Quality of Service Routing

Marília Curado, Edmundo Monteiro

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

1. Introduction

Traditionally, Internet routing is achieved through shortest path protocols that base their decision on the number of hops or administrative metrics. The path computation algorithms belong either to the distance vector or link state families.

Distance Vector protocols have been widely used on the Internet, since the ARPANET and until today. The protocols of the Distance Vector family are used by routers to exchange routing information with their neighbours and to select the shortest paths to all destinations within the network using the Bellman-Ford algorithm, such as in the Routing Information Protocol (RIP) [Malkin1998].

Link-State routing protocol come as an alternative to Distance Vector protocols. Open Shortest Path First is a Link-State Interior Gateway Protocol that is responsible for routing packets within an Autonomous System [Moy1998]. Link state routing protocols use a flooding mechanism to distribute the information about the state of the links of each node to all nodes in the network and apply the Dijkstra Shortest Path First algorithm to compute shortest paths to all destinations in the network.

The routing protocols presented are widely used for Internet routing. However, the inherent shortest-path routing paradigm lacks the potential to support the Quality of Service (QoS) level required by real-time applications. While QoS is provided in actual networks at other levels, such as with queue management and application adaptation mechanisms, QoS routing comes up as the missing piece to achieve an interesting level of QoS.

2. Background

The main goal of QoS routing is to select, based on information about the state of the network, the path that is most suitable according to traffic requirements and thus to contribute to the improvement of traffic performance while maximizing the utilization of network resources [Crawley1998].

The design of a QoS routing protocol needs the definition to three main building blocks: the metrics used to represent the state of the network, the mechanism used for their distribution and the QoS path computation algorithm.

Routing metrics include parameters such as bandwidth, loss rate, delay and jitter. The choice of metrics is one of the main issues that must be addressed in a routing strategy because it determines the characteristics that are offered to traffic and the complexity of the path computation algorithm. Routing metrics are classified according to the composition rule used to compute the value of the metric on a path. The composition rule for additive metrics, such as delay and number of hops, is that the value of this metric over a path is the sum of the values at each hop. The composition rule of multiplicative metrics is obtained by the product of its values in each hop, as it is the case of loss rate or reliability. Bandwidth is a concave metric and its composition over a path corresponds to the minimum value observed in all hops of that path.

The distribution of metrics may be done by flooding, by signalling protocols, or by message probing.

The problem of QoS routing when using two additive or multiplicative metrics, or one additive and one multiplicative metrics is a Non Polynomial complete (NP-complete) problem [Neve2002]. This definition means that there is not an algorithm that is able to find a feasible solution that satisfies both constraints in polynomial time. There are however particular situations concerning the relationship between metrics, the values of metrics, and networks topologies that can be solvable by optimal algorithms in pseudo-polynomial time [Kuipers2003].

The selection of QoS paths subject to multiple constraints is called Multi-Constrained Path (MCP) problem. The MCP problem is to find a path from a source to a destination such that all the QoS constraints are met. The paths that satisfy these constraints are called feasible paths. Since the optimal solution of this type of problems for multiple additive and independent metrics is NP-complete, usually heuristics or approximation algorithms are used. Three of the main solutions for multi-constrained path computation are described next.

First approach - Bandwidth Restricted Path (BRP)

Metric Ordering is the main heuristic used for the solution of the BRP problem. This heuristic requires the identification of the metric that has higher priority and the computation of the best paths according to this metric. Afterwards, it is computed the best path according to the second metric. The algorithms that solve the BRP problem using metric ordering are the Widest-Shortest Path (WSP) [Sobrinho2002] and Shortest-Widest Path (SWP) algorithms [Shen2002]. In these families of algorithms, the width of a path is depicted by the available bandwidth and its length can correspond either to the number of hops or to delay.

A compromise between the two performance objectives of WSP and SWP algorithms, minimize resource consumption and improve load balancing, respectively, is achieved by the All Hops Optimal Path algorithm [Guérin2002]. The AHOP algorithm tries to reduce network cost while achieving load balancing, since a longer path is only used if it has more available bandwidth.

Second approach: Restricted Shortest Path (RSP)

The MCP problem, when two additive metrics are used. All the paths that satisfy the constraint associated with one of the metrics are computed and then the shortest path according to the second metric is selected.

A widely studied case of the RSP is the Delay-Constrained Least Cost problem. The objective of this problem is to find a set of paths that satisfy the delay constraint and then to select the path that minimizes the cost. Even though the DCLC problem is NP-complete, heuristics based algorithms have been proposed, such as the Delay-Cost Constrained Routing [Chen1998] and the Dual Extended Bellman-Ford [Cheng2003].

Third approach: Metrics Combination (MC)

Metrics combination reduces the complexity of the RSP problem by combining both metrics in a single metric and then uses a traditional shortest-path algorithm to compute the path that minimizes the resulting metric. There are two main approaches for metrics composition: linear function composition [Cui2003] and non-linear function composition [Neve2002].

Even tough metrics combination contributes to the simplification of path computation algorithms it prevents the provisioning of guarantees regarding the constraints associated with each one of the metrics involved. In order to overcome this problem, there is the need to define the proper weights used in the combination rule of metrics using, for instance, Lagrange relaxation techniques as in the Binary Search for Lagrange Relaxation algorithm [Korkmaz2000] and in the LAgrange Relaxation based Aggregated Cost algorithm [Jüttner2001].

The combination of metrics in a single metric allows for simple and well known path computation algorithms. However, the rule for combination of the metrics is not always straightforward and the composition of the resulting metric over a path can also be challenging. Therefore, the choice of the type of heuristic must take into consideration factors such as QoS framework, traffic patterns, and traffic engineering objectives.

3. Quality of Service Routing

The deployment of QoS routing protocols raises several problems that must be assessed in order to guarantee an interesting solution in terms of routing behaviour, traffic performance and network utilization. The main issues to be addressed in order to achieve successful QoS routing protocols concern the control of QoS routing overhead, the treatment of routing information inaccuracy, and the avoidance of routing oscillations.

3.1 QoS Routing Overhead

The objectives of QoS routing protocols may be compromised by the additional burden they impose in the network. The load introduced by QoS routing approaches includes the processing overhead due to more complex and frequent computations and the communication overhead caused by the increase in the amount of routing information exchanged within the network. Signalling and additional storage that are needed to support QoS routing protocols are also of importance when assessing the impact of QoS routing on the network.

3.1.1 Metrics quantification and triggering policies

The advertisement of quantified metrics, instead of the advertisement of instantaneous values, is a common approach to avoid the excessive communication cost of dynamic routing protocols [Apostolopoulos1999]. Moreover, instead of distributing updates immediately after a change, the instant of distribution is generally controlled by triggering policies. Triggering policies are classified by the type of trigger used, namely, threshold-based, class-based or time-based.

With threshold-based triggers, a new update is only issued when there is a significant difference between the actual and the previous values of the metrics. Class-based triggers divide the total range of the metrics into intervals and the emission of updates is triggered when the metric crosses the division between classes. In order to limit the instability due to the successive boundary changes, a hysteresis mechanism can be used, and thus, instead of using the actual transposition of a boundary as trigger, the update is only issued when it reaches the middle value of the new class.

As a complement to the policies described, a hold-down timer that ensures a minimum time interval between updates.

Time-based triggers issue routing updates periodically. This type of trigger has the disadvantage of being insensitive to network conditions. On the other hand, it is able to consistently reduce the amount of routing information by a conservative definition of the value of the update period, without influencing traffic patterns.

The solutions described are able to reduce communication and processing overheads, since once routers receive less update messages, they will compute paths less frequently. The utilization of the mechanisms described creates the need for a trade-off between the desired up-to-date state of the network and the burden this imposes in terms of routing overhead [Shaikh2001].

3.1.2 Selective flooding

The flooding mechanism used by Link-State protocols causes a burden that can consume an excessive amount of resources in the network.

The Selective Flooding Protocol uses a combined approach of flooding and probing to avoid the disadvantage of entirely flooding-based approaches [Claypool2001]. Flooding is used to maintain the topological database of every node and probing is performed upon the arrival of a connection request. With this approach, the network overhead is reduced and the selection of paths is made based on up-to-date information collected by the probe packets, however at the expense of connection setup time.

The Restricted Flooding mechanism limits the distribution of updates among border routers of areas or autonomous systems in hierarchical organized networks [Lee2003].

3.1.3 Path pre-computation

Most QoS routing proposals use on-demand path computation in order to obtain paths that guarantee the QoS requirements of traffic based on recent information about the state of the network. However, on-demand path computation has two drawbacks, namely, it introduces some delay before the forwarding of traffic starts and it requires the application of the path computation algorithm for each connection request, introducing additional processing overhead on the routers, especially when the arrival rate of connection requests is high. The pre-computation of paths is the alternative approach to handle the problem of the processing overhead associated with on-demand path computation at the expense of the eventual inaccuracy of the routing decision due to the time frame between path computation and the actual utilization of the path.

Proposals that use path pre-computation to reduce QoS routing overhead include the Multi-constrained Energy Function based Pre-computation [Cui2003], the All Hops Optimal Path [Guérin2002] algorithms, and the Class-based Routing QoS Routing [Curado2003].

Pre-computation is the actual method used for routing in IP networks and thus becomes an interesting option for a smooth transition to QoS routing. Even though pre-computation schemes can reduce the QoS routing processing overhead when the arrival of requests is high, they have two drawbacks. One is related to the utilization of eventually outdated paths, and the other pertains to the need to pre-compute QoS aware paths to all destinations satisfying all possible QoS requests. The first drawback can be overcome by using suitable path computation algorithm triggering mechanisms and the second becomes less important in networks where traffic differentiation has coarser granularity than flow-based, as in the class-based classification style.

3.2 QoS Routing under Inaccurate Information

The main sources of routing inaccuracy are the low frequency of the distribution of link-state updates, the information aggregation in hierarchical networks, propagation delay of routing messages in large networks, and the utilization of estimates about the current state of the network. Due to this wide range of factors, the global state that is kept by each router is just an approximation of the real actual state. When the path computation algorithms use this inaccurate information as if it was exact, their performance can be highly damaged.

One approach to handle routing information inaccuracy is based on algorithms that use probability functions and random variables, such as the following examples.

The Multi-Constrained Path problem under Inaccurate State Information uses a probabilistic function to find a path that that is the most suitable to accommodate a new request, [Mieghem2003].

The Safety-Based Routing combines a probabilistic approach with the limitation of the range the metrics can attain between updates to enhance the performance of QoS routing under inaccurate state information caused by stale link-state information due to large thresholds and hold-down timers [Apostolopoulos1999].

The Bandwidth-Delay-Constrained Path under Inaccurate State Information computes delay and bandwidth constrained paths [Korkmaz2003].

Message probing is another a technique commonly used to deal with imprecise state information. The utilization of probing avoids the staleness of link-state information because the probes gather the most recent state information. The Ticket-Based Probing [Chen1998], the Pre-Computation based Selective Probing Scheme [Lee2000], and the Selective Flooding Protocol [Claypool2001] use probes to collect QoS information about paths in order to solve the DCLC problem under inaccurate state information

3.3 QoS Routing Stability

The stability of QoS routing protocols is a determinant factor of their performance. Instability may occur whenever the responsiveness of the protocol becomes exaggerated introducing unnecessary re-routing of traffic. Specifically, in Link-State protocols, the inappropriate flooding of updates may originate route flaps that will degrade traffic performance. This is particularly problematic when the network is congested, since the additional routing messages consume the already scarce bandwidth resources, and the subsequent application of the path computation algorithm imposes even more load on the router processor.

The problem of routing instability is influenced by several factors, namely, the type of metrics used to compute the best path, the policy that controls the advertisement of the metrics and the path computation algorithm. Network topology and traffic patterns also influence routing behaviour and stability.

3.3.1 Control of metrics distribution

A common approach to avoid instability is the advertisement of metrics that are quantified in some manner, instead of advertising instantaneous values. The use of metric quantification, while contributing to routing stability, reduces the dynamic nature and the adaptation capabilities of the routing protocol, since it increases the period of routing oscillations.

The limitation of the range of values that the metrics can take is an additional method that can be used to control routing oscillations. This is achieved by the definition of a bias term that is a lower bound of the interval where the metrics values can vary. The upper bound of the interval can also be set in order to limit the range of metrics values. A scaling method can be used to fit the measured metrics values on the range defined by the two bounds.

Any mechanism for metric quantification must preserve an adequate trade-off between route adaptability and routing oscillatory behaviour.

3.3.2 Traffic characteristics aware QoS routing

Traffic patterns influence the behaviour of QoS routing algorithms, including its stability capability.

In the Long-Lived-Short-Lived approach, the resources in the network are dynamically shared between short-lived and long-lived flows [Shaikh1999]. The paths for long-lived flows are dynamically chosen based on the load level on the network, whilst the paths for short flows are statically pre-computed. Since the short-lived flows are routed on statically pre-computed paths, the route flaps associated with the dynamic routing of this type of flows are avoided.

In the Differentiated Routing scheme it is used a smoothed metric to compute the shortest path for long-lived flows, and the most recent value of the metrics to compute the shortest path for short-lived flows [Yang2001].

The Enhanced Bandwidth-inversion Shortest-Path the hop count is included in the cost function in order to avoid oscillations due to the increased number of flows sent over the widest-path [Wang2002].

The Proactive Multi-path Routing takes into consideration traffic characteristics and network connectivity [Shen2002]. The path chosen for long-lived flows is the shortest path that satisfies the requirements of the flow, while short-lived flows use multiple alternate paths.

3.3.3 Class-pinning

The dynamic selection of paths may cause routing instability and network oscillatory behaviour. The objective of the class-pinning mechanism is to limit the number of path shifts due to dynamic changes in the state of the network [Curado2005]. The class-pinning mechanism addresses the instability problem by controlling the instant when a traffic class shifts to a new path. Normally, when the state of the network changes, due to events such as the start of a new flow or a traffic burst, routing messages are sent to all nodes, and new paths are computed. Afterwards, traffic of each class will shift to the newer and less congested path, abandoning the path currently used. The next time this process occurs, traffic will eventually go back to the original path, generating an unstable situation. The class-pinning mechanism overcomes this condition by avoiding unnecessary paths shifts. When the class-pinning mechanism is active, new paths are also computed upon the arrival of routing messages. However, after path computation, the weight of the path for each <*destination, class*> pair is compared with the weight of the previously installed path. The new path will only be installed in the routing table if it is significantly better than the path that is currently used by that class.

4. Future Trends

The problem of QoS routing has still many open issues, since there are many trade-offs that must be achieve for a full successful deployment, as was pointed out throughout this discussion. One can expect to see flow-based QoS routing in intra-domain routing, using algorithms satisfying multiple constraints, and class-based QoS routing in inter-domain routing, using algorithms base in a single metric resulting from the combination of multiple metrics regarding the characteristics of the classes used within the QoS framework followed

Moreover, the problem of QoS routing for nowadays and upcoming types of networks such as wireless, sensor, and ambient networks needs further studies due to the specificities of these environments.

5. Conclusion

The need of Quality of Service on the Internet has motivated the development of several mechanisms to evolve actual IP networks. Quality of Service routing is one of these components. QoS routing has as main objective the selection of paths that satisfy the requirements of traffic in the network, while contributing to improved network resource utilization.

The problems addressed show that any QoS routing scheme needs to provide answers to a set of issues that are intertwined and whose implications on the overall performance depend on several and sometimes unpredictable factors.

6. References

[Apostolopoulos1999] G. Apostolopoulos, D. Williams, S. Kamat, R. Guérin, A. Orda, T. Przygienda, “QoS Routing Mechanisms and OSPF Extensions”, Internet Engineering Task Force, Request for Comments 2676, August 1999.

[Chen1998] S. Chen, K. Nahrstedt, “On Finding Multi-Constrained Paths”, Proceedings of IEEE International Conference on Communications (ICC 98), Atlanta, GA, June, 1998.

[Cheng2003] G. Cheng, N. Ansari, “A New Heuristics For Finding The Delay Constrained Least Cost Path”, Proceedings of IEEE GLOBECOM 2003, San Francisco, CA, USA, Dec. 1-5, 2003.

[Claypool2001] M. Claypool, G. Kannan, “Selective Flooding for Improved Quality-of-Service Routing”, Proceedings of SPIE Quality of Service over Next-Generation Data Networks, Denver, Colorado, USA, August 2001.

[Crawley1998] E. Crawley, R. Nair, B. Rajagopalan, H. Sandick, “A Framework for QoS-based Routing in the Internet”, Internet Engineering Task Force, Network Working Group, Request for Comments 2386, August 1998.

[Cui2003] Y. Cui, J. Wu, “Precomputation for Multi-constrained QoS Routing in High-speed Networks”, Proceedings of IEEE INFOCOM 2003, March 30 – April 3, San Francisco, USA, 2003.

[Curado2003] 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.

[Curado2005] Marilia Curado, Edmundo Monteiro, “A Class-Pining Election Scheme for Loop-Free QoS Routing”, Proceedings of IEEE Workshop on High Performance Switching and Routing (HPSR'2005), IEEE Communications Society, Hong Kong, China, May 12-14, 2005.

[Guérin2002] R. Guérin, A. Orda, “Computing Shortest Paths for Any Number of Hops”, IEEE/ACM Transactions on Networking, Vol. 10, No 5, October 2002.

[Jüttner2001] A. Jüttner, B. Szviatovszki, I. Mécs, Z. Rajkó, “Lagrange Relaxation Based Method for the QoS Routing Problem”, IEEE INFOCOM 2001, Anchorage, Alaska, April 22-26, 2001.

[Korkmaz2000] T. Korkmaz, M. Krunz, S. Tragoudas, “An efficient algorithm for finding a path subject to two additive constraints”, Joint International Conference on Measurement and Modeling of Computer Systems, Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems table of contents, Santa Clara, California, United States, 2000.

[Kuipers2002] F. Kuipers, T. Korkmaz, M. Krunz, P. Van Mieghem, “Overview of Constraint-Based Path Selection Algorithms for QoS Routing”, IEEE Communications Magazine, special issue on IP-Oriented Quality of Service, Dec 2002.

[Kuipers2003] F. Kuipers, P. Van Mieghem, “The impact of Correlated Link Weights on QoS Routing”, Proceedings of IEEE INFOCOM 2003, March 30 – April 3, San Francisco, USA, 2003.

[Lee2000] W. Lee, B. Lee, “Pre-Computation based Selective Probing Scheme with Imprecise State Information”, Proceedings of IEEE GLOBECOM 2000, pp 460-464, San Francisco, CA, USA, November 27 - December 1, 2000.

[Lee2003] S. Lee, S. Das, G. Pau, M. Gerla, “A Hierarchical Multipath Approach to QoS Routing: Performance and Cost Evaluation”, IEEE International Conference on Communications 2003 (ICC 2003), Anchorage ,USA, May 11-15, 2003.

[Malkin1998] G. Malkin, “RIP Version 2”, Internet Engineering Task Force, Network Working Group, Request for Comments 2453, November 1998.

[Mieghem2003] P. Van Mieghem, F. Kuipers, T. Korkmaz, M. Krunz, M. Curado, E. Monteiro, X. Masip-Bruin, J. Sole-Pareta, and J. Domingo-Pascal “Quality of Service Routing”, in Quality of Future Internet Services – Michel Smirnof, Jim Roberts and Jon Crowcroft (Editors), Springer, 2003.

[Moy1998] J. Moy, “OSPF Version 2”, Internet Engineering Task Force, Network Working Group, Request for Comments 2328, April 1998.

[Neve2002] H. De Neve, P. Van Mieghem, “TAMCRA: A Tunable Accuracy Multiple Constraints Routing Algorithm”, Computer Communications, Vol. 23, 2002.

[Shaikh1999] A. Shaikh, J. Rexford, K. Shin, “Load-Sensitive Routing of Long-Lived IP Flows”, Proceedings of ACM SIGCOMM’99, Harvard University Science Center, Cambridge, Massachusetts, USA, August 31-September 3, 1999.

[Shaikh2001] A. Shaikh, J. Rexford, K. Shin, “Evaluating the Impact of Stale Link State on Quality-of-Service Routing”, IEEE/ACM Transactions on Networking, Volume 9, Number 2, pp. 162-176, April 2001.

[Shen2002] J. Shen, J. Shi, J. Crowcroft, “Proactive Multi-path Routing”, Proceedings of the International Workshop on Quality of future Internet Services (QofIS 2002), Zurich, Switzerland, October 16-18, 2002.

[Sobrinho2002] J. L. Sobrinho, “Algebra and algorithms for QoS path computation and hop-by-hop routing in the Internet”, IEEE Transactions on Networking, August 2002.

[Yang2001] S. Yang, X. Su, G. de Veciana, “Heterogeneity-aware shortest path routing: Flow Holding Time, User Demand, and Network State”, Proceedings of the IEEE Workshop on High Performance Switching and Routing, Dallas, Texas USA, May 29-31, 2001.

7. Terms and Definitions

Additive metric	A metric whose value over a path is the sum of the values at each hop
Multiplicative metric	A metric whose value over a path is the product of the values at each hop.

Concave metric	A metric whose value over a path is the minimum value observed in all hops of that path
Feasible paths	The paths that satisfy these constraints are called feasible paths
Multi-Constrained Path problem	The selection of QoS paths subject to multiple constraints
Restricted Shortest Path	A simplification of the original MCP problem, when two additive metrics are used
NP-complete	This definition means that there is not an algorithm that is able to find a feasible solution that satisfies both constraints in polynomial time.
Quality of Service routing	The selection, based on information about the state of the network, of the path that can satisfy traffic requirements while maximizing the utilization of network resources