

Traffic Handling in AQUILA QoS IP Network

A. Bak¹, W. Burakowski¹, F. Ricciato², S. Salsano², and H. Tarasiuk¹

¹ Institute of Telecommunications, Warsaw University of Technology,
ul. Nowowiejska 15/19, 00-665 Warsaw (Poland)
{bak, wojtek, halina}@tele.pw.edu.pl
² CoRiTel – v. Anagnina 203, 00040 Roma (Italy)
{ricciato, salsano}@coritel.it

Abstract. The paper describes the traffic handling mechanisms implemented in the AQUILA pilot QoS IP network [10]. The AQUILA project enhances the DiffServ architecture concept [1,2,3] by adding new functionality for admission control and resource management as well as by defining new set of network services. Each network service is optimised for specific type of traffic (e.g. non-reactive and reactive) and has its own traffic handling mechanisms. Exemplary measurement results verifying the effectiveness of AQUILA approach for providing QoS are also included.

1 Introduction

Offering QoS (Quality of Service) in IP networks is of strategic importance for Internet Service Providers. The AQUILA¹ project tackles this challenge, aiming at the definition of a QoS architecture for IP. The problem involves the following aspects: (1) network architecture and traffic control (e.g. definition of QoS signaling, interaction with routing protocols, etc.) (2) traffic handling (e.g. packet scheduling, admission control algorithms, etc.), and (3) management aspects (e.g. user subscription to QoS services, accounting, billing, etc.). AQUILA has focused on a subset of these topics, mainly covering architectural and traffic handling aspects, while, for example, the management and billing issues were not covered.

This paper mainly deals with the traffic handling issues in AQUILA network. A detailed description of the overall AQUILA architecture one can find in [6]. Traffic handling for QoS is rather complex issue involving a set of mechanism operating on different time scales, from milliseconds (packet scheduling) to hours and days (resource provisioning, traffic engineering). AQUILA organizes all the traffic handling aspects in a single vision. Currently, after 15 months from the start of the project, the AQUILA project is running its first trial and it has started the specification for the second project phase. The testbed that has been developed within the project provides a great opportunity to see a running complete QoS architecture [10]. The goal of the paper is to present the specification of the traffic handling mechanisms and

¹ AQUILA (Adaptive Resource Control for QoS using an IP-based Layered Architecture) IST-1999-10077 is an European research project partially funded by the IST programme

to start analysing the project results and the experiences that are being matured in the testbed.

Section 2 introduces us to the AQUILA architecture and investigated concepts, with a special focus on the relationships between the traffic handling mechanisms referring to different time scales. Sections 3 and 4 describe the traffic control mechanisms operating at the packet level and at the flow/aggregate level, respectively. Section 5 presents first experimental results, mainly covering the traffic control at the packet level. Finally, section 6 summarises the paper.

2 Overview of Traffic Handling Approach

This section provides an overview of the AQUILA architecture and concepts. Let us briefly recall the rationale and the architecture of AQUILA. Two important aspects of QoS are QoS *guarantees*, and QoS *differentiation*. In order to provide QoS differentiation, a limited set of Network Services (NS) have been defined in AQUILA project, which represent the services sold by the provider to its customers: Premium Constant Bit Rate (PCBR), Premium Variable Bit Rate (PVBR), Premium Multimedia (PMM), Premium Mission Critical (PMC) and Standard Best Effort (STD). Each NS is meant to support a class of applications with substantially similar *requirements* and *traffic characteristics*. The Network Services are internally mapped by the operator into a set of Traffic Classes. The Traffic Classes use Diffserv based packet handling mechanisms. For providing QoS guarantees the ISP must somehow regulate the volume of traffic submitted to the network, regarded as a limited set of resources. In the AQUILA approach this is accomplished by a distributed layer: the Resource Control Layer (RCL). The RCL components are the End-user Application Toolkit (EAT), located at the end-user site, the Admission Control Agent (ACA), located at the edge of the network (e.g. in the Edge Routers), and the Resource Control Agent (RCA), which is logically a centralized entity within the network itself. The RCL embeds different mechanisms to regulate the traffic at different time-scales – Initial Provisioning, Dynamic Resource Pool and Admission Control. Note that in the first phase of the AQUILA project, the focus is on QoS in a single domain.

Let us now turn our attention on the *traffic handling* mechanisms in AQUILA. Traffic handling is used here as a general term for a set of coordinated mechanisms operating at different time scales:

- **Traffic control** refers to the mechanisms operating at milliseconds time scale like packet scheduling, policing, queue management.
- **Admission control** refers to the algorithms to decide about the acceptance of a new flow in the network, operating at the time scale of seconds to tens of minutes.
- **Resource Pools** refers to the algorithm for short term resource redistribution, to cope with local fluctuations in offered traffic, operating at the time scale of tens of minutes to hours.
- **Provisioning** refers to the algorithm for medium/long term resource allocation and redistribution, operating at the time scale of hours to days.

The relationships among these logical components (Provisioning, Resource Pools, Admission Control and Traffic Control) are briefly described in the following. A very high-level view of the process that enables QoS in the AQUILA architecture is shown in Fig.1, while Fig. 2 presents a simplified pictorial view of the relationships between the different mechanisms. More details are provided in the next sections.

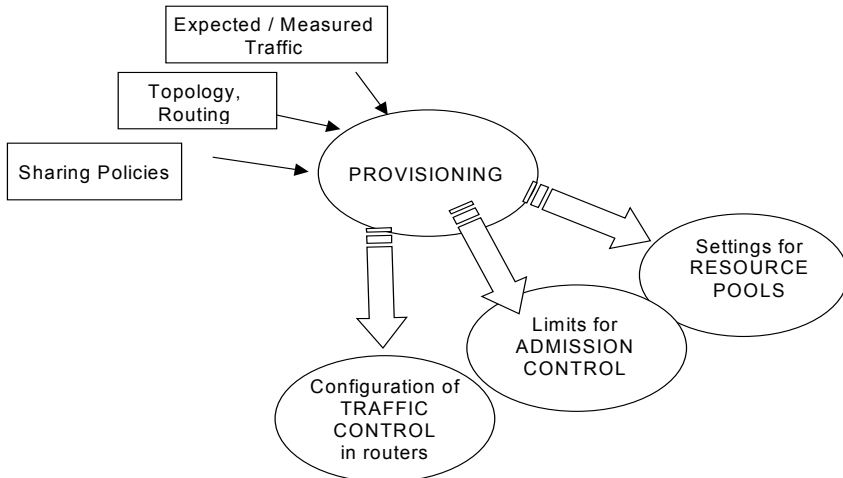


Fig. 1. Enabling QoS in the AQUILA architecture

The Provisioning phase is run off-line before the network operation, and gives the required input to the RCL elements as well as configuration values for setting the router parameters. The initial provisioning algorithm takes as an input global information about the topology, the routing (costs of links), the expected traffic distribution between Edge Routers for each Traffic Class, and any further constraints on the link bandwidth sharing between TCLs. It performs a sort of global computation and produces as output:

- the expected amount of traffic for each Traffic Class on each link, called *provisioned rate*. This is used for the router configuration, i.e. to chose the appropriate setting for the scheduling / queuing parameters (WFQ weights, WRED thresholds) at each router interface.
- the Admission Control Limits for each Traffic Class at each Edge Router. This is used by AC algorithms during the operation phase..
- Definition of the Resource Pools sets (RP will be discussed in section 0).

The Traffic Control mechanisms define how the packets of the different classes are handled by the Edge and Core Routers in the AQUILA network. They includes traffic conditioning (also referred to as *policing*), that is enforced at ingress ERs only, and scheduling / queuing algorithms, implemented at any router interface.

The configuration of the scheduling / queuing mechanism is “static”, i.e. the relevant parameters are configured in the routers at start up. An off-line procedure computes these parameters starting from the provisioned rates produced by the Initial Provisioning algorithm.

Obviously, the configuration of per-flow traffic conditioning parameters at the ingress edge is done run-time according to the admitted requests. Details on the traffic conditioning mechanisms will be given in section 0.

The Admission Control procedure is intended to restrict traffic in order to avoid congestion. The AC procedure is operated on-line, but the AC reference limits, or AC Limits, are computed during the off-line initial provisioning phase and configured during the start-up phase.

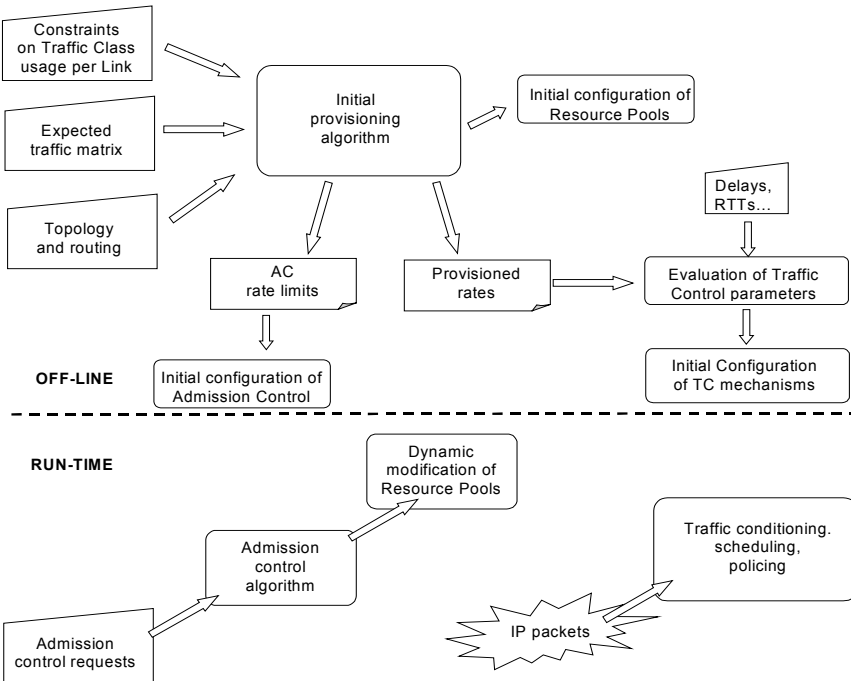


Fig. 2. Initial Provisioning, Traffic Control and Admission Control

The assignment of AC Limits to each Edge Router for each TCL represents a resource assignments to the relevant traffic aggregates. As the AC Limits are computed based on the expected offered traffic at each ER, some deviation can occur during the operation phase between the actual offered traffic and the resource distribution between ERs. The Resource Pools mechanism represent a way to dynamically change the AC Limit to some extent, so as to dynamically track short term fluctuations in traffic requests. Such mechanisms are based on the concept of RP, which are sets of ER that can exchange resources with each other. Such sets are defined during the Initial Provisioning phase.

3 Traffic Handling Mechanisms at Packet Level: The AQUILA Traffic Classes

From the network point of view, the differentiation at the service level into Network Services (NS) naturally introduces a differentiation at the packet handling level. In particular, the implementation of relative priorities between the packets, both for the access to the transmission channel and/or to the buffer space, is exploited to differentiate the delivered end-to-end delays and loss probability. Although the *prioritisation* induced by the applications requirements is a key component of the traffic handling at the packet level, it is not the only reason to introduce packet handling differentiation inside the network: the other factor is the need for *separation* between traffic flows with dramatically different characteristics. As an example, it is well recognized that closed loop flows (typically TCP and in general TCP-friendly flows) should not compete on the same resources with open loop flows (typically UDP). Separation between the two can be achieved e.g. by using different queues served with some bandwidth sharing mechanism (WFQ scheduling).

Prioritisation and separation represent the two main aspects that any design of packet handling mechanisms must take care of. Beyond that, one should also take into account the excess packets treatment and, eventually, those advanced queue management schemes meant to enhance the performances of reactive traffic (typically RED for TCP).

In designing the packet handling mechanisms inside the router, AQUILA has taken into account all of that: AQUILA has defined a set of 5 Traffic Classes (TCL). At each Edge Router (ER), each TCL is assigned a portion of “resource”, i.e. a bandwidth value, which is meant to limit the maximum amount of traffic that the ER can inject into the network for the specific TCL. As this value will be used by the Admission Control algorithm to decide about the acceptance of new flows to the relevant TCL, it will be denoted by AC Limit. The proposed AC algorithms for each TCL will be discussed in section 3.1. Each TCL is associated a different queue in the router output interface, and a bandwidth portion on each link. All the queues except the first one are served by a WFQ scheduler, thus each TCL is associated to a WFQ weight. The queue dedicated to TCL-1 is served with strict priority over the others. Fig. 3 shows the inter-TCL scheduling scheme. It should be considered that such inter-TCL scheduling scheme applies to high speed router interfaces. In fact for low-speed interfaces (few hundreds kilobits per second), which can be found in the access network section, it is not reasonable to statically partition the (little) available bandwidth between the TCLs. To cope with this problem, AQUILA TCLs are implemented in a different way on low-speed interfaces. However for sake of space we will not deal with low-speed interface in this paper. We refer to [7] for further details on this topic.

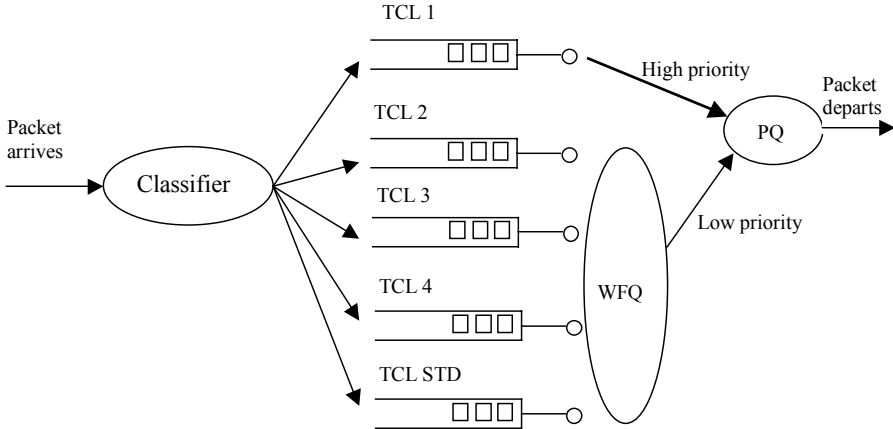


Fig. 3. Design of router output port

TCL-5 is intended to support the Standard Service (STD), i.e. the traditional best-effort traffic. The traffic accessing the STD service is not delivered any QoS and is not regulated by any AC and/or policing function inside the ER. Nevertheless, a non-null amount of bandwidth will be guaranteed to this traffic on each link, accordingly the WFQ weight for TCL-5 will be non-zero.

TCL-1 and TCL-2 are intended to support non-reactive (open loop) traffic with stringent QoS requirements. In particular TCL-1 will be characterized by very high QoS performance (very low delay and very low losses), accomplished by a conservative AC scheme: no statistical multiplexing is achieved within this TCL, and the AC algorithm will exclusively base on the flows peak-rate which is declared by the applications. In the AQUILA architecture TCL-1, which is somehow similar to the EF PHB defined in [2] will exclusively support the PCBR service. Typically, TCL-1 will be entered by flows with small to medium packet size (< 256 B) and not too large peak-rate, as typically originated by real-time streaming applications like VoIP, etc. On the other hand, TCL-2 will deliver a lower QoS level (low delay and low losses) to those streaming application with high emission rate variability and/or large packets: the AC for TCL-2 scheme allows for some degree of statistical multiplexing, thus the mean rate of the flows must be taken into account in the AC algorithm along with the peak-rate. TCL-2 will mainly support the PVBR network service.

TCL-3 and TCL-4 are dedicated to reactive flows (TCP and TCP-like). In particular, TCL-3 will support PMM service and will take long-lived TCP connections (for long file transfers) or other adaptive application flows (audio/video download, adaptive video). These flows are typically greedy, as they continue to expand the emission rate until congestion is reached. TCL-4 instead will support PMC service and will receive non-greedy elastic flows, typically short-lived TCP connections originated by some critical transaction application (e.g. finance) or interactive games. Note that separating long-lived and short-lived TCP connections into different classes, thus avoiding direct competition on the same resources (the WFQ scheduler acts as a sort of arbiter for the bandwidth) prevents the former from starving the latter. Further

details about queuing mechanisms as well as traffic conditioning for each TCL are given below.

It can be noted that there is a one to one mapping between the defined TCLs and the NSs. This should be considered as a sort of coincidence, as there is no one-to-one mapping requirement between NSs and TCLs but the provider has a degree of freedom in this mapping. Nevertheless, for sake of simplicity this degree of freedom has not been exploited in the current specification. It is expected that as the AQUILA approach will evolve, new NSs could be defined while keeping the same set of supported TCLs. On the other hand, there is also the possibility that, based on the ongoing trial experiences, a reduction in the number of TCL could be envisioned: in particular the need of two different TCLs for supporting streaming traffic is under investigation, therefore TCL-1 and TCL-2 could be merged into one single TCL.

Note that the AC function in the network is associated to the TCLs, rather than to the NS. In fact the resources within the network (corresponding to the AC Limits inside the ERs), are assigned on a per-TCL basis.

As a final remark, we are of course aware that a lot of effort is ongoing within the IETF to standardize Per Hop Behaviors (PHB). The scope of AQUILA is not to define new PHBs that will replace the ones defined in the IETF. The main focus of AQUILA is to study the whole QoS picture, from millisecond scale (Traffic Classes) up to long term scale (Provisioning). It was not possible at the AQUILA design time to rely on a stable definition of Diffserv PHBs and it was even farther the possibility to have Diffserv compliant implementation in routers. Therefore it was chosen to design the set of Traffic Classes based on Traffic control mechanism available in the routers. It is possible to re-map the AQUILA TCLs onto the “standard” Diffserv PHBs (EF, AF), but this is out of the scope of this work.

3.1 Traffic Control at Packet Level

This subsection provides further details about the packet scale traffic control mechanisms, i.e. traffic conditioning and queue management, for each TCL.

TCL-1 and TCL-2 are both dedicated to non-reactive (UDP) flows with stringent QoS requirements. According to that, a “severe” purely dropping traffic conditioning at ingress point is associated to both, i.e. all packets exceeding the declared profile are discarded. The traffic profile for TCL-1 is described in terms of a Single Token Bucket, limiting the flow peak rate. The traffic profile for TCL-2 is described in terms of a Dual Token Bucket, controlling both the peak and mean rate of the flow. This is consistent with the definition of AC algorithms: AC for TCL-1 is based only on peak rate, while AC for TCL-2 takes into account the mean rate also to achieve better multiplexing gain (see section 0 for details). For both TCL-1 and TCL-2, queues at router interfaces are simply of FIFO drop-tail type.

TCL-3 is dedicated to long-lived TCP controlled flows. A Single Token Bucket descriptor is used to declare the mean rate only. Traffic conditioning at ingress point is based on 2 colours marking: out-of-profile packets are not discarded but simply marked as such with a different DSCP value. At router interfaces, the TCL-3 queues

employs a WRED management algorithm with two different sets of parameters (minth, maxth, maxp) for in-profile and out-of-profile packets.

TCL-4 is dedicated to short-lived non-greedy TCP-controlled flows with low bandwidth requirements. Dual Token Bucket descriptor is used to declare mean and peak rate. Traffic conditioning and queue management are similar to TCL-3.

Finally, simple FIFO drop-tail queues are used for the STD TCL.

4 Traffic Handling Mechanisms at Flow and Aggregate Level

In the previous sections we have introduced the concept of NS and TCL, which aim at achieving QoS *differentiation* respectively at the service and at the traffic level. This chapter is now focused on the mechanisms that AQUILA implements to deliver QoS *guarantees*. For QoS to be delivered, the amount of traffic entering the network must be somehow regulated. To this purpose, AQUILA considers a combination of three mechanisms at aggregate level: Initial Provisioning, Dynamic Resource Pool and Admission Control. The following of this section discusses these mechanisms in a top-down fashion. These mechanisms are complementary to those discussed above operating at the packet level (e.g. traffic conditioning).

4.1 The Initial Provisioning Phase

For each TCL j each ER i is assigned a certain amount of bandwidth $l_i^{(j)}$, called AC Limit, which is used by the admission algorithm in the ACA responsible for ER i as a reference limit to the traffic it can accept for TCL j . The value of the AC Limit as well as the bandwidth demands for each active reservation are logically stored in the ACA. The initial computation of the set $\{l_i^{(j)}\}$ for each TCL and ER, called the Initial Provisioning phase in AQUILA, can be done off-line and must take into account the following input information:

- a. The complete network topology, included the link capacities and link cost (as used by IGP routing for computation of shortest path).
- b. The long-term expected traffic matrices, i.e. the expected spatial distribution of traffic for each TCL between source/destination router pair.
- c. Some bandwidth sharing policies between TCL, e.g. “no more than 10% of link capacity allocated to TCL-1”, or “no less than 30% of link capacity allocated to best effort traffic (TCL-5)” and similar.

In the simplest scenario classical IGP routing is used within the network (e.g. OSPF, RIP, etc.), and the route for each source/destination pair is unique and constrained to the shortest path. Nevertheless in more advanced scenarios, particularly when MPLS is used for Traffic Engineering purposes, the routing itself is not strictly constrained to shortest paths, and the paths computation could be made jointly with the AC Limits computation. In this case there would be room for some sort of *global optimisation* during the Initial Provisioning phase, whose output would be the set of paths along with the set of AC Limits. This additional perspective has been left out of

the scope of this paper. Also note that fault recovery aspects are not covered by AQUILA.

4.2 The Resource Pool

The scheme described up to now meets the requirement to process each RR locally at the ACA, but at the cost of a high rigidity in the resources distribution. In fact per ER (“horizontal”) as well as per TCL (“vertical”) bandwidth assignment is done statically in the initial provisioning phase. This scheme is not able to dynamically track fluctuations and deviations of the actual offered traffic from the expected one. In order to gain a more dynamical behaviour, it would be desirable to have dynamical sharing of resources between ER (“horizontal” sharing) and/or between TCL (“vertical” sharing) to some extent. AQUILA introduces the concept of “Resource Pool” (RP) to allow for some degree of dynamical sharing of resources between ER (“horizontal” sharing) only. Further introduction of some form of resource sharing between TCLs (“vertical” sharing) is envisioned as a further extension to the model.

In order to explain the RP concept, consider the case that an amount of Resource Reservations for TCL j are being rejected at some ER x due to consumption of the relevant AC Limit, while at the same time some other ER y does not fulfill its resource budget due to lack of demand. In this case it would be desirable to dynamically shift a portion of the resource budget from y to x , by increasing $l_x^{(j)}$ and decreasing $l_y^{(j)}$. In this example we will denote ER y as the “donor”. The concept of RP arises from the general consideration that not any ER can be a meaningful “donor” for any other ER: there are constraints between AC Limits due to the topology and the traffic distribution, mainly related to the presence of *bottlenecks*. Thus, a RP identifies a set of ERs that can be donors to each other: inside a RP the ERs can exchange between each other the assigned bandwidth for a given TCL. Application of the RP concept is straightforward in the case a set of ERs are connected in a star to a core router. The case is depicted in Fig. 4, which will be used to illustrate the dynamical resource distribution algorithm inside the RP.

The above approach can be hierarchically extended to build RP whose elements are not ERs but other RPs. The case is depicted in Fig. 4, where RP “C” is composed by node 9 as root and RPs “A” and “B” as leaves. The hierarchical RP approach can be straightforwardly applied in those networks whose external areas present a hierarchically structured topology, which is expected to be a quite common case in practice.

4.3 Admission Control

In order to control the access to the network, the user applications must “ask for permission” before sending traffic to the network. This permission is granted by a Reservation Request (RR) sent to the network. The task of sending RRs is covered by the EAT module at the user side, which represent a sort of interface between the (legacy) applications and the QoS network. The RR sent from the EAT is processed at

the network side by the ACA, which can accept or reject the flow depending on the *available resources*, i.e. on the profiles of the already admitted flows (for the same TCL, at the same ER) and on the relevant AC Limit.

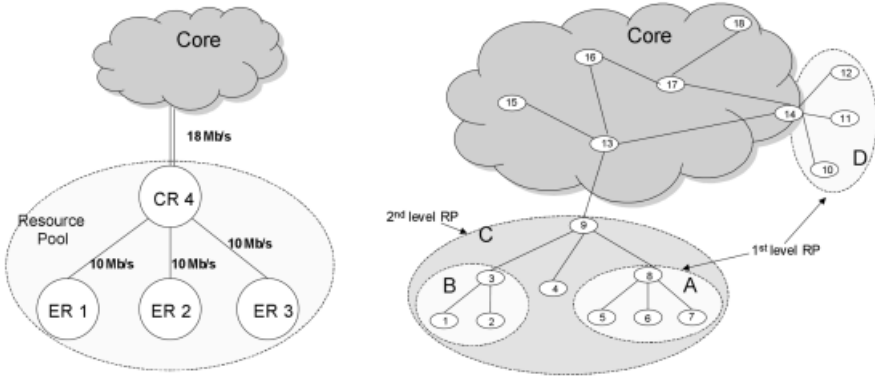


Fig. 4. Resource Pool

The proposed AC algorithms are derived from the results developed in the context of ATM traffic control. The request for network resources is accepted/rejected based on the traffic descriptors provided by the user. In the AQUILA architecture the admission decision is made only at the network ingress and, in some cases, at the egress point. This makes the AC decision more critical, as link-by-link verification of resource availability is not possible. To perform the admission control at the ingress or egress the single link model was considered with capacity C (AC limit) and buffer size B . Furthermore, for simplicity, the isolation between all traffic classes was assumed. In fact, the TCL-1 class has impact on other classes as it is served with the highest priority (see section 3). Whenever below are mentioned parameters C or B , they correspond to the capacity and buffer size dedicated to serve the given traffic class. In the following, details about the AC used in Aquila are provided.

Admission control algorithm for TCL-1

The TCL-1 class uses peak rate allocation scheme [4]. A flow in this class is characterised by the parameters of single token bucket algorithm that correspond to the peak bit rate (PBR) and peak bit rate tolerance (PBRT). In the AQUILA architecture, the TCL-1 traffic is served with the highest priority. Taking this into account, it can be assumed that the TCL-1 streams have negligible packet delay variation [11]. Consequently, the worst case traffic pattern for the superposition of a number of TCL-1 flows takes the form of poissonian stream (with the mean rate equal to the sum of the PBR parameters of the particular flows). Let us assume that the capacity dedicated for TCL-1 class is C_l . In the case, when N_1 flows with $\{PBR_1, PBR_2, \dots, PBR_{N_1}\}$ are currently in progress, a new flow declaring PBR_{new} as its peak rate is admitted if the following condition is satisfied:

$$PBR_{new} + \sum_{i=1}^{N1} PBR_i \leq \rho C_1 \quad (1)$$

Parameter ρ ($\rho < 1$) specifies the admissible load of capacity allocated to the TCL-1 class. The value of ρ is calculated from the analysis of M/D/1/B system taking into account the assumed target packet loss ratio and buffer size [5].

Admission control algorithm for TCL-2

In case of TCL-2 traffic class the REM (*Rate Envelope Multiplexing*) multiplexing scheme is assumed for guaranteeing low packet delay [4]. Therefore, the only QoS parameter that requires concern is the packet loss rate. In the REM multiplexing the buffer (relatively small) has to be dimensioned for absorbing, so called, the packet scale congestion (simultaneous arrival of packets from different sources). For this purpose the N*D/D/1 queuing system analysis is useful. In the TCL-2 class, each flow is characterised by the parameters of the Dual Token Bucket: the peak bit rate (PBR) jointly with the peak bit rate tolerance (PBRT) and the sustainable bit rate (SBR) jointly with the burst size (BSS). It is commonly believed that the worst-case traffic pattern for given values of PBR, SBR and BSS is of ON/OFF type. The proposed admission method for TCL-2 is based on the notion of effective bandwidth. One may find a number of methods for calculating effective bandwidth [4]. For its simplicity, the methods proposed in [12] was chosen for AQUILA. In this method the value of effective bandwidth, $Eff(\cdot)$, is calculated on the bases of PBR, SBR and BSS parameters taking into account the target packet loss rate.

Let us assume that the capacity dedicated for TCL-2 class is C_2 . In the case, when N_2 flows with $\{Eff(1), Eff(2), \dots, Eff(N_2)\}$ are currently in progress, a new flow with $Eff(new)$ is admitted if the following condition is satisfied:

$$Eff(new) + \sum_{i=1}^{N_2} Eff(i) \leq C_2 \quad (2)$$

Admission control algorithm for TCL-3

In the case of TCL-3, each flow is characterised by parameters of single token bucket algorithm that correspond to the sustained bit rate (SBR) and the burst size (BSS). The BSS values for typical flows accessing TCL-3 are expected rather large, allowing for a high variability of submitted traffic. As a difference with other TCLs, the declaration of the PBR parameter is omitted, so that the PBR will be only limited by the access links speed and intrinsic TCP dynamics.

Taking into account that the traffic flows submitted to TCL-3 class are TCP-controlled, only rough QoS guarantees are assumed to be provided. Therefore, an admission control method that maximises the network utilisation can be employed in this case.

Let us assume that the capacity dedicated for TCL-3 class is C_3 . In the case, when N_3 flows with $\{(SBR_1, BSS_1), (SBR_2, BSS_2), \dots, (SBR_{N_3}, BSS_{N_3})\}$ are currently in

progress, a new flow declaring (SBR_{new}, BSS_{new}) is admitted if the following condition is satisfied:

$$SBR_{new} + \sum_{i=1}^{N3} SBR_i \leq C_3 \quad (3)$$

Admission control algorithm for TCL-4

In the TCL-4, a flow is characterised by parameters of dual token bucket algorithm, similarly as in the case of TCL-2 class. The proposed admission control algorithm aims to provide the service rate that will guarantee virtually no packet losses.

Let us assume that the capacity dedicated for TCL-4 class is C_4 . In the case, when $N4$ flows with $\{Eff(1), Eff(2), \dots, Eff(N4)\}$ are currently in progress, a new flow with $Eff(new)$ is admitted if the following condition is satisfied:

$$Eff(new) + \sum_{i=1}^{N4} Eff(i) \leq C_4 \quad (4)$$

The effective bandwidth in this case is calculated by (see [13]):

$$Eff(.) = \max \left\{ SBR, \frac{PBR \cdot T}{B/C + T} \right\} \quad (5)$$

where

$$T = \frac{BSS}{PBR - SBR} \quad (6)$$

5 Experimental Results

This section presents experimental results [10] illustrating quality of service offered by two selected TCLs defined in AQUILA i.e. TCL-1 and TCL-3. These classes are designated for different traffic types, for streaming (PCBR NS) and elastic (PMM NS) traffic.

The experiments were carried out in the AQUILA testbed (see Fig.5), installed in Polish Telecom at Warsaw. The test topology consists of 8 CISCO routers (of different types) connected in the form of the chain for achieving large number of hops. The end terminals are connected to the edge routers by Ethernet ports. The access links (deployed between the edge router and the first core router) are of rather low speed (2 Mbps). The higher capacity links (10 and 155 Mbps) connect the core routers. The following types of routers are installed in the testbed: two edge routers - 1605 (aq1605_2) and 3640 (aq3640_4), 6 core routers - 3640 (aq3640_1, aq3640_2, aq3640_3,) and 7507 (aq7507_1, aq7507_2, aq7507_3). Details about router configuration values can be found in [10].

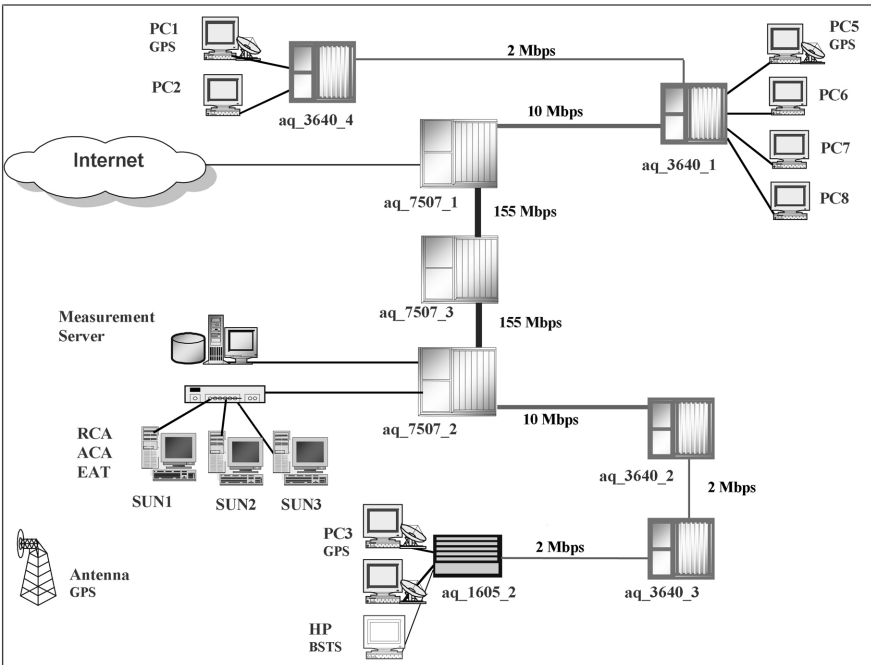


Fig. 5. AQUILA testbed configuration. PC1-8 – terminals, SUN1-3 – SUN Work Stations with implemented RCA, ACA and EAT modules, aq_1605, aq_3640, aq_7507 – CISCO routers

5.1 Results for TCL-1

TCL-1 class was tested assuming that the packet traffic submitted to this class is the maximum traffic allowed by the admission control. The test traffic (TCL-1) was modelled by a Poisson process with constant packet length. Such traffic represents the worst case of the superposition of large number of CBR streams. To take into account the impact of other class traffic on the TCL-1 packets, background traffic was added. The background traffic was sufficient to load the rest of link capacity not dedicated to TCL-1.

In the following experiments, 200 kbps of the access link capacity was reserved for TCL-1. Furthermore, the TCL-1 buffers in the routers were set to 5 packets to guarantee low packet delay requirements. The performance of TCL-1 was validated assuming target packet loss ratio (Ploss) to be 10^{-2} . According to the specified admission control algorithm the maximum admissible load in this case (acc. to the M/D/1 system analysis [4]) is $\rho=0.685$, what is equivalent to 137 kbps.

The foreground traffic was transmitted between PC1 and PC3 terminals (see Fig. 5) while the background traffic was generated only on the access link (between aq_1605_2 and aq_3640_3 routers – see Fig.5). The background traffic was created by a mix of packets with different lengths: 7% with 44 bytes, 21% with 256 bytes and

72% with 1280 bytes. Both foreground and background traffic was transmitted using UDP protocol.

The characteristic of packet loss rate as a function of the TCL-1 traffic volume is reported in Fig.6. One can observe that the measured value of packet loss rate are in the range of 10^{-5} and they are significantly below the assumed target value 10^{-2} even for the load above the admission region (133 kbps). This is rather expected result since the TCL-1 traffic is served with higher priority with the effective service rate of 2 Mbps (instead of 200 kbps as was assumed for the admission control algorithm). Anyway, increasing TCL-1 traffic above the assumed limit (133 kbps) is a bit dangerous since this can degrade the quality experienced by packets carried in low priority classes (e.g. TCL 2). The recommended admissible load of TCL-1 traffic is approximately 10 % of total link capacity.

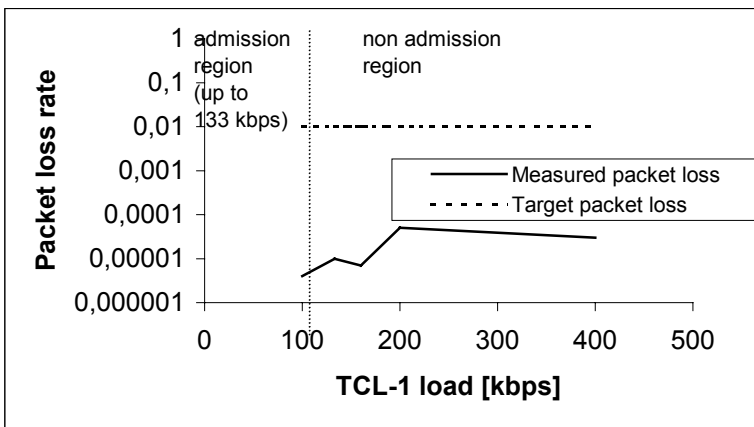


Fig. 6. Packet loss rate vs. TCL-1 traffic load

The characteristics of one-way packet delay as a function of TCL-1 packet length are depicted of Fig.7. These curves were measured assuming that Poisson traffic with the rate equal to 133 kbps (up to admission limit) was submitted to TCL-1. In this case, the background traffic was of ON/OFF type and submitted independently to each intermediate link (see Fig.5), with the peak rates equal to the appropriate link rates. Such type of traffic produces maximum packet delay for the foreground traffic and this is caused by transmission buffer implemented in CISCO routers [10]. The maximum observed delay is below 50 ms. This value is acceptable for applications like for example voice transmission, that tolerates delay in the order of 150 ms, considering that codec and packetisation delay has to be added.




Fig. 7. One-way delay vs. TCL-1 packet length

5.2 Results for TCL-3

The TCL-3 was mainly defined for effective support of TCP-controlled flows. The objective for this service is to guarantee a minimum throughput to the TCP flows. The measurement results reported in this section were obtained assuming that the traffic generated by a number of TCP greedy sources (between PC1 and PC3 – see Fig.5) is submitted to the class in question. The background traffic of CBR type is submitted to other network services with the maximum possible rate corresponding to the assigned capacity for given service. In this way the capacity available for the test TCP connections is limited to the capacity allocated to TCL-3. In the considered case this capacity is 600 kbps with target utilisation factor equal to 0.9 what gives 540 kbps. Notice that the out-of-profile packets in PMM service are not dropped. Consequently, the considered greedy TCP flows can achieve higher rate than requested by the SBR parameters.

The goodput characteristics as a function of best effort traffic for the case of four TCP flows with different declared (and policed) SBR values are depicted on Fig.8. The assigned SBR values were as follows: 135, 135, 70 and 200 kbps. One can observe that the available capacity for TCL-3 is effectively used by the considered TCP connections. When the background traffic exceeds a threshold (in this case 700 kbps) the available capacity is limited to 600 kbps (the capacity dedicated to this service). Moreover this capacity is shared between TCP flows in proportion to their declared SBR rates. For example, the connections with SBR=200 gets 215 kbps while the connection with SBR=70 gets only 95 kbps.

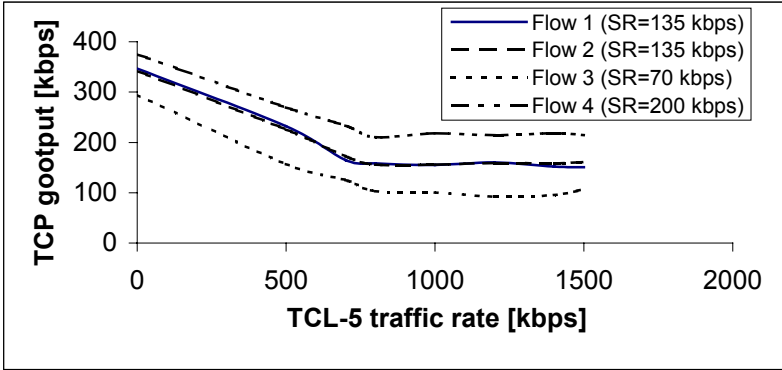


Fig. 8. TCP goodput vs. TCL-5 traffic rate

5.3 Impact of TCL-1 on TCL-3

The impact of TCL-1 traffic on TCL-3 is illustrated on Fig. 9 showing the goodput characteristics of TCP flow 1 and 3 as a function of TCL-1 traffic load. One can observe that increasing the TCL-1 traffic load above the admission limit (133 kbps) decreases the capacity available for TCL-3 and, as a consequence the goodput of TCP flows can be less than the requested rates (SBR values). In the considered case, this effects occurs when TCL-1 traffic is about 230 kbps. This is caused by two factors: (1) there are 100 kbps not assigned to any network service and (2) the target utilisation for TCL-3 is 0.9.

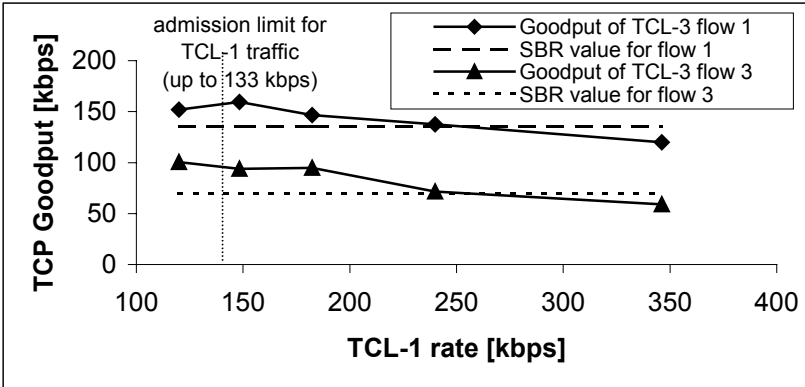


Fig. 9. TCP goodput of flows 1&3 vs. TCL-1 traffic rate

6 Conclusions and Future Work

In the paper the traffic handling mechanisms forced in the AQUILA QoS IP network were described. They correspond to different time scales (from milliseconds to hours or days). The included preliminary measurement results obtained in the AQUILA testbed confirm that by adding new functionalities into the existing IP network there is a possibility to define a set of traffic classes offering differentiated QoS. These experiments were mainly focused on the evaluation of QoS corresponding to the packet and flow level, and were provided for two representative traffic classes (TCL-1 and TCL-3) designated for different types of traffic, non-reactive and reactive.

The following conclusions can be outlined:

- The isolation between different traffic classes can be effectively provided by scheduling mechanisms implemented in the routers.
- The reactive and non-reactive traffic should be definitively submitted to different traffic classes for meeting assumed QoS objectives.
- Different traffic classes require different traffic characterisations and different admission control algorithms.
- The assumed admission control algorithms for the tested traffic classes work according to the expectations.

The currently ongoing work in AQUILA mainly focuses on further trials on the other TCLs (TCL-2 and TCL-4), and on the evaluation of the traffic handling mechanisms at aggregate level (e.g. performances of the Resource Pool dynamics, handling of Reservation Requests).

Future work, for the second phase of the AQUILA project, is to consider inter-domain aspects and to include measurements in the resource control in order to have a feedback from the network performance into the provisioning and resource distribution phase.

References

1. S. Blake et al., "An Architecture for Differentiated Services", Internet RFC 2475, December 1998.
2. B. Davie et al., "An Expedited Forwarding PHB", Internet Draft, draft-ietf-diffserv-rfc2598bis-01.txt, April 2001.
3. Y. Bernet et al., A Conceptual Model for DiffServ Routers, INTERNET-DRAFT, draft-ietf-diffserv-model-02.txt, September 2000.
4. Final report COST 242, Broadband network teletraffic: Performance evaluation and design of broadband multiservice networks (J. Roberts, U. Mocci, J. Virtamo eds.), Lectures Notes in Computer Science 1155, Springer 1996.
5. Final Report COST 257, Impact of New Services on the Architecture and Performance of Broadband Networks (COST 257 Management Committee, Phuoc Tran-Gia, Norbert Vicari eds.), ISBN-Nr. 3-930111-10-1, compuTEAM, Wuerzburg, Germany 2000.

6. Deliverable D1201, System architecture and specification for the first trial, AQUILA project consortium, <http://www-st.inf.tu-dresden.de/Aquila/>, June 2000.
7. Deliverable D1301, Specification of traffic handling for the first trial, AQUILA project consortium, <http://www-st.inf.tu-dresden.de/Aquila/>, September 2000.
8. Deliverable D2301, Report on the development of measurement utilities for the first trial, AQUILA project consortium, <http://www-st.inf.tu-dresden.de/Aquila/>, September 2000.
9. Deliverable D3101, First Trial Integration Report, AQUILA project consortium, <http://www-st.inf.tu-dresden.de/Aquila/>, March 2001.
10. Deliverable D3201, First Trial Report, AQUILA project consortium, Draft version, April 2001.
11. F. Bricet et al., Stochastic ordering and the notion of negligible CDV, Proc. of 15th International Teletraffic Congress, Washington D.C., USA, 1996.
12. K. Lindberger, Dimensioning and design methods for integrated ATM networks, Proc. of 14th International Teletraffic Congress, Antibes, 1994.
13. A. Elwalid et al., A new approach for allocating buffers and bandwidth to heterogeneous, regulated traffic in an ATM node, IEEE Journal on Selected Areas in Communications, p. 1115-1127, 1995