

Deployment of Class-based Routing in Wide Area Networks: Overhead and Performance Assessment

Marília Curado, Orlando Reis, João Brito, Gonçalo Quadros, Edmundo Monteiro

marilia@dei.uc.pt, {oreis, jbrito}@student.dei.uc.pt, {quadros, edmundo}@dei.uc.pt
Laboratory of Communications and Telematics
CISUC/DEI
University of Coimbra
Pólo II, Pinhal de Marrocos, 3030-290 Coimbra
Portugal

Abstract. A Quality of Service (QoS) routing protocol for the Differentiated Services framework is being developed at the University of Coimbra (UC-QoS_R). The main contribution of this paper is the evaluation of the protocol in wide area networks concerning traffic performance and protocol overhead. Many interesting QoS routing proposals have been proposed. However their evaluation has been conducted in majority by simulation. In this paper, we evaluate a QoS routing proposal using a prototype in an emulated Wide Area environment and compare the results with the Local Area scenario. The results show that it is feasible and cost rewarding to use QoS routing in order to improve the performance of QoS aware traffic in the actual existing Internet. The results also show that, besides being feasible, the QoS routing strategy proposed plays a major role in traffic performance both at the local and wide area environment, without imposing excessive performance costs and overhead in the network.

1 Introduction

Quality of Service plays a major role in the deployment of communication system for applications with special traffic requirements, such as video-conferencing or Internet telephony. The need to support these types of traffic has motivated the communication research community to develop new approaches. Some of this work resulted in the Differentiated and Integrated Services architectures proposed by the Internet Engineering Task Force (IETF) [1, 2].

Current routing protocols used in the Internet lack characteristics for QoS provision to support emerging new services. All traffic between two endpoints is forwarded on the same path, even if there are other alternative paths with more interesting properties for the requirements of a specific flow or traffic class. Usually, the shortest path is selected, based on a single static metric that does not reflect the availability of resources. In these situations, congestion easily occurs on the shortest path, with the corresponding degradation of traffic performance, despite the underutilization of network resources on alternative paths. This scenario has motivated the development of QoS aware routing protocols.

The most significant developments on QoS routing are aimed at communication systems where traffic differentiation is done per flow, as in the Integrated Services [1]. The Differentiated Services framework does not explicitly incorporate QoS routing. It is, thus, essential to develop QoS routing protocols for networks where traffic differentiation is done per class. The Quality of Service Routing strategy of the University of Coimbra (UC-QoS_R) was conceived to fulfill this purpose.

The UC-QoS_R strategy selects the best path for each traffic class based on information about the congestion state of the network. This strategy extends the Open Shortest Path (OSPF) routing protocol [3] in order to select paths appropriate for all traffic classes as described in [4, 5].

A prototype of UC-QoS_R was implemented over the Gated¹ platform, running on the FreeBSD operating system [4]. The evaluation of the UC-QoS_R strategy in the Wide Area environment is the main objective of the present paper. The rest of the paper is organized as follows: Section 2 presents some related work; Section 3 describes the UC-QoS_R strategy; test conditions and analysis of results concerning local area networks and wide

¹ <<http://www.gated.org>>

area networks are presented in Section 4; the main conclusions and issues to be addressed in future work are presented in Section 5.

2 Related Work

QoS routing protocols can contribute to improve traffic performance. However, QoS routing introduces additional burden in the network, pertaining to the processing overhead due to more complex and frequent computations and the increased routing protocol overhead. The trade-off between the overhead of QoS routing and its performance is an important issue that was evaluated in some works.

The advertisement of quantified metrics, instead of the advertisement of instantaneous values, is a common approach to avoid the excessive communication overhead of dynamic routing protocols. The quantification rule can be a simple average of the metric values [6] or can use a moving average with configured timescales [7].

Another approach to reduce the communication overhead of QoS routing is the use of trigger policies to control the emission of routing updates [8]. The triggering policies can be classified by the type of trigger used, namely, threshold based, class based and time based. Threshold based triggers control the emission of updates through a threshold. Updates are issued when the relative difference between the last value advertised and the actual value of the metric exceeds the configured threshold. Class based triggers divide the link capacity in classes and trigger the emission of updates when the link capacity changes to a new class. The time based triggers control the emission of updates by some periodic value. An example of such trigger is the hold-down timer used to define the minimum time between updates.

The solutions described above are able to simultaneously reduce communication and processing overhead. This results stems from the fact that once routers receive less update messages, they will not compute paths as often as before. The utilization of the mechanisms described poses the need for a trade-off between the desired updated state of the network and the burden this imposes in terms of routing overhead.

Despite the relevant QoS issues addressed, the proposals for QoS routing analyzed lack the analysis of the applicability to a class-based framework and are only evaluated theoretically or by simulation. The use of a prototype approach limits the dimension of the test-bed, however it introduces processing and communication systems dynamics, being closest to a real situation. Furthermore, there is a lack of the evaluation of the QoS routing strategies proposed in both local and wide area environments.

3 UC-QoS SR Strategy

In this section the main characteristics of the routing strategy UC-QoS SR are briefly described. A more detailed description can be found in previous publications of the authors [4, 5, 9]. The mechanisms used to control the overhead imposed in the network and in routers are presented in detail.

3.1 UC-QoS SR System Model

The UC-QoS SR strategy was designed for hop-by-hop QoS routing in networks where traffic differentiation follows the class paradigm. This strategy is composed of three main components, as follows:

- a) A QoS metric that represents the availability of resources in the network;
- b) Traffic class requirements in terms of QoS parameters;
- c) A path computation algorithm to calculate the most suitable path for each traffic class, according to the dynamic state of the network expressed by a QoS metric.

The availability of resources in the network is measured through a QoS metric that represents the congestion state of routers interfaces. This metric consists of two congestion indexes, one relative to packet delay (*DcI*) and other to packet loss (*LcI*). These indexes evaluate the impact that delay and loss at the router, will have on application performance [10]. The congestion indexes are distributed to all routers in the domain through modified OSPF routing messages (Router Link State Advertisements – R-LSA).

The UC-QoS SR strategy was conceived for communication systems where traffic characterization is based on class sensitivity to delay and loss. Currently, four classes are considered with different delay and loss sensitivities.

The problem of QoS routing when using two additive or multiplicative metrics, or one additive and one multiplicative metrics is a NP-complete problem [11, 12]. Thus, since the congestion indexes are additive metrics, the selection of a path that minimizes both congestion indexes is a NP-complete problem. However, due to their nature, the indexes represent comparable measures, and can be combined in a single metric without loss of information from aggregation of different kinds of units.

The path computation algorithm uses a cost function that combines both congestion indexes, weighted according to delay (δ) and loss sensitivity (λ) of each traffic class. The cost of link l for class i ($c_{l,i}$) results from the combination of the loss congestion index of link l (l_l) and the delay congestion index of link l (l_l), according to:

$$c_{l,i} = \delta_i d_l + \lambda_i l_l. \quad (1)$$

The merging of the congestion indexes origins a value that represents the congestion state of the interface, as it is perceived by traffic belonging to each class. The Dijkstra algorithm is then used to compute the shortest path tree for each traffic class. The UC-QoS SR strategy remains fully compatible with original OSPF because the path selection algorithm is not altered, and because the OSPF configured costs are also advertised in R-LSAs. It is thus possible to establish adjacencies among routers running UC-QoS SR and OSPF.

3.2 Mechanisms for overhead control

QoS routing protocols must contribute to a significant improvement in traffic performance and network resource usage to compensate for the burden they introduce on the network. This overhead is twofold, comprising an increase in the communication load due to routing traffic and a raise in the processing capacity of routers caused by the frequency of path computations. In UC-QoS SR, these overhead are controlled by a policy that controls the emission of link state updates. This policy combines metrics quantification and threshold based diffusion. A similar approach was followed by Apostolopoulos *et al.* but in the flow establishment context [8].

The metrics quantification has two components. In the first step, a logarithmic function is used to smooth extreme instantaneous metric values. In the second step, a moving average of the congestion indexes resulting from Equation 2, with a variable window size (N). The congestion indexes are monitored every second (the lowest time granularity provided by GateD) and the samples are taken continuously. In Equation 2, $MA_d(k)$ is the moving average of N values of the delay congestion indexes at sample k . This function is used to filter the peaks of the QoS metric.

$$MA_d(k) = \sum_{i=k-N}^k \frac{d(i)}{N}. \quad (2)$$

The filtered values, resulting from the application of Equation 2, are then presented to the diffusion control module. In this module, the new value is compared with the one that was previously advertised, and will be diffused only if it significantly different. The decision to issue the advertisements is controlled by the value of a defined threshold.

Besides the link state update policy described above, in UC-QoS SR, OSPF was modified, in order to control even further the protocol overhead and thus increase the possibility of scalability. In original OSPF, the routing messages denominated Network-LSA (N-LSA) identify the routers connected to the network and its diffusion occurs wherever R-LSAs are issued. In the UC-QoS SR strategy, the emission of N-LSAs has been detached from the emission of R-LSAs, because R-LSAs are issued at a higher rate than in OSPF and the information transported in N-LSAs does not change at such a rate. Thus, in the UC-QoS SR, the emission of N-LSA remains periodic and dependent on router connectivity, while the emission of R-LSA is controlled through the threshold of the diffusion module. This strategy allows for a significant reduction of routing messages in the network.

The processing overhead of a routing protocol is due to the application of the path computation algorithm and to the update of the kernel routing table. These two processes are usually linked, as in the original OSPF. This means, that, as soon as the SPF algorithm is applied, the new routes are flushed to the kernel to update the routing table. In the UC-QoS SR the routes are sent to the kernel, only if they are different from the ones that were previously installed. With this approach, the processing overhead associated with the socket communication between the routing daemon and the kernel is significantly reduced.

The policy to control protocol overhead described above contributes also to avoid the number of path shifts that may occur in the network. Combined with these procedures, the UC-QoS SR strategy uses a mechanism named class-pinning, that controls the path shifting frequency of all traffic classes [9].

4 Experimentation

In this section the experimentation made to evaluate the overhead and performance of UC-QoSr in Wide Area and Local Area Networks are presented and its results are analyzed. The WAN environment was emulated by the use of routers running Dummynet².

4.1 Test Conditions

The test-bed used for the experiments presented in this section is depicted in Figure 1. The endpoints 1 to 4 are traffic sources and endpoints 5 to 8 are traffic destinations. Each endpoint only generates or receives traffic of a single class to avoid the influence of endpoint processing on traffic patterns. The routers called Dummynet are responsible for emulating the delays associated with a wide area scenario. These routers run the original OSPF routing protocol to allow for the establishment of adjacencies with the routers running the UC-QoSR protocol.

The routers are PCs with the FreeBSD operating system. The kernel is modified, at the IP level, to include the delay and loss metric modules and to schedule and drop packets according to class sensitivity to these parameters. The monitoring of the delay and loss congestion indexes is needed for the routing decision. The kernel is also modified to interact with the UC-QoSR protocol embedded in GateD. It keeps the routing table with paths for all traffic classes and makes packet forwarding decisions based on destination IP address and Differentiated Services Code Point (DSCP) [2].

The interfaces between endpoints and routers are configured at 100 Mbps. Interfaces between routers are configured at 10 Mbps to introduce bottlenecks. The Wide Area Scenario was emulated through the configuration of Dummynet in order to create a propagation delay of a 10 Mbps connection over 200 Km (a typical distance between two large cities).




Figure 1 – Experimental test-bed

Traffic was generated and measured with the traffic analysis tool Chariot from NetIQ³. Table 1 shows the load distribution that was used in the experiments with 4 traffic classes to attain a total level of load of 24 Mbps. The traffic of all classes is UDP to avoid the influence of TCP flow control in the results. These values were chosen taking into consideration the circulating types of traffic on the Internet⁴, with the majority of the traffic being best-effort traffic (class 1 traffic).

² <http://info.iet.unipi.it/~luigi/ip_dummynet>

³ <<http://www.netiq.com>>

⁴ <<http://www.caida.org>>

	Class 1	Class 2	Class 3	Class 4
Load (Mbps)	9	7	5	3

Table 1 – Load distribution among classes.

The experiments for the evaluation of the UC-QoSRS strategy were done by varying the parameters that control the emission of routing updates, namely, the moving average window size and the threshold that controls the emission of routing updates.

The evaluation of overhead associated with the UC-QoSRS strategy was done by the measurement of protocol indicators. The parameters used to measure the communication overhead of the UC-QoSRS strategy are the following:

- a) Number of routing messages issued (Router-LSA);
- b) Number of routing messages received (Router and Network⁵ LSA).

The processing overhead is evaluated by the following parameters:

- a) Number of times the Shortest Path First (SPF) algorithm is applied; this value evaluates the processing overhead of the routing daemon due to the application of the path computation algorithm for each traffic class;
- b) Number of path shifts; this value evaluates the processing overhead due to the update of the kernel routing table.

The evaluation of traffic performance was made as follows:

- a) Average throughput of all traffic classes – the values shown represent the relation between the obtained throughput and the load generated in each class; this approach was employed to allow for the comparison of the performance of traffic belonging to classes that were generated with different levels of load;
- b) Loss rate of all traffic classes.

Each experiment was carried out for five minutes and was repeated ten times. The results present the averaged values of all tests. The inspection of protocol dynamics was done in all routers using the OSPF-Monitor tool included in GateD. These values were measured by the application Chariot. The plotted results have a degree of confidence of 95%.

4.2 Overhead evaluation in the Wide Area environment

The evaluation of the overhead of the UC-QoSRS strategy was done considering different combinations of the moving average window size used in the quantification rule and threshold of the triggering policy. The results shown pertain to all the routers of the testbed.

Figure 2 shows the results concerning the communication overhead of the UC-QoSRS strategy in the emulated WAN scenario. As could be expected, *Router 2* is responsible for the emission of the highest number of R-LSAs. This stems from the fact that *Router 2* is the bottleneck of the network, suffering from congestion. Another reason for this behavior is due to its three output interfaces. The difference of the number of R-LSA issued by the routers on the paths to the destination endpoints is due to the distribution of the traffic of all classes on the network.

It is clear that the emission of R-LSA is high for small values of the moving average window size (WS), despite of the threshold used to control its emission. However, when a window size of 40 samples is used, the number of R-LSA decreases in all routers and becomes a smaller burden on the network. In *Router 2*, using a WS higher than 80 does not bring any advantage in terms of the number of R-LSA issued. In the other routers of the testbed the reduction is also very small. This leads to the conclusion that the WS can be configured in the range [40, 80] to control the communication overhead of the UC-QoSRS strategy.

Figure 2 also shows the number of LSA received by all routers in the testbed. All routers receive similar numbers of LSA due to the flooding mechanism used by OSPF to distribute routing information.

⁵ The measure used is the total number of LSAs received, that is, the number of Router and Network LSAs received.




Figure 2 – Communication overhead assessment in the Wide Area environment.

Figure 3 shows the results concerning the processing overhead of the UC-QoS strategy in the emulated WAN scenario. The number of times the SPF algorithm is applied is the same in all routers in the testbed due to the link-state nature of OSPF. The SPF indicator for processing overhead drops 50% when the WS used increases from 20 to 40 samples. These results show that, besides controlling the communication overhead, the quantification rules used are able to significantly reduce the processing overhead due to path computation.

The number of path shifts is deployed on the right hand side of Figure 3. Comparing the lines of all routers to the number of SPF on the left, it is possible to see the importance of the detachment between path computation and kernel routing table updating. *Router 2* is the only router that suffers a number of path shifts that is directly related to the application of the path computation algorithm. This is due to the fact that each time the algorithm is applied, a new path is found. Since the other routers only have two output interfaces and the traffic generated is unidirectional, the application of the path computation algorithm seldom produces new paths, and thus the kernel routing table does not need to be updated.

An interesting observation of Figure 3 is that the number of paths shifts starts to rise when a WS of 160 samples is used. This is due to the inaccuracy in routing information that is introduced by the excessive smoothing of the metric. In this situation, important changes in network fail to be considered and distributed and thus routers choose paths based on inaccurate information. This will lead to bad paths that will create congestion and induce the computation of new paths. The problem of QoS routing under inaccurate information was studied, among others, by Chen *et al* and Masip-Bruin *et al*, [13, 14]. The results presented here show that inaccuracy due to metrics quantification can be avoided by a correct parameterization of the quantification rule.




Figure 3 – Processing overhead due to path computation and routing table updating in the Wide Area environment.

4.3 Overhead evaluation comparison in Local and Wide Area environments

The comparison of the overhead of the UC-QoS strategy in LAN and WAN scenarios is the subject of this section. The results depicted on the graphics pertain to *Router 2*, where the congestion bottleneck may occur.

Figure 4 shows the results concerning the communication overhead due to the routing information messages distributed in the network. The number of R-LSA issued by *Router 2* is not significantly different in LAN and

WAN scenarios. The difference is only noticeable for a WS of 20 samples. However, in the WAN scenario *Router 2* receives an important amount of LSA when compared to the LAN scenario. This result is due to the fact that the routers where Dummynet is installed run the original OSPF routing protocol and issue their own R-LSA and N-LSA and participate on the flooding procedure.




Figure 4 – Comparison of the communication overhead in the Wide and Local Area environments.

The processing overhead comparison between LAN and WAN scenarios present some interesting results as depicted in Figure 5. The number of times the path computation algorithm is applied follows a patterns similar to the one observed with the number of R-LSA. This result is as expected, since the application of the SPF algorithm is triggered by the arrival of R-LSA.

The number of path shifts in the LAN scenario shows a behavior with a common pattern for all values of the moving average window size. When the threshold that controls the emission of updates is below 40%, the number of paths shifts is smaller than in the WAN scenario. However, when threshold of 60% or above are used, the number of path shifts increases considerably. The same number of SPF produces an increased number of path shifts because the new paths are better than the old ones and are thus installed on the kernel routing table. The reason for this is the selection of paths based on outdated information and can be avoided by using smaller thresholds, keeping the network updated. Even though the use of high values of the threshold cause instability, the level of instability is controlled by the size of the moving average window.




Figure 5 – Comparison of the processing overhead in the Wide and Local Area environments.

4.4 Traffic performance evaluation

In this sub-section the results concerning traffic performance evaluation are presented. The traffic classes considered are presented in Table 2.

Class	Delay sensitivity	Loss sensitivity	Application
1	Low	Low	<i>Best-effort</i>
2	High	Medium	Internet telephony
3	High	Low	Video real-time
4	Medium	High	Video Training

Table 2 – Traffic classes characterization.

It is clear that the overall performance of all traffic classes is better on the LAN scenario. This is due to the extra delay that is introduced by Dummynet. However, in the WAN environment, class differentiation is done according to class sensitivity to delay and loss, which is a different result from the behavior in the LAN.

Figure 6 shows the results pertaining achieved throughput ratio and loss rate of traffic of the four traffic classes generated in the WAN scenario. Class 1, the class used by best-effort traffic, achieves the lowest throughput ratio and the highest loss ratio. Classes 2 and 4 show the best throughput and the lowest loss rate, according to their delay and loss sensitivity. Class 3 is treated best than best-effort traffic, but is treated worst than classes 2 and 4 because it has a low sensitivity to delay. These results show that the UC-QoS routing strategy is effective in differentiating traffic classes, through scheduling and routing mechanisms, according to their delay and loss sensitivities.

The results also show that the quantification rule is able to contribute to traffic performance, as long as it is correctly configured. When a small moving average window size is used, traffic performance is poor, due to the exaggerated instability in the network caused by frequent distribution of routing messages. On the other side, when the moving average window size is big (over 80 samples), traffic performance is also poor. This is due to the inaccuracy that such a WS introduces in the network. The results show that the moving average window size must be configured in the range [40, 80] in order to achieve the best traffic performance in the Wide Area environment.

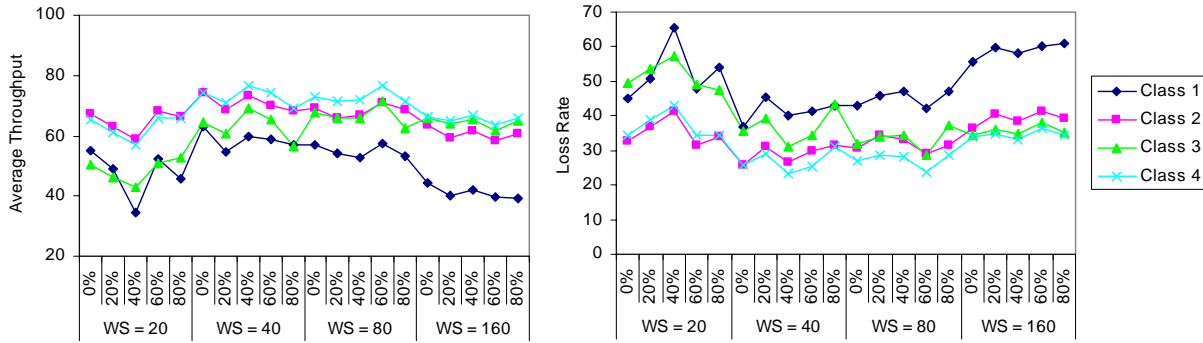


Figure 6 – Performance parameters in the WAN environment.

Figure 7 shows the results pertaining achieved throughput ratio and loss rate of traffic of the four traffic classes generated in the LAN scenario. Then results plotted are better than the ones in the WAN scenario as expected.

The graphics show that traffic differentiation is achieved and is correctly made among classes 2, 3, and 4. However, the results obtained for class 1 traffic (best-effort traffic) are not consistent with class definition, since this is the best treated class according to throughput and loss rate. This behavior is motivated by two factors: the load distribution among classes and the routing oscillations that occur due to the small timescale inherent to the small sized network. Best-effort traffic stays in the same path because it has low sensitivity to delay and loss and does not require a better path, and any class that tries to share the path with class 1, will swiftly shift to another better path. This behavior is corrected when it is used a moving average window size above 80 samples.




Figure 7 – Performance parameters in the LAN environment.

5 Conclusions and Future Work

At the University of Coimbra a QoS routing strategy (UC-QoS SR) was developed to support class-based traffic differentiation. With this strategy the paths suitable for traffic classes are computed based on a QoS metric that evaluates the state of the network in terms of delay and losses. The proposed mechanisms were implemented as an extension to OSPF routing protocol on the GateD platform. Mechanisms for scalability and stability were embedded in the UC-QoS SR strategy in order to overcome the common problems associated with QoS routing.

The behavior of the UC-QoS SR strategy was studied in a local area network testbed and showed effective differentiating traffic while introducing an overhead that was affordable by the communication system. The focus of this paper was the evaluation of the overhead and performance issues of the UC-QoS SR in a wide area environment.

The results showed that the mechanisms in the UC-QoS SR strategy can control the overhead introduced in the network both in LAN and WAN environments. The communication overhead remains a function of the network size and is not significantly affected in the wide area scenario. The processing overhead due to the application of the path computation algorithm is also similar in both environments. However, the overhead due to routing table update is higher in the LAN scenario due to instability created by the short timescale scenario.

In the LAN scenario, where the routing timescale is short, a large window size contributes to stability and to increase traffic performance. The inaccuracy introduced is small because updates are distributed fast to all routers in the network. In the WAN scenario the inherent larger timescale limits the window size that can be used and induces uncertainty in routing information. However this is not an issue, since the best performance can be achieved with window size tuning.

As future work the self-tuning of window size to adapt to network size will be pursued.

Acknowledgements

This work was partially supported by the Portuguese Ministry of Science, Technology and High Level Education (MCES), under program POSI (Projects QoS II and IPQoS) and under the PhD grant PRAXIS XXI/BD/13723/97.

References

- [1] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an Overview", Request for Comments 1633, Internet Engineering Task Force, June 1994.
- [2] S. Blake, D. Black, M. Carlson, E. Davies Nortel, W. Weiss, "An Architecture for Differentiated Services", Internet Engineering Task Force, Request for Comments 2475, December 1998.
- [3] J. Moy, "OSPF Version 2", Internet Engineering Task Force, Request for Comments 2328, April 1998.

- [4] M. Oliveira, J. Brito, B. Melo, G. Quadros, E. Monteiro, "Quality of Service Routing in the Differentiated Services Framework", Proceedings of SPIE's International Symposium on Voice, Video, and Data Communications (Internet III: Quality of Service and Future Directions), Boston, Massachusetts, USA, November 5-8, 2000.
- [5] M. Oliveira, J. Brito, B. Melo, G. Quadros, E. Monteiro, "Evaluation of a Quality of Service Routing Strategy for the Differentiated Services Framework", Proceedings of the 2001 International Conference on Internet Computing (IC'2001), Monte Carlo Resort, Las Vegas, Nevada, USA, June 25-28, 2001.
- [6] A. Khanna, J. Zinky, "The Revised ARPANET Routing Metric", Proceedings of SIGCOMM'89, Austin, Texas, September 19-22, 1989
- [7] B. Lekovic, P. Van Mieghem, "Link State Update Policies for Quality of Service Routing", Proceedings of Eighth IEEE Symposium on Communications and Vehicular Technology in the Benelux (SCVT2001), Delft, The Netherlands, October 18, 2001.
- [8] G. Apostolopoulos, R. Guerin, S. Kamat, S. Tripathi. "Quality of Service Based Routing: A Performance Perspective", Proceedings of SIGCOMM'98, Vancouver, British Columbia, September 1998.
- [9] M. Curado, O. Reis, J. Brito, G. Quadros, E. Monteiro, "Stability and Scalability Issues in Hop-by-Hop Class-Based Routing", Proceedings of the 2nd International Workshop on QoS in Multiservice IP Networks (QoS-IP 2003), Milano, Italy, February 24-26, 2003.
- [10] G. Quadros, A. Alves, E. Monteiro, F. Boavida, "An Approach to Support Traffic Classes in IP Networks", Proceedings of QoFIS'2000 – The First International Workshop on Quality of future Internet Services, Berlin, Germany, September 25-26, 2000.
- [11] Z. Wang, J. Crowcroft, "Quality of Service Routing for Supporting Multimedia Applications", IEEE JSAC, vol. 14, No 7, pp. 1228-1234, September 1996.
- [12] H. De Neve, P. Van Mieghem, "TAMCRA: A Tunable Accuracy Multiple Constraints Routing Algorithm", Computer Communications, 2002, Vol. 23, pp. 667-679.
- [13] S. Chen, K. Nahrstedt, "Distributed QoS Routing with Imprecise State Information", Proceedings of 7th IEEE International Conference on Computer, Communications and Networks (ICCCN'98), pp. 614-621, Lafayette, LA, October, 1998
- [14] X. Masip-Bruin, S. Sánchez-López, J. Solé-Pareta, J. Domingo-Pascual, "A QoS Routing Mechanism for Reducing the Routing Inaccuracy Effects", Proceedings of the 2nd International Workshop on QoS in Multiservice IP Networks (QoS-IP 2003), Milano, Italy, February 24-26, 2003.