

# Multi-Class Measurement Based Admission Control for a QoS Framework with Dynamic Resource Management

Thomas Bohnert<sup>1</sup> and Edmundo Monteiro<sup>1</sup>

*Published online: 28 April 2007*

---

The advent of a Premium Service (PS) featuring Internet is still pending, mainly due its difficult deployment. But do we really need Quality of Service like PS? Likely no. We are therefore investigating an alternative QoS model, which is more flexible to deploy at the expense of looser QoS guarantees. We call it a Better-than-Best-Effort (BBE) service. What distinguishes our approach from traditional ones is dynamic resource management based on measurements and ruled by Perceptual QoS. What yet has been missing was Admission Control (AC) and in this paper we present our latest advance. Based on the existence of measurements taken by the queuing module, we developed a Measurement Based AC algorithm. Design goal was simplicity and general applicability in terms of independence from statistical assumptions. Likely the most interesting finding is that even a very simple design proved to be reasonably effective for our BBE service, mainly due to a cross-layer design, i.e. cooperation between dynamic resource management on queuing level and AC. This is the conclusion a comprehensive simulative performance evaluating.

---

**KEY WORDS:** measurement based admission control; alternative QoS model; perceptual QoS; better-than-best-effort QoS.

## 1. INTRODUCTION

Only a few subjects ever in the history of computer communications have brought in such controversy like the Quality of Service (QoS) issue in the Internet. What actually divides the No-QoS lobby from its counterpart is swiftly written down; persistent technological advance in switching capacity over demand by real-time services. Indeed, recent advances in Ethernet transmission rate surmounted the 100 Gbps barrier and motivated the IEEE standardization body to form the IEEE 802.3 Higher Speed Study Group [1].

---

<sup>1</sup>Laboratory of Communications and Telematics, University of Coimbra, CISUC/DEI, Polo II, Pinhal de Marrocos, 3030-290 Coimbra, Portugal. E-mail: tbohnert@dei.uc.pt, edmundo@dei.uc.pt.

Belief in solving any network inherent quality impairment like packet delay, jitter and loss by pure over provisioning of resources has reached strengthened so much, that G. Bachula, the Vice President of the Internet2 project, has recently stated under what is termed *Net Neutrality*, that there must be no QoS mechanism deployed to secure the evolution of the Internet as barrier-free communication platform [2]. This considerable statement represents the conclusion of a 3-year QoS deployment study carried out by the Internet2 QoS Working Group [3]. In summary, difficulties in providing Differentiated Services (DiffServ) [4] Premium Service (PS), i.e. implementing Expedited Forwarding (EF) PHB (Per-Hop Behavior) [5], outweigh any merit, a similar conclusion as published in [6].

Among other arguments, G. Bachula also supports this conviction by global broadband penetration policy. For example, Europe's declared target is at least 95% broadband territorial coverage by 2008 [7]. It has, however, to be mentioned that alongside pure broadband provisioning, Europe's policy also embraces the advancement of a QoS featuring Internet implemented by founding ambitious research projects like GEANT2, EuQoS, NetQoS or OpenNet only to list a few. One reason is that over provisioning is deemed to be as myopic as Bill Gates' "640 k ought to be enough for anybody" vision in 1981. Just as cellular telephones increase the total number of phone calls including the PSTN,<sup>2</sup> faster Internet access will sooner or later entail new, ever more data-rate greedy services until demand catches-up with technology, and if only when the latter is pushed to its physical limits.

There are ample arguments in support of a QoS featuring Internet but notwithstanding, one cannot ignore the difficulties encountered by an ambitious project on the scale of Internet2. Henceforth, the question is "What are the problems of deploying PS?"

One of the major problems identified by the Internet2 QoS Working Group was insufficient QoS support in routers [3]. This, however, can meanwhile be considered as inappropriate, e.g. see [8, 9], and further indicates the need for QoS support backed up by manufacturer's awareness of future market potentials.

A second matter has its roots in the nature of PS deployment. In short, PS poses an all-or-nothing solution meaning that either all routers are configured respectively or no guarantees can be made at all. In today's Internet configuration, this means that at least access networks have to support PS while the core remains virtually congestion free [10]. The implementation therefore does not only entail technical adaptations but far reaching modifications in business models, which are especially difficult to be implemented on Internet scale. It is obvious that particularly this issue poses a major obstacle and conflicts with the Internet's architecture made of Autonomous Systems (AS).

<sup>2</sup>Public Switched Telephone Network.

On the other hand, many of today's Real-Time (RT) applications adapt to the Best-Effort (BE) Internet. Voice over IP applications, for example, recover some percentage of packet loss using Packet Loss Concealment (PLC) and Forward Error Correction (FEC) and therefore require only limited support from the network [11]. How much support, however, is difficult to quantify and can at best be captured by an application's utility functions [12].

On the basis of this finding, we are investigating a new service model called LCT-QoS. Our objective is a Better-than-BE (BBE) service, which provides looser guarantees but is more flexible to deploy. The key feature is dynamic resource management, so to say on demand, based on the supported application's utility function. How this has been implemented in the context of the DiffServ architecture is described in Section 2.

As in almost any QoS framework, quality assurance means protecting resources from over subscription. Hence, some form of Admission Control (AC) has to be performed and in this paper we present our latest advance towards a Measurement-based Admission Control (MBAC) algorithm for LCT-QoS. Its motivation and formal description is detailed in Section 3 while in Section 4 its simulative performance evaluation is documented. In the same section, we introduce several modifications devised and applied to improve the algorithm. Finally, in Section 5 we conclude on our research and present an outlook on future topics.

## 2. THE LCT-QoS FRAMEWORK

As mentioned before, at our Laboratory of Communications and Telematics (LCT) at the University of Coimbra (UC), we pursue a project whose goal is the investigation of an alternative QoS model. The central idea is to extend the classical BE service to an BBE service by dividing traffic into several classes instead of a single class and to dynamically provide resources based on demand and supported application's utility function. This corresponds to a shift from a *single-class best-effort* paradigm to a *multiple-class better-than-best-effort* paradigm with service specific QoS support.

In order to do so, the LCT-QoS framework is composed of a set of layers where each layer corresponds to a logical module, shown in Fig. 1, Fundamental to all modules is the LCT-UC QoS Metric, which is further detailed in Section 2.1. Placed on top of this module is a queuing module which implements the BBE service referred to before. Classification and forwarding of traffic is in accordance with the DiffServ architecture, see [4], and is therefore called Per-Hop-Behavior D3, which stands for Dynamic Distribution of Degradation (PHB-D3), see Section 2.2. The next higher module is the LCT-UC QoS SR module. This module's intelligence is a set of traffic engineering algorithms to route traffic dynamically within an LCT-QoS domain. It has no direct relation with this paper's topic but the interested




Fig. 1. The modules of the LCT-QoS framework and their location in an LCT-QoS domain.


reader is directed to [13–16]. Finally, on top of the framework is the AC module, which is the principal subject of this article.

## 2.1. Application Utility Based on Perceptual QoS

Traditional IETF QoS assessment focuses on physical parameters like packet delay, loss or jitter, so-called Intrinsic QoS (IQ) [17]. An alternative approach for rating service quality is commonly termed Subjective or Perceived/Perceptual QoS (PQ) [17, 18] and is solely based on human perception or satisfaction regarding service usability. Determining PQ is typically carried out by surveying a set of persons, which participate in a controlled experiment [18, 19]. In contrast to this method, if there are no humans involved and PQ is computed from physical parameters, it is called Objective QoS assessment [18].

In contrast to the ITU, the IETF has not yet adopted the concept of PQ. This is rooted in the nature of the services each standard entity is concerned with. The PSTN's primary service is voice and its quality assessment is a highly subjective matter. Hence, the ITU, as the PSTN's standard body, has been concerned with PQ for more than a century. With the evolution of the Internet as the universal communication platform, however, more and more voice traffic is delivered over Internet infrastructure, henceforth calling for the same quality assessment methodology for voice over IP (Internet Protocol) services. Moreover, it is worth noticing that this reaches beyond Voice over IP (VoIP) services and applies to any audio and/or visual service like for instance IP Television (IPTV) or Videoconferencing.

One of our design goals is therefore to step ahead of the IETF QoS notion and to incorporate PQ. The central role in achieving this goal is held by the LCT-UC QoS Metric. It quantifies quality impairment inflicted by the unreliable communication system and maps it to PQ [20]. In order to do so, it defines degradation and superfluity zones, which in turn define how the impact on PQ varies with physical QoS parameter variations. According to this metric an application's utilization (or PQ) is quantified through a variable named congestion index (CI) where a high CI value means low service quality and vice versa.



**Fig. 2.** Calculation of the delay CI for three traffic classes. With an increasing absolute value of delay the associated CI increases also with a rate depending on the DSlope. Thus, classes sensitivity to delay can be expressed through their DSlopes. according to the metric, Perceived QoS is expressed by the respective CI.

A single IP service is mapped into a service class and for each class there is a CI related to transit delay and a CI related to loss. Solely for the sake of simplicity a linear mapping has been chosen but its generalization to any functional is straightforward as, for instance, demonstrated in [21].

Applying a linear function yields the so-called Degradation Slope (DSlope). It is used by the metric for the definition of classes' sensitivity to delay and loss degradation. A traffic class highly sensitive to PQ degradation for a given QoS parameter will have a high DSlope associated with that characteristic and the opposite applies to low sensitivities. An example for three distinct classes is depicted in Fig. 2. Each class has a different sensitivity to delay degradation (it would be the same if we were talking of loss). Classes with lower DSlope (measured in degrees) are less sensitive so their CI will grow slowly and vice versa. In conclusion, the concept of the LCT-QoS metric can be interpreted as cross-layer QoS assessment in order to express individual application's utility.

The nature of the LCT-QoS metric further allows quantifying fair PQ. In fact, one can say that PQ for all classes is equal if different classes' CIs related to delay and loss are also equal. In other words, the impact on service utility (or quality) inflicted by the current system load is barely the same for all of them, despite the difference in absolute values for delay and loss.

## 2.2. Dynamic Distribution of Degradation (PHB-D3)

Today's Internet traffic composition is highly unevenly distributed. The lion's share, up to 90%, constitutes BE traffic and only the remaining 10% is RT traffic [22]. As a consequence, long-term provisioned resources are in accordance with this distribution. This can lead to the unfavorable situation where in times of high demand resources for one class are short, resulting in flow rejections by AC, while there are still ample resources for another one.

Another issue associated with static long-term resource provisioning is appropriate policing. In a general QoS domain, traffic from multiple ingress points can happen to exit at one egress point and for this reason, the total input of one class must not exceed the maximum output at the weakest egress interface. The Internet2 project has identified this peculiarity as a crucial issue regarding efficient PS deployment [6].

Our LCT-QoS framework overcomes these problems by supporting dynamic resource distribution on the basis of equal CIs. If PQ for one class deteriorates, resources are redistributed such that the impact of the overall increase in traffic load is absorbed by all classes. Although in this way resources are shared by all classes, the system still provides BBE service differentiation due to the different DSlopes.

To integrate this concept a monitor continuously samples average packet delay and loss for class  $c$ ,  $c = \{0, \dots, C\}$  and averages the corresponding congestion indexes, namely  $CI_c^d$  for delay and  $CI_c^l$  for loss. If scheduler weights and buffer sizes are in accordance with classes traffic intensity, corresponding CIs are practically equal for all classes, i.e.  $CI_1^d = CI_2^d, \dots, = CI_C^d$  and  $CI_1^l = CI_2^l, \dots, = CI_C^l$ , meaning equal application utility.

If the traffic intensity of class  $c$  rises, its buffer content will built up which is eventually reflected by a higher  $CI_c^d$  and  $CI_c^l$ . This in turn triggers the redistribution of resources, namely a readjustment of the weights and buffer memory. Although some simple algebra yields the theoretical values for weights and buffer sizes to equalize the CIs, we have to account for the random nature of sampled quantities. Therefore, the adjustment is done by using a Digital Phase Lock Loop to ensure stability and convergence, for details see [23].

Finally, we would like to point out that LCT-QoS shares principles with the Alternative Best-Effort (ABE) service devised by the Internet2 QoS WG; though some time later [24]. Just like LCT-QoS, ABE is essentially a BBE service and is therefore equally flexible and incrementally deployable as LCT-QoS. The major difference, however, is the algorithm for resource distribution. While ABE does also focus on application utility, it obeys a distributed architecture and it is left to applications to choose between a low-delay service on the expense of an increased risk of packet loss. The network is not actively involved but does simply provide a pre-configured service. As we believe that there are profound

conceptual arguments in favour of both models, we refrain from any further comparison.

### 3. MEASUREMENT BASED ADMISSION CONTROL FOR LCT-QoS

A general Admission Control (AC) definition reads as:

A set of actions carried out to decide whether to admit a flow onto a link with a pool of shared resources. The admission is positive if resource demand of the requesting flow superimposed with that of the admitted aggregate can be accommodated *without* violating QoS commitments.

Given this definition, we devised a tailored AC algorithm for LCT-QoS. Since LCT-QoS features measurements to capture resource demand, a natural choice is to leverage these measurements as indicators for resource availability. This means to develop an algorithm of the Measurement Based Admission Control (MBAC) family.

In common with LCT-QoS the power of MBAC lies in its adaptivity, i.e. its ability to operate under changing conditions. That's why MBAC lends itself for such an instationary environment like the Internet [25–27]. Another striking feature is that MBAC does implicitly account for Statistical Multiplexing (SMUX) gain as measurements are taken from the traffic aggregate. This is in contrast to Parameter based AC (PAC)<sup>3</sup> and improves resource utilization considerably.

But MBAC still suffers from a set of open issues. While research in this area was generally called in question in [28], the same authors list in a follow-up [29] a set of yet to be solved problems. What has been found in rigorous, simulative evaluation is that MBAC performance varies considerably for different deployment scenarios and that underlying probability models are frequently too complex, requiring fine-tuning based on operator's experience to achieve reasonable performance. These findings have been further confirmed in [30], where some of the same algorithms have been evaluated in a real system.

One reason for these problems is that general MBAC probability modeling obeys the parametric approach. Parametric means that statistical assumptions are made in advance. In the context of MBAC this applies in particular to the nature of the arrival process' marginal distribution and recently it has become very popular to assume a Gaussian nature, see for instance [27, 31–33]. The motivation is twofold. First, the Central Limit Theorem (CLT) goes well with high-speed access networks and second, Gaussian models are well understood and comply with Long-Range-Dependence (LRD) and Self-Similarity phenomena found in network traffic [34]. Whether, however, this assumption matches reality is rather unclear [35] and the

<sup>3</sup>We intentionally refrain here from using PBAC since this abbreviation is commonly used for Probe Based AC.

only certain conclusion by today is that any assumption shrinks applicability and makes MBAC algorithms sensitive in deployment.

Henceforth, an optimal MBAC should be kept simple and if possible non-parametric to preserve its motivation, general applicability. In a follow-up of this work we have investigated especially this direction, a non-parametric approach for MBAC and our advances have been published in [36].

Inefficiency incurred by parametric modeling can even better be illustrated using Parametric AC (PAC) as example. The principle of this AC family is to model application's resource demand beforehand, for example assuming a simple two-state Markov ON/OFF model for VoIP sources. Given such a well-understood model, the resource requirement for a single source can be computed offline and an aggregate's demand simply the sum of a number of individual sources. The simplest PAC algorithm is the Simple- or Rate Sum (RS) algorithm. Its only parameter is source peak rate and according to this model an aggregate's bandwidth requirement is simply the sum of sources individual peak rates.

The advantage is obviously its simplicity. Moreover, as it is a pure worst-case model, it can be used to provide deterministic QoS. What, however, renders RS as a favorable choice comes at the expense of poor resource utilization since practically no application obeys a so-called Constant Bit Rate (CBR) source, assumed by this PAC. This algorithm therefore poses an intuitive example how assumptions made by parametric modeling can mismatch with reality. Consequently, as there is no account for any SMUX gain, a huge amount of resources are reserved but left unused with an increasing number of sources [37].

Given the advantages and disadvantages of either AC category, it seems natural to make use of the best of each and combine them. In order to do so we developed a two-stage algorithm. In its first stage, called Charging Stage (CS), the algorithm obeys the simple RS concept and the admission decision is solely based on peak rates, not only for its simplicity, but also since this is the only accurate source descriptor at all. This information could be either signaled or pre-configured as in our case. Due to the nature of RS the network is *charged* independently from traffic nature and furthermore, the decision is computed with negligible computational overhead.

So far, we have not accounted for the PHB-D3 and its dynamic resource distribution. If some of a class' resources are temporarily shifted to an other one there is no reason for rejecting flows in charging stage if delay and loss values are still below some thresholds due to a high multiplexing gain. For this reason, we introduce a second stage, which we named Saturation Stage (SS). In SS we make use of delay and loss measurements as indicators for resource availability. Whenever a flow would be rejected in CS, the algorithm switches to SS and checks if delay and loss estimates are below predefined boundaries. If so, the admission decision is positive irrespectively the negative decision in CS.



Recalling the dynamic resource distribution by the PHB-D3 queuing module, this AC hybrid's rationale is to leverage state indicators of the queuing module. This adaptivity does not only allows us to deal with dynamic resource assignment but also to account for SMUX gain. Our algorithm's rationale is therefore in contrast to common AC design where congestion control on AC level is clearly separated from queuing level. Furthermore, we have to point out that we accept packet loss albeit in controlled quantities. This is justifiable given the aim of our BBE service.

Finally, the algorithm's logic can formally be expressed as

$$\chi_k = \begin{cases} \geq 0 & \text{admitt flow } k \\ < 0 & \text{reject flow } k \end{cases} \quad (1)$$

where  $\chi_k$  denotes the cumulative multi-class admission criterion for the requesting flow  $k$ . Its definition reads

$$\chi_k = \min\{\Omega_0, \dots, \Omega_C\} \quad (2)$$

where  $\Omega_c$  is defined as

$$\Omega_c = \begin{cases} \zeta_{n,c} - \left( \left( \sum_{f=0}^F \sigma_{f,c} \right) + \sigma_{k,c} \right) & \text{if } s = n = 0, c = \{0, 1, \dots, C\} \\ \zeta_{n,c} - \hat{\varphi}_{n,c} & \text{if } s = 1, n = \{1, 2\}, c = \{0, 1, \dots, C\} \end{cases} \quad (3)$$

Estimated QoS parameters in (3) are denoted by  $\hat{\varphi}_{n,c}$  where  $n=0$  for bandwidth,  $n=1$  for delay and for loss  $n=2$ . Subscript  $c$  is the class index. Predefined upper bounds for QoS parameters, i.e. admission thresholds, are denoted by  $\zeta$ . Letter  $F$  in (3) denotes the present number of active flows on the link and  $\sigma$  denotes a flow's peak rate. Finally, stages are denoted by  $s=0$  for CS and  $s=1$  for SS.

In line with general Resource Management (RM) frameworks the AC module has been subdivided into two parts, an Admission Policy Sub module (APS) and a Policy Enforcement Sub module (PES) [38]. So far, no signaling has been defined for LCT-QoS and henceforth we apply Implicit AC [39]. This means that arriving flows have to be detected and states have to be maintained for active and rejected ones in order to enforce decisions [37]. These functions are implemented in the PES. Each arriving packet is intercepted by this module and after its classification, the module queries its local data base to see if the packet belongs to a flow with an valid admission token. If this is the case, this flow's activity timer is reset and the packet is forwarded. If not, the packet belongs to a flow which has been previously rejected and a Internet Control Message Protocol (ICMP) message is sent to the source to prevent it from further packet emissions. But if no recorded at all can be found in the data base, it indicates the arrival of an unknown flow and the PES ransoms a APS trigger. The APS in turn, requests state information, i.e. delay and loss estimates, and decides if this flow is admissible. If the decision is positive the




Fig. 3. The LCT-QoS network simulation model and its parameters.

flow id (FID) is associated with a valid admission token plus the before mentioned activity timer and added to the database. Finally, the packet is forwarded onto the link.

#### 4. SIMULATIVE PERFORMANCE ANALYSIS

To carry out a comprehensive performance analysis, we chose the NS-2<sup>4</sup> as simulation framework and the complete LCT-QoS implementation has been based on NS-2's DiffServ module. All results of this study are presented in the sequel of this section.

##### 4.1. Experimental Setup

The network model we chose, depicted in Fig. 3, is a simple LCT-QoS domain and basically consists of two access parts and a network core. Hence, a simplified end-to-end view of today's Internet architecture. Common for such a topology is that access links, connecting edge and core routers, are the bottlenecks and AC is therefore deployed at edge routers.

There are a number of source (client) nodes in the access part and traffic sources (applications) are evenly distributed among them. These traffic sources simulate the applications supported by the LCT-QoS service model. Details are presented in Table I which shows class, delay and loss sensitivity, application and the source model used to simulate application nature. For ST we used either ON/OFF sources with Exponential (EXPOO) or Pareto (POO) distributed ON/OFF times or CBR sources with Random noise (CBRR) influencing packet emission. For ET solely File Transfer Protocol (FTP) traffic over TCP has been used.

The simulated time for each experiment is 3600s which is more than enough for POO aggregates to develop SS and LRD characteristics [40]. To undertake the steady-state analysis, data collected during a 300s warm-up ( $t_w$ ) time has been discarded and source (flow) statistics are only computed for flows which enter

<sup>4</sup>NS-2 Network Simulator <http://www.isi.edu/nsnam/ns/>.

**Table I.** LCT-QoS Service Model and Corresponding Simulation Models

Type:Class	Delay	Loss	Application	FCP	FCE	FEP
ET:0	Low	Low	File Transfer	FTP/TCP	FTP/TCP	FTP/TCP
ST:1	Medium	High	Video Stream	CBRR/UDP	CBRR/UDP	EXPOO/UDP
ST:2	High	Medium	VoIP	POO/UDP	EXPOO/UDP	POO/UDP

and leave the network between  $t_w$  and simulation end. It is widely accepted to model flow arrival pattern by a Poisson process [34] and so did we. In accordance with today’s Internet traffic composition we set  $\lambda = 2.4$  and  $\lambda = 0.6$  for ET and ST. Holding times (flow length) are either Exponential distributed for FTP, CBRR and EXPOO sources or Log-normal for POO. Peak rate of video sources is 0.125 Mbit/s and for VoIP 0.0625 Mbit/s. This corresponds to the rate ON/OFF sources emit packets in ON state.

The configuration of the AC module is as follows. The expected long-term demand for ET is 8.0 Mbit/s and henceforth the target rate in CS, denoted by ( $R_t$ ), is set to this value. The same applies to both streaming sources for which 1.0 Mbit/s has been set for either of them. In total this makes 10 Mbit/s, exactly the capacity of the access links.

For SS we set the delay threshold ( $D_t$ ) to 80 ms for ET and 10 ms as well as 5 ms for ST. The last two values are to incorporate the ITU-T recommendation of 150 ms mouth-to-ear delay for good VoIP quality while assuming 300 ms RTT and 10–15 intermediary nodes to the destination as typical for the Internet. Respectively, the local drop threshold ( $L_t$ ) is set to 10, 5 and 10 percent and was inspired by [41].

The PHB-D3 monitor takes delay and loss measurements with a 10 hz frequency (every 100 ms), like proposed in [23], and temporal peaks are filtered out by a simple Moving Average (MAVG) filter with  $J = 100$ . This value has been found empirically.

$$q_{n,c} = \frac{1}{J} \sum_{j=0}^{J-1} q_{t-j}. \tag{4}$$

To reflect the delay and loss sensitivities listed in Table I, column two and three, DSlopes are set to 37 degree for high, 36 for medium and 33 for low sensitivity. Further, as mentioned in Section 3, admitted flows are assigned an activity timer which is set to 10s by default. If no packets arrive during this period, the flow’s admission is revoked.

Finally, for every simulation scenario (A, B, C, D) we run three simulations in a row, one for each traffic composition. The abbreviation for a composition is

**Table II.** Results Simulation Scenario A for Traffic Composition FCP

Source	$R_t$	$\hat{R}_r$	$\hat{R}_m$	$T_f$	$D_t$	$\hat{D}_m$	$L_t$	$L_r$	$L_w$	$F_c$	$\hat{F}_a$	$\hat{F}_b$
FTP	8.0	33.45	8.09	0.020	0.080	0.058	10.0	<b>18.1</b>	<b>20.3</b>	3873	268	163
CBRR	1.0	1.22	1.08	0.122	0.010	0.0022	5.0	2.7	2.8	258	10	3
POO	1.0	1.00	0.37	0.024	0.005	0.0015	10.0	0.6	0.6	443	16	2

simply the combination of traffic model’s leading letters. For instance, FCP refers to FTP, CBRR, POO and is the abbreviation for the first composition in Table I.

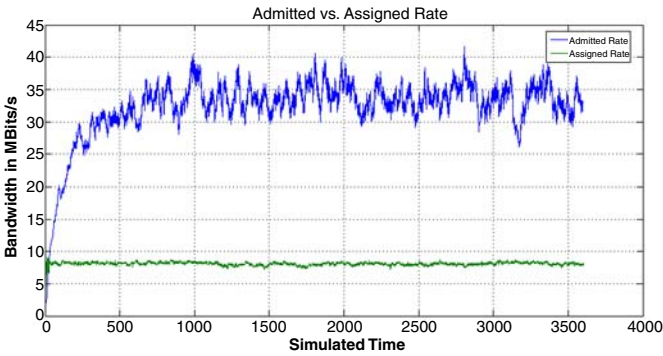
## 4.2. Simulation Scenarios

### 4.2.1. Scenario A

For the leadoff we deployed the AC algorithm in the configuration described in Section 3 and the results are tabulated in Table II. Abbreviations are: CS target rate ( $R_t$ ), mean registered rate ( $\hat{R}_r$ ), measured (mean assigned) rate ( $\hat{R}_m$ ) and per-flow throughput ( $T_f$ ), all in units of Mbit/s. Further, local delay threshold ( $D_t$ ) and mean measured delay ( $\hat{D}_m$ ) in seconds as well as local loss threshold ( $L_t$ ), total drops ( $L_r$ ) at the router, and drops after  $t_w$  ( $L_w$ ) at destinations (retransmissions for TCP sources), each in percent. Finally, the total number of flows completed ( $F_c$ ) between  $t_w$  and simulation end, mean flows admitted ( $\hat{F}_a$ ) and mean number of flows blocked ( $\hat{F}_b$ ).

As to be expected, this simulation scenario confirms interference issues incurred if TCP’s closed-loop congestion control operates over a network which itself features adaptive control. Since the majority of traffic is ET, TCP sources account for most of the CS admission, indicated in the difference between  $R_t$  and  $\hat{R}_r$  and illustrated in Fig. 4. Due to TCP’s nature, admitted sources try to adapt to the network and resource distribution on queuing level keeps almost constant simply because demand does not increase sharply enough. If eventually enough TCP sources are admitted onto the link and their cumulative probing for resources is sufficient, the total number of drops crosses the threshold, see  $L_t$ ,  $L_r$  and  $L_w$ . In this state, while system goodput is reasonable high, micro scale performance is poor as single TCP connections starve for bandwidth, see  $T_f$  and also Fig. 5. This confirms what has been found earlier in [39], also in the context of MBAC, namely that TCP sources require a minimum of bandwidth for reasonable performance.

Concluding on this findings, we restrict SS admission to RT only. Since this observation applies to any traffic compositions, we refrain from presenting results for the other two experiments.



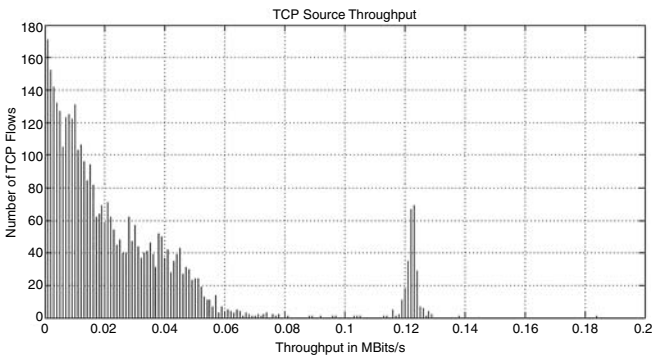
**Fig. 4.** Elastic traffic admitted (top) and actually assigned bandwidth (down). As indicated, offered load is much higher than the available capacity.

4.2.2. Scenario B

Using the same experimental setup as for the previous runs, we evaluated the performance of the modified algorithm. The results for all three experiments are tabulated in Table III.

Like in the preceding experiment, the difference between  $R_t$  and  $\hat{R}_r$  is a quantitative indicator for CS admission gain. However, with the new modification in place, it's almost only class 1 who gains, up to 4 times  $R_t$  for traffic pattern FEP (see Table I for class mapping). Comparing  $R_t$  and  $\hat{R}_m$  also acknowledges the PHB-D3 principle, service differentiation by priority and temporal demand.

In contrast to this benefit, with the only exception of class 2's delay target for the FEP composition, both QoS targets were missed for all simulations. In Table III,



**Fig. 5.** Elastic traffic per-flow throughput frequency. The majority of the flows suffer too low bandwidth because of link over subscription.

**Table III.** Results Simulation Scenario A for Traffic Composition FCP

Source	$R_t$	$\hat{R}_r$	$\hat{R}_m$	$T_f$	$D_t$	$\hat{D}_m$	$L_t$	$L_r$	$L_w$	$F_c$	$\hat{F}_a$	$\hat{F}_b$
FTP	8.0	33.45	8.09	0.020	0.080	0.058	10.0	<b>18.1</b>	<b>20.3</b>	3873	268	163
CBRR	1.0	1.22	1.08	0.122	0.010	0.0022	5.0	2.7	2.8	258	10	3
POO	1.0	1.00	0.37	0.024	0.005	0.0015	10.0	0.6	0.6	443	16	2

these deviations are emphasized using bold font. We explain this by the length of the measurement interval and the lack of prediction in this interval. More precisely, delay and loss estimates are computed every 100 ms and these estimates are valid for the entire interval. All admission requests within this interval are evaluated based on these estimates and thus equal criteria. This admission linearity implies that if an estimated QoS criteria is very close but yet below its threshold, a burst of flow arrivals would be admitted but likely results in an over subscription and if so leads to QoS target violations. This inherent issue of any measurement-based approach is further magnified if time required to capture the impact of a flow admissions is large. In other words, if much time elapses until parameter estimates reflect the additional load of recent admissions.

Another performance relevant issue is finding the optimal configuration for flow (in)activity detection. As explained in Section 3, the PES monitors flow activity as there is no explicit signaling informing the AC algorithm about flow departures. In order to release resources right in time, the PES therefore maintains an activity timer for each admitted flow. If a flow has been inactive for the value of this timer, its departure is assumed, admission revoked and registered resources are released. If this value is kept small the algorithm's dynamics are preserved but the risk of communication interruption might be high. To the other extreme, a large value incurs an increasing risk of over subscription. Inactive flows are reflected in low delay and loss estimates, which in turn lead to positive admission decisions in CS. If, after a time of inactivity, a number of flows become simultaneously active again, competing for resources with the additionally admitted flows during their inactivity, the PHB-D3 module responds accordingly with a redistribution of resources, which in turn is likely to result in missed QoS targets.

#### 4.2.3. Scenario C

After the disclosure of this issue, we further modified the algorithm towards better robustness. We refined SS admission policy such that each time a flow has been rejected all subsequent flows are also rejected until a flow departure is detected since this implies released resources. As before, the efficiency of this modification has been evaluated under exactly the same simulation conditions. The results are presented in Table IV.

**Table IV.** Results Simulation Scenario C for Traffic Composition FCP, FCE, and FEP, Top Down

Source	$R_t$	$\hat{R}_r$	$\hat{R}_m$	$T_f$	$D_t$	$\hat{D}_m$	$L_t$	$L_r$	$L_w$	$F_c$	$\hat{F}_a$	$\hat{F}_b$
FTP	8.0	8.0	6.52	0.089	0.080	0.073	10.0	7.8	8.6	499	64	143
CBRR	1.0	2.80	2.60	0.118	0.010	0.004	5.0	<b>5.8</b>	<b>6.0</b>	82	22	4
POO	1.0	1.03	0.41	0.019	0.005	0.001	10.0	0.6	0.6	122	17	4
FTP	8.0	8.0	6.51	0.087	0.080	0.073	10.0	7.7	8.6	529	64	141
CBRR	1.0	2.72	2.53	0.118	0.010	0.048	5.0	<b>5.6</b>	<b>5.6</b>	85	22	4
EXPOO	1.0	1.23	0.49	0.025	0.005	0.002	10.0	0.8	0.8	78	20	4
FTP 0	8.0	8.0	6.68	0.090	0.080	0.072	10.0	7.6	8.4	497	64	141
EXPOO	1.0	4.36	2.34	0.068	0.010	0.004	5.0	<b>5.7</b>	<b>5.8</b>	127	35	3
POO	1.0	1.30	0.52	0.019	0.005	0.002	10.0	0.8	0.8	139	21	4

The results show that with this modification in place, the algorithm's reactivity is limited and aligned to flow departure events. Juxtaposing the results, however, discloses some intriguing facts. While for scenario A the AC algorithm exhibited great inaccuracy, this modified version does meet any delay ( $D_t$ ) targets and approaches loss targets  $L_r$  and  $L_w$  with a much lower deviation, however, still in the order of 10 to 20 percent. As a consequence of the applied modification, a more conservative admission policy is to be expected and is indeed confirmed and quantified by  $F_c$ ,  $\hat{F}_a$  and  $\hat{F}_b$ . However, the effects do not severely affect  $F_c$  and  $\hat{F}_a$ . The more measured admission policy is also reflected in  $\hat{R}_r$ .

As a means to compare the performance with and without the last modification enabled, we defined a performance index.


$$I_p = (R_r * L_t) / (R_t * L_w). \quad (5)$$

A larger  $I_p$  indicates a better SS admission gain to QoS target violation ratio in comparison to a smaller value. The result of the comparison is illustrated In Figs. 6 and 7 in which left bars in a couple represent scenario B and right bars scenario C. According to this metric the latest modification improves the performance in the sense of a better tradeoff between high resources utilization and meeting QoS targets.

#### 4.2.4. Scenario D

Contrary to previous practice, the modification presented in this scenario does not tackle and any empirically found issue but is to investigate an alternative onset, namely a non-parametric probability model motivated by facts elucidated in Section 3.


In its current state, the algorithm estimates mean values of packet delay and loss and these, in turn, serve as indicators for QoS and resource availability. A



**Fig. 6.** Class 1 performance index comparison for traffic pattern FCP-FEP, left to right.

mean value, however, does not reflect the fluctuation of the underlying process. As a consequence, for two distinct delay (or loss) processes with equal mean but different variance and distribution function the algorithm performs equal, irrespective their disparate natures. Therefore, the MBAC is so far tailored to *average* requirements.

This finding motivates another modification, whose target is to distinguish processes by variance and distribution in order to provide an empirical, probabilistic upper bound for delay and loss. In order to do so, we apply a (non-parametric) histogram to capture the marginal distribution of delay (and loss) processes. The choice for a histogram based approach results from its general applicability and



**Fig. 7.** Class 2 performance index comparison for traffic pattern FCP-FEP, left to right.



the uncertainty associated with the statistical characteristics of today’s Internet traffic as briefly discussed in Section 3 and extensively elaborated in [36].

As in its initial design, delay and loss values are sampled but now the algorithm also computes a histogram there from. Let  $n$  denote the sample size,  $B_k = [t_k, t_{k+1})$  the  $k$ -th bin where  $t_{k+1} - t_k = h$ , for all  $k$ . The number of observations falling into the  $k$ -th bin is represented by  $v_k$ . The histogram at a point  $x$  is then defined as

$$\hat{f}(x) = \frac{v_k}{nh} = \frac{1}{nh} \sum_{i=1}^n I_{B_k}(X_i) \tag{6}$$

where  $I_{B_k}(X_i)$  is the indicator function, which reads

$$I_{B_k}(X_i) = \begin{cases} 1, & X_i \text{ in } B_k \\ 0, & X_i \text{ not in } B_k \end{cases} \tag{7}$$

Using a histogram enables us to estimate the probability with which delay ( $\hat{p}_d$ ) and loss ( $\hat{p}_l$ ) estimates exceed a given threshold ( $t_d, t_l$ ), see (8) for an example for delay.

$$\hat{p}_{1,c} = P\{P_{1,c} \geq t_{1,c}\} = 1 - \sum_{x=0}^{t_d} \hat{f}_{1,c}(x). \tag{8}$$

In this onset (3) translates to

$$\Omega_c = \begin{cases} \zeta_{n,c} - \left( \left( \sum_{f=0}^F \sigma_{f,c} \right) + \sigma_{k,c} \right) & \text{if } s = n = 0, c = \{0, 1, \dots, C\} \\ \theta_{n,c} - \hat{p}_{n,c} & \text{if } s = 1, n = \{1, 2\}, c = \{0, 1, \dots, C\} \end{cases} \tag{9}$$

with  $\theta_{n,c}$  denoting the probability threshold.

For the evaluation we set the sample size  $n$  to 2000 delay (loss) measurements, corresponding to a 200s memory while keeping the measurement frequency as before. The histogram binwidth  $h$  is dynamically calculated at admission request as the maximum measured value divided by the sample size.

$$h = \frac{\max\{\hat{\varphi}_{n,c}\}}{n} \tag{10}$$

Simulation parameters are equivalent with those of the previous experimental setups, except that each QoS parameter is now associated with a probability of  $\theta_{n,c} = 10^{-2}$ . The results for this set of simulations are presented in Table V.

The new model is clearly more conservative when compared with its counterpart, indicated by comparing  $F_a$  in Tables IV and V. The results also indicate some sensitivity to traffic with possible LRD characteristic since not a single sources has been admitted in SS for the FCP composition and only a few for FEP, see the difference between  $R_t$  and  $R_r$  for all classes. An analysis of the estimated loss

**Table V.** Results Simulation Scenario D for Traffic Composition FCP, FCE, and FEP, Top Down

Source	$R_t$	$R_r$	$R_m$	$T_f$	$D_t$	$D_m$	$L_t$	$L_r$	$L_w$	$F_c$	$F_a$	$F_b$
FTP	8.0	8.0	8.1	0.109	0.080	0.057	10.0	6.0	7.0	518	64	138
CBRR	1.0	1.0	0.96	0.123	0.010	0.002	5.0	2.0	1.5	41	8	4
POO	1.0	0.99	0.40	0.021	0.005	0.001	10.0	0	0	101	16	4
FTP	8.0	8.0	6.66	0.089	0.080	0.070	10.0	8.0	8.0	518	64	141
CBRR	1.0	2.58	2.40	0.118	0.010	0.003	5.0	6.0	6.0	69	21	4
EXPOO	1.0	1.20	0.48	0.025	0.005	0.001	10.0	1	0	74	16	4
FTP	8.0	8.0	8.54	0.112	0.080	0.054	10.0	6.0	6.7	496	64	139
EXPOO	1.0	1.02	0.56	0.070	0.010	0.003	5.0	2.0	2.1	34	8	4
POO	1.0	0.99	0.40	0.020	0.005	0.001	10.0	0	0	119	16	4

values, however, disclosed an interesting fact. Since for all three patterns the loss percentage for POO traffic is close to zero (values  $\leq 10^{-2}$  are considered zero) it can't be the influence of POO traffic which prevents SS admissions. In fact, we found that it is almost solely due to ET.

Basically, there are two explanations for this finding. The first is simply ET volume. Recall, that ET accounts for around 80 percent of the total traffic volume and delay and loss violation probability is likely to increase with the total amount of packets. A second explanation is based on a complex coherence between AC and the dynamics of PHB-D3. Highly variant POO traffic triggers frequent redistributions which in turn leads to an increasing risk of packet loss, especially for TCP sources which need time to readjust transmission rates to network performance. This would also explain the similarity of pattern FEP in scenario C and D. Indeed, for almost all performance parameters, both models perform equal. Even loss target violations experienced by class one are grossly the same. We explain this by the smoothing phenomena of EXPOO traffic. Albeit the same average number of admitted flows for EXPOO and POO sources, compare  $F_a$  in Table V, the former model is by far smoother. Thus, redistribution of resources happens less frequent and TCP sources do not have to readjust transmission rates so frequently what prevents packet loss and therefore keeps the algorithm more daring. As a quantitative indicator see the difference between  $R_t$  and  $R_r$  for class 2 in pattern FCE in Table V.

## 5. CONCLUSION

At our laboratory we pursue the investigation of an alternative QoS model for the Internet. As one result of the yet to happen advent of QoS support in the Internet, many of today's applications implement features to deal with the unreliable nature

of the Internet. These applications therefore do not require much support from the network and a service Better-Than-BE (BBE) might be far enough. This is the principle of LCT-QoS, which does provide less strict QoS guarantees in favour of more flexible deployment.

Key to LCT-QoS is dynamic resource distribution based on delay and loss estimates, computed from real-time measurements and mapped to application utility. While this proved to be an efficient concept for our BBE model, any protection from overwhelming network load in form of AC was yet missing. This was subject of this article.

The presented AC algorithm is simple, fast and reasonable accurate. This has been achieved by taking the best of features of PAC and MBAC and combination them. As it turned out, some simple modifications could improve the performance of AC considerably. However, simulations also indicate that AC based on pure measurements is doomed to limited precision. With a lack of prediction it is difficult to gain perfect control about resource management.

In a slightly different onset, we made a step ahead and presented a modified algorithm based on a non-parametric probability model. While, the initial focus of LCT-QoS is set on applications with average requirements, we enhanced the algorithm to provide probabilistic upper bounds. Without requiring any change in the PHB-D3 module, simulations showed that the modified algorithm is able to meet the performance expectations while evaluated on set of heterogeneous traffic patterns.

Finally, simulations provided valuable insight in complex coherence between PHB-D3 and MBAC performance. The dynamic distribution of resources introduces another dimension of complexity compared to common fields of MBAC application. However, results also showed that a close cooperation between the queueing and AC module can be beneficial, especially in order to keep AC design simple and flexible.

## ACKNOWLEDGMENTS

This work was supported by the Portuguese Ministry of Science and High Education (MCES), and by European Union FEDER under program POSI (project QoSMap).

## REFERENCES

1. IEEE Higher Speed Study Group, sep 2006. <http://ieee802.org/3/hssg/public/index.html>. Accessed on 07.01.2007.
2. G. R. Bachula, Testimony of G. R. Bachula, Vice President, Internet2 Before the United States Senate Committee on Commerce, Science and Transportation: Hearing on Net Neutrality. Technical report, Internet2, February 2006.

3. B. Teitelbaum and S. Shalunov, Why Premium IP Service Has Not Deployed (and Probably Never Will). Technical report, Internet2, May 2002.
4. S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, *An Architecture for Differentiated Services*. RFC2475. IETF, December 1998.
5. B. Davie, A. Chany, F. Baker, J. C. R. Bennet, K. Benson, J. Y. Le Bedec, A. Chui, W. Courtney, S. Davari, V. Firoiu, C. Kalmanek, K. Ramakrishnan, and D. Stiliadis, *An Expedited Forwarding PHB (Per-Hop Behavior)*. RFC 3246. IETF, March 2002.
6. C. Fraleigh, F. Tobagi, and C. Diot, Provisioning IP backbone networks to support latency sensitive traffic. In *INFOCOM 2003*, IEEE, Vol. 1, pp. 375–385. April 2003.
7. J. Cornu and G. Hughes, digital divide and broadband territorial coverage. Technical report, eEurope Advisory Group, June 2004.
8. J. Evans and C. Filsfil, Deploying DiffServ at the network edge for tight SLAs, Part 1. *IEEE Internet Computing*, pp. 61–65, January 2004.
9. J. Evans and C. Filsfil, Deploying DiffServ at the network edge for tight SLAs, Part 2. *IEEE Internet Computing*, pp. 61–69, March 2004.
10. V. Padmanabhan, L. Qiu, and H. J. Wang, Server-based inference of Internet link lossiness. In *IEEE INFOCOM 2003*. IEEE, 2003.
11. J. H. James, B. C. Chen, and L. Garrison, Implementing VoIP: A voice transmission report. *IEEE Communications Magazine*, Vol. 42, No. 7, July 2004.
12. D. Miras, Network QoS needs of advanced Internet applications a survey. Technical Report, Internet2, nov 2002.
13. M. Curado, O. Reis, J. Brito, C. Quadros, and E. Monteiro, Deployment of class-based routing in wide area networks: Overhead and performance assessment. In *6th IEEE International Conference on High Speed Networks and Multimedia Communications (HSNMC 2003)*, Vol. 2601 of *Lecture Notes in Computer Science*, Estoril, Portugal, Springer, July 2003.
14. M. Curado, O. Reis, J. Brito, G. Quadros, and E. Monteiro, QoS routing for differentiated services: simulations and prototype experiments. In *Internetworking 2003 International Conference*, San Jose, California, USA, IFIP, June 2003.
15. M. Curado, O. Reis, J. Brito, G. Quadros, and E. Monteiro, Stability and scalability issues in hop-by-hop class-based routing. In *2nd International Workshop on QoS in Multiservice IP Networks (QoS-IP 2003)*, Vol. 2720 of *Lecture Notes in Computer Science*, Milan, Italy, Springer, February 2003.
16. M. Oliveira, J. Brito, G. Quadros, and E. Monteiro, Simulation analysis of the UC-QoS routing strategy. In *5th World Multiconference on Systemics, Cybernetics and Informatics (SCI 2001)*, Orlando, USA, July 2001.
17. J. Gozdecki, A. Jajszczyk, and R. Stankiewicz, Quality of service terminology in IP networks. *IEEE Communications Magazine*, Vol. 41, No. 3, pp. 153–159, March 2003.
18. A. Takahashit, H. Yoshino, and N. Kitawaki, Perceptual QoS assessment technologies for VoIP. *IEEE Communication*, Vol. 42, No. 7, pp. 28–35, July 2004.
19. A. Watson and M. A. Sasse, Measuring perceived quality of speech and video in multimedia conferencing applications. In *ACM Multimedia*, ACM, pp. 55–60, October 1998.
20. E. Monteiro, F. Boavida, G. Quadros, and V. Freitas, Specification, quantification and provision of quality of service and congestion control for new communication services. In *Proceedings of the 16th AFCEA Europe Brussels Symposium & Exposition, AFCEA/IEEE ComSoc*, Les Pyramides, Brussels, Belgium, January 1995.
21. T. M. Bohnert and E. Monteiro, On autonomic qos provisioning for VoIP services. In *4th IEEE Consumer and Communications and Networking Conference, CCNC'07*, Las Vegas, USA, IEEE, January 2007.

22. M. Fomenkov, K. Keys, D. Moore, and K. Claffy, Longitudinal study of Internet traffic in 1998–2003. In *Winter International Symposium on Information and Communication Technologies (WISICT)*, January 2004.
23. G. Quadros, A. Alves, E. Monteiro, and F. Boavida, An approach to support traffic classes in IP networks. In *QoITS*, pp. 285–299, 2000.
24. P. Hurley, J.-Y. Le Boudec, P. Thiran, and M. Kara, ABE: Providing a low-delay service within best effort. *IEEE Network*, Vol. 15, No. 3, pp. 60–69, June 2001.
25. T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido, A nonstationary poisson view of internet traffic. In *Infocom*, Hong Kong, Hong Kong, IEEE, March 2004.
26. J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun, On the nonstationarity of Internet traffic. In *ACM SIGMETRICS '01*, New York, USA, pp. 102–112, ACM Press, July 2001.
27. M. Grossglauser and D. N. C. Tse, A framework for robust measurement-based admission control. *IEEE/ACM Transactions on Networking*, Vol. 7, No. 3, pp. 293–309, 1999.
28. L. Breslau, S. Jamin, and S. Shenker, Measurement-based admission control: what is the research agenda. In *IWQOS*, London, UK, pp. 3–5, March 1999.
29. L. Breslau, S. Jamin, and S. Shenker, Comments on the performance of measurement-based admission control algorithms. In *INFOCOM*, Tel Aviv, Israel, pp. 1233–1242, March 2000.
30. A. W. Moore, An implementation-based comparison of measurement-based admission control algorithms. *J. High Speed Networks*, Vol. 13, No. 2, pp. 87–102, 2004.
31. S. Georgoulas, P. Trimintzios, and G. Pavlou, Joint measurement-and traffic descriptor-based admission control at real-time traffic aggregation. In *ICC '04: IEEE International Conference on Communications*, 2004.
32. J. Qiu and E. W. Knightly, Measurement-based admission control with aggregate traffic envelopes. *IEEE/ACM Transactions on Networking*, Vol. 9, No. 2, pp. 199–210, 2001.
33. T. K. Lee and M. Zukerman, Practical approaches for connection admission control in multiservice networks. In *Proceedings of IEEE ICON '99*, Brisbane, Australia, pp. 172–177, May 1999.
34. V. Paxson and S. Floyd, Wide-area traffic: The failure of poisson modeling. In *SIGCOMM*, London, United Kingdom, pp. 257–268, August 1994.
35. J. Kilpi and I. Norros, Testing the Gaussian approximation of aggregate traffic. In *Internet Measurement Workshop*, ACM, pp. 49–61, 2002.
36. T. M. Bohnert, E. Monteiro, Y. Koucheryavy, and D. Motchanov, Non-parametric and self-tuning measurement-based admission control. In *IFIP Networking 2007*, Atlanta, USA, IFIP, May 2007.
37. T. M. Bohnert, Measurement-based admission control for a class-based QoS framework. Master's thesis, University of Coimbra, Coimbra, Portugal, January 2005.
38. *Resource and Admission Control Sub-system (RACS); Functional Architecture*. Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN). ETSI, Sophia Antipolis Cedex, France, June 2006.
39. R. Mortier, I. Pratt, C. Clark, and S. Crosby, Implicit admission control. *IEEE Journal on Selected Areas in Communications*, Vol. 18, No. 12, pp. 2629–2639, December 2000.
40. T. M. Bohnert and E. Monteiro, A comment on simulating LRD traffic with pareto ON/OFF sources. In *CoNEXT'05: ACM Conference on Emerging Network Experiment and Technology*, Toulouse, France, pp. 228–229, ACM Press, October 2005.
41. T. J. Kostas, M. S. Borella, I. Sidhu, G. M. Schuster, J. Brabiec, and J. Mahler, Real-time voice over packet-switched networks. *IEEE Network*, Vol. 12, pp. 18–27, January 1998.

**Thomas Michael Bohnert** received a diploma degree in Computer Engineering in 2005 from the University of Technology Mannheim in Germany. Since then he is pursuing PhD studies in Admission Control at the University of Coimbra, Portugal. His scientific interests are General Network Modelling and Performance Analysis as well as Perceptual and Intrinsic Quality of Service provisioning. He is also chairing the ESF COST290 Action WiMAX Study Group.

**Edmundo Monteiro** is Associate Professor at the University of Coimbra, Portugal from where he graduated in 1984 and got a PhD in Electrical Engineering-Computer Speciality, in 1995. His research interests are Computer Communications, Quality of Service, Mobility, Routing, Resilience and Security.