On local CAC schemes for scalability of high-speed networks

Javier Aracil, Senior Member IEEE, Jose A. Hernández, Antonio J. Elizondo, Raúl Duque, Óscar Gonzalez de Dios*
Universidad Autonoma de Madrid, EPS, C/ Fco. Tomás y Valiente, 11, 28049 Madrid
** Telefónica I+D, Emilio Vargas, 6, 28043 Madrid*
Tel: +34 91497 2272, Fax: +34 91 497 2235, e-mail: javier.aracil@uam.es

ABSTRACT

Next generation networks are required to provide bandwidth on-demand for fine granularity sessions. In this sense, centralized CAC (Connection Admission Control) approaches could suffer from scalability problems if the number of requests for connections were very high. In this paper we investigate local CAC schemes where the decisions are performed at the network edges, based on pre-calculated admission quotas.

Keywords: CAC, distributed control, high-speed networks.

1. INTRODUCTION

Call admission control is an essential functionaly in networks supporting Quality of Service (QoS). In this paper, we report about distributed CAC schemes that serve to alleviate the scalability and responsiveness problems of their centralized counterpart. Actually, currently proposed Call Admission Control methods for most next generation networks are centralized, thus complicating matters for scalability and robustness. On the other hand, distributed CAC techniques show a better behaviour in terms of reliability. However, the CAC becomes more sophisticated and difficult to implement and manage, primarily because resource brokerage is now carried out on an individual basis in many different CAC units. Consequently, such CAC units must be either notified about the availability of resources network-wide, so as to make a decision in a coordinated, yet decoupled, manner or must be allocated a certain share of the resources to be managed in a completely independent fashion. In this paper we follow the latter approach, proposed in [1] by the MUSE IST project, in which the CAC units are assigned a certain share of the resources beforehand, which are updated on-demand. In what follows, we shall refer to this technique as a “local” CAC.

This study provides a performance evaluation of the local CAC in terms of capacity planning for the network control plane and bandwidth efficiency. On the other hand, we also assess the responsiveness of such local CAC to the input traffic demands. It follows immediately that the more frequent the resource allocation updates are, the more closely the CAC controller follows the input demand. However, this may lead to undesired oscillations of network resources, together with a significant signalling overhead. In order to tackle these issues we first evaluate the dynamic behaviour of a single CAC controller. Our aim is to find the rate of the resource allocation updates such that the input demand is satisfied with a certain probability.

1.1 Application scenarios

Local CAC is originally thought to be used at the border of multi-service networks that must deal with a huge number of service requests that require QoS guarantees, as a centralized approach for handling these requests may show important scalability problems.

The MUSE IST project (see [1]) proposes the usage of local CAC in the (aggregation network) access nodes (DSLAMs, OLTs, etc) for a subset of services that have stringent QoS requirements such as IPTV. Note that the implementation of a local CAC requires being locally able of handling service signalling in order to be able to make the user aware of the CAC decision result (acceptance or rejection). In this way, local CAC can easily handle multicast traffic (e.g. IPTV), as it is possible to be locally aware of IGMP messages in most of current IP DSLAMs.

The proposed technique is also suitable to be employed in the delivery of VoD services. Within this scenario, there are unicast connections established between the video server and each one of the end users. So, with the application of the “local” CAC in the video server, it is possible to make a local decision about the network resource availability for the desired route according the assigned quota for the corresponding destination, stored in the CAC module of the video server.

For other services such as VoIP, local CAC may be used provided that the access nodes were able to participate in the signalling process. In this way, the proposal of distributing SBC capabilities ([2]) into the access nodes would allow to locally take the CAC decisions.

A long term application scenario consists of local CAC at the edge of an Optical Burst Switching (OBS) Network. In this light, edge nodes would have pre-established quotas per destination and class of service, assuring that the load inside the network is controlled, avoiding harmful congestion.

2. Analysis

We provide capacity planning rules for the local CAC controller, based on the input demand (traffic matrix). Secondly, we perform an experiment using real traffic traces from the Spanish NREN, in order to show the suitability of the results in a real network scenario.

2.1 Capacity planning for the local CAC controller (packet case)

Let us first assume a packet network case in which traffic is expressed as a Gaussian fluid in the graph (V,E) where V is the set of nodes and E the set of links. The input demand is expressed as a matrix **T**, with V rows and V columns and the uplink and downlink traffic to a certain node *i* is expressed as the sum of the entries of row *i* and column *i* respectively. In what follows we will derive the capacity planning rule (quota assignment) for local CAC *i* assuming that the average traffic intensity is *I* in bits per second and the standard deviation is given by σ. The capacity planning problem is stated as follow, find *C* such that

$$P(X > C) < \varepsilon$$

where X is the offered traffic, resulting in the rule $C = I + \sigma Z_{1-\varepsilon}$, where $Z_{1-\varepsilon}$ is the $1-\varepsilon$ percentile of a standard Gaussian random variable. The following figure shows the value of C, n times the average input rate, for a coefficient of variation σ / I equal to 1 (bursty traffic) and equal to 1/3 (highly multiplexed traffic).

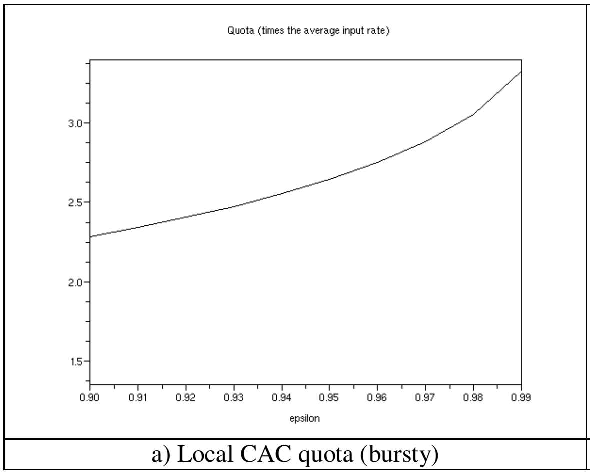

a) Local CAC quota (bursty)

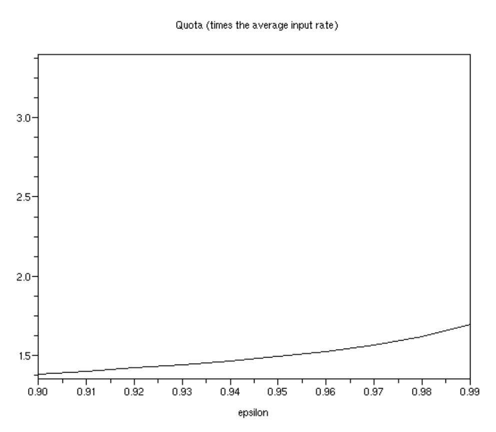

b) Local CAC quota (highly multiplexed)

It turns out that in a bursty case (exponential variability) the local CAC quota should be three times the input demand, for a desired quality of service level larger than 97% approximately. On the other hand, the above results show that the *increase of the quota is linear with the demand increase, if the coefficient of variation remains constant.* For example, if the ratio (standard deviation)/mean is constant and the traffic demand doubles, then the quota simply doubles. If, on the other hand, the standard deviation is left constant, then the quota decreases, as the traffic becomes less bursty. The following graph shows a case for constant standard deviation, and a traffic intensity that doubles.

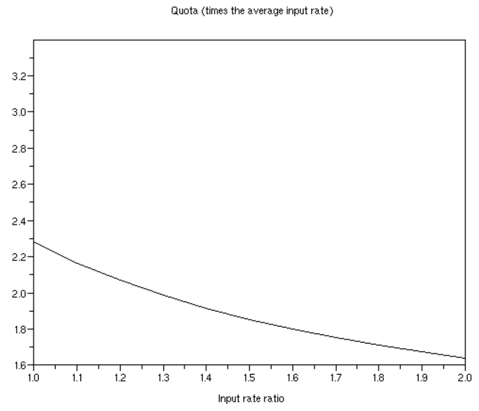

a) Local CAC quota with traffic intensity that doubles and $\varepsilon = 0.9$. Variance is left constant.

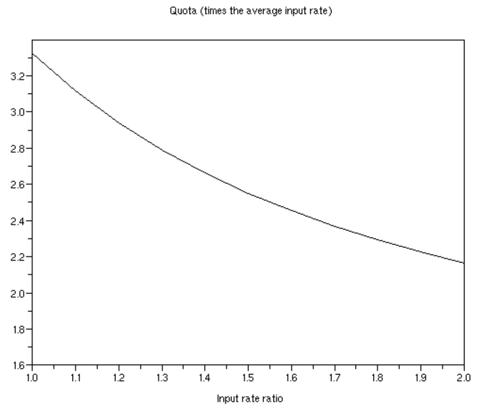

b) Local CAC quota with traffic intensity that doubles and $\varepsilon = 0.99$. Variance is left constant.

3. A case study with real data

In this section, we consider a case study with real data from the Spanish academic network. We collect the busy hour traffic in the access links of four universities, two of them being small and the other two large. The following figure shows the topology of the Spanish academic network, along with the experimental setup.

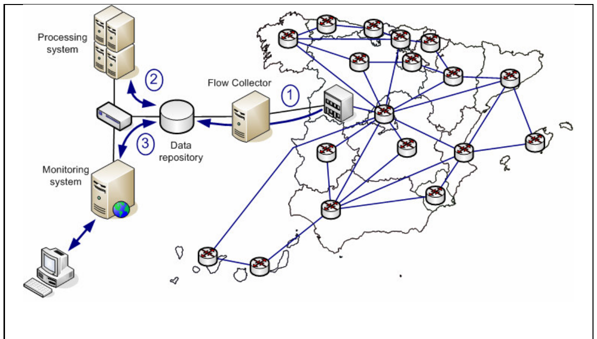

Measurement testbed in the Spanish academic network

We consider that each of the access links is controlled by a local CAC and proceed with the Gaussian capacity planning rule. The results are shown in the following table.

University	Mean (Mbps)	Standard deviation	Quota ($\varepsilon = 0.99$)
U1	32	5.6	45.02
U2	8	2.6	14.04
U3	206	7.6	223.68
U4	206	35.6	288.81

It turns out that the variability of the busy hour is relatively small in a real case, thus making the local CAC scheme amenable for use, as the resulting quotas are not too large in comparison with the real demand.

4. CONCLUSIONS

In this paper we have analyzed a case of local CAC. The analysis has been performed both from an analytical and from an experimental point of view, including real traffic traces from the Spanish NREN. Our results show the advantages of the local CAC mechanism, with resulting quotas that are close to the average traffic demand per CAC controller.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the BONE Network of Excellence, partially funded by the European Union Seventh Framework Programme, the MUSE project, partially funded by the EU Sixth FP, and the RUBENS project as part of the EUREKA CELTIC initiative.

REFERENCES

[1] Elizondo et al, MUSE QoS Architecture White Paper (https://www.ist-muse.org/Abstracts/abstract_whitepaper_QoS.htm)
[2] De Vos et al, MUSE deliverable DTF1.7 Multimedia Support in Access Architecture (https://www.ist-muse.org/Abstracts/abstract_DTF1.7.htm)
[3] John T. Kent and Nicholas T. Longford, An eigenvalue decomposition for first hitting times in random walks, *Probability Theory and Related Fields*, Springer Berlin,Volume 63, Number 1, march 1983