

On local CAC schemes for scalability of high-speed networks

Javier Aracil, José Alberto Hernández

Dept. Ing. Informática, Universidad Autónoma de Madrid, Spain

Email: {javier.aracil,jose.hernandez}@uam.es

Antonio J. Elizondo, Raúl Duque, Óscar González de Dios

Telefónica I+D, Spain

Email: {ajea, ogondio}@tid.es, rduque@entel.es

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

Index Terms—Connection Admission Control; Distributed control; High-speed networks.

I. INTRODUCTION

Call Admission Control (CAC) is an essential functionality for networks supporting Quality of Service (QoS). In this paper, we study distributed CAC schemes that serve to alleviate the scalability and responsiveness problems observed in 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 on many different CAC units. Consequently, such CAC units either need to be 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, initially proposed in [1] by the MUSE IST project¹, and later followed in RUBENS 4, 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 “local” CAC.

This article has been extended from the original: “On local CAC schemes for scalability of high-speed networks” by the same authors, which was presented on the International Conference on Transparent Optical Networks (ICTON) held in Athens, Greece, on June 22-26, 2008

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.

¹Multiservice Access Everywhere Project: Website: <http://www.ist-muse.org/>

This study further provides a performance evaluation of such a local CAC strategy 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 occur, 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.

A. 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, it may show important scalability problems.

The MUSE IST project (see [1]) proposes the usage of local CAC in the (aggregation network) access nodes, such as the Digital Subscriber Line Access Multiplexers (DSLAMs) or the Optical Line Terminals (OLTs) for a subset of services that have stringent QoS requirements such as IPTV traffic. 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. IP Television), as it is possible to be locally aware of Internet Group Management Protocol (IGMP) messages in most of current IP DSLAMs.

The proposed technique is also suitable to be employed in the delivery of VoD services. In 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 Session Border Control (SBC) capabilities [2] into the access nodes would allow to locally take the CAC decisions.

RUBENS proposes to take into account several parameters to accept the request of a new application (it can be a game session, VoD service, etc). One of these inputs to make such decision is the availability of bandwidth for a session of such application, which may be performed in a scalable way by means of the local CAC.

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.

II. ANALYSIS

We provide capacity planning rules for the local CAC controller, based on the input demand (traffic matrix). Second, we analyze the timescale for updates. This is the timescale at which the capacity allocation (quota) per CAC controller should be updated. Third, a reliability analysis is presented. Finally, we perform an experiment using real traffic traces from the Spanish National Research and Education Network, in order to show the suitability of the results in a real network scenario.

A. 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 over 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. The goal is to derive the capacity planning rule (quota assignment) for local CAC i assuming that the average traffic intensity is I in bits per second and its standard deviation is given by σ . The capacity planning problem is stated as follow, find C such that

$$P(X > C) < \epsilon \tag{1}$$

where X is the offered traffic, resulting in the rule $C = I + \sigma$, where is the 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).

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

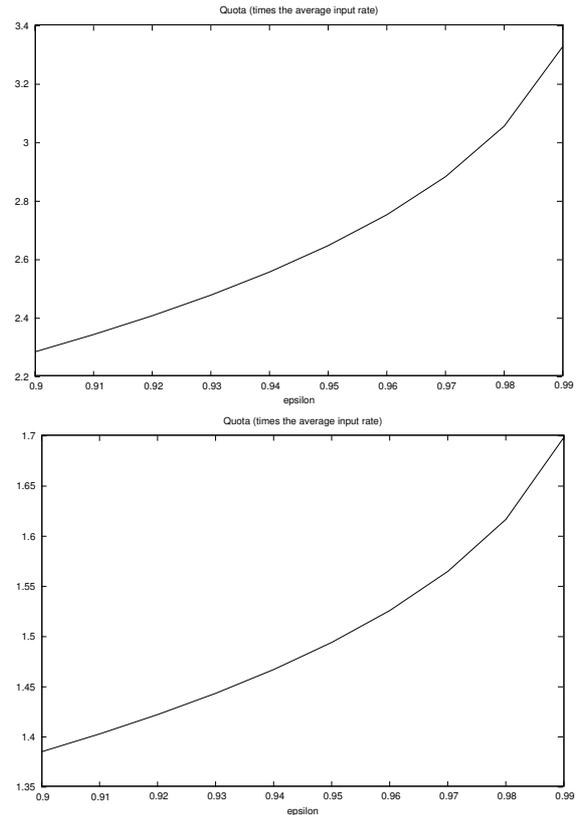


Figure 1. Top: Local CAC quota (bursty); Bottom: Local CAC quota (highly-multiplexed)

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.

B. Timescale for updates

As a first approximation, we assume that the input demands follow a Poisson arrival process with rate λ arrivals per time unit and consume a resource unit (for example, a lightpath). Such resource unit is held for an exponential time, with mean $1/\mu$ in time units per resource unit. The local CAC controller is provided with N resource units. We further assume that there is no queueing of demands, i.e. if the N resource units are already occupied, the CAC controller does not take any subsequent demand arrival, which are consequently dropped. Note that the CAC controller can be modelled by a $M/M/N/N$ birth-death process. Let $X(t)$ denote the number of busy resource units. Provided that the CAC controller initial state is empty, namely $X(0) = 0$, we wish to find the value of , such that

$$P(X(t_0 + t_b) = N | X(t_0) = 0.9N) < 0.01 \tag{2}$$

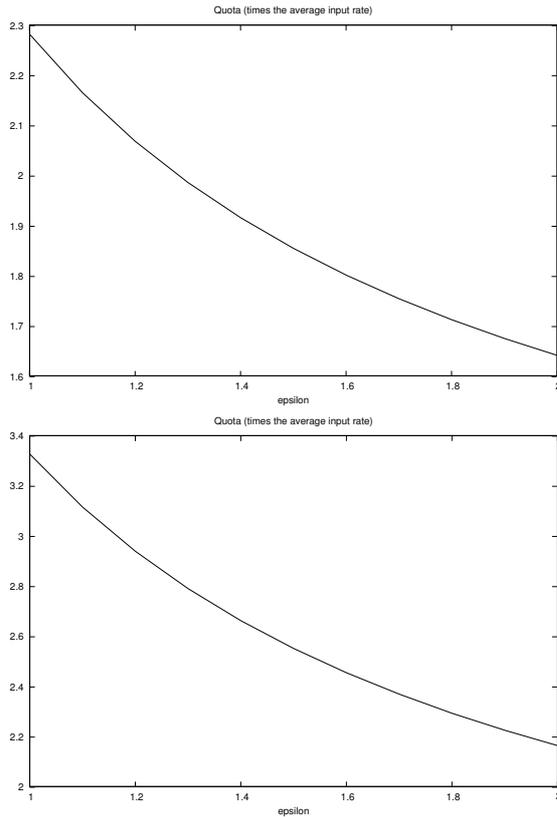


Figure 2. Top: Local CAC quota with traffic intensity that doubles and $\epsilon = 0.9$; Bottom: Local CAC quota with traffic intensity that doubles and $\epsilon = 0.99$

The equation can be explained as follows. Let us assume that t_0 is the hitting time to state $0.9N$ of the $X(t)$ process. Namely, the CAC controller boots at $t = 0$ and it is empty. Then, at time $t = t_b$ the CAC reaches an occupancy of 90%. This is an indication that the resources should be updated at the CAC. The point now is *How much can the CAC wait for an update?* Note that the time t_b represents such maximum waiting time for an update, because at time $t_0 + t_b$ the CAC will be fully booked and new demand arrivals will thus be blocked.

In order to evaluate such a conditional probability we consider the infinitesimal generator of the birth-death process of eq. 3, where state $N + 1$ is absorbent and let $Q(t) = e^{-Pt}$. The parameter $I = \lambda/\mu$ is the traffic intensity in Erlang.

Then,

$$P(X(t_0+t_b) = N | X(t_0) = 0.9N) = Q(t_b)[0.9N, N+1] \tag{4}$$

where $[a, b]$ means selection of the row a and column b in the matrix. For simplicity, we assume that $0.9N$ is an integer. There is a direct eigenvalue decomposition for hitting times of this sort [3]. However, we provide a numerical solution to the matrix exponential. The computation is performed by first block-diagonalizing the matrix and then applying a Pade approximation on each block.

In what follows, t_b is denoted by “Time to saturation”. The following figures compares the distribution of the time to saturation from state zero and state nine for

a local CAC controller with ten resource units, and 60% utilization (traffic intensity of 6 Erlang offered to the 10 circuits). The x-axis represents “average session duration”. As shown, average session durations in the range of minutes demand more frequent resource updates than session durations in the range of days.

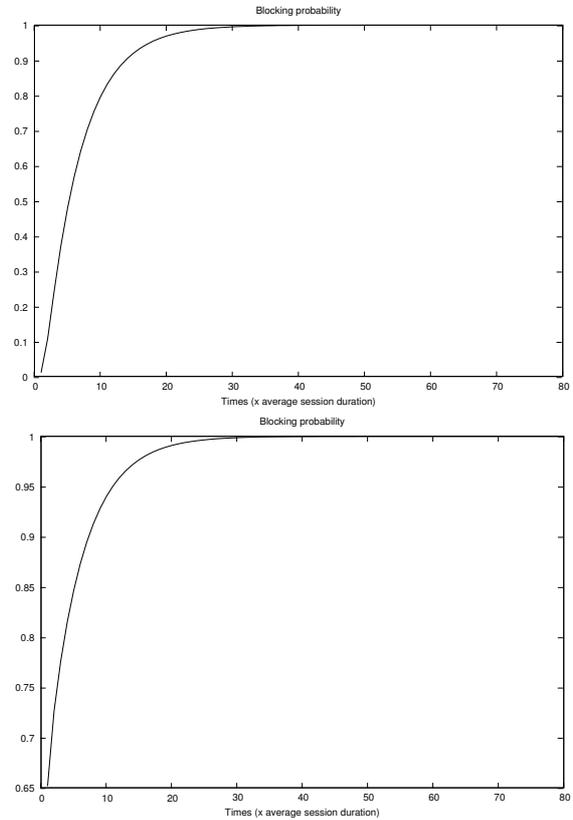


Figure 3. Top: Distribution of t_b (from empty CAC controller); Bottom: Distribution of t_b CAC (from CAC loaded 90%)

We note that the time to saturation is smaller in the loaded CAC, which matches our intuition. Clearly, a loaded system is more likely to saturate than an empty system, in the same time interval.

On the other hand, it is worth pointing out that the timescale for saturation can be relatively small, especially if the session duration is short. The following figure shows the lower percentiles of the distribution of the time to saturation, for the loaded CAC controller.

It turns out that in order to achieve the targeted 1% probability of not reaching saturation the updates must be as frequent as 0.001 times the average session duration. Let us assume that a worst case of a resource unit per video clip, with an average duration of 2.7 minutes². Then, the CAC controller requires an update every 0.162 seconds. Assuming that such updates comprise 1000-bit management packets, then the control plane requires 1620 bps per CAC controller.

However, *this is not comparable to the centralized CAC case in terms of signalling load in the control plane, which*

²<http://www.itworld.com/Tech/2987/070913onlinevid/>

$$\mathbf{P} = \begin{pmatrix} -I & I & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 1 & -(1+I) & I & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & 2 & -(2+I) & I & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & 3 & -(3+I) & I & 0 & \dots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & -N & N & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \end{pmatrix} \quad (3)$$

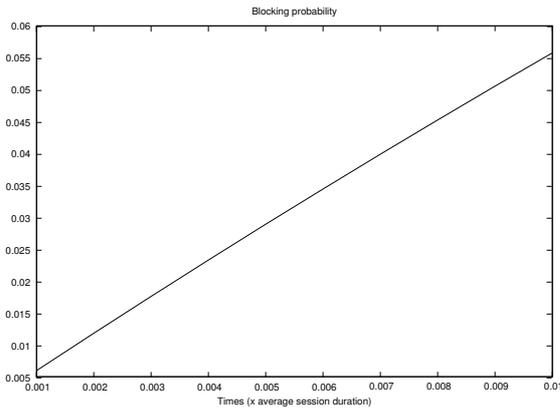


Figure 4. Time to saturation, loaded system (90%)

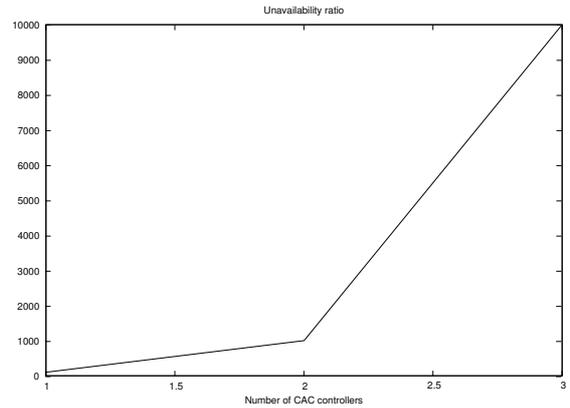


Figure 5. Unavailability ratio

is much higher. In that case, we have a resource allocation request per session which implies a much higher rate.

C. Reliability analysis

In what follows, let us assume that the failure probability of a CAC, either local or centralized is p , which can be obtained as the ratio between the MTBF and MTTR.

Let us now compare the reliability ratios obtained by the local and centralized CAC. We consider the probability that resources are unavailable, i.e, no bandwidth allocation request can be satisfied in the network. This implies failure of all CACs in the local case and failure of the centralized CAC in the centralized case, which can be calculated using a Binomial distribution.

The following graph shows the ratio between such probabilities (resource unavailability) versus the number of local CAC units, for a very large failure probability of 10%.

The results show that the unavailability ratio is as high as 10000 times as much in the centralized case, for a small network of 3 CAC controllers and very high failure rate. For smaller CAC failure rates, the unavailability ratio is even higher.

III. 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 measurement collection infrastructure.

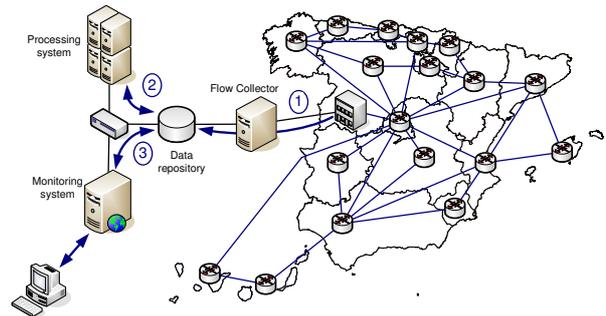


Figure 6. Measurement testbed in the Spanish academic network

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)	Std. deviation (Mbps)	Quota ($\epsilon = 0.99$)
U1	32	5.6	45.02
U2	8	2.6	14.04
U3	206	7.6	223.68
U4	206	35.6	288.81

TABLE I. RESULTS IN A CASE STUDY WITH REAL DATA

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.

IV. 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] E. et. al., "MUSE QoS Architecture," white Paper.
- [2] D. V. et. al., "MUSE deliverable DTF1.7 Multimedia Support in Access Architecture," available at https://www.ist-muse.org/Abstracts/abstract_DTF1.7.htm.
- [3] J. T. Kent and N. T. Longford, "An eigenvalue decomposition for first hitting times in random walks," *Probability Theory and Related Fields*, vol. 63, no. 1, March 1983.

Javier Aracil received the M.Sc. and Ph.D. degrees (Honors) from Technical University of Madrid in 1993 and 1995, both in Telecommunications Engineering. In 1995 he was awarded with a Fulbright scholarship and was appointed as a Postdoctoral Researcher of the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley. In 1998 he was a research scholar at the Center for Advanced Telecommunications, Systems and Services of The University of Texas at Dallas. He has been an associate professor for University of Cantabria and Public University of Navarra and he is currently a full professor at Universidad Autónoma de Madrid, Madrid, Spain. His research interest are in optical networks and performance evaluation of communication networks. He has authored more than 50 papers in international conferences and journals.

José Alberto Hernández completed the five-year degree in Telecommunications Engineering at Universidad Carlos III de Madrid (Madrid, Spain) in 2002, and the Ph.D. degree in Computer Science at Loughborough University (Leics, United Kingdom) in 2005. After this, he joined the Networking Research Group at Universidad Autónoma de Madrid (Spain), where he actively participates in a number of both national and european research projects concerning the modeling and performance evaluation of communication networks, and particularly the optical burst switching technology. His research interests include the areas at which mathematical modeling and computer networks overlap.

Antonio J. Elizondo obtained his Master's degree in Telecommunications Engineering at the Universidad Politécnica de Madrid (Madrid, Spain) and his Master's degree in Economics at the Universidad Nacional de Educación a Distancia (Spain). From 1998 he has been working for Telefónica I+D on Traffic Engineering and traffic controls for IP and ATM networks,

Quality of Service differentiation and scheduling algorithms, and planning and techno-economical evaluations of new technologies for access networks and advanced services. From 2004 to 2008, he worked in the Department of Advanced Networks Planning division, being mainly involved in IST projects like NOBEL or MUSE, and leading the AGAVE project. Currently he works as a technology expert for the Direction of Networks and Services Platforms, where he contributes to organize the Telefónica I+D research agenda on networks and platforms.

Raúl Duque completed the five-year degree in Telecommunications Engineering at Universidad Carlos III de Madrid (Madrid, Spain) in 2007. In 2006, he joined the New Network Technologies Division at Telefónica I+D where he developed his degree dissertation about a new network resource management mechanism based on the Nominal Route concept. He is currently working in the Photonic Network Division at Telefónica I+D, where he is participating in internal research projects for the Telefónica group. His research interests include network performance, optical network planning, resilience techniques, multicast transport algorithms for IPTV and access control strategies.

Óscar González de Dios received a Master Degree from the University of Valladolid in 2000 in Telecommunications Engineering. In addition to his work in Telefónica, he is at the same time pursuing his PhD degree in TCP/IP performance over OBS networks in the University of Valladolid. In 2000, he joined Telefónica I+D, where he worked for several years in the development and testing of telephony applications and interactive voice-response platforms. In 2005 he joined the Advanced Network Planning department in Telefónica I+D, where he is currently participating in R & D European projects, such as NOBEL II and e-Photon/One+, concerning analysis and performance evaluation of optical networks.