Analysis of Interdomain Smart Routing and Traffic Engineering Interactions

A. Fonte*†, M. Pedro*‡, E. Monteiro*, F. Boavida*

*Departament of Informatics/CISUC, University of Coimbra, Polo II, Pinhal de Marrocos, 3030-290, Coimbra, Portugal

Email: {afonte,macpedro,edmundo,boavida}@dei.uc.pt

†Polytechnic Institute of Castelo Branco, ESTCB, Av. do Empresario, 6000-767, Castelo Branco, Portugal

‡Polytechnic Institute of Leiria, ESTG, Campus 2, Morro do Lena - Alto Vieiro, 2411-901, Leiria, Portugal

***Abstract*—Multihoming Smart Routing plays a significant role in improving the performance of Internet access of IP Network Customers through dynamic path switching. On the other hand, Inter-domain Traffic Engineering becomes indispensable for IP Network Providers, in order for them to meet their traffic objectives. Unfortunately, this ability of both parties to choose their own routing policies does not necessary lead to the best routing in the Internet. A major challenge of the research on Traffic Engineering and Internet Routing architectures is to accommodate this tussle. However, little is known about the interactions between Inter-domain Smart Routing and Traffic Engineering.**

In this paper, we analyse such interactions through intuitive descriptions and a simulation model. In our evaluations, we observed that the overall stability can benefit from combining Inter-domain Adaptive Smart Routing and Traffic Engineering. However, we also observed that the effectiveness of Inter-domain Traffic Engineering can be negatively affected by interdomain routing changes performed by Smart Routing.

***Index Terms*— Multihoming, Routing, Traffic Engineering.**

I. Introduction

Multihoming is a well-known technique for improving performance and reliability of Internet accesses. It consists of the increasing of Internet connectivity by contracting multiple broadband lines (e.g., Business DSL or T1) from two or three INPs (IP Network Providers), common referred as ISPs (Internet Service Providers). It estimates that stub ASs (Autonomous Systems) employing multihoming experience a potential performance improvement in at least 40% [1]. However, even though this benefit of multhoming, ID (Inter-Domain) routing is still dependent of BGP (Border Gateway Protocol), which do not provide efficient route control [2].

Smart Route Controllers (SRC) are, therefore, being increasingly used by multi-homed stub ASs, as they provide stub ASs a holistic way to solve their traffic challenges through shifting some traffic between INPs. SRCs as they were introduced are not new concept. Many enterprises have been devoting efforts to research and development of SRC products [3], [4]. In the research community, some papers have devoted attention to the design and stability of SRCs [5], [6], [7]. In this paper, we focus on a collateral issue: on the potential interactions between ID Smart Routing and Traffic Engineering (TE) in ASs of INPs. Henceforth, we refer stub ASs simply as IP Network Customers (INC).

Compared to INCs, INPs have more complex routing policies because they usually operate transit ASs and run also internal BGP (iBGP). In addition, INPs should engineer the ID traffic properly. One common TE problem is the BGP Egress Router Selection (BERS) problem, i.e., *how to assign the incoming traffic to multiple egress points, such the traffic objectives are met (e.g., to minimize the maximum link utilization (min-MLU) or Load-balancing (LB))*? The biggest difficulty with BERS is as far as the number of traffic aggregates, objectives and egress point choices increase it becomes harder to solve it. Several studies, such as [8], [9], [10], constitute the base framework for most of the current proposals.

The consequence of both INPs and INCs having distinct business models and traffic challenges is that ID Smart Routing and TE tools are selfish by nature in that each party seeks to control its traffic according only to its own goals without considering the effects over the traffic or network of the other party. This raises two open questions that we seek to answer in this paper: *Does the ability of both INCs and INPs to make greedy choice of ID routes lead to the best routing of IP packets in the Internet? Why and how can these tools interact?*

Most of significant related work used a game-theoretic-based analysis. In theoretical thread, a precursor study compared the global Internet performance obtained through *selfish routing* to the optimum achieved through global routing coordination [11]. Specifically, the authors proved *'that if the latency of each edge is a linear function of its congestion, then the total latency of the routes chosen by selfish network users is at most* $4/3$ *times the minimum possible total latency*'. Besides this study can give us some insights about the interaction mechanisms, it seems not feasible to extend this result to our case. Afterwards, in [12] the authors found that *selfish overlays* can significantly reduce the effectiveness of intra-domain TE (i.e., MPLS optimization), but taking a somewhat more realistic model through using of measured ISP topologies. Motivated by this study in [13] the authors deeply modelled this kind of interaction.

Hence, it is still unclear *if* and *how* these penalties can also occur in inter-domain environments, which motivate this work. We believe that understanding the potential interactions between ID Smart Routing and Traffic Engineering can provide some guidelines for future work in the area, such as research on cooperative TE [14].

1930-529X/07/$25.00 © 2007 IEEE

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2007 proceedings.

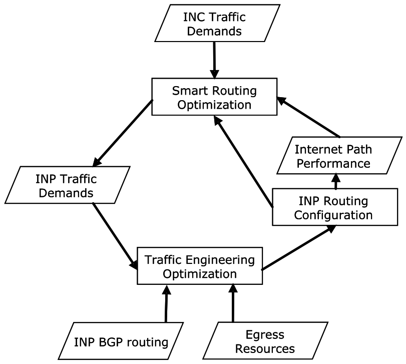

Fig. 1. Model of the Interactions between ID Smart Routing and Traffic Engineering.

The remaining sections of this paper are organized as follows. In Sect. II, we give further motivation and analysis of the interactions through intuitive descriptions. Then, in Sect. III, we describe the SRC system and the genetic ID TE algorithm used in our evaluations. In Sect. IV, we present our evaluations. Finally, we conclude the paper in Sect. V and present some directions for future work.

II. Background and Interaction Descriptions

Figure 1 presents a comprehensible model of the potential interactions between ID Smart Routing and Traffic Engineering. By using this model and an example, this section seeks to answer the question: *Why and how can these tools interact?*

A. Background

Let us first briefly describe a typical ID TE process in an INP network with a set of ingress points I, a set of egress points E and a set of reachable prefixes P. The inputs of an ID TE process are the incoming traffic demands (TD), the egress point choices for each prefix $p \in P$ (given by BGP) and the egress point capacities. Its output is the optimal ID routing to problem. At a given timescale t this process is repeated to return the network at the optimal performance regime as the traffic demands might fluctuate over time due to changes on INC's routing or applications. In this paper, we focus on fluctuations due to routing changes in INCs.

More formally, we represent the predicted TD over the INP $k \in K$ for the slot time t by the matrix $D_{(k,t)}$, where K is the set of INPs and each entry $D_{(k,t)}(i,p)$ is the predicted TD over ingress point i of INP k for prefix p. If each INC $h \in H$ have T data transfers at rates $x_h(p)$ to distribute over its INPs, then each entry of $D_{(k,t)}$ is defined as:

$$D_{(k,t)}(i,p) = \sum_h \sum_p x_h(p).r_{(h,k,t)}(i,p),$$

where H is the set of INCs and $r_{(h,k,t)}(i,p) \in \{0,1\}$ is an indicator function to select whether the route from INC h to prefix p via ingress point i of INP k is active (i.e., True (1), False(0)). On the other hand, we represent the ID egress resources of INP k as C_k, where each entry $C_k(e)$ is the capacity at each egress point e. And, we represent the ID routing from INP k as $\epsilon_{(k,t)}$, where each entry $\epsilon_{(k,t)}(i,e,p) \in \{0,1\}$ is an indicator function that tells whether the $D_{(k,t)}(i,p)$ is assigned to the egress point e.

For concreteness, our BERS problem is: *How to assign each entry of traffic demands $D_{(k,t)}(i,p)$ to an egress point e so as to optimize a certain traffic performance objective.* In this paper, we encoded two typical objectives (i.e., Min-max link utilization and load-balancing) in BERS problem to ensure that egress link utilizations are at lowest levels and thereby to minimize congestion, giving us two BERS versions. Other objectives can be also encoded in BERS (e.g., min-cost routing) [8]. Before proceeding, let first define the link utilization of e for a routing $\epsilon_{(k,t)}$ as:

$$U_e = \sum_i \sum_p \frac{\epsilon_{(k,t)}(i,e,p).D_{(k,t)}(i,p)}{c_k(e)}, \epsilon_{(k,t)}\ is\ a\ routing.$$

Objective 1 - Minimizing the Maximum Link Utilization (min-MLU). One possibility for the BERS problem would be to minimize the maximum link utilization, i.e.,

$$min\ max\ U_e, \forall e \in E, \epsilon_{(k,t)}\ is\ a\ routing. \quad (1)$$

Objective 2 - Load-balancing (LB). Load-balancing is closer to min-MLU objective. Both objectives can be used as interchangeable, but load-balancing is more sophisticated and stringent. For this case the BERS problem would be,

$$min\ |U_i - U_j|, \forall i,j \in E, \epsilon_{(k,t)}\ is\ a\ routing. \quad (2)$$

Both objectives are subject to constraints:

$$\sum_i \sum_p \epsilon_{(k,t)}(i,e,p).D_{(k,t)}(i,p) \le C_k(e), \forall e \in E \quad (3)$$

$$with,\ \sum_e \epsilon_{(k,t)}(i,e,p) = 1, \forall i \in I \quad (4)$$

B. Interactions between INCs and INPs

The interactions between INCs and INPs concerns the effects of changing ID routing at one party over the traffic (or network) of the other party (see Fig.1). In the first interaction, ID routing changes at INC networks achieved by Smart Routing interact with INPs by changing the TD over INPs. In the Internet, in Figure 2, suppose each dual-homed INC, INC_1 and INC_2, has exactly 4 data transfers (i.e., $T = 4$) for prefixes $\{p_1, p_2, p_3, p_4\}$ at rate $x(.,.) = 10$ units of traffic. Then, suppose that at the beginning of the first time slot (i.e., $t = 0$) INC_1 selects the path $INP_2 - INP_1$ for all prefixes and in turn INC_2 selects $INP_3 - INP_1$. This implies that the predicted TD matrix $D_{(1,0)}$ of the target INP, INP_1, for the first TE cycle (0) is the matrix presented in the left side of Table I. Following, imagine that at an instant of time $t^{'} \in]0,1]$ SRCs in INC_1 and INC_2 decide to move traffic between INPs through switching some paths based on latency measures. For instance, suppose INC_1 decides to use the path $INP_3 - INP_1$ for prefixes p_1 and p_2 and in turn INC_2 the

TABLE I
ILLUSTRATION OF ID TD MATRIX CHANGES

	$t=0$				$t^{'} \in]0,1]$			
	p_1	p_2	p_3	p_4	p_1	p_2	p_3	p_4
$D_{(1,t)}(1,.)$	10	10	10	10	0	0	10	10
$D_{(1,t)}(2,.)$	10	10	10	10	10	10	10	10
$D_{(1,t)}(3,.)$	0	0	0	0	10	10	0	0

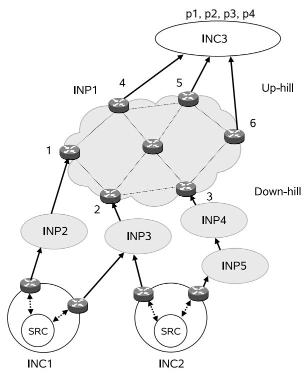

Fig. 2. An Internet Example.

path $INP_5 - INP_4 - INP_1$ for these prefixes. As a result, the (real) TD matrix of INP_1 changes to the matrix presented in the right side of Table I.

This example shows that as far as SRCs distribute traffic across INPs, several entries of TD matrices over INPs might be changed. Consequent on, if the overall change in the TD matrices is far beyond a given tolerated margin, the initial conditions of the ID TE algorithms assumed by INPs to compute optimal ID routing can be broken (e.g., $|D_{(1,t')}| > |D_{(1,t)}| + \sigma_1^2$, where σ_1^2 is the tolerated fluctuation in the TD over INP_1). The practical result of this change is that it can lead the networks of INPs becoming far away from optimal performance regimes. In short, this example shows that Smart Routing can reduce the ID TE effectiveness of INPs as well as imposing additional burden over ID TE to return the network to optimal regime.

In turn in the second interaction, upon finding a new solution for the BERS problem INPs interact with INCs by changing the end-to-end performance of ID traffic leaving the INCs networks. In effect, INPs might decide to shift some ID traffic to others egress points. As a result, a new set of paths is then provided to INCs, which might have different performance characteristics (e.g., latency). In response, SRCs at INCs adapt their ID routing to these changes, and so the first interaction may be repeated.

III. A CLOSER LOOK AT ID SMART ROUTING AND TRAFFIC ENGINEERING

A. Smart Routing Controller Systems

The key ideas of SRCs are decoupling the routing control part from the BGP-based infrastructure and operation in shorter timescales. This way, most of the complexity needed to cope with BGP inefficiency is set apart from BGP. Figure 3

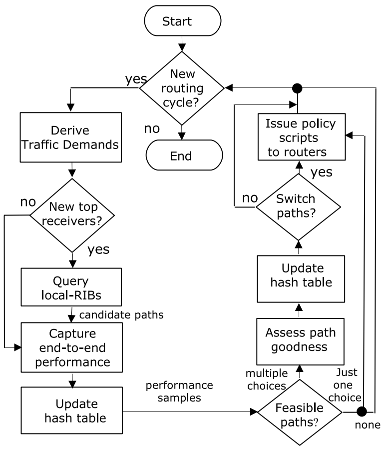

Fig. 3. Smart Routing Controller logical presentation.

presents a logical diagram of our design. It attempts to capture the common functionalities from BGP-based SRCs [7], [6]. We would like to explore other design alternatives, but many technical details about SRC products are unknown [3], [4].

1) Exploiting path diversity and TE capabilities of BGP: The key technique used by SRCs to get better end-to-end performance is exploiting the AS-level path diversity by choosing the best next-hop INP to forward packets. For this purpose, a SRC first collects the routes from local BGP RIBs (Routing Information Bases) to construct a comprehensive view of ID routing. Then, every routing cycle it does an online processing to compute the path changes needed according to the performance of the available path choices (see point *3)*). Finally, it issues command scripts to routers with the corresponding BGP tweakings. In our design, SRCs tweak the LOCAL-PREFERENCE attribute of routes to indicate their ranks.

2) Path Monitoring with Active Probing: One critical issue is to remove the synchronization between the probes [6]. Otherwise, it can lead to path oscillations. In our design, we removed synchronization between the probes thru adding randomization in the sampling times (rather than adding in path switching). For this purpose, we used a pseudo-Poisson sampling process with N_i samples uniformly distributed over a slot of time t_i (in seconds) [15]. Controlling probing overhead and burden over SRCs is also critical. To achieve this goal, our SRCs uses a conservative-enough frequency of probes, while keeping efficient routing. More precisely, $f_i = N_i/t_i = 8/36$ Hz, which is less demanding than the 10Hz used in [6]. In addition, SRCs only focus on the traffic for top receivers, so called popular prefixes, which are a small fraction (i.e., about 10%) of the total number of receivers. The later feature comes from the property of traffic demands being consistent with Zipf-like distributions [17].

3) Dynamic Path Switching: Typically SRCs are multi-objective driven. This implies that if multiple paths obeying to multiple QoS constraints are available for a prefix, a SRC has to select the best path to allocate the traffic based on the QoS

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2007 proceedings.

measures. This problem is called the Multi-Constrained Path (MCP) routing problem, which is NP-hard [16]. In our design, the solving of the MCP routing problem, is through the Metrics Combination (MC) heuristic, by combining latency and spare bandwidth in a single metric, i.e., $\alpha_1.latency_t + \alpha_2.\frac{1}{abw_t}$, where $latency_t$ is the median of the measured RTT (Round-Trip Time) on the path and abw_t is the average spare bandwidth in the target egress link of the INC network at time slot t. As our SRC design focus is on network performance, we only combined network factors affecting QoS. Latency was chosen because it suggests the response time of Internet. In turn, spare bandwidth was chosen to contribute for the load-balancing at INC networks. Then, a SRC picks the path that has the smallest value of the metric in the time slot t_i.

When using the MC heuristic, the biggest difficulty to face is finding the right scale factors $\alpha_i, i = 1, 2$, while keeping efficient and stable routing. In our analysis, we thus used two variations of smart routing as terms of comparison, depending of the approach used to tune $\alpha_i, i = 1, 2$: (1) Adaptive Smart Routing (ASR), where $\alpha_i, i = 1, 2$ are self-tuned as network conditions change, and (2) Conventional Non-Adaptive Smart Routing (NASR), where values for $\alpha_i, i = 1, 2$ are heuristically assigned. Further details of the ASR framework can be found in [7] and the references therein.

B. ID Traffic Engineering Heuristics

As described so far, our BERS goal is to minimize traffic objectives such as (1) or (2). The capacity constraint (3) ensures that the total resource requirements of the traffic flows assigned to each egress point do not exceed the available contracted capacity. The assignment constraint (4) guarantees that each traffic flow is assigned to exactly one egress point e. Our BERS algorithm is based on a genetic single objective version of [10]. The algorithm, belonging to the class of evolution strategies used in optimization, resembles the process of biological evolution, where each individual is described by its genetic code, called a chromosome. On the other hand each chromosome is composed of individual genes. In the problem in hand, a gene is the assignment of a single aggregate traffic flow to an egress point of the INP, and an individual (i.e., a chromosome) is a potential solution. The basic algorithm steps are presented in Figure 4. It starts with the creation of the initial generation, where the individuals are created randomly. Then, an evaluating step based on the proposed objective function (1) or (2) follows. After that, and for a number of generations, a new generation of children is created and compared with the corresponding generation of parents. From this comparison the better elements will compose the next generation of parents. The last of the generations is the aimed TE solution.

IV. Analysis of the Interactions

This section presents the results of a pool of simulation tests performed to assess the ID smart routing and traffic engineering interactions.

```
Create the initial parent generation;
Evaluate the generation;
For a number of generations;
    Create the child generation;
    Evaluate both generations together;
    Rank both generations together;
    Replace worst parents with better children;
End;
```

Fig. 4. Basic steps of the genetic ID Traffic Engineering algorithm.

A. Methodology

In our simulation tests, we evaluated the performance of Adaptive Smart Routing (ASR), Non-Adaptive Smart Routing (NASR), default BGP routing and BGP routing with ID TE for either the min-MLU objective or the LB objective - abbreviated by TE.mMLU and TE.LB - for the possible combinations depending of whether SR or TE mechanisms being switched ON/OFF, giving us eight different simulation configurations. The simulation tests were carried out using the J-Sim simulator [18]. The SRCs were developed on top of the BGP implementation available in this platform. In turn, the genetic ID TE algorithms were coded in the MATLAB language. To emulate the iterative process described in Sect. II, we also implemented a coordination mechanism between both environments. In our simulation model, Smart Routing plays at a timescale 50 times smaller than ID TE.

1) Simulation Setup: The AS-level topology used in the experiments is similar to the example of Fig. 2. The simulated network represents an Internet core composed of 100 Tier-2 INPs divided into two portions: down-hill and up-hill. Both topology portions were built using BRITE [19] according to Waxman's model with (α,β) set to $(0.15, 0.2)$, and a ratio of ASs to links of 1:3. Then, both portions were interconnected by the target INP, a full-meshed Tier-1 INP composed of 8 POPs with a ratio of POPs to links of 1:3. During the tests 300 SRC sources located at down-hill portion send traffic to popular prefixes located at up-hill portion through the target INP. Each traffic aggregate is composed of a variable number of multiplexed Pareto flows (i.e., VoIP flows) according to a Zipf distribution (i.e., a Weibull distribution with shape 0.4). Flow arrivals are described by a Poisson process. We ensured that the overall traffic load can lead the network to the most cost-effective regime in case of a perfect traffic distribution across the network. In other words, we do not simulate greatly over or under provisioned links to avoid silly situations (e.g., stable path selections or persistent congestion). Finally, we configure SRCs to observe the ITU-T's G.114 recommendation to maintain high voice quality, which suggests a one-way delay bound of 150ms, giving us a RTT bound of 300ms.

2) Performance Metrics and Objectives: We use three metrics to evaluate the performance of ID Smart Routing and TE: number of path shifts, latency and ID TE performance ratio. The number of path shifts is obtained by adding the number of path changes that is needed such the performance of the aggregates meet the RTT bound. Latency is defined as the average of RTTs measured for traffic aggregates. The ID

1930-529X/07/$25.00 © 2007 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2007 proceedings.

TE performance ratio of a ID routing ϵ_t at INP and traffic demands $D_{t'}$ over the INP is defined as the real performance of the INP's network divided by the performance in the optimal regime, i.e., $R(\epsilon_t, D_{t'}) = \frac{L(\epsilon_t, D_{t'})}{Op(\epsilon_t, D_t)}$, where $L(.)$ is the maximum utilization or load-balancing across egress links of the target INP on a given $D_{t'}$ and ϵ_t, and $Op(.)$ is the optimal utilization or load-balancing on a given D_t and ϵ_t. The ID TE performance ratio $R(.)$ measures how far the routing ϵ_t is from being optimal on a given $D_{t'}$. Higher values of $R(.)$ (> 1) implies that the routing ϵ_t is farther away from optimal.

The first objective of this simulation study is to evaluate the potential impact of ID TE on Smart Routing stability. To achieve this goal, we compared the average number of path shifts that occurred at SRC when ID TE is switched ON/OFF. The second objective is to assess the implications of ID Smart Routing and TE interactions on the traffic performance. To achieve this goal, we compared the end-to-end latency for all simulation configurations. The third objective is to evaluate the potential impact of Smart Routing on ID TE performance. To achieve this goal, we compared the ID TE performance ratios for mMLU and LB objectives for all simulation configurations.

B. Evaluation Results

Figure 5 and 6 show, respectively, the number of path shifts and the cumulative number of path shifts, measured at the end of each ID TE cycle. In turn, Figure 7 shows the traffic latency. These results revealed some observations. First, comparing both smart routing schemes, what we see is, there is an *overall stability benefit from combining ASR and Traffic Engineering*. In effect, this combination roughly needs no more than 32% (mMLU) or 37% (LB) of the total number of path shifts needed by the NASR and ID TE combination to meet the RTT bound. These results show similar trends of the previous study in [7]. That is, using ASR the oscillations can be drastically reduced, on average, with similar traffic performance. Second, the growth of path shifts number is almost linear, except in the case of ID TE being switched OFF. This is expected because the overall load of network is moderate and flows arrivals are poissonian. On the other hand, when ID TE is switched OFF, the higher latency observed for traffic reveals that default BGP routing is unable to provide enough traffic distribution across INP egress links and avoiding congestion. This is expected because BGP decisions are mainly based on AS path length criterion. Furthermore, this higher latency for traffic summed up to a similar number of path shifts observed for both schemes ASR and NASR is a clear sign that even SRCs cannot avoid congestion if there is not enough capacity or in other words if paths are greatly overlapped. Third, we also observe that if ID TE uses LB objective both schemes ASR and NASR perform worse. This is an expected result because LB objective is more stringent than mMLU objective (as we observe next).

Figures 7 and 8 show the values of statistics related to ID TE performance ratios for both objectives mMLU and LB. These results also revealed some important observations. First, *Smart Routing changes at INCs can reduce the effectiveness of ID Traffic Engineering.* As we can observe, best ratios, i.e., averages and medians closer to 1, were gathered if Smart Routing is switched OFF; and for the more stable combination we registered a TE performance ratio of 3.9 units. This implies that although ID TE is able to adapt efficiently the ID routing to traffic demands fluctuations due to sources, it is unable to accommodate stronger traffic changes due to Smart Routing. Second, as expected we observed that if ID TE is switched OFF the network has the worst performance. The higher values in averages and medians of ID TE performance ratios reveals significant traffic unbalances across INP egress links. Third, although ID TE has positive impact over ASR, we also observe that ASR scheme has episodes of significant interaction with ID TE (as it shows the bigger excursions of ratios). Furthermore, given the better stability of ASR this might be a sign that this interaction could involve big aggregates. However, the medians and averages for ratios gives a sign of lesser interaction in steady state. Fourth, if the ID

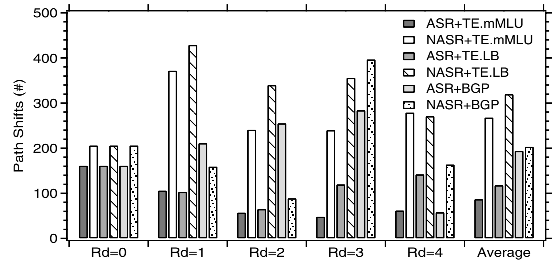

Fig. 5. Number of Path Shifts measured in each TE cycle.

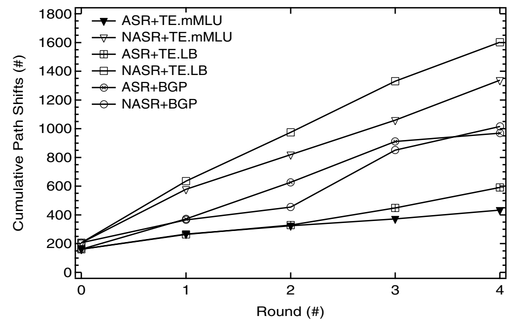

Fig. 6. Cumulative Number of Path Shifts measured in each TE cycle.

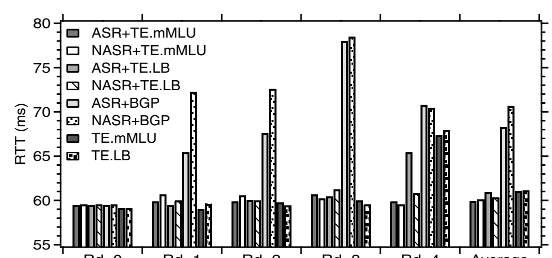

Fig. 7. Average latency measured in each TE cycle.

1930-529X/07/$25.00 © 2007 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2007 proceedings.

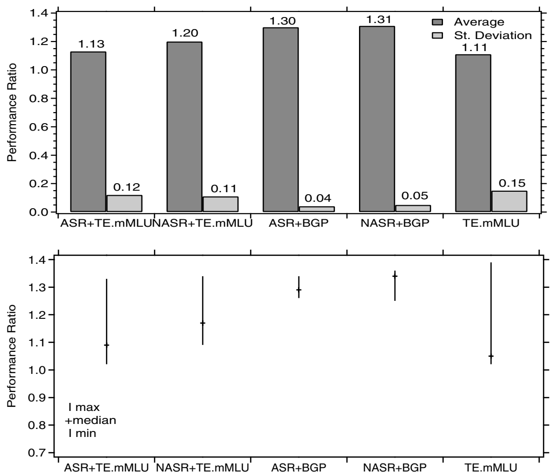

Fig. 8. ID TE.mMLU performance ratios.

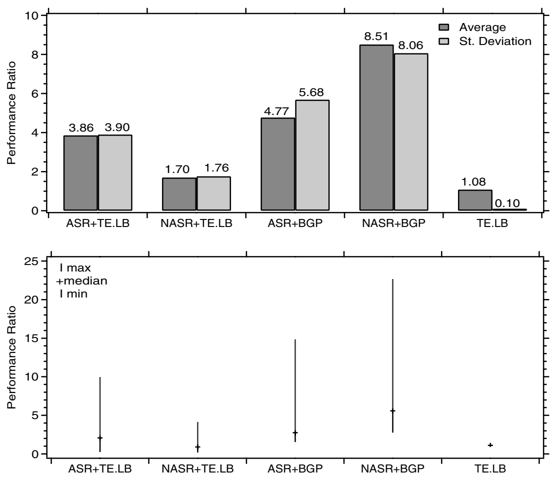

Fig. 9. ID TE.LB performance ratios.

TE uses LB and NASR is ON, the INP network experience better performance at least in average, but at the expense of much larger number of path shifts. Fifth, the higher values in all statistics for the LB objective reveals that stringent objectives can intensify the interactions. Moreover, the high values of standard deviations gives a clear sign of high risk of interaction. This result perfectly matches the fact of smart routing being more instable when the LB objective is used.

V. Conclusion

Multihoming Smart Routing provides a way for IP Network Customers to improve performance of their Internet access through dynamic path switching. On the other hand, IP Network Providers are increasingly employing Inter-domain Traffic Engineering. In this paper, we have shown that interactions occur between these techniques due to inter-domain routing changes. We analyzed the interactions through intuitive descriptions and through a simulation model. In our evaluations, we used two variations of Smart Routing - Adaptive and Conventional non-adaptive - and a genetic Inter-domain Traffic Engineering algorithm aiming at meeting two similar objectives (i.e., min-MLU and LB). From the IP Network Customers perspective, we observed that the overall stability can benefit from combining an adaptive smart routing scheme and Inter-domain Traffic Engineering. However, we have also observed that the effectiveness of Inter-domain Traffic Engineering can be negatively affected by routing changes performed by Smart Routing. We also observed this risk is greater when more stringent Inter-domain Traffic Engineering objectives are used. Yet two open issues remain. First, we plan to investigate how the granularity and timescale of traffic rule these interactions. Second, we also plan to model in detail this kind of interaction though the Game Theory.

Acknowledgment

This work was partially funded by the European Commission through Network of Excellence CONTENT (FP6-0384239)

References

[1] A. Akella et al., A Measurement-Based Analysis of Multihoming. In the Proc. of ACM SIGCOMM 2003, Karlsruhe, Germany, 2003.
[2] S. Savage et al., The end-to-end effects of Internet path selection. In Proc. of ACM SIGCOMM 99, Cambridge, MA, August 1999.
[3] Cisco Systems, Inc., Optimized Edge Routing.
[4] Internap Networks Inc., Flow Control Platform.
[5] F. Guo et al., Experiences in Building A Multihoming Load Balancing System. In Proc. of IEEE INFOCOM 2004, Hong Kong, March 2004.
[6] R.Gao et al., Avoiding Oscillations due to Intelligent Route Control Systems. In the Proc. of IEEE INFOCOM 2006, Barcelona, Spain, April 2006.
[7] M. Yannuzzi et al., On the Advantages of Cooperative and Social Smart Route. In Proc. of ICCCN 2006, Arlington, Virginia, USA, October 2006.
[8] T. Bressoud and R. Rastogi., Optimal Configuration for BGP Route Selection. In Proc. of IEEE INFOCOM 2003, San Francisco, USA, April 2003.
[9] S. Uhlig, A multiple-objectives evolutionary perspective to interdomain traffic engineering. International Journal of Computational Intelligence and Applications, World Scientific Publisher, 2005.
[10] M. Pedro et al, An Approach to Off-line Inter-domain QoS-Aware Resource Optimization. In Proc. of IFIP Networking 2006, Coimbra, Portugal, May 2006.
[11] T. Roughgarden and E. Tardos, How bad is selfish routing, J. Ass. Comput. Mach., vol. 49, no. 2, pp. 236-259, 2002.
[12] L. Qiu et al., On Selfish Routing in Internet-Like Environments. In Proc. of ACM SIGCOMM 2003.
[13] Y. Liu et al., On the Interaction Between Overlay Routing and Underlay Routing. In Proc. of IEEE INFOCOM 2005.
[14] G. Shrimali et al., Cooperative Interdomain Traffic Engineeing using Nash Bargaining and Decomposition. In Proc. of IEEE INFOCOM 2007, Anchorage, Alaska, May 2007.
[15] G. Almes et al., A one-way delay metric for IPPM, IETF, Request for Comments 2679, September 1999.
[16] F. Kuipers et al., Overview of Constraint-Based Path Selection Algorithms for QoS Routing, IEEE Communications Magazine, Vol 40, n12, Dec 2002.
[17] W. Fang and L. Peterson, Inter-as traffic patterns and their implications. In Proceedings of the 4th Global Internet Symposium, December 1999.
[18] J-Sim Homepage, http://www.j-sim.org.
[19] A. Medina et al, BRITE: An Approach to Universal Topology Generation. In Proc. of MASCOTS, August 2001.