Service Level Agreement Enforcement for Differentiated Services

Paulo Rogério Pereira
INESC ID, Rua Alves Redol, 9. 1000-029 Lisboa, Portugal.
Phone: +351-213100345. Fax: +351-213145843. Email: prbp@inesc.pt

I. INTRODUCTION

This paper describes a hierarchical architecture of active policies [1] that performs the management of a differentiated services (DiffServ) network. These policies monitor quality of service (QoS) parameters and dynamically optimize some aspects of the existing equipment and services to provide the best possible QoS to users. This helps the enforcement of Service Level Agreements (SLA) between provider and users.

The results presented were obtained with NS2 *network simulator* [2], for a 15-node network, with three traffic classes: Expedited Forwarding (EF), Assured Forwarding (AF) and Best Effort (BE). The results show that the use of active policies improves the service offered to users, by constantly adapting to the network state, and helping to fulfill SLAs with minimum costs to the service provider.

II. SYSTEM ARCHITECTURE

Figure 1 shows the active policies deployed and how they are related to each other.

Basically, active policies act over the network, when certain thresholds on the QoS are crossed, to avoid SLA violation. Examples are: modifying Random Early Detection (RED) parameters for increasing packet discard, especially for out-of-profile packets; restricting the admission of new flows into the system; activating extra lines or increasing the bandwidth available for some paths; and selectively degrading the class of certain flows.

On the other hand, if the QoS is better than specified in the SLA, the reverse actions are possible to improve the service provider profits. For instance, allowing some more flows will allow the provider to charge more. To ensure stability, the actions in the direction of improving the QoS are more intense than the actions in the direction of degrading the QoS.

The different levels in figure 1 correspond to different abstraction levels, allowing a refinement of the QoS requirements. At the Equipment level, policies are local to each device. The Network level is sub-divided in two sub-levels to ease implementation and improve scalability. The Service level is the first that sees end-to-end QoS parameters.

As is it difficult to specify the exact thresholds where the policies should operate, a planning policy makes long-term adjustments to ensure a more precise SLA fulfillment.

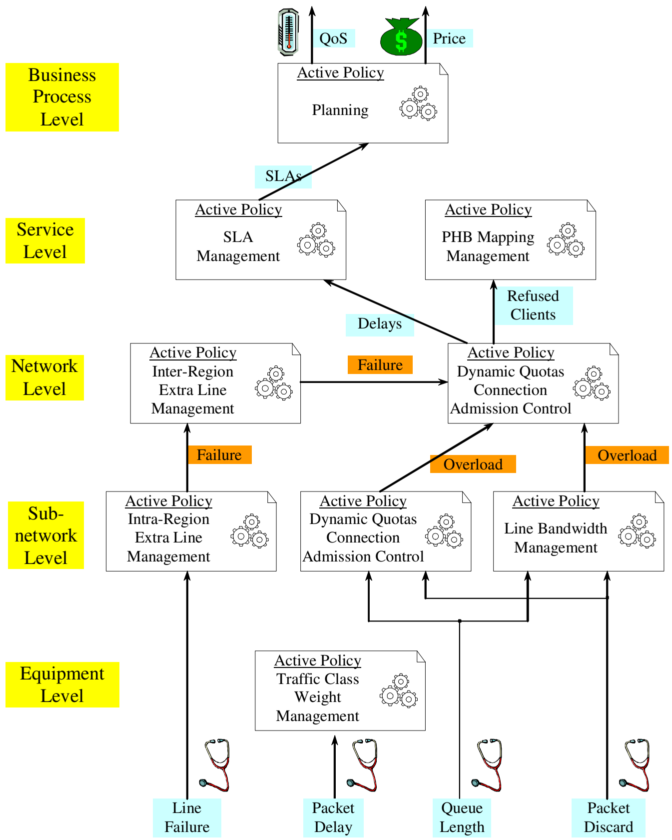

Fig. 1 Active Policies used

A. SLA Management without Planning

The SLA management active policy, at service level, monitors the end-to-end delays for the EF and AF traffic classes, acting when these delays approach the SLA limit, by adjusting parameters from lower level policies that do not have this end-to-end view.

Figure 2 shows the average end-to-end delays obtained with the following QoS objectives:

SLA(*delayEF*) = 60 ms and *SLA*(*delayAF*) = 120 ms, and all policies operating, except the planning policy. These results show it is possible to enforce the SLAs with some error. The error from the SLA objective is about 12-32% without the policies and about 7-23% with the policies. This shows it is difficult to set the exact policies operation points.

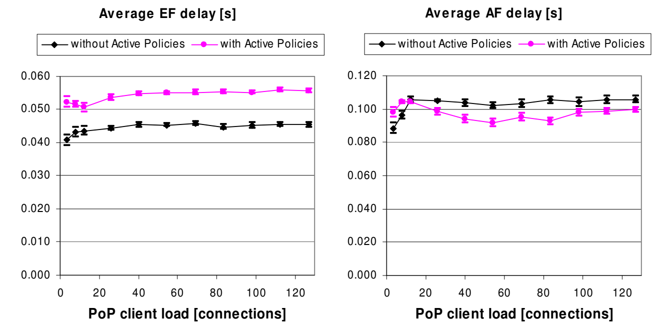

Fig. 2 Average EF and AF delay

B. Planning Policy

The top level planning active policy adjusts the SLAs of the SLA management active policy to obtain better long-term results.

Now, the previous delay objectives, as required in the SLA, are used to set the *SLA*(*delayEF*) and *SLA*(*delayAF*) to be used by the SLA management active policy. Both are modified by an integral feedback system as shown in figure 3. The integrator gain is set to 0.25, to make slight changes in the SLAs for each sampling period. The sampling period is large as compared to the network delays, but much smaller than the SLA verification period. An anti-windup mechanism was added to limit the SLA values to a range of 0.25 to 4 times the required SLA.

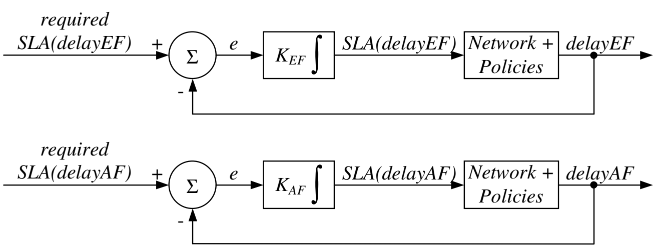

Fig. 3 SLA control system

Figure 4 shows the average end-to-end delays obtained with this policy, as compared with those without policies. These graphics show that the static error from the SLA objective is greatly reduced by the control system to about 1-8%, so that the EF delay is much nearer the required value of 60 ms, and the AF delay of 120 ms. There were no packet drops both for EF and in-profile AF traffic.

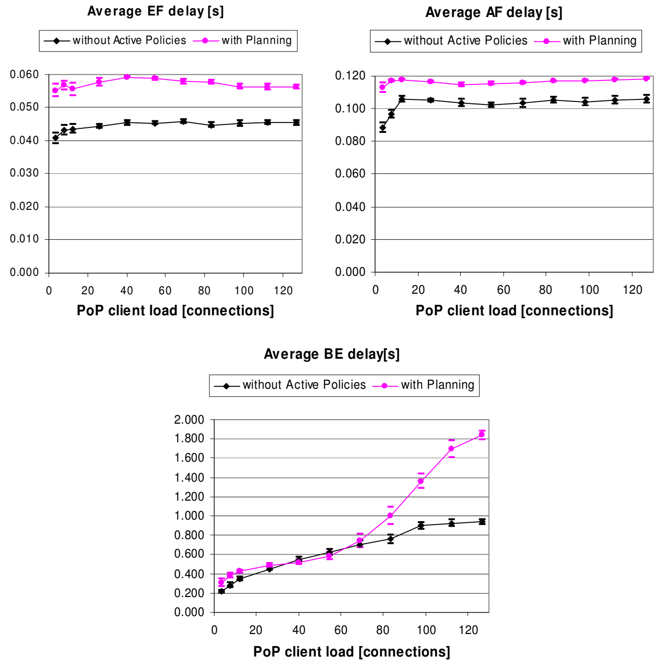

Fig. 4 Average delay per traffic class

III. CONCLUSION

It can be concluded that the QoS specified in SLAs can be provided within certain limits by using active policies that dynamically adjust network parameters to optimize network performance.

The planning policies provide more precise QoS results circumventing the difficulty in setting policy operating points.

The active policies can be reused for different situations, just by changing the variables they control and their configuration parameters.

IV. REFERENCES

[1] P. Pereira, D. Sadok, P. Pinto, "Service Level Management of Differentiated Services Networks with Active Policies", in *ConfTele'2001 proceedings*, Figueira da Foz, Portugal, April 2001, pp. 542-546. ISBN: 972-98115-2-0.

[2] UCB/LBNL/VINT Network Simulator (version 2). http://www.isi.edu/nsnam/ns/