Analytical approach and verification of a DiffServ-based priority service

Christos Bouras, Afrodite Sevasti

Research Academic Computer Technology Institute-CTI, 61 Riga Feraiou Str., 26221 Patras, Greece
Department of Computer Engineering and Informatics, University of Patras, 26500 Rion, Patras, Greece
{bouras, sevastia}@cti.gr

Abstract. The provision of Quality of Service (QoS) in a seamless way over the dominating internetworking protocol of our times (IP), has been a challenge for many researchers in the past years. Strict qualitative guarantees have proven difficult to provide in a way that has discouraged efforts in the area. The lack of a coherent provisioning methodology has been identified as the main reason for this. In this work, we are attempting an analytical yet straightforward approach to the provisioning methodology proposed for premium service of high-quality demanding traffic in the wide-area. Our approach is based on a series of well-known results of queuing theory but is proven to provide good approximations to experimental results as well as worthwhile qualitative guarantees.

1. Introduction

The DiffServ framework has been designed for the provision of QoS services in large-scale environments, where the extensive aggregation of flows does not allow solutions for QoS provisioning on a per-flow scale. At the same time, DiffServ is a framework and as such, it provides guidelines rather than strictly defined service models. Network designers have, thus, at their disposal multiple different individual mechanisms and alternatives for QoS services' implementation. This results in a variety of solutions, the compatibility and relevant importance of which can rarely be pinpointed. Our work attempts to approach analytically and evaluate a set of guidelines for the provision of a high-priority service to quality-demanding traffic in wide-area networks, while at the same time ensuring deterministic upper bounds to critical metrics, such as one-way delay.

Dimensioning and provisioning QoS on the DiffServ framework basis has been the topic of many research initiatives in the last years. Some of them focus in DiffServ-based QoS service models and definitions ([4],[6],[9]) while others perform experimental studies on individual mechanisms and modules of DiffServ-based QoS services ([5],[7],[8]).

Also a number of analytical approaches have been presented in the latest years, some of which are briefly presented here. In [13], the authors are proposing the use of a statistical rather than deterministic analysis for the estimation of loss rate and delay guarantees along a multistage network for packets served in aggregates by a non-preemptive priority service. The main reason for this is that deterministic analysis seems to be too pessimistic while statistical analysis allows for better approximation of quality metrics in the expense of statistical guarantees. The results of this work that are relevant to the analysis made here is the validity of the M/D/1 queuing model in the analysis of queues serving aggregates of flows with non-pre-emptive priority. Delays and buffer saturation probabilities are estimated using a Poisson stream of MTU-sized packets for each such flow aggregate and end-to-end delay guarantees are estimated with respect to the waiting time in queues of packets along a multistage network.

In [14], the authors are proposing a closed loop edge-based framework for flow control having as an ultimate goal, that of providing QoS. Their approach is also focused on the accumulation of packets at each router attributed to each flow for a sequence of routers. However, this is in the framework of devising a flow control mechanism based on the operation of TCP and aiming at preserving queue occupancy (and thus queuing delays) to desired levels for a certain level of quality.

In this work, we are aiming at the analytical modeling of the situation imposed when a high-priority service is introduced to a best-effort IP network. We are also combining a service definition with a set of mechanisms, towards an integrated solution for QoS provisioning to quality-demanding traffic, with bounded quality guarantees. With respect to the related work done so far, the contribution of this work is two-fold: we initially provide an analytical model for dimensioning and provisioning a high-priority service to IP multimedia traffic over a backbone topology. In the sequel, we use simulation in order to verify our analytically supported provisioning model.

In section 2 of this paper, the model used for both the analytical and simulation work is provided. In section 3, the analytical approach of DiffServ-based QoS service is provided. In section 4, a set of scenarios verifying the analysis of section 3 and demonstrating the performance of the service are provided. This paper concludes with our intended future work on this topic and conclusions drawn from the work carried out so far.

1.1. The Model

The case that will be further investigated in this work is that of dimensioning and providing a high-priority, low latency QoS service for aggregated traffic in a wide-area network. This service, referred to as 'Gold' service from now on, is built according to the Expedited Forwarding Per-Hop-Behavior ([1]) of the DiffServ framework. The Gold service aims at offering the equivalent of an end-to-end virtual leased line (VLL) service at the IP layer across multiple domains. In general, traffic requesting for such a service is sensitive to delay, jitter and packet losses, so the proposed service and provisioning method will be evaluated against these metrics.

For each of the customers (C_i) an appropriate SLA that specifies the characteristics (traffic envelope) of the marked as Gold traffic injected by C_i into the network is required. The Gold service then guarantees a specific rate of service for each aggregate, a specific bounded end-to-end delay (D), bounded jitter and minimal packet losses to all the legitimate aggregates served by it. Legitimate aggregates are those the traffic envelope of which is examined (and enforced if necessary) by the use of a policer for each one of them, as depicted in Figure 1.

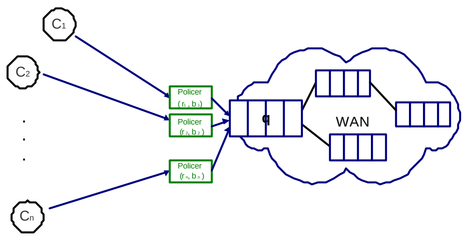

Fig. 1. Enforcing token bucket profiles to Gold service customers

One of the most common policers used and supported by commercial network equipment is that of a token bucket (r, b), that imposes conformance to an average rate and maximum burst size to the incoming traffic. Gold traffic aggregates are served by priority queues (like q of Figure 1) along the wide area network. The aim of the analytical and experimental study that follows is to investigate rules and principles for the Gold service provisioning and the quality achieved for Gold traffic.

2. Analytical Study

For the Gold service provisioning, it is proposed to use appropriate configuration of all the routers comprising the network in order to ensure the existence of proper buffering space and for the Gold traffic to be served with the required quality.

The basic principles on which the service dimensioning is based are:

- The configuration of routers so that Gold traffic will receive at least a strictly bounded minimum service rate at each router, regardless of the load imposed on the router
- The policing of the total of Gold traffic entering each Point-of-Presence (PoP) is such that the arrival rate of the aggregated Gold traffic entering the PoP router, created by the merging of all Gold aggregates of customers attached to a specific PoP does not exceed the minimum rate of service that Gold traffic receives at each egress interface of the router (according to [1]).
- Based on these principles, the goal of Gold service provisioning would be:
- To assure that each flow participating in a Gold traffic aggregate originating from a customer attached to a PoP demonstrates a throughput equal to or larger than that of entering the network as it exits the network.
- Delay and jitter for the packets of this flow are both bounded between the points where the packets enter and exit the network, while the flow perceives negligible or zero packet losses

2.1. Gold Service Dimensioning

As explained extensively in [15], Gold service provisioning is based on resources' over-provisioning and therefore the provision of capacity guarantees can be obtained by ensuring that for the total capacity C_q appointed to each router queue serving Gold traffic the aggregated Gold traffic injected into this queue is only a percentage of C_q. As it was shown in [15], this means that

$$\sum_i r_i \le kC_q \Rightarrow \sum_i l_i \times r_{access} \le kC_q \Rightarrow \sum_i l_i \le \frac{kC_q}{r_{access}} . \quad (1)$$

Factor k determines the over-provisioning factor for the implementation of the Gold service on an outbound interface of an ingress router while r_{access} determines the capacity of the interface through which each customer is attached to the WAN for each customer C_i (assumed to be constant for all customers for simplicity), r_i the token bucket profile rate with which the Gold aggregate injected by customer i to the WAN is policed and l_i the percentage of r_{access} that can be used for Gold traffic for each customer C_i. It is important to stress out at this point that the value of l_i or r_i is one upon which the customer can be charged for receiving the service.

If S_q is the set of customers that are able to inject a Gold traffic aggregate through ingress router queue q, there is also a limitation to the value of C_q, which should not exceed a percentage of the capacity of the WAN backbone links. This means that

$$C_q \le m \times C, \text{ with } m \in (0,1) . \quad (2)$$

so that under heavy load, best effort traffic will not face denial of service. From (1) and (2) it follows that

$$\sum_{S_q} l \times r_{access} \le k \times C_q \le k \times m \times C \Rightarrow l \le \frac{m \times C}{|S_q| \times r_{access}} \times k . \quad (3)$$

where l_i has again been considered equal to l for all customers for simplicity. Thus, the percentage of Gold capacity on customers' access links to the WAN is limited for each customer according to k, the over-provisioning factor of Gold traffic on the backbone. In the experimental analysis to follow it will be demonstrated how k affects perceived quality by the traffic served with the Gold service and will provide relevant guidelines. Recommendations from related work ([10]) and other experimental evidence ([11]) lead to the selection of small values for k, in the range of $(0.05, 0.2)$. After selecting the over provisioning factor, the analysis above can be used to determine the values of l_i's in a Gold service provisioning scenario.

Apart from transmission rate guarantees, bounded end-to-end delay can be ensured for the whole of Gold traffic by selecting the token bucket policing profile for each aggregate i served by a ingress router so that the worst-case delay that Gold traffic packets will face during their stay in each router j will not exceed a maximum value $D_{\max j}$. Again here the over-provisioning factor seems to play a crucial role. As outlined again in [15], the guaranteed bounded delay that a packet can face from the moment it enters until it exits the WAN is denoted by:

$$D_{tot} \leq \sum_{j} D_{\max j} + \sum_{n} D_{prop_n} + \sum_{n} D_{trans n} . \quad (4)$$

where D_{trans_k} is the transmission delay and D_{prop_k} is the propagation delay of the packet on transmission line n. Equation (4) adds up all delays that a packet faces on each router j and every transmission line it crosses. The following analysis deals with the value of $D_{\max j}$, since the values of D_{trans_k} and D_{prop_k} cannot be affected by the implementation of a service over existent equipment/transmission lines. Also for clarity purposes, the analysis that will follow deals only with the metric of end-to-end delay. However, the experimental approach that will follow anticipates for jitter as well. Under the assumption of Poisson traffic (exponential inter-arrival times) for both Gold and best-effort (BE) aggregates, assuming a general service time distribution and a non-preemptive priority scheduler providing Gold traffic with the highest priority, $D_{\max j}$ on router j for Gold packets can be expressed as:

$$D_{\max j} = \frac{R}{1-k_j} . \quad (5)$$

where k_j is the service rate for Gold traffic in j (arrival rate to service rate ratio, see also (1)) and R is the mean residual time in the router, given by:

$$R = \frac{1}{2}\, l_{Gold} \overline{X^2_{Gold}} + \frac{1}{2}\, l_{\mathrm{BE}} \overline{X^2_{\mathrm{BE}}} . \quad (6)$$

where l_{Gold} is the arrival rate of Gold traffic to the Gold traffic queue ($\sum_i r_i$ for an ingress router in (1)) and $\overline{X^2_{Gold}}$ is the second moment of service time for traffic entering the router (corresponding values for best-effort traffic are denoted similarly, but for simplicity without losing generality one can assume that $\overline{X^2_{Gold}} = \overline{X^2_{BE}} = \overline{X^2}$).

One immediate observation from (6) is that arrival to service ratio (k_j) for the Gold queue has to be sufficiently limited (a strong indication for over-provisioning requirements) so as to make $D_{\max j}$ limited on the end-to-end path and thus achieve the quality required in terms of end-to-end delay. Actually, arrival to service ratio (k_j) and the arrival rate ($\sum_i r_i$) are the only parameters in the disposal of the Gold service designer in order to achieve bounded queuing delay. Based on a known service time distribution (e.g. deterministic distribution for VoIP traffic comprising of constant size IP packets), the value of $\overline{X^2}$ to be used in (6) can be determined. The Gold service designer can proceed in determining R from (6) and $D_{\max j}$ in (5) and finally providing an end-to-end guarantee via (4).

From the analysis made above it has become apparent that the selection of the policing profile (especially the value of r_i) of each Gold aggregate injected by a customer to a WAN PoP and the provisioning factor or utilization at each router queue serving Gold traffic can affect the delay guarantees provided by the WAN to its Gold service customers. In Figure 2, an indicative comparison of the queuing delay perceived by an experimental set-up and the corresponding theoretical bound for a single router queue is provided.

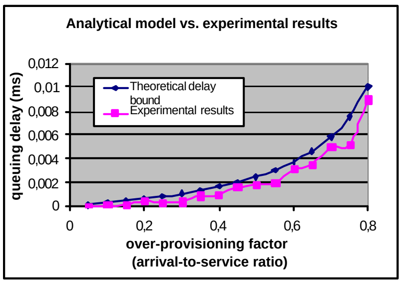

Fig. 2. Comparative presentation of theoretical and experimental queuing delays

Inversely, provided that specific need for bounded delay exists, e.g. 40ms for VoIP traffic crossing the network one can determine the token bucket profiles that can be offered to SLAs with customers.

3. Experimental Study

In this section, a series of experiments that validate the service provisioning principles already outlined and investigate the different alternatives for Gold service provision are described. These experiments aim at:

- Demonstrating how appropriate provisioning of the Gold service affects the quality guarantees provided to eligible traffic and providing useful guidelines wrt the overprovisioning required

- Studying how well the theoretical bounds are approximated by experimental environments set up in a variety of topologies and scenarios. Guaranteed delay has been our main concern here

The components of the simulation environment that have been used are analytically presented in [12]. The experiments carried out aim at investigating the Gold service provisioning on the backbone of a WAN, comprised by a number of PoPs, to which a number of customers with different access capacities are attached. Gold traffic on the backbone links co-exist with background traffic served as best effort and inserted to the topology as cross-traffic ([2],[3]). An introductory work of the experiments presented here are provided in [15]. For all experiments the value of MTU has 1500 bytes. The scheduling mechanism simulated at the routers has been that of Modified Deficit Round Robin of Cisco GSRs.

3.1 Determining the Over-provisioning Factor for Gold Service

In this section a series of experiments that investigate the efficient over-provisioning factor (parameter k in (1)) values through different scenarios are conducted. This work aims at adding to related existent work, by considering k as an input parameter to the Gold service provisioning methodology and examining how well the experimental results approach the theoretical ones for different values of k. For the experiments presented, Poisson modeling of traffic has been used. The experimental setup was this of Figure 3.

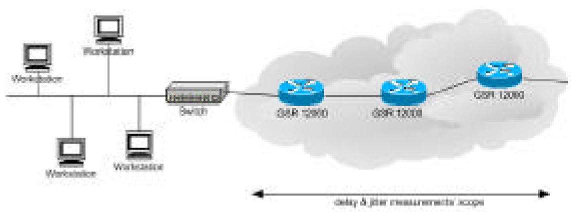

Fig. 3. Over-provisioning factor monitoring testbed

The workstations were configured to transmit VBR traffic over an Ethernet LAN and the aggregated traffic was then injected to a backbone of Cisco GSRs, using MDRR-ALT scheduling. The configuration imposed on each GSR was such that the over-provisioning factor k in $\sum_i r_i \leq kC_q$, where r_i are the token bucket rate policers configured in the ingress interface of the entering backbone GSR for each one of the workstations ranged from 97% to 30%, for the different cases tested. In all of the three cases, it holds that $\sum_i r_i < C_q$, so that Gold traffic is always over-provisioned with a varying over-provisioning grade. In Figure 4, the throughput measured between the measurement points of the simulated topology is presented. It is obvious how over-provisioning factors above the level of 80% affect the throughput perceived by Gold traffic.

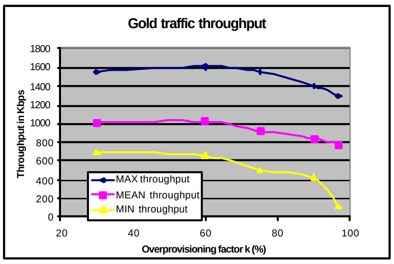

Fig. 4. Throughput measurements for the topology of Figure 3

A second experiment with different levels of aggregation has also been conducted according to the topology of Figure 5. Here, the same traffic is inserted to the backbone cloud, however, not a single point of entry is used. All of the routers are configured so that $\sum_i r_i \leq kC_q$ at each router, however in this experiment, $\sum_i r_i$ is different for each router and therefore the capacity reserved for the Gold traffic C_q has been adjusted accordingly.

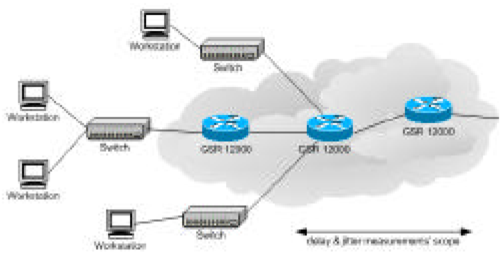

Fig. 5. Over-provisioning factor monitoring testbed with different aggregation levels

In this experimental setup, the points of throughput measurement are shifted to a scope where all of the Gold traffic flows are aggregated (i.e. the last hop of the packets' route in the backbone network). As can be seen from Figure 6, the throughput achieved is now more smoothed wrt the different over-provisioning factors used and this is explained by the fact that traffic is injected gradually to the network and the router policing mechanisms provide a better statistical spreading of Gold traffic.

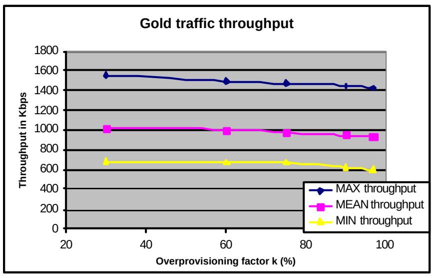

Fig. 6. Throughput measurements for the topology of Figure 5

3.2. Verification of Theoretical Delay Guarantees

In this section, a comparison of the theoretical model with experimental results is presented for the case of delay guarantees. The results of the previous two section, as far as the over-provisioning factor k is concerned are taken into consideration. However, the focus here is on how well delay guarantees are achieved, always in comparison with the theoretical expected values, for more realistic scenarios than that indicatively presented in Figure 2 of section 3.1.

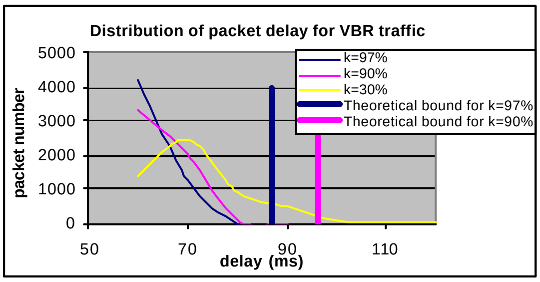

Fig. 7. The effects of MDRR-ALT scheduling in the delay perceived by VBR traffic

The measurements presented here have been performed over the topology of Figure 3. As can be seen from Figure 7, the experimental results for the case of high over-provisioning are confirmed by the theoretical bounds calculated by (4),(5),(6).

4. Future Work and Conclusions

Most of the work presented here has been based on Poisson traffic modeling. Part of our future work will consist of applying the methodology proposed to different traffic mixes, taking into more recent findings, such as that of self-similar traffic. Our future work on the Gold service will also concentrate on an analytical approach for the jitter guarantees that can be provided for a service such as that of Gold. We also intend to further deploy and evaluate the proposed algorithm for the dimensioning of the Gold service that, according to the analysis made here, will be useful for network administrators in order to introduce the service to their networks.

The work presented in this paper aims at providing an analytical estimation of the qualitative guarantees that can be provided by a high-quality service to IP traffic over a wide area network. We also provide guidelines for dimensioning the service by estimating worst-case bounds on quality metrics. A brief presentation of a simulation set-up used to simulate Gold service provisioning to quality demanding traffic together with how that verified the preceding analysis has been provided.

References

[1] V. Jacobson et al., An Expedited Forwarding PHB, IETF RFC 2598 (1999)
[2] C. Dovrolis, D. Stiliadis and P. Ramanathan, Proportional Differentiated Services: Delay Differentiation and packet Scheduling, IEEE/ACM Transactions in Networking (2002)
[3] M. Yuksel, B. Sikdar, K.S. Vastola, B. Szymanski, Workload generation for ns Simulations of WANs and the Internet, Proceedings of Communication Networks and Distributed Systems Modelling and Simulation Conference, San Diego, CA (2000)
[4] M. Goyal, A. Durresi, C. Liu and R. Jain, Performance Analysis of Assured Forwarding, Internet Draft (draft-goyal-diffserv-afstdy-00) (2000)
[5] I. Yeom, N. Reddy, Impact of marking strategy on aggregated flows in a differentiated services network, proceedings of IWQoS workshop (1999)
[6] S. Shalunov, B. Teitelbaum, Qbone Scavenger Service (QBSS) Definition, Internet2 Technical report, Proposed Service Definition, (2001)
[7] M. May, J-C. Bolot, A. Jean-Marie and C. Diot, Simple performance models of tagging schemes for service differentiation in the Internet, Proc. Infocom '99, New York (1999)
[8] TF-TANT Task Force, Tiziana Ferrari, Editor, Differentiated Services: Experiment report, Phase 2 (2000)
[9] M. Campanella, Implementation Architecture specification for the Premium IP service, Delivrable D2.1-Addendum 1, SEQUIN Project (IST-1999-20841) (2001)
[10] Y. Le Boudec and P. Thiran, 'Network Calculus: A Theory of Deterministic Queuing Systems for the Internet, LNCS 2050, Springer Verlag (2001)
[11] C. Bouras, M. Campanella, M. Przybylski, A. Sevasti, 'QoS and SLA aspects across multiple management domains: The SEQUIN approach, Journal of Future Generation Computer Systems, 939, Elsevier Publishing, pp. 1-14 (2003)
[12] C. Bouras, D. Primpas, A. Sevasti, A. Varnavas, Enhancing the DiffServ architecture of a simulation environment, Sixth IEEE International Workshop on Distributed Simulation and Real Time Applications ((DS-RT 2002), Texas, USA, (2002)
[13] T. Bonald, A. Proutiere, J. Roberts, Statistical Performance Guarantees for Streaming Flows using Expedited Forwarding, Proc. Infocom 2001, pp. 1104-1112 (2001)
[14] D. Harrison, Y. Xia, S. Kalyanaraman, K. Ramachandran, An edge-based Framework for Flow-Control, submitted for publication
[15] C. Bouras, A. Sevasti, Deployment of a DiffServ-based priority service in a MAN/WAN environment, The International Conference on Advanced Information Networking and Applications (AINA), Xi'an, China, (2003)