Resilient Inter-Carrier Traffic Engineering for Internet Peering Interconnections

Stefano Secci, *Member, IEEE*, Huaiyuan Ma, Bjarne E. Helvik, *Member, IEEE*, and Jean-Louis Rougier

***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 occur 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.**

***Index Terms*—Routing resiliency, Internet reliability, peering management, game theory, BGP, inter-domain routing, multipath routing.**

I. Introduction

RESILIENCY in Internet routing is usually associated with the level of path diversity and with the ability to perform load balancing and multipath routing. More generally, this stands for distributed routing protocols, in any context, from sensor networks to optical networks, and Border Gateway Protocol (BGP) networks. However, path diversity in distance (path) vector routing protocols, such as BGP, is particularly difficult to achieve at an acceptable level, especially when local policy filters are applied. In this article, we present a framework that increases path diversity and resiliency in Internet routing. We propose a multipath routing policy meant to be applied on peering interconnections, which in fact represent the real bottleneck of current Internet.

Manuscript received July 30, 2010; revised February 8, 2011 and April 25, 2011. The associate editor coordinating the review of this paper and approving it for publication was G. Pavlou.

S. Secci is with the LIP6, Université Pierre et Marie Curie (UPMC), 4 place Jussieu, 75005 Paris, France (e-mail: stefano.secci@lip6.fr).

H. Ma and B. E. Helvik are with the Q2S Center of Excellence, Norvegian University of Science and Technology (NTNU), O.S. Bragstads Plass 2E, 7491 Trondheim, Norway (e-mail: {huaiyuan, bjarne}@q2s.ntnu.no).

J.-L. Rougier is with the Institut Telecom, Telecom ParisTech, LTCI CNRS, 23 avenue d'Italie, 75013 Paris, France (e-mail: rougier@telecom-paristech.fr).

This work has been partially supported by the Euro-NF (Anticipating the Network of the Future - From Theory to Design) Network of Excellence and the ETICS (Economics and Technologies for Inter-Carrier Services) Integrated project of the Seventh Framework Programme of the European Union.

A preliminary version of this article has been presented at the 2010 IEEE Global Communications Conference (GLOBECOM 2010) [1].

Digital Object Identifier 10.1109/TNSM.2011.110311.100064

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 [2] [3]. 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 BGP, a multipath routing mode has not been standardized; recommendations are however given in [4], and some vendors now offer forms of BGP multipath (see, e.g., [5] and [6]). To our knowledge, these extensions are however seldom used [7].

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 [8] [9].

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, often 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 [10] model the peering coordination problem with non-cooperative game theory. A peering game - called "ClubMED game"- can be built using the MED attribute of BGP to disseminate routing and congestion costs, and is conceived to be applied only for inter-peer critical flows. When selecting more than one equilibrium of the ClubMED game, one obtains a multipath routing solution across multiple peering links. In [11], 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 the adopted link failure model and reports simulation results. Section VI reviews related work. Section VII concludes the paper.

II. Peering Equilibrium MultiPath Framework

In the following, we first recall how to model the routing across peering interconnection with non-cooperative game theory (defining the ClubMED game), then we present possible multipath routing coordination strategies for this framework (called Peering Equilibrium MultiPath, PEMP, strategies).

A. The ClubMED peering game

The ClubMED game modeling is characterized in detail in [10]. 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

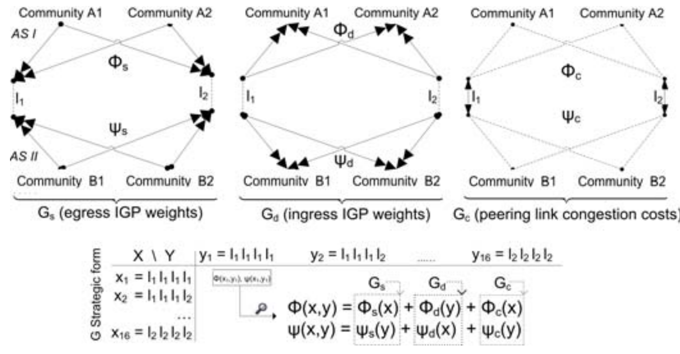

Fig. 1. Multi-pair 2-link ClubMED game composition example.

could represent critical Internet flows that deserve careful peer routing, because, e.g., they produce high bit-rate aggregates, or 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., with respect to 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).

Mathematically, it is a particular potential game, in which the equilibria correspond to the minima of a potential function and vice-versa. Fig. 3 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. 3, 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. The single Pareto-efficient profile, thus belonging to the Pareto-frontier, is (l_1,l_3), which is not an equilibrium. 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 [11]).

Fig. 2. Reference timeline scenario example.

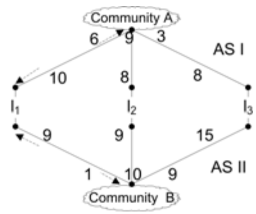

I\II	l_1	l_2	l_3
l_1	$(16,10)^2$	$(19,10)^2$	$(13,16)^8$
l_2	**(14,19)**^{0}	**(17,19)**^{0}	$(11,25)^6$
l_3	***(14,18)***^{0}	**(17,18)**^{0}	$(11,24)^6$

Fig. 3. 3-link example.

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 [15], 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 with respect to 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. 3, 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.

The four PEMP strategies[1] were tested and compared in [11]. The evaluation on the bilateral routing cost, the peering link load, the route stability and the time complexity, shows that the Unselfish-Jump strategy (in turns relying on the NEMP strategy solution as above-mentioned) offers reasonable performance trade-offs. In the following of this paper, we present how the PEMP behaviour can be affected by the occurrence of network impairments and how the equilibrium solution can be reinforced. The resilience motivation is to prevent from peering disconnections that could cause worldwide outages of the Internet.

[1]It is worth noting that forms of inter-AS Equal Cost Multi Path (ECMP) routing do not apply to the PEMP routing context. Informally, we can state that the PEMP is the rational extension of ECMP, where the path cost equality condition of ECMP is extended to the equilibrium equivalence condition of PEMP (for peering routes only).

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 [12] [13] - i.e., anomalies like traffic matrix variations (see [8]) or link/node failures (see [9]). 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. The PEMP peering context is in fact an "overlay " routing between ClubMED nodes, on top of the underlying IGP. The two routing layers have, however, a tight coupling due to the transit IGP path cost adopted in the ClubMED game setting. The PEMP routing network is composed of egress routers (i.e., the ClubMED nodes), border routers interconnected via peering links, and selected inter-peer flows routed between ClubMED nodes across peering links (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),
(ii) controlling how intermediate PEMP solutions are applied after the occurrence of network impairments (reactive step).

In the following, we detail the execution and impairment management policies we propose for resilient peering management.

A. PEMP execution policy

In [11], 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 [14]).

In such a scenario, 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.

It is worth noting that a router failure can entail many link failures. In Section III-B, we detail precisely how these network impairments modify the ClubMED game setting.

Fig. 2 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. They rely, instead, quite often on local router interconnections at Internet eXchange Points (normally, private IXPs for top-tier peering carriers). These interconnections may be reestablished rapidly (empirical support of this assumption is for instance found in [32], where the authors observed that peering link failures are typically transient with short downtimes). As an acceptable simplification, we thus assume that intra-AS link failures can persist after the next IGP-WO while peering link failures 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 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. The best case would correspond to a very low reduction of the starting equilibrium set, while the worst case to an empty intersection at the first impairment.

B. Impairment management

Let us detail precisely how the impairments are managed in the policy. Anomalies in the PEMP framework can be peering link failure, inter-peer traffic matrix variations - i.e., peering route and bit-rate variations for the PEMP-managed critical flows - and intra-AS link failures. As explained hereafter, the first two impairments are managed with transient rerouting, while the latter is managed proactively with robust equilibrium computation and reactive transient rerouting.

1) Peering link failures(transient rerouting):: As already mentioned, peering link failures are expected to be transient and not to last until 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 consist of:

- the update of all the G_c congestion game components corresponding to the multipath strategies with that link selected for at least one flow, i.e., if link l fails: $\phi_c(x) = \infty,\ \forall x \in X | l \in x,\ \ \psi_c(y) = \infty,\ \forall y \in Y | l \in y$
- the computation of the new equilibrium set;
- the application of the routing solution corresponding to the non-empty intersection with the previous set, or the new set if the intersection is empty;
- the restoration of the previous equilibrium set when the peering link failure is restored.

2) Inter-peer traffic matrix variations(transient rerouting):: 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 before the next IGP-WO, it may be worth considering them immediately to avoid congestions at peering links. As for the peering link failure reaction, transit rerouting shall be applied, recomputing the congestion cost components of G_c, 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: transient rerouting and robust equilibrium set computation: Intra-AS link failures have 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. Their frequent occurrence is in fact one of the major issues behind BGP routing instability and converge issues. In our policy, we manage them with two steps: transient rerouting upon their occurrence, and proactive robust equilibrium computation using link failure forecasting models.

Let us better detail how intra-AS link failures affect the routing game composition and how they can be managed by the resilient PEMP policy. As explained in Sect. II-A, the selfish game and the dummy game are built using the transit IGP path costs[2]. 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

[2]It is worth stressing that, as in BGP hot-potato routing, intra-AS path cost variations can lead to egress link change. In PEMP routing this may also happen, even if such changes are subject to rational and strategically-justified bilateral decisions (routing equilibria).

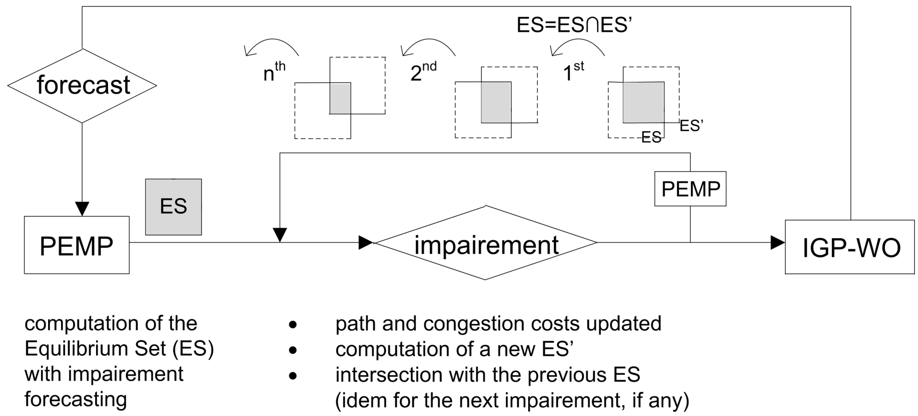

Fig. 4. PEMP execution policy chart.

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.,

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

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

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, which in fact can be done during the impairment forecasting phase introduced in Section III-A before the main PEMP execution (see Fig. 4). A possible link failure model is described in Section IV-A. 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 [14] 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:

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

And the dual computation stands for the second peer.

In order to take into account these path cost variations, the two peers should exchange these parameters (e.g., coded as an extension of the MED field); this operation may, however, give to the peer an excessive insight on its network, and moreover generate too much signaling overhead.

Instead of announcing the $\widetilde{\delta}$ variations, each peer could thus announce a global directional path cost error, one for the egress direction and one for the ingress one. Just an error, instead than explicit per-path cost variations, can sufficiently abstract critical intra-domain routing information. Let ϵ^I and ϵ^{II} be the egress cost errors for AS I and

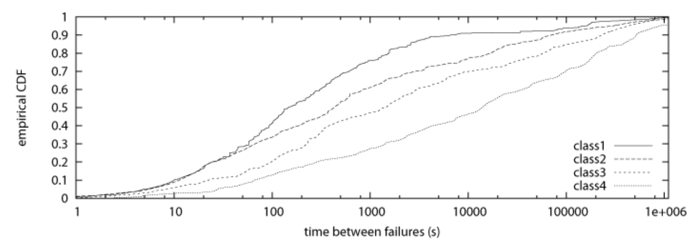

Fig. 5. Per-class CDFs of the time-between-failures distribution, [14].

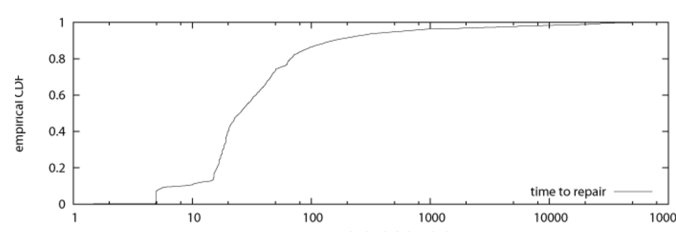

Fig. 6. CDF of the time-to-repair distribution, [14].

AS II (resp.). Being aware that path costs variations may assume quite different values, an optimistic computation correspond to the median error (that discards extreme errors), $\epsilon^I = \text{median}_{(\mathrm{i,j})} \{\delta_{\mathrm{i,j}}/\mathrm{c}^{\mathrm{I}}_{\mathrm{i,j}}\}$, 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. The game can be easily extended to take into account these error margins. They define a *potential threshold* under which a profile becomes an equilibrium. More precisely, the minimum potential strategies are found, then the other profiles that have a potential within the minimum plus a threshold (τ_P) are considered as equilibria too. Each potential difference ΔP 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:
$a_I(x_1, x_2) = \epsilon^I \cdot (\phi_s(x_1) + \phi_s(x_2))$
$a_{II}(y_1, y_2) = \epsilon^{II} \cdot (\psi_s(y_1) + \psi_s(y_2))$.

Under this setting, an optimistic threshold can be:

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

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.

It is worth noting that this procedure can be done jointly with an IGP path cost errors estimation due to routing strategy changes after two main PEMP executions, as explained in [10], which may in fact further increase the potential threshold. Indeed, such errors are likely to be relevant for low-dimensioned networks in which discrete load variations can have a relevant impact on the IGP weight. For simplicity, in the following we will consider that intra-AS networks are lightly loaded and thus that such errors are negligible.

IV. Performance Evaluation

We emulated a realistic 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 employed topology is depicted in Fig. 7.

The time horizon is 3600 hours with 200 18h-spaced samples of the real traffic matrix from the two networks were used. We used 252 successive traffic samples from [22] for Geant2 and from [23] for Internet2. The TOTEM toolbox [21] was used to run a IGP-WO heuristic, with a maximum IGP weight of 50 for both ASs. The original link capacity was scaled down by 10 to create an intra-domain congestion risk (and hence to observe significant transit path cost variations after IGP-WO). The inter-cone routing generates additional traffic for the traffic matrices. We used a random inter-cone traffic matrix such that flows are balanced with 200 Mb/s per direction, which corresponds to $2/3$ of the total available peering capacity; with respect to the intra-AS traffic, the cone-to-cone connection requests have an average roughly corresponding to the global (two-side) average of intra-AS ones. To evaluate the effectiveness of the congestion game we considered peering links with 100 Mb/s per direction.

A. Link Failure Model

With the aim to simulate our proposal on realistic instances, we need a link failure model for operational IP backbone networks. An investigation on the intra-AS IP link failure behavior in a backbone carrier network can be done as explained in [14] (results for the Sprint network). The authors show how to estimate link failure rate, time-between-failures and time-to-repair probability distributions.

In our simulations, we adopted the distributions of [14], which to our knowledge is the only link failure model pertinent for our framework in the literature. Link failures are classified into planned failures and unplanned failures. It is shown that 20% of all failures are planned failures. The unplanned failures are further classified into individual link failures and shared link failures (router-related or optical related). Individual link failures account for 70% of the unplanned failures. An interesting observation is that 55% of all individual link failures are caused by 2.5% of the links. These are denoted *high failure links*. All high failure links are backbone links, while low failure links are mainly access links. Since we deal only with core IP links, we are interested in the distributions of high failure links.

1) Individual link failure distribution: The individual link failure distribution is analytically identified in [14]. As mentioned above, the observed failure frequencies are highly heterogeneous, with the high failure links being affected by most of the failures. Most of high failure links are inter-POP links, and half of them share a router with another high failure link. Once classified in the two classes of high failure and low failure links, two power law regimes could be identified. After ordering links k from the most failing one to the least failing one, the failure probability p_k as function of the ranked position $r(k)$ follows a power-law distribution $p_k \propto r(k)^{-0.73}$ (while for low failure links, the power-law distribution is instead more steep, $p_k \propto r(k)^{-1.35}$). For simulations, ASs' links need to be ranked according to their propensity of failure; in our simulations, we rank first the link connected to nodes with higher degree, and among those with equal degree the longer links are ranked first.

2) Time-between-failure distribution: Empirical Cumulative Distribution Functions (CDFs) of the time-between-failures for high failure links are given in [14]. A corresponding analytical approximation could not be found. It is shown that a subset of the high failure links experiences very bursty failures, i.e., most of failures occurs over a short time period. On the other hand, some other subsets of high failure links exhibit failures patterns that persist over the entire time period. Therefore, the time-between-failure process presents quite heterogeneous cases; however, for our simulations, it was possible to qualify four different subsets of high-failure-links sharing similar distributions, and to assign the links to each class using the already mentioned link ranking adopted to assign the p_k. The empirical CDFs for each subsets are reported in Fig. 5, for the sake of presentation.

3) Time-to-repair distribution: The empirical CDF of the time-to-repair for the unplanned link failures are also investigated in [14]. The CDFs for the different type of unplanned failures - i.e., high failure, low failure and shared (router and optical) failures - are significantly different from each other. For example, the ratio of links having a time-to-repair less or equal than 25s is roughly 50% for high failure links, 30% for low failure links, and only 6% for shared failures. All in all, the time-to-repair for high failure links (whose distribution is used in our simulation) is generally much shorter than for the other types of unplanned failures. The adopted time-to-repair CDF for high failure links is reported in Fig 6, for the sake of presentation.

B. Numerical results

In our simulations, we performed impairment forecasting according to the individual link failure distribution, and built the intra-AS link failure scenario according to the other two distributions above. For peering links, no failure model exists

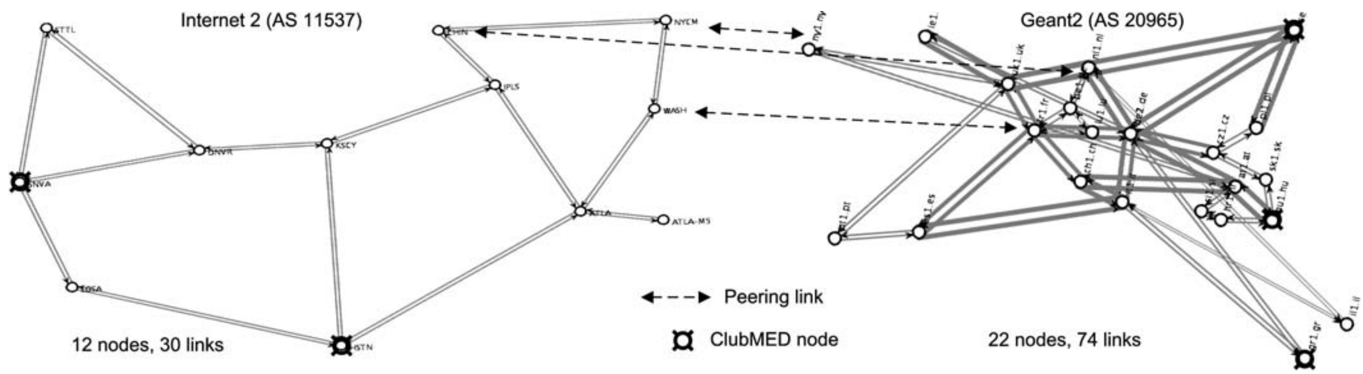

Fig. 7. Internet2 - Geant2 peering scenario with 3 peering links.

in the literature as of our knowledge; we used the same time-to-repair distribution scaling it down by a factor of ten (thus assuming a much shorter time for peering links). Moreover, for peering links, we used the fourth class of the time-between-failure distribution (thus a smaller occurrence than for intra-AS links), disregarding concurrent peering link failures. 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 as it appeared as the best one with respect to performance and trust-enforcement criteria; its choice is, however, not binding and other strategies, also different than the four ones presented in Section II-B, can be adopted since the policy is totally transparent to the adopted PEMP set selection algorithm.

1) Impact of impairments on route deviations: In Fig. 8, we show the path diversity of the equilibrium solution, 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 occur after the previous IGP-WO occurrence. 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 is 12 (6 best-path flows from Internet2 to Geant2, and 6 in the opposite direction), and the maximum 36 (i.e., full multipath - 18 paths per direction). 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 solution showing that often robust singleton solutions appear (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 because of a higher occurrence of null intersections.

2) Equilibrium set intersection dynamics: Also to better investigate the characteristic noted above (iii), we report in Fig. 9 the statistics about the intersection ratio for the equilibrium set size after impairments, i.e., the percentage of previous multipath equilibria kept in the new PEMP set after a failure. We remember that in the case there is a null intersection, the whole new set is used; such cases are considered as 0% intersections. We can assess that:

(i) with respect to the case with no resilience management, we obtain a high median ratio of intersection, above 30% also for intervals with a high number of impairments (more than nine);

(ii) the trend is slightly decreasing with the number of successive impairments, with still a first quartile always bigger than 20% and a third quartile bigger than 55%.

These results confirm that the simple method adopted for impairment forecast is appropriate, and thus that the approximation made for (3) of a single failure is acceptable in this framework (even if in the simulations a non negligible number of concurrent multiple failures could be observed).

3) 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. 10 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

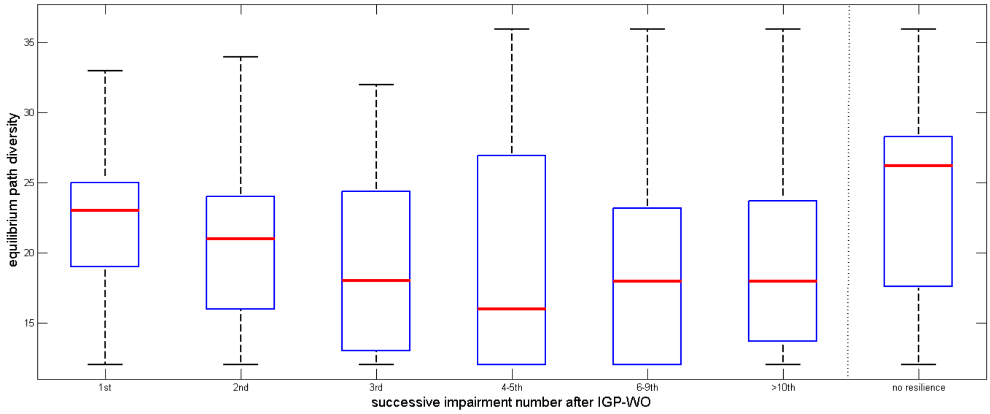

Fig. 8. Equilibrium path diversity (boxplot statistics).

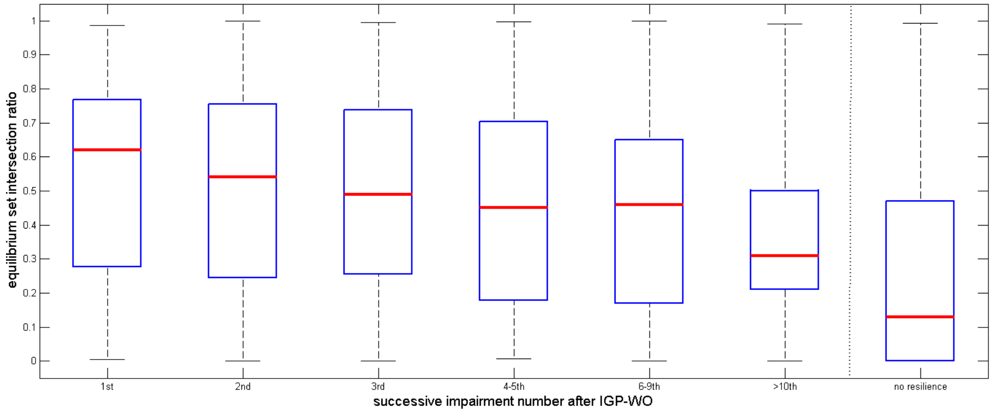

Fig. 9. Equilibrium set intersection dynamics (boxplot statistics).

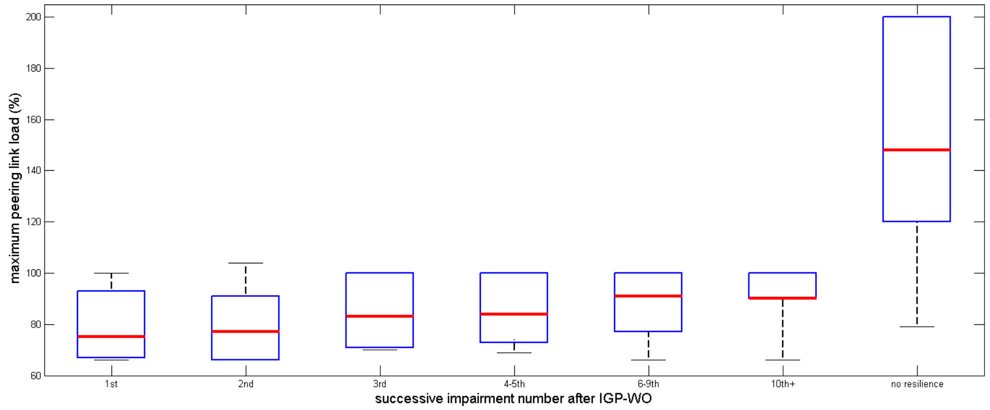

Fig. 10. Maximum peering link load (boxplot statistics).

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. Implementation and Practical Aspects

As any network and service management proposition, the resilient execution policy shall be implementable with a few or no change to existing protocols for the sake of scalability and practical feasibility. About the aspects that merely relate to the building of the ClubMED game, to the computation complexity of PEMP strategies, and to the BGP-PEMP interactions, we address the reader to the corresponding implementation section in [11]. With respect to the practical issues that relate to the resilient execution policy proposed in this article, the following aspects may arise:

- *peering link failure duration assumption*: in the policy requirement description, we assumed that peering link failures persist less than intra-AS ones, with the justification that peering links are typically short-distance links installed in dedicated facilities. This is in fact confirmed by some measurement studies, as [32]. In exceptional cases, namely when the peering agreement involves content providers (e.g., Amazon, Yahoo![3]) that have simplified transport infrastructures, this assumption may need to be relaxed. From a practical perspective, the important aspect is how link failures are modeled in the different steps of the policy, as described in Section III-B. Whether peering link failures show to be persistent rather than transient, and peering link failure prediction information is available, the starting robust equilibrium set may also be computed taking them into account. This implies including estimated G_c component variations for the robust PEMP set computation, while keeping the described transient rerouting immediately upon failure for fast restoration.

- *IGP-WO interval timescale assumption*: we assumed that intra-AS link downtime can be longer than the IGP-WO execution interval. This assumption is indeed dictated by the fact that intra-AS backbone links are typically long-haul links with potentially very long restoration times. In exceptional cases, namely when the peering agreement involves content providers, or low-tier regional providers with non-stringent service level agreements on QoS performance, IGP weight reconfigurations may be performed on very long time scales (e.g., months). With such long latencies, the pertinence of the link failure forecasting model is still more important, in order to obtain very robust routing equilibrium sets. Otherwise, more null PEMP set intersections, hence transient path diversity increase, can be expected during transient rerouting.
- *transit path cost errors*: in Section III-B3, we assumed that it is strategically pertinent to allow the exchange of abstracted path cost errors. Different levels of error visibility can be conceived: bidirectional, per-direction, per-path or per-link, and we assumed that per-direction visibility is appropriate since bidirectional ones would be unrealistic due to asymmetry of Internet routing, and per-path and per-link visibilities would give too much insight in the operational network. In peering scenarios where the level of trust among peers is very high (which is likely impossible among tier-1s, but more possible among non-competing lower tier ASs), per-path errors may be exchanged with possibly little improvement on the overall performance. Such improvements could appear, however, negligible as then “only" the potential sensibility is considered to build the equilibrium set: the different errors finally summed together may produce a negligible net result difference.

Our resilient peering management framework has been primarily designed to meet the requirements of top-tier ASs that (i) regularly perform IP traffic engineering, (ii) need to coordinate the routing across peering links because of routing instabilities and frequent performance loss (as most of top-tier peering todays), yet (iii) are willing to keep high levels of routing independency with no binding agreements. Whether particular conditions characterizing the peering settlement modify the scenario so that (i)-(iii) holds only partially, the practical aspects above described may arise and tackled as suggested.

[3]indeed, large content providers typically establish peering agreements with carrier providers for improving content distribution, so that the carrier is able to better control the traffic load and the content provider is able to offer better quality-of-experience to users [24].

VI. Related Work on Internet Routing Resiliency

Resilience management in Internet routing is commonly adressed at two main scopes: intra-domain and inter-domain routing.

At the intra-domain scope, the network domain is managed by a single autonomous authority, which has thus a large degree of operations for its own network. In this framework, very high levels of resilience can be guaranteed thanks to a variety of network technologies, ranging from best-effort ones at the IP network layer, to connection-oriented ones involving also lower communication layers. A thorough survey of major propositions on this domain can be found in [16].

Because of the ability of guaranteeing high level of reliability within domain boundaries, nowadays, when discussing about Internet routing resilience issues one is immediately faced with the important lack of reliability in inter-domain routing. Indeed, at the borders between networks managed by different autonomous authorities, the spoken language is only IP and BGP at the network layer, and technical collaboration incentives are not straightforward. This is particularly relevant at peering interconnections where there is no customer-provider relationship supporting interconnection reliability requirements.

At the inter-domain (inter-AS) scope, routing resilience is commonly chased by conceiving forms of multi-path selection in interdomain routing protocols. The reason is the need of higher routing visibility for alternative path selection, after failure or quality degradation along the current best path. In the short term, the lack of resilience and path diversity of BGP has lead to overlay service routing architectures to the Internetwork layer via host-based routing, and to multihoming-based inter-domain traffic engineering. However, many measurements campaigns, such as [17], have shown that both overlay routing and multihoming do not allow a significant and long-term improvement of Internet routing reliability. A more far-sighted view suggests to pursue Internet reliability where it can be more natively provided, i.e., in the network layer.

At the network layer, the reliability objective can be basically chased either adopting forms of inter-domain source-routing - see, e.g., [18],[19] and [20], which are able to consider policies and path diversity contraints (more or less explicitely tending toward connection-oriented routing) - or sticking with the more scalable BGP-like path-vector best-effort routing paradigm and defining evolutionary improvements.

Many recent propositions at the state of the art fall in the latter class of path-vector multi-path inter-domain routing protocols. The simplest way would consist in allowing BGP routers to announce not only their chosen best path, but also additional alternative paths. A number of propositions have been discussed in standardization fora on this possibility; a review can be found in [25] and [26] together with a thorough comparison of all possible dissemination strategies of additional paths that remain compliant with the BGP decision process. Surely, more accurate filtering of the announced paths can be performed, e.g., as proposed in [28] considering also confidentiality issues. Another solution is to let two ASs negotiating diverse downstream route and establish them with tunnelling parallel to BGP routing, as done in [29], which dictates however tunnel state management requirements to the Inter-network layer.

As already mentioned, there is no work at the state of the art that specifically characterizes inter-domain (peering) link failures. Such failures may, however, be partially inferred from BGP transient routing failures, whose occurrence can be deterministically identified as explained in [27]. As argued in the introduction, a significant part of inter-domain reliability risk todays lays on top-tier, likely peering, inter-domain links. Despite the world of network operations is almost daily faced with peering link failures, disconnections or congestions, almost no effort has been devoted to resilient peering management in the networking literature. The work in [10] poses the basis of the resilient peering management we describe in this paper.

Finally, little attention has been given to the possibility to implement new resilient routing policies that account for the physical interconnection infrastructure among Internet carriers, which in fact may dispose of available resources that are not visible at the IP routing layer because of policy routing. An interesting recent work in this direction is [30], where the authors argue with results related to an experimental analysis of the current Internet interconnection topology, namely including Internet eXchange Points, that the overall Internet resiliency might be significantly increased by relaxing peering agreements and similar policies to transit or sibling routing in case of major disasters to the infrastructure. Another recent work in this direction is [31], where special mutual transit agreement are formalized as a new model for interconnection policy in between paid-peering and sibling agreements.

VII. 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 for 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 opportunely using the MED attribute of BGP [11]; this 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.

With respect to the state of the art on the subject, our work opens a new and pragmatic unexplored path to improve the Internet routing resilience. Many contributions about Internet routing resiliency improvement concentrate on mechanisms that indiscriminately increase the path diversity at the global routing scope. Our work more pragmatically stresses that the more serious resilience concerns are nowadays concentrated at peering interconnections, where Internet carriers exchange traffic aggregates with an almost absent routing coordination. We proposed a resilient routing policy executed over peering equilibria; our framework models the routing interaction among peering carriers and motivates the rationale adoption of coordinated multipath equilibria with proven resilient behavior after peering networks' equipment failures.

References

[1] S. Secci, H. Ma, B. E. Helvik, J.-L. Rougier, "A resilient routing policy for peering management," in *Proc. IEEE GLOBECOM 2010*.

[2] R. Teixera, K. Marzullo, S. Savage, G. M. Voelker, "In search of path diversity in ISP networks," in *Proc. IMC 2003*.

[3] J. He, J. Rexford, "Towards Internet-wide multipath routing," in *IEEE Netw. Mag.*, vol. 22, no. 2, Mar. 2008.

[4] A. Lange, "Issues in Revising BGP-4," draft-ietf-idr-bgp-issues-05, Aug. 2011.

[5] "Configuring BGP to select multiple BGP paths," JUNOS documentation.

[6] "BGP best path selection algorithm," Cisco documentation.

[7] E. Elena, J.-L. Rougier, and S. Secci, "Characterization of AS-level path deviations and multipath in Internet routing," in *Proc. NGI 2010*.

[8] P. Casas, L. Fillatre, and S. Vaton, "Multi hour robust routing and fast load change detection for traffic engineering," in *Proc. IEEE ICC 2008*.

[9] A. Sridharan and R. Guérin, "Making IGP routing robust to link failures," in *Proc. Netw. 2005*.

[10] S. Secci, J.-L. Rougier, A. Pattavina, F. Patrone, and G. Maier, "ClubMED: coordinated multi-exit discriminator strategies for peering carriers," in *Proc. NGI 2009*.

[11] S. Secci, J.-L. Rougier, A. Pattavina, F. Patrone, and G. Maier, "Peering equilibrium multipath routing: A game theory framework for Internet peering settlements," *IEEE/ACM Trans. Netw.*, vol. 16, no. 2, pp. 419–432, Apr. 2011.

[12] B. E. Helvik, *Dependable Computing Systems and Communication Networks; Design and Evaluation*. Tapir Academic Publisher, 2007.

[13] B. E. Helvik, "Perspectives on the dependability of networks and services," *Telektronikk*, vol. 100, no. 3, 2004.

[14] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. Chuah, Y. Ganjali, and C. Diot, "Characterization of failures in an IP backbone," *IEEE/ACM Trans. Netw.*, vol. 16, no. 4, 2008.

[15] R. B. Myerson, *Game Theory: Analysis of Conflict*. Harvard University Press.

[16] P. Cholda, A. Mykkeltveit, BE. Helvik, OJ. Wittner, and A. Jajszczyk, "A survey of resilience differentiation frameworks in communication networks," *IEEE Commun. Surveys & Tuts.*, vol. 9, no. 4, 2007.

[17] J. Han, D. Watson, and F. Jahanian, "An experimental study of Internet path diversity," *IEEE Trans. Dependable Secure Comput.*, vol. 3, no. 4, 2006.

[18] H. T. Kaur, S. Kalyanaraman, A. Weiss, S. Kanwar, and A. Gandhi, "BANANAS: an evolutionary framework for explicit and multipath routing in the Internet," in *Proc. ACM SIGCOMM Workshop Future Directions Netw. Archit.*, 2003.

[19] K. Argyraki and D. R. Cheriton, "Loose source routing as a mechanism for traffic policies," in *Proc ACM SIGCOMM Workshop Future Directions Netw. Archit.*, 2004.

[20] S. Secci, J.-L. Rougier, and A. Pattavina, "AS-level source routing for multi-provider connection-oriented services," *Comput. Netw.*, vol. 54, no. 14, pp. 2453–2467, Oct. 2010.

[21] J. Lepropre, S. Balon, and G. Leduc, "Totem: a toolbox for traffic engineering methods," in *Proc. INFOCOM 2006*.

[22] S. Uhlig, B. Quoitin, S. Balon, and J. Lepropre, “Providing public intradomain traffic matrices to the research community,” *Comput. Commun. Rev.*, vol. 36, no. 1, 2006.

[23] By courtesy of Y. Zhang, Abilene topology and traffic dataset. Available: http://www.cs.utexas.edu/Ÿyzhang/research/AbileneTM.

[24] G. Goth, “ New Internet economics might not make it to the edge,” *IEEE Internet Comput.*, vol. 14, no. 1, pp. 7–9, Jan./Feb. 2010.

[25] V. Van den Schrieck, P. Francois, and O. Bonaventure, “BGP add-paths: the scaling/performance tradeoffs,” *IEEE J. Sel. Areas Commun.*, vol. 28, no. 8, pp. 1299–1307, Oct. 2010.

[26] J. Uttaro, V. Van den Schrieck, P. Francois, R. Fragassi, A. Simpson, and P. Mohapatra, “Best practices for advertisement of multiple paths in BGP,” draft-ietf-idr-add-paths-guidelines-01, May 2011.

[27] F. Wang, J. Qiu, L. Gao, and J. Wang, “On understanding transient interdomain routing failures,” *IEEE/ACM Trans. Netw.*, vol. 17, no. 3, 2009.

[28] F. Wang and L. Gao, “Path diversity aware interdomain routing,” in *Proc. IEEE INFOCOM 2009*.

[29] W. Xu and J. Rexford, “MIRO: multi-path interdomain routing,” in *Proc. SIGCOMM 2006*.

[30] C. Hu, K. Chen, Y. Chen, and B. Liu, “Evaluating potential routing diversity for internet failure recovery,” in *Proc. INFOCOM 2010* (Mini-conference).

[31] Y. Liao, L. Gao, R. A. Guérin, and Z.-L. Zhang, “Safe inter-domain routing under diverse commercial agreements,” *IEEE/ACM Trans. Netw.*, vol. 18, no. 6, pp. 1829–1840, Dec. 2010.

[32] O. Bonaventure, C. Filsfils, and P. Francois, “Achieving sub-50 milliseconds recovery upon BGP peering link failures,” *IEEE/ACM Trans. Netw.*, vol. 15, no. 5, Oct. 2007.

Stefano Secci (S’05-M’10) received the M.Sc. and Ph.D. degrees in communications engineering from Politecnico di Milano, Milan, Italy, in 2005 and 2009, and a Ph.D. degree in computer science and networks from Telecom ParisTech, Paris, France, in 2009.

He is an Associate Professor at the LIP6, Pierre and Marie Curie University (UPMC - Paris 6 - Sorbonne Universités), Paris, France. He worked as Post-Doctoral Research Associate with Telecom ParisTech, and NTNU, Norway, in 2009-2010, and was visiting researcher at George Mason University, USA, in 2010. Before receiving the Ph.D. degrees, he worked as a Network Engineer with Fastweb Italia, Milan, Italy, and as a Research Associate with Ecole Polytechnique de Montréal, Montréal, QC, Canada, and with Politecnico di Milano. He actively participated in the ANR ACTRICE, FP7 EuroNF, FP7 ETICS projects, and CELTIC TIGER2 projects, and a US ONR project on future Internet. He is the author of more than 20 articles in major international journals and conferences. His works space from optical network planning and switching to IP routing optimization and traffic engineering. His current research interests are about future Internet routing resiliency, mobility, and policy.

Dr. Secci is the recipient of the NGI 2009 Best Paper Award. He has been member of the Technical Program Committee of IEEE ICC 2011-12, IEEE GLOBECOM 2010-11, NGI 2010-11, and CFIP 2009, TPC co-chair of NoF 2011, and has been reviewer on a regular basis for major international journals and conferences in networking.

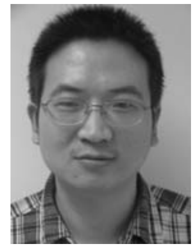

Huaiyuan Ma works at EMGS Norway as a senior software engineer. He had been a PhD student in Q2S Research Centre, NTNU, Norway, from Oct. 2007 to Dec. 2010. His major research topic is Internet traffic and Internet routing scalability. From Oct. 2005 to Aug. 2007, he worked in Huawei Technologies Co., Ltd as a system engineer. He was involved in the development of carrier-level routers, as development responsible of many modules. From Aug. 2003 to Aug. 2005, he studied Dependable Computer Systems in Chalmers University of Technology, Gothenburg, Sweden, receiving his MSc degree. From Aug. 1995 to June 2003 he worked in the 13th Research Institute of Ministry of Electronics Industry of China as a software programmer, test engineer and quality control engineer.

Bjarne Emil Helvik (1952) received his Siv.ing. degree (MSc in technology) from the Norwegian Institute of Technology (NTH), Trondheim, Norway in 1975. He was awarded the degree Dr. Techn. from NTH in 1982. He has since 1997 been Professor at the Norwegian University of Science and Technology (NTNU), the Department of Telematics. He is principal academic at the Norwegian Centre of Excellence (CoE) for Quantifiable Quality of service in Communication systems (Q2S). Since August 2009, he has been Vice Dean with responsibility for research at the Faculty of Information Technology, Mathematics and Electrical Engineering at NTNU. He has previously held various positions at ELAB and SINTEF Telecom and Informatics. In the period 1988-1997 he was appointed as Adjunct Professor at the Department of Computer Engineering and Telematics at NTH.

His field of interests includes QoS, dependability modelling, measurements, analysis and simulation, fault-tolerant computing systems and survivable networks as well as related communication system architectural issues. His current research focus is on distributed, autonomous and adaptive fault-management in telecommunication systems, networks and services. He has authored/co-authored a number of research papers

Professor Helvik has been active in RACE, ACTS, IST and EURESCOM collaborations, and in a number of program committees. Among them committees for the International Symposium on Fault-Tolerant Computing (FTCS), the European Dependable Computing Conferences (EDCC), the ITC Specialist Seminars, the International Teletraffic Congress (ITC), GlobeCom, the International Conference on Dependable Systems and Networks (DSN), International Telecommunications Network Strategy and Planning Symposium (NETWORKS), International Conference on Dependability (DEPEND), Design of Reliable Communication Networks (DRCN), Reliable Networks Design and Modeling (RNDM) and has co-chaired the Euro-NGI 2007 conference on Next Generation Internet Networks.

He has developed courses and is teaching dependability at NTNU on PhD and MSc level. In connection with the latter he has (co-)authored two textbooks/lecture notes on the topic.

Jean-Louis Rougier received his engineering diploma in 1996 and his Ph.D. in 1999 from Télécom ParisTech (formerly called Ecole Nationale Supérieure des Télécommunications). He joined the computer science and networks department of Télécom ParisTech in 2000 as an associate professor. His research interests are performance evaluation and algorithms, traffic engineering, and routing in networks.