicit and implicit AC scenarios, widening the diversity of services supported.

The contents of this paper is organized as follows: current related work on SLS definition and management is debated in Section 2; an overview of the proposed QoS and SLS management architecture and operation is given in Section 3; the specification of the service-oriented AC criteria, including the proposed AC rules, is included in Section 4; the proof-of-concept and performance results are provided in Section 5; finally, the main conclusions are summarized in Section 6.

2 SLS definition and management

An SLS defines the expected service level, QoS related parameters and traffic control issues. The definition of a standard set of SLS parameters and semantics, apart from being a key aspect for QoS provisioning, is crucial for ensuring end-to-end QoS delivery and for simplifying interdomain negotiations [1]. Several working groups are committed to SLS definition and management [2,3]. Although a large combination of quality, performance and reliability parameters is possible, service providers are expected to offer a limited number of services. To define SLSs for quantitative and/or qualitative standard services adapted to different application types is, therefore, a major objective [4, 5].

For SLS management and control, solutions based on Bandwidth Brokers (BBs) (RFC 2638) centralize service information required to perform control tasks such as AC, removing them from the network core. However, this involves a large amount of information to manage and a functional dependence on a single entity. To improve reliability and scalability in large networks, several approaches consider distributed service control tasks with variable control complexity depending on the QoS guarantees and predictability required. To provide guaranteed services, existent proposals tend to require significant network state information and, in many cases, changes in all network nodes [6]. To provide predictive services, control tasks based on network measurements performed node-by-node [7,8] and end-to-end [9,10] have deserved special attention. These solutions lead to reduced control information and overhead, but eventually to QoS degradation. To control elastic traffic, for efficient network utilization, implicit strategies i.e., without requiring explicit signaling between applications and the network, have also been defined [11].

In these proposals, detailed in [12], aspects such as the trade-off between service assurance level and network control complexity for a scalable and flexible support of distinct service types and corresponding SLSs, intra and interdomain, are not covered or balanced as a whole. This grounds the motivation for the present proposal.

3 Multiservice QoS and SLS management architecture

The proposed self-adaptive architecture for managing multiple service levels involves different tasks related to service definition, monitoring and control, interrelated as illustrated in Fig. 1.

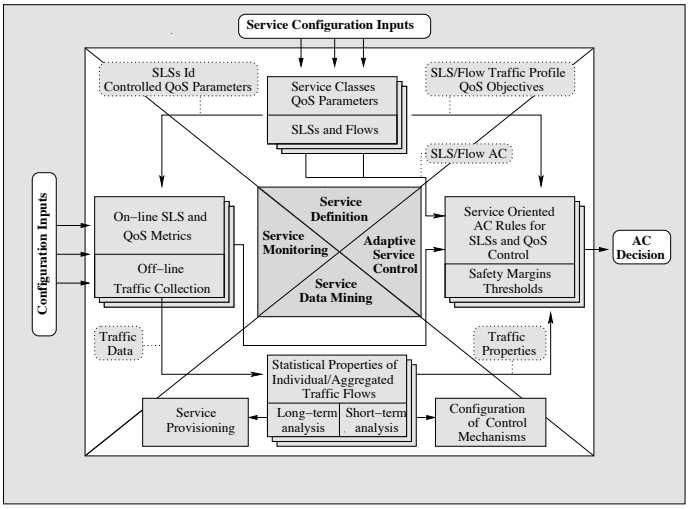

Fig. 1. Multiservice management architecture

Service definition involves the definition of basic service classes oriented to application with different requirements, the definition of relevant QoS parameters to control within each service type and the definition of SLSs' syntax and semantics. *Service monitoring*, performed on-line, keeps track of QoS and SLS status in the domain through a set of well-defined service metrics, providing feedback to drive *Self-adaptive service control* tasks such as AC. Traffic aggregates may also be collected for subsequent off-line analysis and characterization. *Service data mining* allows to determine the statistical properties of each class as a result of traffic aggregation so that more realistic service-oriented control mechanisms and service provisioning can be established. The knowledge resulting from interrelating these areas provides the basics for defining a multiservice management architecture and corresponding *AC decision criteria.*

In order to pursuit design goals such as flexibility, scalability and easy deployment, the service control strategy illustrated in Fig. 2 comprises: (i) distributed control between edge nodes; (ii) no control tasks within the network core; (iii) reduced state information and control overhead; (iv) measurement-based self-adaptive behavior regarding network dynamics. This model, oriented to accommodate multiple services, may perform AC irrespectively of the applications' ability to signal the network.

A primary idea of the AC strategy is to take advantage of the need for on-line QoS and SLS monitoring in today's networks and use the resulting monitoring information to perform distributed AC. This monitoring process, carried out on a per-class and edge-to-edge basis, allows a systematic view of each service class load, QoS levels and SLSs utilization in each domain, while simplifying SLSs' auditing tasks. An additional and crucial characteristic of the devised

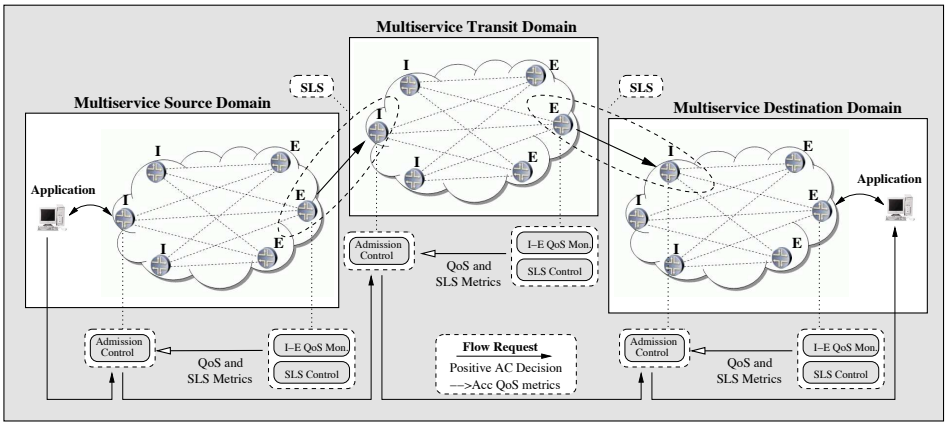

Fig. 2. Distributed monitoring-based AC approach

AC strategy is to consider a service-dependent degree of overprovisioning in order to achieve a simple and manageable multiservice AC solution. These levels of overprovisioning, embedded in the AC rules, allow to relax the AC process widening the range of service types covered by a monitoring-based AC solution.

As shown in Fig. 2, in the model's operation, while ingress nodes perform implicit or explicit AC resorting to service-dependent rules for QoS and SLS control (see Section 4), egress nodes collect service metrics providing them as inputs for AC. When spanning multiple domains, collecting and accumulating the QoS measures available at each domain edge nodes will allow to compute the expected end-to-end QoS. This cumulative process is consistent with the cascade approach for the support of interoperator IP-based services, which has the merit of being more realistic, i.e., in conformance with the Internet structure and operation, and more scalable than the source-based approach [1].

4 Specifying the multiservice AC criteria

For controlling both the QoS levels in the domain and the utilization of existing SLSs, the following rules have been defined: (i) rate-based SLS control rules; (ii) QoS parameters control rules; (iii) end-to-end QoS control Rules. The specification of these rules, following the notation in [13], is presented in Table 1. The conformance of the defined rules determine the acceptance of a new flow F_j. Note that Eq. (3) is not flow dependent, i.e. it is checked once during Δt_i to determine $AC\_Status_{\Delta t_i}$. An $AC\_Status_{\Delta t_i}$ =`accept` indicates that the measured QoS levels for SC_i are in conformance with the QoS objectives and, therefore, new flows can be accepted. An $AC\_Status_{\Delta t_i}$ =`reject` indicates that no more flows should be accepted until the class recovers and restores the QoS target values, which will only be checked at Δt_{i+1}. For a service class SC_i under *im-*

plicit AC, as flows are unable to describe r_j, traffic flows are accepted or rejected implicitly according to the value of $AC\_Status_{\Delta t_i}$.

Table 1. Control rules summary

Type of Rule	Description
SLS Rate Control Rules	Verify upstream and downstream SLSs utilization
$\tilde{R}_{i,(I_n,*)} + r_j \leq \beta_{i,I_n} R_{i,I_n}$ (1)	$\tilde{R}_{i,(I_n,*)}$ is the current measured rate of flows using SLS_{i,I_n} independently of the egress nodes E_m involved; r_j is the rate of the new flow F_j; $0 < \beta_{i,I_n} \leq 1$ is a service-dependent safety margin defined for the negotiated rate R_{i,I_n} of SLS_{i,I_n}.
$\tilde{R}^+_{i,(*,E_m)} + r_j \leq \beta^+_{i,E_m} R^+_{i,E_m}$ (2)	$\tilde{R}^+_{i,(*,E_m)}$ is the current measured rate of flows using SLS^+_{i,E_m}, considering all ingress-to-E_m estimated rates of flows going through E_m; r_j is the rate of the new flow F_j; $0 < \beta^+_{i,E_m} \leq 1$ is the service-dependent safety margin for the rate R^+_{i,E_m} defined in SLS^+_{i,E_m}.
QoS Control Rules	Verify the conformance of QoS levels in the domain
$\forall (P_{i,p}, \beta_{i,p}) \in P_{SC_i} : \tilde{P}_{i,p} \leq T_{i,p}$ (3)	$\tilde{P}_{i,p}$ is the ingress-to-egress measured QoS parameter; $\beta_{i,p}$ is the corresponding safety margin; $T_{i,p}$ is the parameter's upper bound or threshold, given by $T_{i,p} = \beta_{i,p} P_{i,p}$, used to set the acceptance status for Δt_i.
End-to-end Control Rules	Cumulative computation and verification of e2e QoS
$\forall P_{j,p} \in P_{F_j}$: $(\mathtt{op}_1\ (P^{acc^-}_{j,p}, P_{i,p}))\ \mathtt{op}_2\ (\gamma_{j,p} P_{j,p})$ (4)	$P_{j,p}$ is a flow's QoS parameter, allowing a tolerance factor $\gamma_{j,p}$; $P^{acc^-}_{j,p}$ is the parameter's cumulative value when crossing upstream domains; $P_{i,p}$ the corresponding target value in present domain.

In order to increase the scalability of the control strategy, it is the QoS parameter target value for the service class that bounds the corresponding SLS's expected QoS value and respective flows. Depending on the semantics of each parameter, $P_{i,p}$ can either be an upper or lower bound. Embedding the expected SLS parameters values in the respective class parameter target values is of paramount importance as QoS and SLS control in the domain is clearly simplified.

5 Proof-of-concept

To evaluate the performance of the service control strategy regarding its ability to manage multiple service commitments in a multiclass environment, a simulation prototype was set using NS-2. This prototype implements three functional interrelated modules - Automatic Source Generation Module, AC Decision Module, and QoS and SLS Monitoring Module. Fig. 3 presents a simplified diagram of the simulation model architecture, including the relation between these modules and the main underlying functions and variables. The two recursive modules represented in gray are responsible for the dynamic behavior of traffic source generation and monitoring.

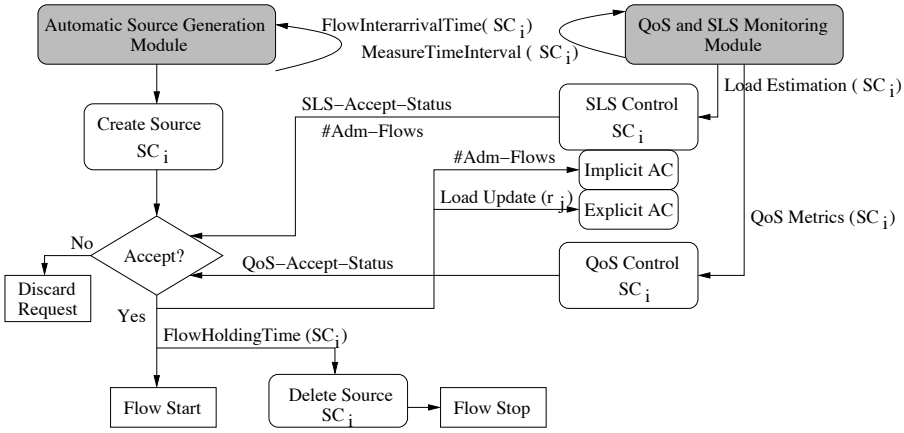

Fig. 3. Simulation model diagram

Taking into consideration current service configuration guidelines [14], three initial service classes have been defined. Table 2 summarizes the service classes implemented, highlighting AC and QoS monitoring parameters used to configure the AC rules specified in Table 1. Three downstream SLSs have been considered, one per service class, with a negotiated rate (R^+_{i,E_m}) defined according to the traffic load share intended for the corresponding class in the domain. As shown, the parameterization of the AC rules is service-dependent and larger safety margins β^+_{i,E_m} and tighter thresholds $T_{i,p}$ are defined for more demanding classes. For instance, a $\beta^+_{i,E_m} = 0.85$ corresponds to impose a safety margin or degree of overprovisioning of 15% to absorb load fluctuations and optimistic measures.

Table 2. Configuration of service classes SC_i

SC_i	Serv. Type	AC Type	R^+_{i,E_m}	β^+_{i,E_m}	$P_{i,p}$	$T_{i,p}$	Example	Traffic Src
SC1	guaranteed (hard-RT)	explicit and conservative	3.4Mbps (10% share)	0.85	IPTD ipdv IPLR	35ms 1ms 10^{-4}	VoIP Cir.Emulation Conv. UMTS	Exp.or Pareto on/off (64kbps,pkt=120B on/off=0.96/1.69ms)
SC2	predictive (soft-RT)	explicit and flexible	17Mbps (50% share)	0.90	IPTD IPLR	50ms 10^{-3}	audio/video streaming	(256kbps,pkt=512B on/off=500/500ms
SC3	best-effort	implicit	13.6Mbps	1.0	IPLR	10^{-1}	elastic apps.	FTP (pkt=512B)

The network domain consists of ingress routers I_1, I_2, a multiclass network core and an edge router E_1. I_2 is used to inject *concurrent* or *cross* traffic (referred as CT-I2), allowing to evaluate concurrency effects on distributed AC and to assess the impact of cross traffic on the model performance. The test scenarios with cross traffic allow to evaluate the presence of unmeasured traffic within the network core. This type of traffic, likely to occur in real environments, impacts

on domain's QoS and load without being explicitly measured by E_1 SLS rate control rule (Eq. (2)). The domain routers implement the three service classes according to a hybrid Priority Queuing - Weighted Round Robin (2,1) scheduling discipline, with RIO-C as AQM mechanism. The domain internodal links capacity is 34Mbps, with a 15ms propagation delay. Δt_i is set to 5s.

5.1 Simulation results

This section intends to assess the self-adaptive behavior and effectiveness of the proposed service control strategy in keeping QoS levels and SLSs share consistently and efficiently controlled. For this purpose, Section 5.1-A presents results illustrating the solution's ability to ensure domain QoS levels in presence of concurrent and cross traffic, highlighting also its capacity to self-adapt to new QoS thresholds; Section 5.1-B illustrates the significance of the results facing the high utilization levels achieved.

A - Ensuring domain QoS levels Fig. 4 illustrates the obtained IPTD and IPLR for service classes SC1, SC2 and SC3. As shown, the classes exhibit a stable behavior regarding the pre-defined QoS levels: (i) SC1 is very well controlled, with IPTD kept almost constant throughout the simulation period. The mean ipdv assumes a low value (0.1ms) as a result of small variations, bounded by a well-defined maximum and minimum values ($\pm$0.4ms). With concurrent traffic no loss occurs; (ii) SC2 and SC3 IPLR evolution tends to the defined IPLR thresholds of 10^{-3} and 10^{-1}, respectively. For SC2, IPLR has a less regular behavior as it results from occasional loss events, converging to the defined threshold; (iii) the percentage of packets exceeding delay QoS thresholds is very small: for SC1 only 0.007% of packets exceed the IPTD threshold (35ms) and 0.0005% the ipdv threshold (1ms). For SC2, 2.95% of packets exceed the delay threshold (50ms). Note that, exceeding a QoS threshold does not necessarily imply a service QoS violation, as the defined concept of threshold comprises a safety margin to the QoS parameter target value (see Eq. (3)). These results are particularly encouraging attending to the high network utilization obtained (see Section 5.1-B).

Impact of cross traffic and adaptation to new thresholds This new test scenario intents to illustrate the model's ability to self-adapt to distinct QoS thresholds, in particular, to control new delay and loss bounds. In addition, the traffic conditions are now more demanding as the traffic submitted to ingress I2 is cross traffic. Fig. 5 presents a multimetric analysis of the IPTD and IPLR obtained for each service class in each Δt_i.

As shown, when a tighter IPTD threshold of 35ms is set for SC2, AC is effective in bringing and maintaining IPTD controlled around that value. Simultaneously, considering a new IPLR threshold of 0.05 for SC2 and SC3 (more relaxed and tighter than the previous one of 10^{-3} for SC2 and 10^{-1} for SC3), it is notorious that the control strategy has been able to bring IPLR to the new

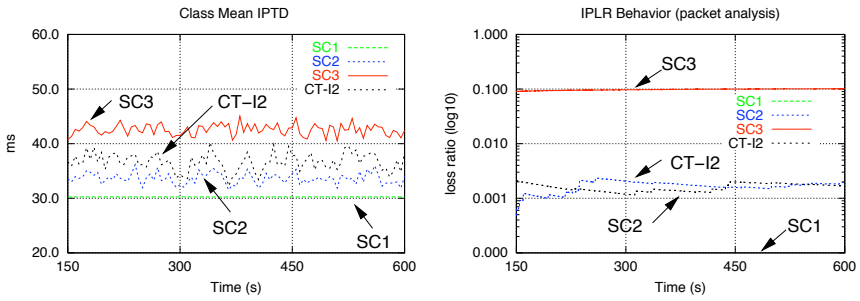

Fig. 4. IPTD and IPLR evolution

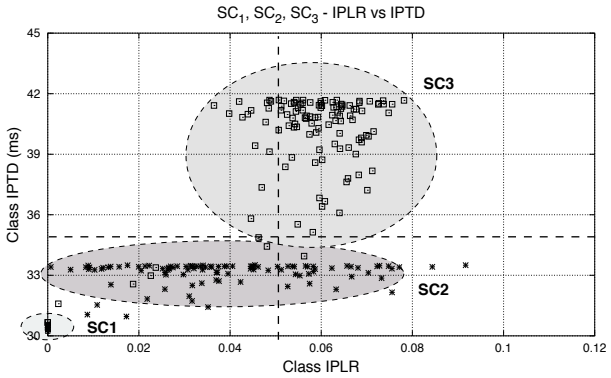

Fig. 5. Adaptation to new QoS thresholds in presence of cross traffic

value. It is also evident that IPLR is more difficult to keep tightly controlled, however, a consistent behavior around 0.05 is achieved. In the presence of cross traffic, the main rule determining AC decisions is the QoS control rule Eq. (3), with $AC\_status_{\Delta t_i} = reject$ activated mostly by IPLR threshold violations. This rule by itself maintains the QoS levels controlled, as shown in Fig. 5.

From these set of experiments, the relevance of the defined AC rules becomes evident for assuring service commitments in the domain. While the rate control rule (Eq. (2)) assumes a preponderant role for service classes SC1 and SC2 to control the traffic load and indirectly QoS, particularly in situations involving concurrent traffic, the QoS control rule (Eq. (3)) is decisive to assure the domain QoS levels in presence of unmeasured cross traffic. In real environments, where the two type of situations are likely to occur simultaneously, the two AC rules will complement each other to increase the domain capabilities to guarantee service commitments.

B - Controlling SLSs share Fig. 6 (a) illustrates the obtained share of each class under the concurrent traffic test scenario. Note that, for the safety margins (β^{+}_{i,E_m}) and the SLSs rate share (R^{+}_{i,E_m}) defined in Table 2, the utilization target for SC1, SC2 and SC3 is 8.5%, 45% (SC2+CT-I2) and 40%, respectively. As shown, the share configured for each class and SLS is well accomplished and the global utilization is kept very high. SC2 and CT-I2 obtain a similar behavior and share and SC3 exceeds its share slightly. This occurs due to the adaptive nature of traffic within SC3, the more relaxed implicit AC criterion and the work conserving nature of the scheduling algorithm in use, which allows SC3 to take advantage of unused resources. The per-class and global utilization with cross traffic decreases slightly comparing to the concurrent case. This decrease is a consequence of the effect of cross traffic on queues occupancy increasing loss events and triggering the QoS control rule more frequently. However, the rate share of each class is also well accomplished and the global utilization very high.

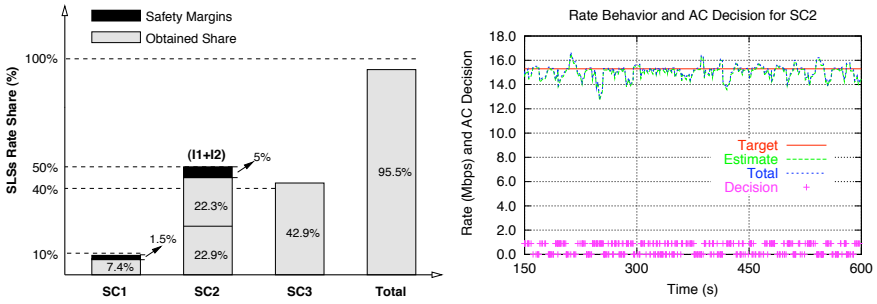

Fig. 6. (a) Rate share for each service class; (b) Rate and AC behavior for SC2

When observing the behavior of the AC rules for SC2 and the resulting AC decision along time (Fig. 6 (b)[1]), it can be seen that, although the rules are effective in blocking new flows when QoS degradation or an excessive rate is sensed, the effect of previously accepted flows may lead to some episodes of rate estimates above the target line. In fact, traffic fluctuations reflecting a low estimation in Δt_{i-1} may lead to over acceptance during Δt_i. Despite the fast time reaction of the control rules, these situations evince the advantage of using protective safety margins.

[1] In Fig. 6 (b), *Target* line represents the value $\beta^{+}_{i,E_m} R^{+}_{i,E_m}$ above which AC rejection occurs, *Estimate* line represents the estimated rate or load of SLS^{+}_{i,E_m}, i.e., $\tilde{R}^{+}_{i,(*,E_m)}$, and *Total* line reports to the previous estimate by adding the new flow rate r_j. *Decision* dots represent a positive (dots above the x-axis) or negative (dots overlapping the x-axis) AC decision, considering also the QoS control rule evaluation.

6 Conclusions

In this paper, a self-adaptive service management strategy has been proposed to improve QoS guarantees and enforce SLSs fulfillment in multiservice networks. The solution relies on service-dependent AC rules which allow a versatile and consistent control of QoS levels and SLS usage both intra and interdomain. The proof-of-concept has demonstrated that the self-adaptive behavior inherent to on-line measurements combined with the proposed AC rules is effective in controlling the quality levels of each service class. Under demanding traffic conditions, the relevance of the two defined AC rules became evident complementing each other to increase the domain capabilities to guarantee service commitments.

References

1. Georgatsos, P. et al.: Provider-level Service Agreements for Inter-domain QoS delivery. Fourth International Workshop on Advanced Internet Charging and QoS Technologies (ICQT04) (September 2004)
2. Morand, P. et al.: Mescal D1.2 - Initial Specification of Protocols and Algorithms for Inter-domain SLS Management and Traffic Engineering for QoS-based IP Service Delivery and their Test Requirements. Mescal Project IST-2001-37961 (January 2004)
3. Diaconescu, A., Antonio, S., Esposito, M., Romano, S., Potts, M.: Cadenus D2.3 - Resource Management in SLA Networks. Cadenus Project IST-1999-11017 (May 2003)
4. Seitz, N.: ITU-T QoS Standards for IP-Based Networks. IEEE Communications Magazine **41**(5) (June 2003)
5. D. Miras et al.: A Survey of Network QoS Needs of Advanced Internet Applications. Internet2 Working Document (November 2002)
6. Stoica, I., Zhang, H.: Providing Guaranteed Services Without Per Flow Management. In: ACM SIGCOMM'99. (October 1999)
7. Jamin, S., Danzig, P., Shenker, S., Zhang, L.: A Measurement-Based Call Admission Control Algorithm for Integrated Services Packet Networks (Extended Version). IEEE/ACM Transactions on Networking (February 1997) 56–70
8. Breslau, L., Jamin, S.: Comments on the Performance of Measurement-Based Admission Control Algorithms. In: IEEE INFOCOM'00. (March 2000)
9. Cetinkaya, C., Kanodia, V., Knightly, E.: Scalable Services via Egress Admission Control. IEEE Transactions on Multimedia **3**(1) (March 2001) 69–81
10. Elek, V., Karlsson, G., Rnngren, R.: Admission Control Based on End-to-End Measurements. In: IEEE INFOCOM'00. (2000)
11. Mortier, R., Pratt, I., Clark, C., Crosby, S.: Implicit Admission Control. IEEE Journal on Selected Areas in Communication **18**(12) (December 2000) 2629–2639
12. Lima, S.R., Carvalho, P., Freitas, V.: Admission Control in Multiservice IP Networks: Architectural Issues and Trends. IEEE Computer Communications Magazine **45**(4) (April 2007) 114–121
13. Lima, S.R., Carvalho, P., Freitas, V.: Self-adaptive Distributed Management of QoS and SLSs in Multiservice Networks. In: IEEE/IFIP International Conference on Integrated Management (IM 2005), Nice, France, IEEE Press (May 2005)
14. Babiarz, J., Chan, K., Baker, F.: Configuration Guidelines for Diffserv Service Classes. RFC 4594 (August 2006) Informational.