This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE Globecom 2010 proceedings.

A Resilient Routing Policy for Peering Management

Stefano Secci^a, Huaiyuan Ma^b, Bjarne Emil Helvik^b, Jean-Louis Rougier^c
^{a}LIP6, Université Pierre et Marie Curie (UPMC - Sorbonne Universités). E-mail: stefano.secci@ieee.org
^{b}Q2S Center, NTNU, Norway. E-mail: {huaiyuan, bjarne}@q2s.ntnu.no
^{c}Institut Telecom, Telecom ParisTech, LTCI CNRS, France. E-mail: rougier@telecom-paristech.fr

***Abstract*—We present a novel resilient routing policy for controlling the routing across peering links between Internet carriers. Our policy is aimed at offering more dependability and better performance to the routing decision with respect to the current practice (e.g., hot-potato routing). Our work relies on a non-cooperative game framework, called Peering Equilibrium MultiPath (PEMP), that has been recently proposed. PEMP allows two carrier providers to coordinate a multipath route selection for critical flows across peering links, while preserving their respective interests and independence. In this paper, we propose a resilient PEMP execution policy accounting for the occurrence of potential impairments (traffic matrix variations, intra-AS and peering link failures) that may occurr in both peering networks. We mathematically define how to produce robust equilibrium sets and describe how to appropriately react to unexpected network impairments that might take place. The results from extensive simulations show that, under a realistic failure scenario, our policy adaptively prevents from peering link congestions and excessive route deviations after failures[1].**

I. INTRODUCTION

Internet multipath routing is gaining much interest in the networking research community. It is generally considered beneficial for traffic engineering in IP networks as it allows higher path diversity and load-balancing on equivalent routes [1] [2]. Multipath and load-balancing can be implemented in both internal and external routing. With link state Internal Gateway Protocols (IGPs), load-balancing can be performed on multiple equal cost paths, or on arbitrary paths in traffic-engineered networks. For external routing protocol, i.e., with the Border Gateway Protocol (BGP), a multipath routing mode has not been standardized; recommendations are however given in [3], and some vendors now offer forms of BGP multipath (see, e.g., [4] and [5]). To our knowledge, these extensions are however seldom used [6].

Nowadays, the common practice is to over-dimension carrier networks for both operational simplicity and reliability purposes, which brings to observed utilizations largely under critical levels. It is well know that, in fact, most providers upgrade internal links with more capacity when the mean utilization gets greater than an arbitrary low threshold. Under these circumstances, *intra-AS multipath* routing is not as useful as desired because, in absence of network impairments (failures, traffic matrix variations), congestions become very rare events. Congestions after impairments may also be avoided by opportunely optimizing the IGP link weights [7] [8].

Therefore, congestion for intra-AS links is no longer a critical issue. It is instead moving at the edges of carrier providers, more specifically at peering points - where normally providers exchange (the respective clients') traffic for free. At peering points, links may not be upgraded in accordance to traffic growth, particularly in case of tensions between peers (e.g., due to traffic asymmetry). Furthermore, the free-transit relationship does not provide any incentive for peers to coordinate their routing strategies (e.g., following the preferences of the neighbor) which alleviates the potential congestion issues. The routing across peering links follows the basic BGP routing with no MED (Multi-Exit Discriminator) signaling, which would represent a trial of coordination. BGP routing is instead guided by the so-called hot-potato and tie-breaking routing (the first purely selfish, the second artificial and inefficient), which may lead to inefficient configurations.

Aiming to an improved routing management for peering settlements, authors in [9] model the coordination problem with non-cooperative game theory. A peering game is built using the MED attribute of BGP to disseminate routing costs, and is to be applied only for inter-peer critical flows. When selecting many equilibria of the game, one obtains a multipath routing solution across multiple peering links. In [10], Peering Equilibrium MultiPath (PEMP) routing strategies are defined to fine-select the equilibrium solution set. The bilateral routing cost and the number of route deviations can so be significantly decreased, and peering link congestion can be avoided.

In this paper, we propose a resilient execution policy for the PEMP routing framework accounting for network impairments, i.e., intra-AS failures, peering link failures and inter-peer traffic matrix variations. The main objective is to control the number of route deviations, and the risk of peering link congestions, due to the occurrence of such network impairments. We pragmatically consider that traffic engineering operations, modifying the IGP transit costs in accordance to intra-AS traffic matrix variations, are regularly scheduled for both the peering networks, and that network impairments can happen in between affecting the peering game cost components and changing the PEMP routes. Our resilient policy is based on two steps: (i) proactively, it first computes equilibrium sets that are robust against possible impairments; (ii) in reaction to some impairments, it reduces the multipath equilibrium set size, intersecting each new equilibrium set with the previous one. Hence, the peering path diversity (defined as the overall number of inter-AS paths used for the peering flows), corresponding to the robust equilibrium set, is adaptively decreased after impairments without deviating to new paths but restricting the number of pre-selected paths. By extensive simulation of realistic topologies and failure scenarios, we show that our policy correctly prevents route deviations and peering link congestions.

The structure is the following. Section II resumes the routing game modeling. Section III presents the resilient execution policy. Section IV presents simulation results for a realistic scenario. Section V concludes the paper.

---

[1] This work has been partially supported by the EU EuroNF Network of Excellence and the FP7 ETICS IP project.

978-1-4244-5638-3/10/$26.00 ©2010 IEEE

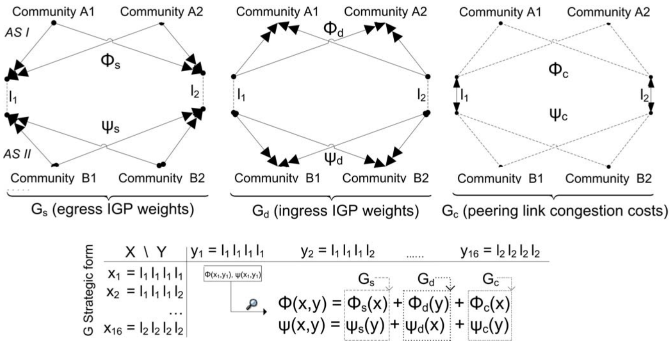

Fig. 1. Multi-pair 2-link ClubMED game composition example.

II. PEERING EQUILIBRIUM MULTIPATH FRAMEWORK

In the following, we recall the routing game modeling and the PEMP routing coordination strategies.

A. The ClubMED peering game

The ClubMED ("Coordinated MED") game modeling is characterized in detail in [9]. The idea is to re-use the MED as the means to exchange loose routing and link congestion costs between peering networks for a subset of customers' destination prefixes. The scheme relies on a non-cooperative game-theoretic modeling where each peer is represented as a rational player that can take benefit by routing accordingly to a cost game. The principle is to make the peering routing decision following efficient equilibrium strategy profiles of the game, thus allowing better collaboration between carriers. The result possibly encompasses multipath routing across the available peering links. The peering game is defined to allow a careful routing for some *destination cones* grouping a subset of customers' destination prefixes. The flows among these destination cones could represent critical Internet flows that deserve careful peer routing, because, e.g., they produce high bit-rate flow aggregates, have particular QoS or reliability requirements.

Each destination cone is reachable behind a single AS border router not at the peering border (called "ClubMED node"), and each peering AS can manage several cones. The inter-cone flows are supposed to be equivalent, e.g., w.r.t. their bandwidth, so that their path cost can be fairly compared and their routing coordinated. Practically, a destination cone can be identified by a BGP 'community id' tag in order to give to the decision process the means to identify the ClubMED routing scope. The game is to be built only at the ClubMED nodes connecting the destination cones; its 'solution' relies on a coordinated peering equilibrium indicating at least one egress peering link for each inter-cone flow.

As depicted in Fig. 1, the peering game is composed of three games: a selfish game G_s built upon the egress IGP path costs (from the ClubMED node toward the peering links), a dummy game G_d built upon the ingress IGP path cost (inverse direction), and a congestion game G_c built upon congestion costs assigned to peering links. The IGP path costs can be coded with little primitive extensions via a composite MED attribute in BGP announcements. To build the congestion game, the bit-rate of each inter-cone flow should be known by each ClubMED node (e.g., via Netflow).

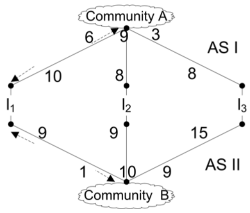

I\II	l_1	l_2	l_3
l_1	$(16,10)^2$	$(19,10)^2$	$(13,16)^8$
l_2	$(\mathbf{14,19})^0$	$(\mathbf{17,19})^0$	$(11,25)^6$
l_3	$(\boldsymbol{\mathit{14,18}})^0$	$(\mathbf{17,18})^0$	$(11,24)^6$

Fig. 2. 3-link example.

Mathematically, it is a particular potential game, in which the equilibria are minima of a potential function and vice-versa. Fig. 2 is an example with G_s and G_d only (the superscripts are the potential values). It is possible to have a single equilibrium, or many as in Fig. 2, and the equilibria may be Pareto-inefficient such as (l_3,l_1), (l_2,l_2) and (l_3,l_2) - remembering that a profile is Pareto-efficient if no other profile decreases a player's cost without increasing the other player's cost. To capture potential variations of the path costs, e.g., due to regular IGP Weight Optimization (IGP-WO), some incertitude parameter needs to be added to the game cost components. This leads to a very important enlargement of the game equilibrium sets, and rises the need for a coordination policy to fine-select a few equilibria.

B. PEMP coordination strategies

The following PEMP strategies can be implemented to fine-select a multipath equilibrium solution (see [10]).

1) Nash Equilibrium MultiPath (NEMP): it is the one-shot strategy to which to coordinate; it selects the equilibria of the Nash set, only the Pareto-superior ones if any.

2) Pareto-frontier: in an infinitely repeated context, it selects the profiles of the Pareto-frontier. Indeed, from "folk-theorem"-like results [13], this strategy is an equilibrium of the repeated game and grants a maximum gain for the players in the long-run. Nevertheless, computing these strategies is very complex combinatorially. Moreover, the unilateral trust for such a strategy could decrease whether in a short period of analysis the gains reveal to be in favor of a single peer.

3) Unselfish-Jump: to guarantee balance in gains in the short term (helping to keep a high level of reciprocal trust), in this strategy, after shrinking the Nash set w.r.t. the Pareto-efficiency, for each equilibrium the ASs might agree to make both a further step towards the best available strategy profile such that the loss that one may have moving from the selected equilibrium is compensated by the improvement upon the other AS. One AS may unselfishly sacrifice for a better bilateral solution. This strategy makes sense only if the other AS is compensated with a bigger improvement, and returns the favor the next times.

4) Pareto-Jump: the jump is constrained toward a Pareto-superior profile only (not necessarily in the Pareto-frontier), hence avoiding unselfish sacrifices. E.g., in the example of Fig. 2, we would jump from the Pareto-superior Nash equilibrium (l_3, l_1) to the Pareto-superior profile (l_1, l_3). We would not have this jump for the Unselfish-Jump policy, that would prefer instead (l_1, l_1) with a global gain of 6 instead of "just" 3 with (l_1, l_3). Finally, note the jump strategies are not binding: it would be enough to associate them with the menace to pass to one of the more selfish choices. Also note that MEDs from different ASs should be normalized to the same IGP weight scale in order to be comparable.

III. RESILIENT PEERING MANAGEMENT

A resilient routing framework generally aims to produce a dependable network state that does not suffer from the occurrence of network impairments [11] - i.e., anomalies like traffic matrix variations (see [7]) or link/node failures (see [8]). In an AS IP network, with a link-state IGP routing protocol, *robust* routing algorithms pro-actively compute IGP link weights accounting for the anomalies and aiming at some *performance goal* on the service level to guarantee after failures (typically, congestion avoidance or delay bounds). PEMP is completely agnostic on the implementation of robust algorithms in IGP routing to prevent from anomalies such as intra-AS link failure and intra-AS traffic matrix variations. Our peering context is in fact an "overlay"routing between ClubMED nodes, on top of the underlying IGP. The two layers have, however, a tight coupling due to the transit IGP path cost adopted in the game setting. The PEMP network is composed of ClubMED nodes, border routers interconnected via peering links, and selected inter-peer flows (see Fig. 1). The transit IGP path costs from and to ClubMED nodes and peering routers compose the ClubMED game, whose equilibria are used for PEMP routing. In this context, the *robustness performance goals* we target for PEMP routing are:

1) reduction of PEMP route deviations;
2) prevention from peering link congestions.

They can be pursued by (i) computing robust PEMP solutions just after IGP-WOs (proactive step), and (ii) controlling how intermediate PEMP solutions are applied after the occurrence of network impairments (reactive step).

A. PEMP execution policy

In [10], the PEMP strategies were assumed to be executed just after IGP-WO, one at each side, each one considering updated traffic matrix. Aiming to defining a resilient execution policy for the original PEMP routing framework, additional dependability assumptions can be made. In fact, IGP-WO operations may be executed quite rarely, and between two executions the game cost components may change significantly after network element failures (which should not be considered as unusual events in carrier networks [12]). A PEMP solution, adhering to the previous ClubMED game and IGP-WO settings, may be not consistent with the current network state. Between two IGP-WO operations, the following impairments could modify the ClubMED game setting:

- intra-AS link failures and restorations: even if the single IGP link weights do not change, transit IGP path costs (hence G_s and G_d) could be modified since some IGP paths will change to circumvent the failure;
- peering link failures: the corresponding unilateral congestion cost components of G_c is set to infinity;
- significant inter-peer flow traffic matrix variations: the cost components of G_c are updated.

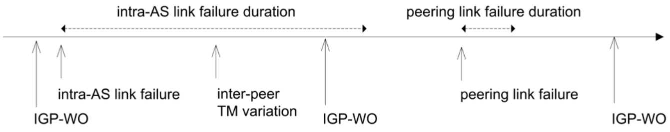

Fig. 3. Reference timeline scenario example

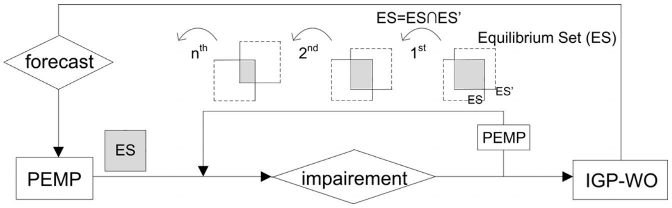

Fig. 4. PEMP execution policy chart

Fig. 3 depicts a reference scenario example. Between two planned IGP-WOs, the inter-peer flow traffic may change significantly, and intra-AS and peering link failures may occur. The frequency, time-to-repair and time-between-failures distributions of link failures can vary (but can be estimated as described in Section IV-A). We assume that intra-AS link failures can last much more than peering link failures, because peering links are rarely trenched as lengthy direct physical connections between two carriers; they rely, instead, quite often on local router interconnections at Internet eXchange Points (normally, private IXPs for top-tier peering carriers) that can be reestablished rapidly. As an acceptable simplification, we thus assume that intra-AS link failures can persist after the next IGP-WO while peering link failure do not.

The resilient execution policy we design, is summarized in the chart of Fig. 4. After each IGP-WO, a new *robust PEMP equilibrium set* is computed anticipating future network impairments (see the next section). Then, when a network impairment occurs, a new set is computed, and the retained set for PEMP is the intersection between this set and the previous one. This process is repeated until the next IGP-WO. If an intersection is empty, then the whole new set is applied.

When a network impairment occurs, the intersection should be retained first, instead of the whole new set, to avoid excessive route deviations. The intersection will in fact induce a reduction on the equilibrium path diversity, i.e., withdraws of some paths from the current multipath solution. Subsequent intersections will tend toward a best-path routing solution. Nevertheless, if the impairment forecast is reliable, the equilibrium set intersections will often be comprehensive (hence, with a small number of path withdraws) and rarely be empty.

B. Impairment management

Let us detail precisely how the impairments are managed in the resilient PEMP execution policy.

1) Peering link failures: As already mentioned, peering link failures are expected to be transient and not to last untill the next IGP-WO. Therefore, this type of impairment does not need to be considered in the impairment event forecasting, and shall be managed with *transient-rerouting*.

Peering transient-rerouting in PEMP should:

- update the G_c components corresponding to the strategies with that link for at least one flow, i.e., if link l fails: $\phi_c(x) = \infty,\ \forall x \in X|l \in x,\ \ \psi_c(y) = \infty,\ \forall y \in Y|l \in y$
- compute the new equilibrium set;
- route according to the non-empty intersection, or the new set if the intersection is empty;
- restore the previous set when the failure is restored.

978-1-4244-5638-3/10/$26.00 ©2010 IEEE

2) Inter-peer traffic matrix variations: Changes of the inter-peer flow bitrates are normally already considered at each new IGP-WO. However, if those changes exceed an arbitrary alert threshold, it may be worth considering them immediately to avoid congestions at peering links, recomputing the congestion cost components of G_c (see [10]), thus the new equilibrium set, then applying the intersection solution. Unlike the previous type, this routing change is obviously not transient.

3) Intra-AS link failures: This kind of failures has a potentially broader impact on the ClubMED game setting in that it can change the IGP cost (from and to ClubMED nodes and peering routers) for many transit paths. As explained in Sect. II-A, the selfish game and the dummy game are built using the transit IGP path costs. Reminding the notation in Fig. 1, let $\phi(x, y)$ and $\psi(x, y)$, and X and Y, be the cost functions and the strategy sets for the first and the second peer, respectively, with $x \in X$ and $y \in Y$. The cost functions are decomposed in their selfish ($\phi_s(x)$, $\psi_s(y)$), dummy ($\phi_d(y)$, $\psi_d(x)$) and congestion ($\phi_c(x)$, $\psi_c(y)$) game cost components.

Only the selfish and dummy components are affected by the occurrence of the intra-AS link failures. More precisely, let $\delta^k_{i,j}$ be the integer IGP path cost variation induced on the path from the ClubMED node i toward the peering router j by a failure at intra-AS link l_k, and ${\delta'}^k_{j,i}$ the dual from the peering router j toward the the ClubMED node i. A failure on link l_k will thus produce cost changes for all the strategies that included that link for at least one flow, i.e.,

$$\phi_s(x)' = \phi_s(x) + \sum_{i,j} \delta^k_{i,j}, \quad \forall x \in X \ s.t. \ l_k \in x \tag{1}$$

$$\phi_d(y)' = \phi_d(y) + \sum_{i,j} {\delta'}^k_{j,i}, \quad \forall y \in Y \ s.t. \ l_k \in y \tag{2}$$

where $\phi_s(x)'$ and $\phi_d(y)'$ are the new cost function values. $\psi_s(y)'$ and $\psi_d(x)'$ are similarly computed. The δ cost variations can be pre-computed during the impairment forecasting phase introduced in Section III-A before the main PEMP execution. We propose to use the pre-computed δ cost variations, together with a probability estimation of the intra-AS link failure, to compute a robust PEMP equilibrium set.

Let p_k be the probability that a failure on link l_k will occur before the next IGP-WO, say for the first peer. It is possible for a carrier provider to estimate such a probability distribution, as argued in [12] and explained in Section IV-A. Under the assumption that having simultaneous link failures in the same network is a stochastically negligible event, and that thus an IGP path cost variation is given by a single failure, one can compute the expected IGP path cost variations as:

$$\widetilde{\delta_{i,j}} = \sum_k p_k \delta^k_{i,j}, \quad \widetilde{\delta'_{j,i}} = \sum_k p_k {\delta'}^k_{j,i}, \ \forall i, j \tag{3}$$

And the dual computation stands for the second peer. To take into account these path cost variations, the two peers should exchange them; this operation may, however, give to the peer an excessive insight on its network, and generate too much signaling. Instead, each peer could announce a global directional path cost error, one for the egress direction and one for the ingress one. Let ϵ^I and ϵ^{II} be the egress cost errors for AS I and AS II (resp.). The path costs variations may assume quite different values; an optimistic computation is the median error, $\epsilon^I = \mathrm{median}_{(i,j)} \left\{\delta_{i,j}/c^I_{i,j}\right\}$, where $c^I_{i,j}$ is the shortest path cost for the ClubMED node i to the peering router j of AS I. Similarily for ϵ^{II} and the ingress errors. Extending the game to take into account these errors, one gets a *potential threshold* (τ_P) under which a profile becomes an equilibrium. In fact, each potential difference from (x_1, y_1) to (x_2, y_2) can be increased of $a_I(x_1, x_2) + a_{II}(y_1, y_2)$, where $a_I(x_1, x_2) = \epsilon^I \cdot (\phi_s(x_1) + \phi_s(x_2))$ and $a_{II}(y_1, y_2) = \epsilon^{II} \cdot (\psi_s(y_1) + \psi_s(y_2))$. An optimistic threshold can be:

$$\tau_P = \min_{x_1, x_2 \in X} \{a_I(x_1, x_2)\} + \min_{y_1, y_2 \in Y} \{a_{II}(y_1, y_2)\} \tag{4}$$

Denoting with $P(x_0, y_0)$ the potential minimum, all strategy profiles (x, y) such that $P(x, y) \leq P(x_0, y_0) + \tau_P$ will be considered as equilibria. This operation can also allow escaping selfish (endogenous) solutions mainly guided by G_s + G_c, introducing Pareto-superior profiles in the Nash set.

IV. PERFORMANCE EVALUATION

We emulated the real peering scenario between the Geant2 and Internet2 research networks with 3 peering links, 5 destination cones, 2 in Internet2 and 3 in Geant2, i.e. a total of 6 flows per direction to be routed between Internet2 and Geant2. The time horizon is 3600 hours with 200 18h-spaced samples of the real traffic matrix from the two networks were used. This is the same datasets as [10], which is referred to for details omitted here due to space limitations. The inter-peer PEMP flows occupy 2/3 of the peering capacity at each side.

A. Link Failure Model

To simulate our proposal on realistic instances, we adopted the backbone link failure model for operational IP networks presented in [12]. The authors show how they estimated link failure rate, time-between-failures and time-to-repair probability distributions. These distributions give, to our knowledge, the only link failure model in the literature pertinent for our framework. For the link failure rate distribution, power-law regimes are available. For the time-between-failures and time-to-repair probability distributions, they could not find an analytical approximation, hence we sampled the available experimental comulative distribution functions.

B. Numerical results

We performed impairment forecasting according to the link failure distribution, and built the intra-AS link failure scenario according to the other two distributions. No failure model for peering links exists in the literature as of our knowledge; hence we used the same time-to-repair distribution scaling it down by a factor of ten. For inter-peer traffic matrix variations, we considered [-30%, 30%] uniformly distributed variations, with four of such impairments positioned in a time instant uniformly distributed over each IGP-WO interval. We adopted the Unselfish-Jump PEMP strategy (see II-B).

1) Impact of impairments on route deviations: In Fig. 5, we show the equilibrium path diversity, defined as the overall number of paths used to route all the flows, for the new equilibrium solutions after the occurrence of impairments between two IGP-WO. The graphic is drawn using a classical boxplot statistics format representing the minimum, first quartile, median, third quartile and maximum values. We illustrate the result obtained for the successive impairments that

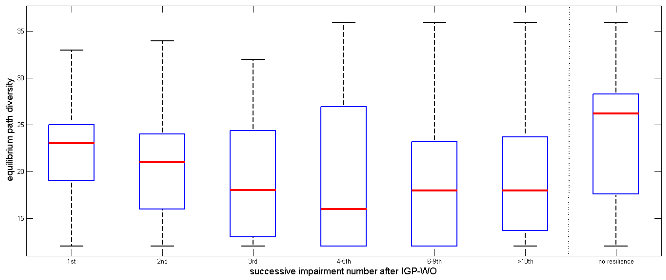

Fig. 5. Equilibrium path diversity (boxplot statistics)

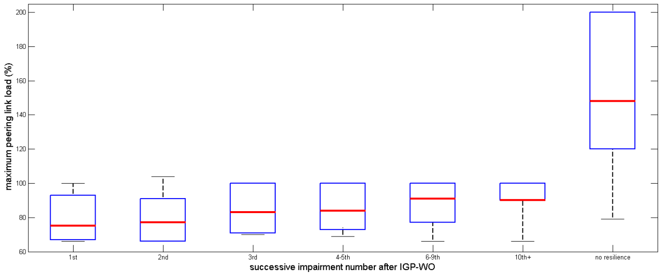

Fig. 6. Maximum peering link load (boxplot statistics)

occur after the previous IGP-WO occurence. For the sake of comparison, the last boxplot has been added, representing the results that would be obtained re-executing PEMP without the resilient policy framework, but with intermediate equilibrium set intersections. It is worth noting that the minimum path diversity (6 best-path flows from Internet2 to Geant2, and 6 in the opposite direction), and the maximum 36 (i.e., full multipath). We can assess that:
(i) the implementation of our resilient policy manifests with a decreasing equilibrium path diversity, symptom of positive set intersections and progressive pruning of those paths that are no longer appropriate after impairments;
(ii) the high median of the solution without our resilient policy shows the utility of the impairments forecast phase that allows the selection of a robust PEMP starting solutions;
(iii) both the median and the first quartile decrease for the first 5 impairments, and the median starts increasing smoothly afterwards. In particular, for the 4-9th impairments the first quartile remains equal to the minimum. This clearly indicates that as many successive impairments occur, the multipath routing solution degenerates towards a best-path showing that often robust singleton solutions exist (we expect, however, this not to manifest so often with more peering links and more flows, with the resulting larger number of strategy alternatives);
(iv) the median increases after the 4th impairment likely because of a higher occurrence of null intersections.

2) Impact of impairments on peering link loads: We measure the quality of the transient-rerouting solution after peering link failure and inter-peer traffic matrix variations, looking at the maximum peering link load after these events. Results are reported in Fig. 6 with and without our resilient PEMP policy. Without the resilient PEMP policy, we assume that the flows are rerouted over the remaining paths (excluding the failed peering link) or across the shortest path, if no remaining PEMP path exists. It is seen that we will have periods with severe overload of the peering links in the non-resilient case. Remembering that the best peering link load remains 66% because 2/3 of the peering capacity is at least used by the inter-peer flows, we can assess that:
(i) peering link congestion after impairments can be practically avoided, thanks to the appropriate use of the congestion game;
(ii) the median utilization smoothly increases with the occurrence of impairments, which is likely to be due to the decrease of the equilibrium set size. In fact, a lower number of equilibria induces less load-balancing and hence a worse filling of the available peering capacity.

V. SUMMARY

In this paper, we propose a resilient framework for peering management. As a matter of fact, peering interconnections among top-tier carriers are becoming the real bottleneck of the Internet. A research challenge is thus to define appropriate routing frameworks to manage the routing of critical flows across peering links. An interesting solution consists in performing peering multipath routing following the equilibria of a routing coordination game built by peering carriers by opportunely using the MED attribute of BGP [10], which in fact allows important performance improvements, decreasing the occurrence of route deviations and peering link congestions.

We rise the important issue of dealing with network impairments (link failures, traffic matrix variations) in the peering management framework. We propose a resilient execution policy of multipath equilibrium routing, which consists of a first block of robust peering equilibrium selection accounting for intra-AS link failures, and a second block defining consistent rerouting procedures in case of impairments between scheduled traffic engineering operations.

By extensive simulations on a realistic scenario, we show that peering link congestion after failure can be avoided, and that the route deviations are significantly reduced. In this resilient framework, route deviations manifest as successive path withdraws from the multipath solution, with a resulting decrease of the path diversity used for the routing solution.

REFERENCES

[1] R. Teixera et al., "In Search of Path Diversity in ISP Networks", in *Proc. of ACM Internet Measurement Conference*, Oct. 2003.
[2] J. He, J. Rexford, "Towards Internet-wide multipath routing" in *IEEE Network magazine*, March 2008.
[3] A. Lange, "Issues in Revising BGP-4", draft-ietf-idr-bgp-issues (2003)
[4] "Configuring BGP to Select Multiple BGP Paths", JUNOS document.
[5] "BGP Best Path Selection Algorithm", Cisco documentation.
[6] E. Elena, J.-L. Rougier, S. Secci, "Characterization of AS-level Path Deviations and Multipath in Internet Routing", in *Proc. of NGI 2010*.
[7] P. Casas, L. Fillatre, S. Vaton, "Multi hour robust routing and fast load change detection for traffic engineering", in *Proc. of IEEE ICC 2008*.
[8] A. Sridharan, R. Guérin, "Making IGP Routing Robust to Link Failures " in *Proc. of Networking 2005*.
[9] S. Secci et al., "ClubMED: Coordinated Multi-Exit Discriminator Strategies for Peering Carriers", in *Proc. of NGI 2009*.
[10] S. Secci et al., "PEMP: Peering Equilibrium MultiPath routing", in *Proc. of GLOBECOM 2009*.
[11] B. E. Helvik, *Dependable Computing Systems and Communication Networks; Design and Evaluation*, Tapir academic publisher, 2007.
[12] A. Markopoulou et al., "Characterization of Failures in an IP Backbone", *IEEE/ACM Transactions on Networking*, Vol. 16(4), 2008.
[13] R.B. Myerson, *Game Theory: Analysis of Conflict*, Harvard Univ. Press.