Self-Adaptive Routing: The Inevitable Evolution of Interdomain Route Optimization Tools

M. Yannuzzi, X. Masip-Bruin,
S. Sánchez-López, J. Domingo-Pascual
Advanced Broadband Communications Center
Technical University of Catalonia
Avgda. Víctor Balaguer, s/n – 08800 Vilanova i la Geltrú
Barcelona, Catalonia, Spain
Email: {yannuzzi, xmasip, sergio, jordid}@ac.upc.edu

A. Fonte[1,2], M. Curado[1], E. Monteiro[1]
[1]Laboratory of Communications and Telematics
University of Coimbra
Pólo II, Pinhal de Marrocos, 3030-290 Coimbra, Portugal
[2]Polytechnic Institute of Castelo Branco, Av. Pedro Álvares
Cabral, nº12, 6000-084 Castelo Branco, Portugal
Email: {afonte, edmundo, marilia}@dei.uc.pt

***Abstract*—Multihoming in combination with route optimization tools working in short timescales is becoming a common practice in order to improve end-to-end QoS at the edge of the Internet. Nevertheless, the stability repercussions of these selfish practices under massive utilization are absolutely unpredictable. Therefore, self-adaptive mechanisms will be inevitable if masses of completely independent and uncoordinated multihomed stub Autonomous Systems are allowed to simultaneously change their interdomain traffic distributions seeking only for the best of their own purposes in short, and even very short timescales. In this paper we first investigate the main requirements and expected functionalities of the future self-adaptive route optimization tools from a practical viewpoint. Then, based on this analysis we propose the first self-adaptive route optimizers for an interdomain routing framework that we developed in a previous work. Our main contribution is that these self-adaptive route optimizers are able to exploit the capabilities of BGP with the aim of enhancing the performance of delay-sensitive applications, while providing significant improvements in terms of overall network stability.**

Keywords-component; Interdomain Edge Routing, Multihoming, BGP, Stability, Self-Adaptive.

I. INTRODUCTION

Multihomed stub Autonomous Systems (ASes) are claiming for cost-effective mechanisms to manage and distribute their interdomain traffic in order to exploit their multi-connectivity and hence improve their performance. For this reason, solutions generally known as route optimizers are starting to become widely deployed at the edge of the Internet [1-3]. These solutions are targeted to the largest fraction of ASes in the Internet, which crowds together mostly medium and large Enterprise Customers (ECs), Content Service Providers (CSPs), and small Network Service Providers (NSPs). Route optimizers are attractive for multihomed stub ASes since they are able to operate in very short timescales. However, the implications of rearranging interdomain traffic in such timescales are completely unpredictable in terms of global stability, especially, when the number of sources simultaneously injecting these selfish perturbations to the network grows large. Given that no entity will exist to control and coordinate these massive perturbations, self-adaptive mechanisms will become unavoidable in the near future [4].

Consequently, the goals of this paper are, first, to identify the main requirements for self-adaptive route optimization tools and the functionalities that could be integrated in such tools, from a practical point of view. This is surveyed in Section II. Our second goal is to show that none of the currently proposed solutions fulfills all those requirements, especially from a stability standpoint. Thus, Section III analyses related work and based on this study we emphasize the motivation behind this work. Our final and most important goal is then to propose and test a self-adaptive route optimization scheme, which fulfills all previous requirements and helps BGP to improve end-to-end QoS in short timescales. The proposal and its evaluation are presented in Sections IV and V respectively. Finally, Section VI points out some future directions and concludes the paper.

II. REQUIREMENTS FOR THE FUTURE SELF-ADAPTIVE ROUTE OPTIMIZERS: A PRACTICAL APPROACH

In this section we identify six main requirements for the future self-adaptive route optimizers. We have assumed that these optimizers will still be capable of managing interdomain traffic at the edge of the Internet in short timescales, and in a selfish manner. We have also assumed a pragmatic position, which is supported by the cited references. Last but not least, we have tried to highlight some expected functionalities that could be easily, and advantageously, integrated in the future route optimization tools. The main set of requirements is:

1. Exploit multihoming as much as possible in order to improve end-to-end performance

Multihomed stub ASes could largely benefit from cost-effective tools that rapidly, and incessantly, let them make a better use of their multi-connectivity to the Internet. These kinds of solutions are not applicable to large transit ASes, such as Tier1 or Tier2 Internet Service Providers (ISPs). This is because the effects of greedily managing large amounts of interdomain traffic in short timescales are unpredictable. In fact, most ISPs will keep relying on rather static Traffic Engineering (TE) provisioning techniques. Accordingly, this is

This work was partially funded by the Spanish Ministry of Science and Technology under contract FEDER-TIC2002-04531-C04-02, the Catalan Research Council under contract 2001-SGR00226, and the European Commission through E-NEXT under contract FP6-506869.

a strong requirement but for multi-connected ASes located at the edge of the Internet, which typically gathers together CSPs, medium/large ECs, and also small NSPs.

2. Exploit BGP capabilities

Compelling recent work demonstrates that it is not necessary to circumvent BGP routing in order to improve end-to-end performance [5]. In fact, whereas several multihomed stubs are starting to use route optimization tools, it remains to be seen if the overlay proposals that one can find in the literature will ever become deployed at the AS-level. Thus, we foster incremental approaches such as route optimizers, which smartly exploit the capabilities of BGP to get better performance. However, tweaking interdomain traffic in short timescales by means of BGP is only feasible for the outbound traffic from the AS. This is because tuning BGP in short timescales for the inbound traffic may lead to BGP damping [6]. Nevertheless, small NSPs, and especially, CSPs and ECs could take full advantage of these BGP-based solutions.

3. Provide transparency

Network administrators of multihomed stub ASes are not willing to adopt complex mechanisms to manage their interdomain traffic. Moreover, they do not want to get into the details of how and when their traffic should be rearranged. They simply want to take plain decisions, and they expect that these decisions last in time. Based on these demands, we consider that the future self-adaptive route optimizers must provide transparency to the traffic reallocation decision process. Fig. 1 depicts the advantage of self-adaptive mechanisms from this perspective. A self-adaptive mechanism may gather QoS information (using passive and/or active measurements) and adapt/hide the QoS dynamics from the traffic reallocation decision process. This supplies an appealing solution, since administrators may decide how conservative or opportunistic they want to be in terms of their traffic reallocations by making a simple decision. For instance, if a certain threshold is met and the network conditions are adequate then allow the reallocation of traffic. In such a case, network administrators would simply need to select and configure this fixed threshold. It should become clear that how conservative or opportunistic an ASes will actually be strongly depends on the network dynamics, and so this may vary over time. In contrast, the use of self-adaptive mechanisms in combination with the selection of fixed thresholds allows these ASes to straightforwardly decide how opportunistic they are willing to be, and this decision will last in time.

4. Global stability

The relevance of a self-adaptive tool is in its strengths in terms of supplying local and global stability. Under highly changing network conditions, it is imperative that each route optimizer counts with a self-adaptive mechanism allowing it to learn from those dynamics, and diminish or even prevent the number of path shifts until the network conditions are once again stable. Such network conditions will certainly occur if masses of non-self-adaptive route optimizers are allowed to simultaneously perturb the network, or they could even occur during link flaps, or routing misconfigurations. Indeed, multihomed stub ASes not using self adaptive mechanisms may find that the number of traffic reallocations they are actually allowing could be much higher than the expected. In other words, under highly changing network conditions even a conservative opportunistic approach may lead to network instability. As an alternative, self-adaptive mechanisms are able to adapt themselves to those changing conditions, so that they could be able to reflect the choices made by network administrators independently of the network dynamics.

5. Control only a reduced set of the available paths

Compelling recent studies like [7] demonstrate that the problem of tracking and controlling most of the traffic of multihomed stub ASes turns out to be impracticable. This is because the large variability of the topological characteristics of interdomain traffic, in addition to the limited aggregation of this traffic, indicates that the number of paths to be tracked and controlled is not only highly variable, but also very large. Despite these variability and lack of aggregation issues, several recent studies also show that a very small number of invariant paths are still responsible for a significant fraction of the existing traffic [7, 8]. By invariant we mean that these paths are stable, i.e., they are, typically, permanently present in the BGP tables and hence are not affected by the variability issues mentioned before. For instance, the measurements conducted in [7] reveal that only six invariant AS paths carried about 36% of the one-month total traffic of a real multihomed stub AS, which is indeed a significant fraction of the overall traffic. Thus, a realistic approach is that it is possible to track and optimize a significant portion of the existing traffic by simply controlling a reduced set of stable paths (typically 6-10 paths).

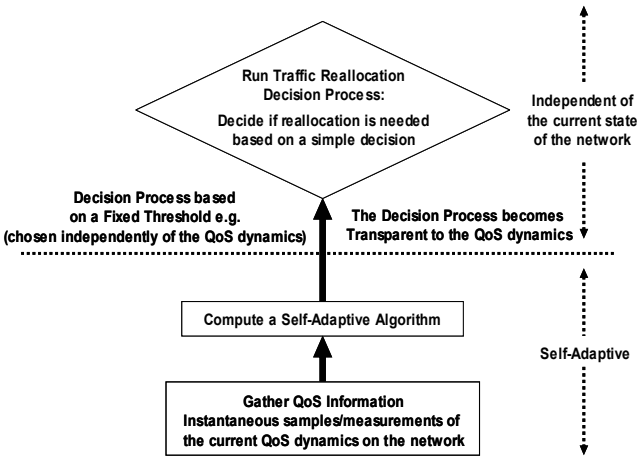

Fig. 1. A Self-Adaptive algorithm provides transparency to the traffic reallocation decision process.

An important issue is that multi-connectivity to the Internet does not necessarily guarantee improved end-to-end path diversity. This is due, first, to the fact that BGP only advertises the best path it knows, so BGP considerably prunes the total number of available paths between distant ASes, and second, to the topological characteristics of the Internet at the AS-level. Even though several studies have addressed these scarce path diversity issues [9, 10], recent studies like [5] demonstrate that in practice multihoming in combination with route optimizers are powerful tools to improve end-to-end performance. A sound explanation for this is that the Internet core is in effect an over-provisioned network.

6. Cooperation among remote ASes

This is indeed expected for a number of reasons. First, recent studies show that interdomain traffic is mainly exchanged among ASes that are not directly connected. Instead, they are typically 2, 3 and 4 AS hops away [8], and this applies as well to the reduced set of paths to be tracked and optimized in 5. Second, cooperation between remote ASes is expected for both performance and economical reasons. Indeed, such scenario is perfectly suitable, for example, for medium and large ECs with multi-connected premises using different ASes. In such cases, self-adaptive route optimizers primarily need to improve the performance but for the traffic exchanged among the enterprise sites. Once again, the number of relevant paths to be tracked and controlled per-site will be typically small. For this reason, strong manufacturers such as Cisco have recently announced that they will go in this direction [2].

III. Related Work and Motivation

Several research efforts are being carried in order to improve end-to-end performance at the interdomain level. However, none of the current proposals is able to provide end-to-end QoS in short timescales for multihomed stub ASes, and cope at the same time with the global stability issues described in Sections I and II. Indeed, none of the available proposals simultaneously fulfills all the requirements presented in Section II. These proposals could be primarily divided into two different groups.

1. In-band solutions

This first group gathers together solutions trying to enhance BGP with new capabilities, such as TE and QoS extensions. All these in-band solutions, i.e., solutions intrinsically supported and signaled using BGP, are able to supply significant improvements in terms of rather static QoS or TE provisioning. However, they are inadequate to handle the rearrangement of interdomain traffic in short timescales. The reasons for this are that: (a) BGP is a slow reacting protocol [11]; (b) such approaches would significantly increase the number of messages exchanged between BGP peers, which may lead to network instabilities [6]. Thus, these solutions are not able to cope with the current demands of multi-connected stub ASes, and mainly fail to fulfill requirements 1 and 4 in Section II.

2. Out-of-band solutions

This second group is conformed by solutions tending to decouple the complexity of QoS or TE provisioning from BGP devices. These out-of-band solutions, i.e., solutions not intrinsically supported and signaled using BGP can be in turn divided into two different types of approaches.

2.1 Overlay networks

Overlays are able to improve end-to-end QoS in several ways at the cost of circumventing BGP routing. However, they are not capable of tackling down the central issues of QoS provisioning and overall network stability at the AS-level. In effect, not even the most compelling proposals in the area, such as [12, 13], have been implemented in real practice at the AS-level.

Akella et al [5] have recently compared overlay routing to multihoming route control using route optimizers, and their results clearly show that it is not necessary to add the complexity of bypassing BGP routing to achieve good end-to-end performance. Furthermore, the stability repercussions of massive deployment of overlay networks at the AS-level are absolutely unpredictable.

Thus, current overlay proposals are not suitable as route optimization tools at the AS-level, since they are not capable of fulfilling requirements 2 and 4 in Section II.

2.2 Route optimizers

Two types of route optimizers are emerging, namely, DNS-based optimizers [1], and BGP-based optimizers [2-3]. The DNS-based solutions are addressed to small organizations which are not willing to deal with the difficulties of BGP peering and management, and are out of the scope of this work.

Given the limitations of the aforementioned related work and the potential strengths of BGP-based route optimizers [4, 5], we foster the utilization and deployment of these solutions, but highlighting that the perturbations they cause in the network should be timely controlled. Current route optimizers are quite limited. Above all, they lack of mechanisms that allow them to learn from the network dynamics and "socially" deal with the global stability issues of greedily managing interdomain traffic in short timescales. In addition, none of the available route optimizers fulfills requirements 3, 4, and 6 in Section II.

Thus, the main motivation for this work is that even though commercial route optimizers are becoming widely deployed at the edge of the Internet, none of the state-of-the-art proposals is able to fulfill the main requirements and expected functionalities of these route optimizers. As a consequence, new self-adaptive routing mechanisms need to be developed for these route optimizers.

IV. A SELF-ADAPTIVE EDGE ROUTING SCHEME

In [14] we proposed a distributed architecture in which a pair of Smart Routing Managers (SRMs)[1] within two non-peering multi-homed stub ASes were able to: (a) exchange a Service Level Specification (SLS) regarding the traffic among them; (b) examine the compliance with this SLS; (c) and accurately configure on-the-fly BGP to avoid link failures, or service degradation for a set of Classes of Service (CoSs). The foremost motivations for influencing traffic in this way are rooted on requirements 1, 2, 5 and 6 in Section II. The essence in this approach is that the QoS perception between a pair of remote ASes in our scheme is basically the one that the SRMs have of each other.

This scheme has several advantages. First, with only a very small number of route optimizers (the SRMs), but located at strategically selected remote multihomed ASes, is enough to control a significant portion of the traffic of an AS (requirement 5 in Section II). Second, no SRMs are needed in any transit AS connecting the remote ASes in our scheme. Our approach is that a SRM within a source AS dynamically manages BGP in order to control the allocation of its outbound traffic towards a remote AS in our scheme, depending on the network conditions and QoS constraints for each CoS. This allows tweaking BGP even in very short timescales given that no BGP messages will be ever spawned.

Despite the strengths presented above, the current SRMs lack of mechanisms that allow them to learn from the network dynamics, and adapt themselves in order to address the stability issues presented so far. Therefore, in this Section we propose a novel self-adaptive SRM with the aim of covering requirements 3 and 4 of Section II. This self-adaptive SRM is especially designed to improve the performance of delay-sensitive applications, such as Voice of IP (VoIP) or video applications. We took this approach because these are major applications which are perfectly suitable for collaborative edge routing schemes, such as medium/large ECs with multihomed premises in different locations. Therefore, we choose the One-Way Delay (OWD) [15] as the end-to-end QoS information that the SRMs will collect from the network (Fig. 1).

These new SRMs are designed to perform opportunistic routing management, so they will take full advantage of multihoming allowing the reallocation of traffic for a given CoS each time a sufficiently better end-to-end path exists. This opportunistic behavior imposes a trade-off in the timescale of the traffic reallocations, since an excessively conservative approach may under-utilize network resources, whereas frequently shifting traffic may lead to overall network instability. To cope with this problem our opportunistic SRMs have in-built self-adaptive features.

In order to gather the OWD information (Fig. 1), the SRMs are endowed with mechanisms to spawn small probes targeting the reduced set of remote ASes in our scheme. To fix ideas about the burden behind this practice, typically, each SRM will only need to probe 6-10 remote SRMs (requirement 5 in Section II), for 2-3 CoSs, and through 2-3 egress links of the AS.

Following the recommendations in the literature, we used a Pseudo-Random Poisson Process to generate the probes [15]. In addition, we set the size and frequency of the probes to correlate the measurements with the class of traffic (application) being controlled.

1. A Self-adaptive cost metric for the Smart Routing Managers

In order to decide which is the best path to reach any remote AS in our scheme, each SRM collects QoS information and using this information computes the cost to reach that destination. This QoS information provides two components to that cost. The first component consists of end-to-end QoS information, which is based on processing and filtering the OWD inferred from the probes. We call this a Smoothed OWD (SOWD)[2]. The second component consists of local QoS information, and is based on collecting the Available Bandwidth (ABW) from the egress links of the AS.

Equation (1) presents the cost metric to be used by each SRM in our edge routing scheme. In terms of notation, M_{ij} represents the cost to reach a distant SRM through the i_{th} egress link of the source AS for traffic of class j. The SOWD and the ABW components described before are represented as S_{ij} and ABW_i respectively. In addition, the non-negative parameters α_j and β_j are self-adaptive weights which endow the SRMs with self-adaptive capabilities. The bound $\overline{D_j}$ represents the maximum OWD tolerable to reach a remote AS for traffic of class j in our scheme. This constraint is specified in the SLS exchanged between the SRMs, using the SRM protocol [14]. On the other hand, the bound B_i represents the minimum acceptable ABW in the i_{th} egress link of the AS. This is an important constraint which supplies a minimum bandwidth guarantee per-egress link avoiding best-effort traffic starvation.

$$M_{ij} = \begin{cases} Floor\left[\alpha_j S_{ij} + \dfrac{\beta_j}{ABW_i}\right] & if \quad (S_{ij} \le \overline{D_j}) \ \wedge \ (ABW_i \ge B_i) \\ \infty & if \quad (S_{ij} > \overline{D_j}) \ \vee \ (ABW_i < B_i) \end{cases} \quad (1)$$

The main motivations for the selection of this additive cost metric can be summarized as follows:

- This self-adaptive cost is not only able to reflect the present QoS dynamics, but also to evolve with them providing transparency to the traffic reallocation decision process (requirement 3 in Section II). Based on the computation of this cost the SRMs are able to tweak BGP in short timescales for the outbound traffic of the AS. In fact, network administrators of any multihomed stub AS using a SRM will only need to select and con-

[1] In [14] we called these mangers Overlay Entities (OEs). We rename these optimizers/managers here to emphasize that they do not circumvent BGP.

[2] Our aim is to avoid frequent changes of the QoS information, so the SRMs use a SOWD instead of instantaneous values of the OWD collected.

figure a fixed threshold, so that if this threshold is met then the reallocation of traffic is allowed. The self-adaptive capabilities mean that if the network conditions are adequate to reallocate the traffic, then the threshold will be more frequently met. However, if the conditions are inadequate, then larger variations in the QoS conditions are needed to meet the same threshold value.

- Equation (1) allows to flexibly design the cost unbalancing the sensitivity of the weights α_j and β_j. The motivation for this is to magnify the relevance of the OWD over the local ABW. However, the additive term of ABW in (1) preferably allows the allocation of traffic over links with more ABW when OWD conditions are similar. This allows an overall better traffic distribution for the outbound traffic from the AS.

2. Qualitative description of the self-adaptive traffic reallocation process

In order to compute the cost in (1) the SRMs smooth the OWD samples gathered from the probes using two filters in cascade:

1. The first filter corresponds to the median OWD through a sliding window. Our motivation for choosing the median is that it is widely accepted as the best estimator of the OWD that user's applications are actually experiencing. The sliding window is heuristically chosen in order to get a good trade-off between responsiveness and correlation between the measurements and the applications data being controlled.
2. The second filter is applied to the median OWD and works like an A/D converter with a self-adaptive pace of conversion. This pace is constantly adapted depending on the QoS conditions on the network. If the conditions are smooth the pace is small, and more traffic reallocations are allowed. In such conditions a SRM is able to take full advantage of the multi-connectivity of its AS to the Internet. However, if the conditions may lead to instability the pace increases and the traffic reallocations allowed by the SRMs are diminished or even stopped until the network conditions become smooth once again.

Fig. 2 helps to understand these cascade filters. The bound $\overline{D_j}$ is the maximum OWD tolerable for the class of traffic j. The outcome of the second filter is the term S_{ij} in (1), and as it is indicated in (1), if during any interval $S_{ij} > \overline{D_j} \Rightarrow M_{ij} = \infty$. This means that the egress link i of the AS should be removed from the list of available links for output traffic of class j as long as a violation to the SLS exists.

The reallocation of traffic of class j is triggered by the SRM whenever the difference of the costs between the current egress link and any alternate egress link reaches the threshold R_j^{th}. This threshold is preconfigured in the SRM for each class of traffic by the network administrator of the AS. This parameter basically defines the degree of conservativeness of the AS in terms of the number of and frequency of path shifts that the AS is willing to permit.

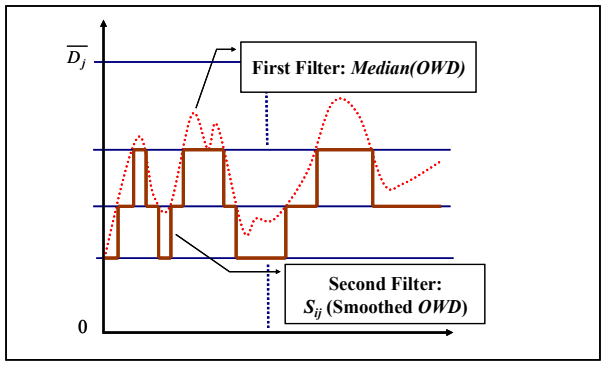

Fig. 2. Self-adaptive filtering processes for the QoS information collected

V. Evaluation Results

We have conducted several experiments to study the performance of the proposed self-adaptive SRMs. The first objective of these experiments is to study how the SRMs are able to adapt their-selves depending on the mid and long-term QoS dynamics. This is shown through its ability to exploit the available paths in order to reduce the OWDs of packets composing QoS traffic aggregates. In addition, we assess the traffic transfer efficiency to examine the traffic performance of each QoS service. This efficiency parameter is given by $Ef_{nj} = C_{nj} / C_{sj}$, where C_{nj} is the throughput at a given destination n, and C_{sj} is the overall throughput at the source domain s for class of traffic j. The second objective is to study how the self-adaptive SRMs contribute to the network stability under variable QoS dynamics, and when different SLSs are used between the remote domains. In this case, as performance indicator we use the number of path shifts needed to meet the agreed QoS constraints for each service for the different SLSs.

Our simulations were performed using the J-Sim [16] simulator with the BGP Infonet suite [17] in which we have implemented all the functionalities of the SRMs. We have introduced a set of QoS extensions to BGP in the Infonet suite for QoS provisioning, such as we did in [14]. The network model for our simulations is illustrated in Fig.3, which is part of the GÉANT European Academic Backbone [18]. The simulated network aims to be a multi-service network with BGP routers supporting the standard Per-Hop Behaviors (PHB): EF, AF21, AF11 and Best-Effort [19, 20].

Our traffic mix consists of Voice over IP (VoIP) calls, video calls, prioritized data downloading, web browsing and email downloading. New voice and video connections uniformly distributed arrive at users' border routers throughout the simulation time. Additionally, the data connections are active during almost all the simulation running time. The source models are based on the work done by Alshaer and Horlait [21] as follows: (i) EF Voice traffic is marked with

the EF code-point. The EF source is an ON-OFF VoIP generator. In the ON state, the EF voice source generates traffic at a peak rate of 64kbps; (ii) video traffic is marked with the AF21 code-point. The AF21 source is a video traffic generator characterized by Pareto ON-OFF. In the ON state the AF21 source generates video traffic at a peak rate of 200 kbps; (iii) finally the prioritized data traffic is marked with the AF11 code-point. Both AF11 traffic and BE connections are characterized by Poisson processes. AF11 and BE traffic sources generate traffic at 350Kbps and 500Kbps, respectively.

The frequency and size of the probe packets are selected in correlation with the corresponding traffic DSCP. These parameters were empirically tuned after several simulation runs.

In our experiments we run the same simulations using first self-adaptive SRMs and then non-adaptive SRMs. In both experiments, we collect the data results considering three different SLSs previously exchanged between remote stubs domains based on maximum packet OWD for Voice, video and prioritized data traffic. These maximum OWDs tolerated per-service were heuristically chosen to represent reasonable values of the kinds of traffic sources considered, and to allow a moderate probability of SLS violation events, as depicted in Table 1.

The end-to-end traffic performance is contrasted in Fig. 4 and 5 when self-adaptive and non-adaptive SRMs are in use. In terms of latency, Fig. 4 shows that the self-adaptive SRMs generally outperform the non-adaptive SRMs. In particular, for the case of AF21 traffic, where the delay sensitive traffic is classified, the self-adaptive SRMs clearly improve the traffic performance. Therefore, the self-adaptive capability of adjusting the cost metric based on the mid and long-term network conditions, improves the route optimization decisions that the SRMs are able to take. In addition, the contribution of the self-adaptive SRMs to reduce the latency of delay-sensitive applications does not lessen the traffic transfer efficiency as shown in Fig 5. Indeed, both kinds of SRMs provide almost the same traffic transfer efficiency.

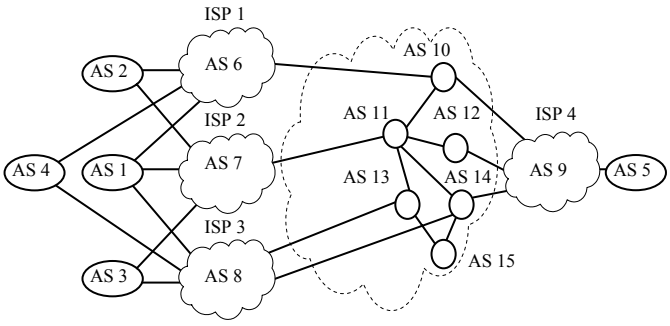

Fig. 3. Simulation network topology

TABLE 1.
SLS DESCRIPTIONS

SLS ID	Voice (EF) [ms]	Video (AF21) [ms]	Prioritized (AF21) [ms]
1	45	55	75
2	55	65	85
3	65	75	95

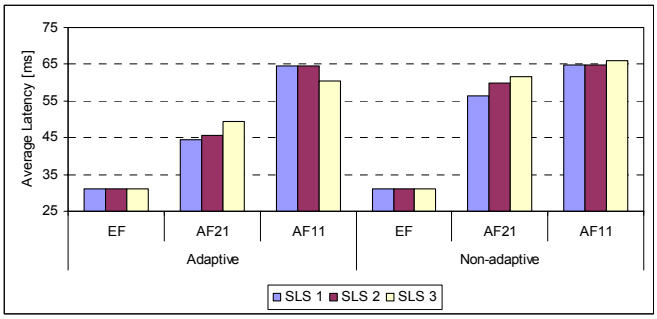

Fig. 4. Average latency of QoS sensitive traffic when using self-adaptive and non-adaptive SRMs

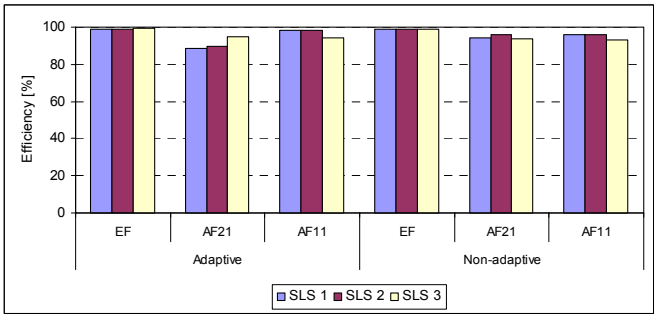

Fig. 5. Traffic transfer efficiency of QoS sensitive traffic when using self-adaptive and non-adaptive SRMs.

The contribution of the self-adaptive SRMs to the stability of the network is evaluated by the total number of path shifts that occur in all the multihomed stub ASes. This is assessed for different values of the degree of conservativeness, i.e. the threshold R_j^{th}, as depicted in Fig. 6. From the analysis of this figure, it is clear that even though the selfish selection of paths is prone to situations of instability, the self-adaptive SRMs are able to restrain this undesired behavior. Specifically, the results show that the number of path shifts introduced by the self-adaptive SRMs is clearly much smaller than those that occur for the non-adaptive SRMs. Indeed, in some cases the number of path shifts is reduced by nearly one order of magnitude. The ability of the proposed self-adaptive SRMs to foresee situations of SLS violation, and to take adequate measures by choosing alternative paths before congestion becomes excessive, is capable of improving network stability. As a complementary mechanism, the activa-

tion of the degree of conservativeness (the threshold R_j^{th}) is able to reduce the number of path shifts, and therefore to contribute to an additional improvement of network stability.

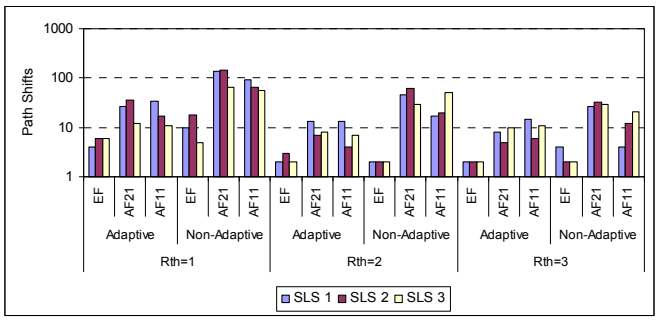

Fig. 6. Number of path shifts when using self-adaptive and non-adaptive SRMs for different normalized thresholds.

VI. Conclusions and Future Directions

In this paper, we analyzed the problems that will arise in the Internet if the current trend of selfishly optimizing traffic at the edge of the network becomes widely deployed. Modern optimization tools lack of mechanisms that allow them to learn from the network dynamics and "socially" deal with the global stability issues of greedily managing interdomain traffic in short timescales.

Once identified this problem, we have investigated the main requirements and expected functionalities of the future self-adaptive route optimization tools from a realistic viewpoint. Based on this analysis we have proposed and evaluated the first self-adaptive route optimizers for an interdomain routing architecture that we developed in a previous work. Simulation results show that our self-adaptive Smart Routing Managers are not only able to greedily improve end-to-end performance, but also to drastically reduce the number of path shifts needed to accomplish the desired performance. We are currently working towards testing our SRMs in a real testbed. In addition, we have plans to develop a stability model for our edge routing architecture based on the application of physical similes.

References

[1] A. Akella, S Seshan, and A. Shaikh, "Multihoming Performance Benefits: An Experimental Evaluation of Practical Enterprise Strategies," USENIX Annual Technical Conference 2004, Boston, MA, USA

[2] Cisco Optimized Edge Routing, http://www.cisco.com/

[3] Internap Network Services Corporation, http://www.internap.com/

[4] M. Yannuzzi, X. Masip-Bruin, E. Monteiro, "Towards Self-Adaptive Inter-Domain Edge Routing," accepted for publication in the IEEE Infocom Student Workshop, Miami, USA, March 2005.

[5] A. Akella, J. Pang, B. Maggs, S. Seshan and A. Shaikh, "A Comparison of Overlay Routing and Multihoming Route Control," in Proceedings of ACM SIGCOMM04, Portland, USA, August 2004.

[6] C. Villamizar, R. Chandra, R. Govindan, "BGP Route Flap Damping," Internet Engineering Task Force, Request for Comments 2439, November1998.

[7] S. Uhlig, V. Magnin, O. Bonaventure, C. Rapier and L. Deri, "Implications of the Topological Properties of Internet Traffic on Traffic Engineering," Proceedings of the 19th ACM Symposium on Applied Computing, Special Track on Computer Networks, Nicosia, Cyprus, March 2004.

[8] B. Quoitin, S. Uhlig, C. Pelsser, L. Swinnen, O. Bonaventure, "Interdomain Traffic Engineering with BGP", IEEE Communications Magazine, May 2003.

[9] A. Akella, B. Maggs, S. Seshan, A. Shaikh, R. Sitaraman, "A Measurement-Based Analysis of Multihoming," in Proceedings of ACM SIGCOMM 2003, Karlsruhe, Germany.

[10]. C. de Launois, B. Quoitin, and O. Bonaventure, "Leveraging network performances with IPv6 multihoming and multiple provider-dependent aggregatable prefixes," 3rd International Workshop on QoS in Multiservice IP Networks (QoSIP 2005), Catania, Italy, February 2005.

[11] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet routing convergence," in Proc. ACM SIGCOMM, 2000.

[12] D. G. Andresen, H. Balakrishnan, M. F. Kaashoek, R. Morris, "Resilient Overlay Networks", in Proceedings, 18th of ACM SOSP, 2001.

[13] S. Agarwal, C. Chuah, R. Katz "OPCA: Robust interdomain policy routing and traffic control," IEEE Openarch, April 2003.

[14] M. Yannuzzi, A. Fonte, X. Masip, E. Monteiro, S. Sánchez, M. Curado, J. Domingo, "A proposal for Interdomain QoS routing based on distributed overlay entities and QBGP," in Proceedings of the First International Workshop on QoSR (WoQoSR), co-located with QoFIS'04, Barcelona, Spain, October 2004.

[15] G. Almes, S. Kalidindi, M. Zekauskas, "A One-way Delay Metric for IPPM", Internet Engineering Task Force, Request for Comments 2679, September 1999.

[16] J-Sim Homepage, http://www.j-sim.org.

[17] Infonet Suite Homepage, http://www.info.ucl.ac.be/~bqu/jsim/

[18] GÉANT Website, http://www.dante.net/server/show/nav.007

[19] V. Jacobson, K. Nichols and K. Poduri, "An Expedited Forwarding PHB," RFC 2598, IETF, June 1999.

[20] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski, "Assured Forwarding PHB Group," RFC 2597, IETF, June 1999.

[21] H. Alshaer and E. Horlait, "Expedited Forwarding Delay Budget Through a Novel Call Admission Control," 3rd European Conference on Universal Multiservice Network (ECUMN'2004), Porto, Portugal.