

A FRAMEWORK FOR COOPERATIVE INTER-DOMAIN QoS ROUTING

Alexandre Fonte,^{1,3} Edmundo Monteiro,¹
Marcelo Yannuzzi,² Xavier Masip-Bruin,² and Jordi Domingo-Pascual,²

¹*University of Coimbra,
Laboratory of Communications and Telematics,
CISUC/DEI, Polo II, Pinhal de Marrocos, 3030-290, Coimbra, Portugal*
afonte,edmundo@dei.uc.pt

²*Universitat Politècnica de Catalunya,
Departament d'Arquitectura de Computadors,
Avgda. Victor Balaguer, s/n-08800 Vilanova i la Geltru, Barcelona, Catalunya, Spain*
yannuzzi,xmasip,jordid@ac.upc.es

³*Polytechnic Institute of Castelo Branco,
Av. Pedro Alvares Cabral, n-12, 6000-084, Castelo Branco, Portugal*

Abstract Currently, a straightforward way to design BGP-based Traffic Engineering (TE) tools for stub Autonomous Systems (AS) is to rely on selfish routing mechanisms. Although TE tools can find an optimal solution, this optimum represents only a local optimum for outbound traffic. Indeed, this is one of the main limitations of the selfish routing approach. This approach makes the TE tools unaware of the effects of their route choices on transit AS throughout the chosen paths, due to uncoordinated routing decisions, and congestion can occur on distant intra- or inter-domain links. Thus, cooperation among AS is the key to avoid the performance degradation and routing instability caused by the selfish routing approach and it would be fundamental for the future QoS-aware Internet. With these objectives in mind this paper presents and discusses a framework for coordinated Inter-domain QoS Routing (QoSRS) decisions among stub and downstream AS taking into account multiple traffic QoS constraints and routing preferences. The paper includes a description of the main mechanisms and algorithms that integrate the framework, and finally a discussion of the implementation issues.

Keywords: Inter-domain QoS routing, Cooperative inter-domain routing.

1. Introduction

Nowadays Internet Autonomous Systems (AS) are largely operated by commercial entities. Part of the inter-domain traffic exchanges are governed by Service Level Specification (SLS) reflecting AS business inter-relationships. Stub AS (AS that don't carry transit traffic) demand robust SLS from transit AS to satisfy the Quality of Service (QoS) requirements of user applications. However, mainly due to economical reasons, transit AS delay the deployment of mechanisms to support the provision of multi-level QoS-aware connectivity services.

On the stub AS side, an emergent approach to either protect traffic aggregates from QoS degradation or in general optimize traffic performance and business goals (e.g. transit costs) is to deploy on-line Traffic Engineering (TE) mechanisms. These TE mechanisms can be located at specialized Smart Routing Managers (SRM) entities (also called Overlay Entities - OE) [Rekhter, 1995; Yannuzzi, 02] working on top of the Border Gateway Protocol (BGP). SRMs are out-of-BGP-band autonomous routing agents able to perform intelligent path selections satisfying multiple QoS constraints (e.g., latency and throughput). They operate on shorter timescales and exploit as much as possible AS multihoming and BGP capabilities. Specifically, SRMs are able to change BGP attributes (e.g. LOCAL-PREFERENCE) on-the-fly, reflecting the required traffic exchanges.

The benefits of such SRM-based approaches, operating on shorter timescales, are evident from the technical and economical perspectives and so they are also being developed as commercial products [Dai, 2003; Internap, 2005; Cisco, 2005]. However, these tools also rely on the selfish routing approach and are focused on outbound traffic optimization. They allow stub AS to perform outbound traffic optimization but the optimum is only a local one since the routes (or transit providers) are greedily selected.

Selfish routing approaches make SRMs unaware of the effects of their route choices on transit AS (e.g. on tier-3 or tier-4 providers) throughout the chosen exit paths, due to uncoordinated routing decisions. Namely, congestion can occur on distant intra- or inter-domain links, since it is very difficult to generate accurate forecasts of traffic demands to support resource distribution optimization [Feldmann, 2001]. Recent theoretical results by T. Roughgarden and E. Tardos [Roughgarden, 2002] confirm these limitations. The study revealed that important performance losses can be introduced by local optimization mechanisms comparing to global optimization. The cited work also shows that when general link latency functions are used the price of anarchy (that is the worst-case ratio between the total latency for the selfish case and the one for global optimization) is unbounded.

Poor routing stability is another significant shortcoming of the selfish approach. Uncoordinated routing reactions attempting to face congestion could produce cycles of influence between AS and thus routing instability. In the worst case, selfish routing mechanisms can cause extensive route flapping that can seriously affect the performance of large parts of Internet [Griffin, 2002]. Furthermore, the currently available techniques to enable AS to react and repair the congestion problems by means of inbound traffic control are limited to BGP in-band techniques. These techniques can assign new MED (Multiple-exit Discriminator) values or prepend the AS paths to indirectly signal congestion by influencing their neighboring link choices. Unfortunately, those techniques require external BGP updates and operate typically on TE management cycles of several minutes or hours. For that reason, they reveal to be more effective for persistent congestion problems. Moreover, some studies have documented problems and limitations such as MED oscillation and coarse-grained traffic control, arising by using these techniques [Griffin2, 2002; Chang, 2005].

Being aware that uncoordinated routing decisions, that are inherent to selfish approach, are on the origin of the above mentioned problems, a solution to help SRMs to alleviate the congestion caused by their selfish actions is widely desirable, namely, a solution to coordinate routing decisions among SRMs on behalf of stub and transit AS. Ideally, one global authority (in other words a global coordination mechanism) would be designed to support and guide SRMs on finding global optimums. Unfortunately, given the size and the complexity of today's Internet, this is inconceivable. Consequently, alternative approaches must be sought to cope with selfishness that don't rely on a global authority or on a global policy repository, even at some additional costs on signaling. Moreover, it could be expect that those approaches would be simpler as well as involve considerably lesser deployment costs as global ones.

To address the above issues, this paper presents a conceptually simple out-of-BGP-band cooperative approach able to support coordinated inter-domain QoS routing decisions among stub and transit AS for the context (but not limited) of the future QoS-aware Internet. The particularity of this approach is that it allows finding 'social' optimums by honoring both individual stub and transit AS preferences (or constraints) regarding the carrying (or admission of new) traffic amounts of certain traffic class (TC) aggregates. To summarize, the proposed approach is motivated and well-suited to:

- 1 Protect stubs from SLS violations of TC aggregates;
- 2 Reduce selfish costs (i.e. performance losses and instability);
- 3 Improve the traffic exchanges predictability;
- 4 Improve network resources utilization;

The remaining sections are organized as follows. In Sect. 2, a brief analysis of the related work is given. Then, in Sect. 3, the proposed cooperative approach is described, including the basic concepts and the complete algorithm for coordinated routing decisions. In Sect. 4, some implementation considerations are discussed. Finally, in Sect. 5 the paper is concluded and some directions for future (already ongoing) work are discussed.

2. Related Work

TE state-of-art is primitive [Feamster, 2003]. In previous work concerning intelligent routing schemes willing to control inter-domain traffic exchanges, two sets of TE mechanisms coupled with BGP have been proposed. The first set includes proposals for on-line and off-line BGP route optimization techniques. Among these proposals there are only few papers dealing with the design of algorithms for multi-objective (e.g. performance and cost) route optimization [Akella, 2003; Goldenberg, 2004]. The second set of mechanisms was designed to meet another central TE issue that is the attainment of smooth traffic distributions on egress links. In [Uhlig, 2004] this problem is addressed by using an evolutionary TE algorithm. One important aspect common to the cited works is that the proposed mechanisms operate on relatively large timescales of a few minutes to avoid frequent BGP route changes and the consequent BGP updates advertisements. This means that these techniques are not able to handle real-time metrics (e.g. in the order of magnitude of the network Round-Trip Time). In parallel to the cited scientific work, commercial products are also appearing. These products are globally known as smart (or optimized) edge routing and are able to operate on short timescales [Internap, 2005; Cisco, 2005]. In general commercial solutions try to select the most cost-effective routes or providers. In contrast with research tools, their internal details are unknown. For instance, their effective performance improvements and their impact on BGP performance is unclear. Another important common aspect is that these techniques behave as selfish routing schemes. On the other hand, solely two works attempt to coordinate route selections by using a negotiation strategy to avoid resource policy violations among AS [Mahajan, 2005; Winick, 2002]. However, both works are focused on the scenario of a pair of multi-connected ISPs (Internet Service Providers).

The proposal of the present paper differs from the above described techniques in the following aspects 1) it is based on a cooperative approach as part of (but not limited to) an overall solution to address inter-domain QoSR issues; 2) it is able to handle SLS violations by the automatic reconfiguration of BGP routers parameters; 3) it is focused on coordinated path selections; 4) it takes into account both individual stubs and transit AS preferences and constraints.

3. Cooperative Approach

As stated previously the inter-domain QoSR framework proposed in this paper is based on SRMs working on top of BGP. Due to scalability concerns the proposed framework focuses on the control of routing between pairs of multi-homed stub AS [Yanuzzi, 02]. The key idea is that the peering SRMs belonging to remote stub AS, would like to co-operate in a reflective (mirroring) manner to share QoS measurement data and to manage traffic exchanges between both AS, such as the traffic goals of TCs en route to the remote stub can be fulfilled. Examples of such goals for a TC are the latency minimization and throughput maximization.

This cooperative approach for inter-domain QoSR allows the incremental deployment because it does not require cooperation with intermediate AS and most of the introduced complexity is located in the edges of the Internet. Unfortunately, similarly to the existing scientific and commercial SRM-based mechanisms, it suffers from the selfish shortcomings, described so far. This section presents an extension of this cooperative approach able to perform coordinated QoSR decisions.

3.1 Basic Concepts

As mentioned in Sect. 1, rather than assuming that a global authority or a global policy repository exists, our approach is based on a simpler route coordination strategy for finding 'social' optimums, that honors both individual stub and transit AS preferences (or constraints). The design of a routing mechanism relying on this approach is achieved by introducing extensions to SRM's route controller mechanisms to enable traffic adaptation making use of information on both QoS measures of TCs and the corresponding acknowledged transit AS routing preferences. For sure, it is assumed that stub AS are willing to sacrifice themselves for the benefit of the overall routing environment. Figure 1 illustrates a possible scenario for the cooperation among remote SRMs. The figure also illustrates the exchange of signaling messages carrying transit AS preferences. For clearness the SRMs mechanisms located in transit AS are omitted.

Let us assume that each multi-homed stub subscribes among K providers, several connectivity services for transit traffic classified into M TCs. Similarly to DiffServ (Differentiated Services) framework, we use the concept of TC, denoted as TC_n , where n is the class index and $n = 0, \dots, M - 1$ for M TC, to divide the network traffic (e.g., premium, gold, silver and bronze) [Blake, 1998]. Moreover, we do not consider any special notation to distinguish among the paths used to route packets classified in each one of M TCs. We simply denote a path as $P_k, k = 1, \dots, K$.

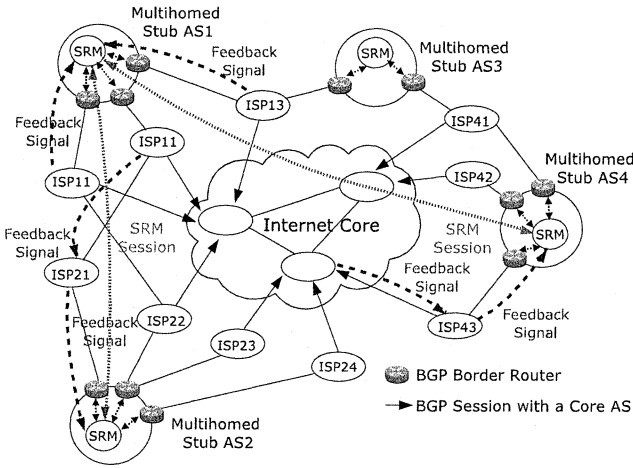


Figure 1. Illustration of the proposed approach

Assuming that all QoS interconnections are established and also that default active paths to target destinations are inserted into local RIBs (Routing Information Base) (i.e., Local-RIB [Rekhter, 1995]) of stub's BGP routers, let us describe our route coordination strategy. Figure 2 outlines a useful model of the route coordination strategy. For simplicity, consider the case that an SLS violation or else that a significant QoS change occurs on a certain active path P_k carrying the traffic classified into a certain TC_n . Upon detecting one of these events, an ingress SRM performs the following steps:

- 1 First, it obtains all available QoS feasible paths $P_{k'}$ for the traffic aggregate affected, alternative to the active one P_k . In the descriptions we assume that a $P_{k'}$ is feasible if $\forall c_i(P_{k'}) < C_i, i = 1, 2, \dots, m$, where c_i is the i_{th} additive QoS constant assigned to $P_{k'}$ and $C_i > 0$ the corresponding i_{th} constraint. We denote the set containing all $P_{k'}$ as A .
- 2 It then requests to the SRMs of transit AS throughout each path $P_{k'}$ their preferences regarding the admission of the affected TC aggregate, TC_n .
- 3 Next, after obtaining all preferences from transit AS it joins them together with the corresponding $c_i(P_{k'})$ constants, and map them into the parameters of the route rank function, $rank()$. The $rank()$ basically implements the SRM's route decision process and manages the SRM's route ranking tables for each TC_n to target prefix destinations.
- 4 Finally, after $rank()$ outputs the 'social' optimum path, it sets the local BGP routers to select this path. Of course, if the outputted path differs

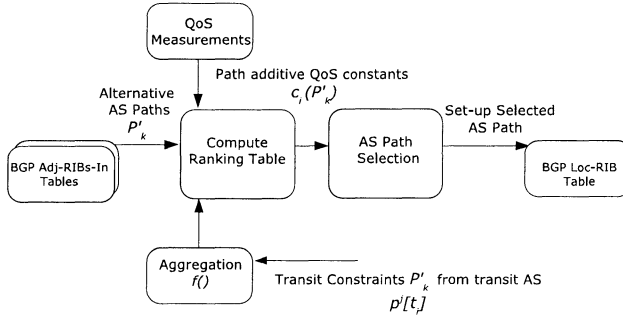


Figure 2. Model of the Route Coordination Strategy

from the current active path, it withdraws that path from BGP Local-RIBs. In principle, at this stage, the IP packets classified in TC_n are routed into the best QoS interconnection.

This process is continuously repeated, step-by-step, for each traffic aggregate affected by degradation. Concerning transit AS (as it will be detailed in next subsection) and besides giving feedback about preferences to stub SRMs, SRMs of transit AS have the job of coordinating link change decisions among neighbor AS as well.

There are some basic concepts concerning the overall approach and the design of this kind of coordinated routing mechanisms as follows.

AS Routing Preferences: For the effectiveness of our route coordination strategy, we assumed that the transit AS members of all available paths at stub (i.e., from all P_k and $P_{k'}$), are willing to co-operate by reporting their routing preferences. An AS routing preference is data structure similar to a cost, denoted as p^j , which represents the cost reported by j_{th} transit AS throughout a path P_k . It may reflect either the estimated overall congestion state of an AS on carrying a new traffic amount for a certain TC or (after admission) an estimation of the current AS's congestion state regarding the carrying the traffic aggregate. In practice, this concept easily allows transit AS to influence how traffic enters into their networks targeting to both maximize of their network utilization and to fix congestion problems.

Coordination Metric: We introduce the concept of a coordination metric to support transit AS on periodical generation and expression of their AS routing preferences. The coordination metric values are computed by an ordinal function $M()$. Thus, a given preference could be expressed as $p^j[t_i] = M(v^j)$. In terms of this notation $p^j[t_i]$ is the routing preference instance reported by the j_{th} transit AS at instant t_i for traffic classified in a certain TC; and v^j represents a given vector of internal information of that peer being mapped into p^j . As example, a natural $M()$ function is the Excess of Bandwidth Demanded

(EBD) for TC_n for all individual customers, denoted as E_n . Precisely, if at the provider j the sum of the bandwidth demanded for TC_n is D_n and S_n is the bandwidth supplied by the provider for a certain TC_n . So thus, $v^j = (D_n, S_n)$ and, then, the EBD reported to customers is given by $E_n = D_n - S_n$.

Furthermore, the coordination metric is an essential mechanism to handle the business constraints on the disclosure of internal information (e.g. internal configurations or even policy details). We suggest the use of functions $M()$, for which the problem of finding its inverse $M()^{-1}$, is computationally unfeasible.

Preference Aggregation and Feasible Outputs: At stub AS, SRMs consider that the gathering operation of the AS routing preferences along a certain QoS feasible path is well-succeeded when every preference from all AS members is gathered. The result of all well-succeeded gathering operations is a subset of alternative paths, i.e., $R \subset A$. It is assumed that all collected preferences $P^{k'} = \{p^1, \dots, p^j, \dots, p^n\}$, for a $P_{k'}$, where p^j is the AS routing preference from the j_{th} transit AS, are aggregated using a function $f()$. In case of additive routing preferences, the result of the aggregation denoted as $\zeta^{k'}$ and defined as $\zeta^{k'} = f(P^{k'})$, is equal to the sum of the p^j values of $P^{k'}$. For nonadditive routing preferences, $\zeta^{k'}$ can be the minimum (or maximum) among the p_k elements in $P^{k'}$. Alternatively, the cumulative routing preferences can be calculated hop-by-hop across transit AS throughout a $P_{k'}$. This has the advantage of reducing signaling and local processing cost. However, this complicate the location of congestion points. In our model, the subset R (including the associated $\zeta^{k'}$ values and $c_i(P_{k'})$) is the set of feasible outputs of the coordination mechanism. In other words, these paths are the elements of the SRM's route ranking table for TC_n to the target prefix destination.

Path Selection Triggering and Routing Optimization Problem: The following questions regarding the dynamic traffic adaptation process must be answered during the coordination mechanism design: *What is the approach followed for the path selection triggering? What is our routing problem? How can we solve it?*

As described so far, SRMs manage routing patterns by composing and updating multiple ranking tables. To achieve these objectives, first SRMs actively check QoS bounds and transit AS constraints from paths in BGP Adj-RIBs-In tables, whether they are installed into BGP Local-RIBs or not. Then, whenever a significant variation on any constraint value (e.g., on $c_i(P_{k'})$ constant or on a p^j value) is registered or in response to abnormal network operation events, the $rank()$ function is called and thus the path decision process is triggered to remove or insert a path into the ranking or simply to update the order of the paths.

Although, the proposed approach imposes higher CPU and memory consumption in the SRMs, as well as additional load on the network due to SRMs signaling messages, it is characterized by reduced response times, since all




Figure 3. SRM recovering a QoS affected Path

alternative paths are "pre-computed". To bound the implementation costs according to TE guidelines SRMs should focus on engineering the paths to a small fraction of the prefix destinations (e.g., popular prefix destinations or on stable traffic volumes) [Feamster, 2003].

Naturally, SRMs must be aware of traffic goals, of the precise routing problem and how to solve it. In other words, this means that SRMs must be able to select the feasible path $P_{k'} \in R$ that maximize its degree of satisfaction regarding *multiple independent objectives* (i.e., $c_i(P_{k'})$ constants and $\zeta^{k'}$), among the set of feasible outputs R . However, this MCP (Multi-constrained Path) problem is known as being NP-complete [Garey, 1979]. To cope with the NP-completeness of m -constrained problems, researchers had focused on the development of polynomial time heuristics [Kuipers, 2002]. An SRM's route controller dealing the MCP problem should follow a similar approach, that is, to integrate a computationally efficient optimization algorithm.

3.2 Interactions between Smart Routing Managers

Rather than assuming an identical arrangement of AS independently of the existing inter-relationships (like BGP does), our out-band approach allows our interdomain QoSR model to adopt a different strategy to improve stability and scalability. We propose two levels of interactions between the SRMs. In the first level SRMs perform coordinated path change decisions, as described so far. Contrary, in the second level the SRMs perform coordinated link changes decisions to reallocate TC aggregates among links shared by pairs of AS. Figure 3 illustrates these interactions.

The clustering for the first level of interactions between SRMs is obtained by identifying the AS, which have previously agreed to perform coordinated

path change decisions. On the other hand, the clustering for the second level interactions between SRMs is obtained by identifying pairs of dense AS, which have previous agreements to perform coordinated link change decisions. The later arrangement is motivated by recent tomography studies showing that transit AS are usually very dense and normally share among them multiple links [Map, 2005]. In both cases, the addresses of SRMs and their capabilities can be encoded in a new BGP attribute (e.g., an extended community attribute).

When coupling both levels of interactions, instead of immediately changing an affected traffic aggregate to an alternative path, an SRM explicitly spawns a QoS degradation warning message on the current path. Consequently, this requires the pairs of SRMs with second level relations to seek for alternative links able to improve the current offered QoS. This feature has the advantage of keeping the current path in the case of a successful recovery operation pre-required, avoiding thus unnecessary path shifts and the corresponding BGP updates. This process is supported by the BGP path concept to be agnostic of any detail about AS interconnections.

3.3 Coordinated Link Changes

In the coordinated link changes case, the ingress SRM adapts the TC aggregates among the alternative routing options (i.e. the alternative links) following a similar process to the path change case. In contrast, this is accomplished by exchanging of link preferences between the egress SRM and its peer.

Among all available routing options, the ingress SRMs must select the ones that don't hurt downstream peers of the egress AS. As depicted in Fig. 4, the routing option that implies the use of different egress points at the egress AS could be selected as the solution. Therefore, congestion therefore could occur in downstream peers of the egress AS along the select path the target prefix destination. This in turn could trigger a subsequent cascade of coordinated link decisions processes before achieving convergence. To avoid these effects and to alleviate BGP as well, one solution is the egress SRM to signal its peer to use only the equivalent links (i.e., the ones whose the traffic changes uses the same egress point) [Winick, 2002]. Alternatively, to produce the same effect and reduce signaling costs, we prefer to apply a condition to the alternative links (similar to the use of a drain plug) to obtain the subset of routes composed only by the routes which keep the current egress point. Finally, only the corresponding routing preferences are reported by the egress SRM to the ingress SRM. We call this subset the Null Effect Route Subset (NERS). In practice, NERSs are obtained simply passing the routes by a tuned filter.




Figure 4. Illustration of NERS paths

3.4 Algorithm for Coordinated Routing Decisions

The replacement process of routes impacted by performance degradations is accomplished in several steps. As described above, after path selection is triggered, SRMs start a new routing cycle. Upon that, this cycle finishes after a 'social' optimum path is computed and established. If no 'social' optimum path can be found we assume that SRMs can perform greedy route selections as a last chance to improve the QoS. Figure 5 shows the complete algorithm for coordinated routing decisions. The modularity and independency of exact implementation details about coordination metric computation, aggregation functions, traffic objectives and optimization algorithms and finally signaling, which are part of our mechanism, are fundamental requirements. Therefore, the presented algorithm mainly focuses on the flow of SRMs functions execution and on interactions between SRMs, themselves and BGP routers.

4. Implementation Considerations

This section contains some considerations related to the SRMs functions needed to support the proposed cooperative approach. The section includes considerations related to QoS measurements, gathering of alternative paths, setting-up the selected path and signaling.

QoS Measurements: We assume that both remote SRMs have the knowledge of the existing SLSS, namely the QoS requirements for traffic exchanges among their stubs. To support path evaluation and SLS violation detection, we adopt a strategy based on active end-to-end measurements of QoS parameters. So thus, SRMs must incorporate efficient measurement methods following the recommendations of recent standardization efforts [Shalunov, 2004; Almes, 1999]. In the transit AS side, we assume that there is a general purpose measurement system to provide local traffic characteristics and periodically monitor the usage of the resources previously allocated to each service. This data is then used to compute AS routing preferences to be reported within feed-

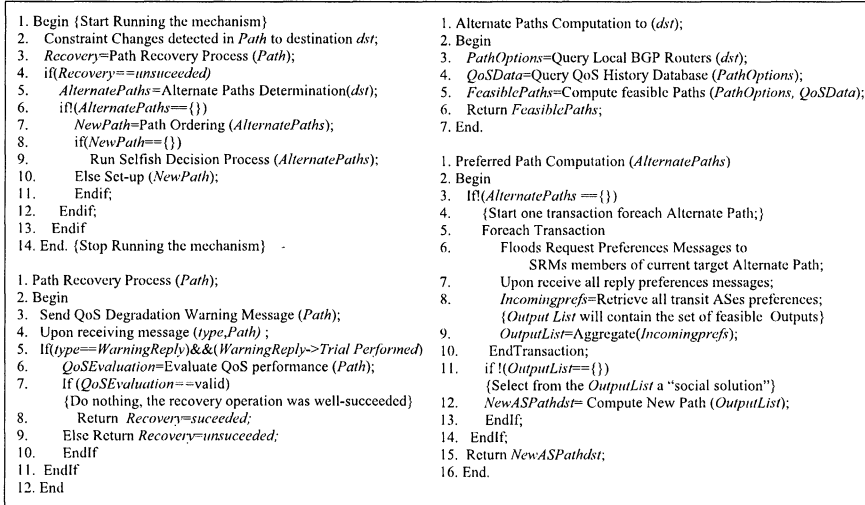


Figure 5. Coordination Mechanism Algorithm

back signaling messages sent to stub SRMs. Recent efforts on developing of ISP-based capturing tools are found in [IPMON, 2005; ATT, 2005].

Gathering of an alternative path set: A fundamental requirement to deploy our out-band approach is that SRMs must have administrative control over BGP speakers and thus full access to the Routing Information Base (RIB), namely to the Adj-RIBs-In and the Loc-RIB databases. A precise definition of the terms Adj-RIBs-In and Loc-RIB can be found in [Rekhter, 1995]. Furthermore, in the QoS context the RIB must be extended to hold all P_k for all TCs into the QoS extended BGP Adj-RIBs-In table from peer at k_{th} -provider. The gathering of alternative paths able to accommodate a traffic aggregate affected by a strong QoS degradation is an essential function. These paths and the corresponding next-hops are thus retrieved from those QoS extended Adj-RIBs-In, depending on the QoS measurements history.

Setting-up the Selected Path: The final step of the proposed mechanism is to set-up the selected path as the active path. This can easily be done by installing the route into the QoS extended BGP Loc-RIB table. Rather than this, our proposal is to enable SRMs to modify the IP forwarding tables directly. This enables to create a soft state routing allowing SRMs to rollback routing decisions and to avoid overload BGP routers during instability episodes. The new routes are inserted into the BGP Loc-RIB only when they are considered as stable.

Signaling: Communication between SRMs is asynchronous and of limited network bandwidth. A signaling protocol is required. This must include requests, replies and acknowledgments to gather AS routing preferences, QoS

data sharing or other actions. Messages are sent to individual or groups of SRMs. In addition, recent recommendations in the proposal of an IP signaling protocol with QoS extensions, should be taken into account in the deployments [Hancock, 2004].

5. Conclusions and Future Work

In this paper, we have proposed an cooperative framework for coordinated inter-domain QoS routing decisions. The particularity of our framework is to rely on out-BGP-band routing mechanisms and signaling to support stubs on finding 'social' optimums by honoring both individual stub and transit AS preferences. We believe that this kind of approach is the only way to allow predictable inter-domain traffic exchanges between AS and to support robust SLSS in the Internet environment. As discussed in the paper, two main set of open issues are still part of our research agenda. First, we intent to design a specific coordination metric, and efficient traffic optimization algorithms to find 'social' optimums. Secondly, to ensure scalability, it is essential to build hierarchical SRMs organizations and to define their relations in order to design a signalling protocol for SRMs data sharing and actions' requests and acknowledges. Finally, in a simulation environment based on J-Sim [Jsim, 2005] and Infonet BGP suite [Infonet, 2005], we plan extensive simulations contrasting the behavior of the described cooperative approach with current in-band BGP mechanisms. The results collected will enable the evaluation of the strengths and limitations of the contributions and will lead to refinements.

Acknowledgment

This work was partially funded by the European Commission through Network of Excellence E-NEXT (contract FP6-506869) under SATIN Grant *Study of Coordination Mechanisms and Signaling Protocols for Inter-domain Quality of Service Routing in a Distributed Overlay Entities Architecture*.

References

- A. Akella, B. Maggs, S. Seshan, A. Shaikh, and R. Sitaraman. *A measurement-based analysis of multihoming*. In Proc. of ACM SIGCOMM 2003, August 2003
- G. Almes, S. Kalidindi, M. Zekauskas, *A one-way delay metric for IPPM*, IETF, RFC 2679, September 1999
- Network Measurement Tools. <http://www.research.att.com/>. Web page accessed at August 2005.
- S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang and W. Weiss, *An Architecture for Differentiated Services*, RFC 2475, December 1998.
- Cisco Systems, *Optimized Edge Routing*. <http://www.cisco.com/>. Web page accessed at August 2005.
- R. Dai, D. Stahl, and A. Whinston. *The economics of smart routing and QoS*. In Proc. of the Fifth Inter. Workshop on Networked Group Comm. (NGC'03), 2003.

- Internap Network Services. *Internap Flow Control Platform*. <http://www.internap.com/>. Web page accessed at August 2005.
- R. Chang and M. Lo, *Inbound Traffic Engineering for Multihomed ASs Using AS Path Prepending*. IEEE Network, March/April 2005.
- N. Feamster, J. Borkenhagen, and J. Rexford. *Guidelines for interdomain traffic engineering*. ACM SIGCOMM Comput. Commun. Rev., 33(5) : 19 to 30, 2003.
- A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, *Deriving traffic demands for operational IP networks: Methodology and experience*. IEEE/ACM Trans. Networking 9, June 2001.
- M. Garey and D. Johnson, *Computers and Intractability, A guide to the theory of NP-Completeness*. Freeman, San Francisco, 1979.
- D. Goldenberg, L. Qiu, H. Xie, Y. Yang, and Y. Zhang. *Optimizing cost and performance for multihoming*. In Proc. of ACM SIGCOMM 2004, August 2004
- T. Griffin. *What is the Sound of One Route Flapping?* IPAM talk, 2002
- T. Griffin, and G. Wilfong. *Analysis of the MED Oscillation Problem in BGP*. In Proc. of the 10th IEEE International Conf. on Network Protocols (ICNP'02), 2002
- R. Hancock, G. Karagiannis, J. Loughney and S. Bosh. *Next Steps in Signalling*. IETF, Internet-draft, November 2004
- Infonet Suite Homepage. <http://www.info.ucl.ac.be/bqu/jsim/>. Web page accessed at August 2005.
- IP Monitoring Project (IPMON). <http://ipmon.sprint.com/>. Web page accessed at August 2005.
- J-Sim Homepage. <http://www.j-sim.org>. Web page accessed at August 2005.
- F. Kuipers, P. V. Mieghem, T. Korkmaz, and M. Krunz, *An Overview of Constraint-Based Path Selection Algorithms for QoS Routing*, In IEEE Communications Magazine, December 2002
- R. Mahajan, D. Wetherall and T. Anderson. *Negotiation based routing between neighboring domains*. In Proc. of NSDI, May 2005.
- The Internet Mapping Project. <http://research.lumeta.com/ches/map/>. Web page accessed at August 2005.
- Y. Rekhter, T. Li. *A Border Gateway Protocol 4 (BGP-4)*. IETF, RFC 1771, March 1995
- T. Roughgarden and E. Tardos. *How bad is selfish routing?* Journal of ACM, 49(2):236 to 259, 2002.
- S. Shalunov, B. Teitelbaum. *One-way active measurement protocol (OWAMP) requirements*. IETF, RFC 3763, April 2004
- S. Uhlig, O. Bonaventure, *Designing BGP-based outbound traffic engineering techniques for stub ASes*. ACM SIGCOMM CCR, October 2004
- J. Winick, S. Jamin, and J. Rexford. *Traffic engineering between neighboring domains*. <http://www.research.att.com/jrex/papers/interAS.pdf>, July 2002.
- M. Yannuzzi, A. Fonte, X. Masip, E. Monteiro, S. Sanchez, M. Curado, J. Domingo. *A proposal for inter-domain QoS routing based on distributed overlay entities and QBGP*. In Proc. of WQoS'04, LNCS 3266, October 2004