Towards Self-Adaptive Inter-Domain Edge Routing

Marcelo Yannuzzi
Advisors: Xavier Masip-Bruin, Edmundo Monteiro[†]
Advanced Broadband Communications Lab, Technical University of Catalonia (UPC)
Vilanova i la Geltrú, Barcelona, Spain
{yannuzzi, xmasip}@ac.upc.edu, [†]edmundo@dei.uc.pt

Abstract— **At present, greedily managing inter-domain traffic at the edge of the network is becoming a common practice. However, the stability implications of these practices under massive utilization are completely unknown. Given that global stability is a must for the current and future Internet, self-adapting mechanisms will be necessary if masses of absolutely independent and uncoordinated multi-homed stub Autonomous Systems (ASs) are allowed to simultaneously change their traffic patterns seeking only for the best of their own purposes in short, and even very short timescales. From our perspective, multi-homing in combination with routing edge optimizers is a powerful tool in which stub ASs can rely on in order to improve their end-to-end QoS [1]. In this sense we foster this kind of inter-domain Traffic Engineering (TE) approach, but highlighting that the set of perturbations introduced by these devices to the network should be timely controlled. Given that no centralized entity will ever exist to control and coordinate these massive perturbations, self-adapting tools will become unavoidable. This paper surveys the main advantages of self-adapting tools at the edge of the Internet, especially when traffic is reallocated in short timescales. This is part of ongoing work. Our first results show that self-adaptive algorithms not only provide significant improvements in terms of end-to-end QoS, but they are also able to achieve better traffic distribution and obtain better overall link utilization, while timely limiting the number of AS path shifts.**

Keywords-component; Multi-homing; Stability, Self-Adaptive, Inter-Domain Edge Routing

I. MOTIVATION: LOCAL AND GLOBAL STABILITY

It is widely accepted that the main problem with end-to-end QoS provisioning is on the very foundations of the current inter-domain network paradigm. The central issue is that the Border Gateway Protocol (BGP) has not inbuilt QoS capabilities given that it was designed with very different goals in mind by the early nineties [2]. Although some researchers have proposed to replace BGP, in practice, only incremental approaches are realistic and will have chance to become deployed. From this perspective, the initial inter-domain heuristics mostly tended to add QoS and Traffic Engineering (TE) extensions to BGP, but quite recently some researchers and manufacturers have started to avoid new enhancements to the protocol and propose to decouple part of these tasks from BGP devices [3-7]. While heuristics extending BGP are only able to improve end-to-end performance for internets under low routing dynamics, the latter result much more effective, especially, when routing changes occur more frequently. The main difference between these two approaches is that the latter decouples part of the policy control portion of the routing process from BGP devices. Hence, the two approaches basically differ in how routing policies are controlled and signaled. In-band techniques, that is, those inherently supported by BGP can feasibly operate over long timescales which means that they are appropriate for rather static or pseudo-static TE provisioning. On the other hand, out-of-band techniques, that is, those decoupled from BGP are in fact able to operate in much shorter timescales so they result perfectly appropriate for dynamic or even highly dynamic TE provisioning. Even so, the stability implications of rearranging inter-domain traffic in very short time scales is not yet understood. Indeed, the effect of managing large amounts of inter-domain traffic in this way is completely unpredictable. Thus, these kinds of solutions are definitively not applicable, for example, to large transit Autonomous Systems (ASs) such as Tier-1 or Tier-2 Internet Service Providers (ISPs). Additionally, the rearrangement of small fractions of inter-domain traffic in short timescales, but magnified by the number of sources simultaneously injecting these perturbations to the network may also result unpredictable in terms of global stability.

Above all, multi-homed stub ASs are those which could benefit the most from novel mechanisms providing them with dynamic TE capabilities in medium or short timescales. This particular fraction of ASs crowds together mostly medium and large enterprise customers, Content Service Providers (CSPs), and small Network Service Providers (NSPs), which altogether actually represent more than 60% of the total number of ASs in the Internet. Therefore, the blast of multi-homed stub ASs in the last few years has gained huge interest in both research and commercial fields, and that's why several optimized edge routing proposals are starting to appear as commercial products.

II. SELF-ADAPTIVE EDGE ROUTING

Two major advantages can be foreseen in using self-adaptive inter-domain TE tools. First of all, it is widely known that network administrators of multi-homed stub ASs are not willing to adopt complex mechanisms to control how their traffic is managed in medium and short timescales. 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 such decisions last in time. Indeed, NSPs, CSPs, and medium/large enterprise customers are eagerly claiming for straightforward mechanisms that allow them to opportunistically manage their inter-domain traffic in short timescales depending on the existing end-to-end performance. Thus, a major advantage behind self-adaptive mechanisms is that they are perfectly suitable for this kind of opportunistic and selfish demands. The next figure depicts the significance of self-adaptive mechanisms from this perspective, given that they are able to hide the QoS dynamics from the traffic reallocation decision process. This supplies an appealing solution for multi-homed stub ASs, since they may decide how conservative or opportunistic they want to be by simply selecting a fixed threshold for example, without worrying about the stability implications of their decision. It should become clear that how conservative or opportunistic such ASs 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 ASs to straightforwardly decide how opportunistic they are willing to be, and this decision will definitively last in time.

In second place, the significance of a self-adaptive tool is in its strengths in terms of guaranteeing local and global stability. Under highly unpredictable network conditions, such as link flaps, or routing misconfigurations it is imperative that each edge optimizer counts with a self-adaptive mechanism which allows it to learn from these dynamics and prevent or diminish the number of AS path shifts until the network conditions are once again stable. Indeed, multi-homed stub ASs 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 aggressive network dynamics even a conservative opportunistic approach may lead to network instability caused by the excessive number of path shifts allowed. Thus, the assertion of being “conservative” strongly depends on the network dynamics. As an alternative, self-adaptive algorithms/metrics are able to adapt themselves to those changing conditions, so that they could be able to reflect the choices made by a network manager independently of the network dynamics.

To conclude, it should become clear that these self-adaptive tools will in fact become suitable in the near future not only to reallocate flows of inter-domain traffic for a set of IP prefixes, but also sets of inter-domain tunnels, or sets of inter-domain IP/MPLS Label Switch Paths (LSPs) in shorter timescales.

III. Ongoing Work

In [5] we have designed a completely distributed routing edge framework and architecture, including a set of heuristics in order to supply local and global stability based on contention, prioritization and penalization/relaxation. Our goals at this step are, firstly to design self-adaptive QoS cost metrics and self-adaptive QoS Routing (QoSR) and TE techniques for multi-homed stub ASs, which will help BGP to improve end-to-end QoS in a selfish, but globally stable manner. Secondly, we are working on the development of a stability model for the proposed architecture. We are applying physical similes, such as thermodynamic models for open systems, aiming at counting with an energy model which could be treated by means of the Lyapunov criterion.

References

[1] A. Akella, et al., “A Comparison of Overlay Routing and Multihoming Route Control,” in Proceedings of ACM SIGCOMM04, Portland, USA, August 2004.

[2] Y. Rekhter, T. Li , “A Border Gateway Protocol 4 (BGP-4),” Internet Engineering Task Force, Request for Comments 1771, March 1995.

[3] L. Subramanian, I. Stoica, H. Balakrishnan, R. H. Katz, “OverQoS: Offering Internet QoS using overlays,” ACM SIGCOMM Computer Communications Review, vol. 33-1, January 2003.

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

[5] M. Yannuzzi, et al., “A proposal for inter-domain QoS routing based on distributed overlay entities and QBGP,” in Proceedings of the First International Workshop on QoSR (WoQoSR), co-located with QoFIS’04, LNCS 3266, Barcelona, Spain, October 2004.

[6] Internap Network Services Corp., Flow Control Platform, http://www.internap.com/

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

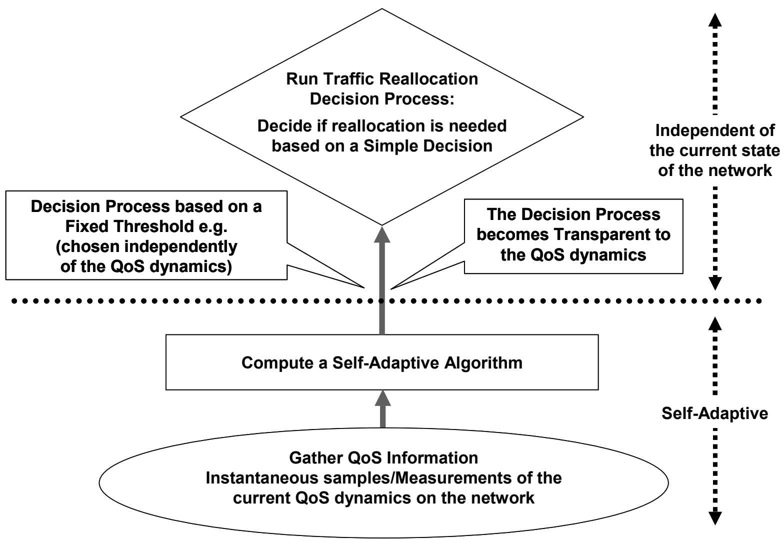

Figure 1. A Self-Adaptive algorithm provides transparency to the traffic reallocation decision process