Peering Games for Critical Internet Flows

Stefano Secci[a,b], Jean-Louis Rougier[a], Achille Pattavina[b], Fioravante Patrone[c], Guido Maier[b]
[a]Institut Telecom, Télécom ParisTech, LTCI CNRS, France. E-mail: {secci, rougier}@telecom-paristech.fr
[b]Politecnico di Milano, Italy. E-mail: {secci, pattavina, maier}@elet.polimi.it
[c]Università di Genova, Italy. E-mail: patrone@diptem.unige.it

I. CONTEXT[1]

We propose to model the routing of critical Internet flows across peering links with the theory of non-cooperative games. The goal is to pro-actively tackle major issues in current Internet routing: (i) lack of coordination between IGP and BGP routing (ii) (whose symptom are) frequent BGP route deviations (iii) (which can cause) sudden congestions at inter-AS links. Since peering settlements furthermore suffer from a low operation coordination (peering links less straightforwardly upgraded than transit links, BGP multi-exit discriminator, MED, signaling rarely used there), we believe that the peering settlements represent a practical implementation scope for the proposed framework.

II. THE CLUBMED PEERING GAME

The ClubMED (Coordinated MED) routing framework is characterized in detail in [1]. Within it the MED signaling between peering ASs is modeled as a non-cooperative peering game that can allows the peers to coordinate towards rational, efficient and stable multipath routing solutions.

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 game-theoretic modeling of the load sharing problem. Each peer is represented as a rational player that can take benefit by routing accordingly to a cost game built upon routing and congestion costs. The principle is to take the peering routing decision following efficient equilibrium strategy profiles of the game - in its one-shot form or repeated form - 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 across peering links 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 or have particular QoS or reliability requirements, or have similar characteristics.

Each ClubMED destination cone is reachable behind a single AS Border Router (ASBR) not at the peering border (called "ClubMED node"), and each peering AS can manage several destination cones. The inter-cone flows are supposed to be equivalent, for instance w.r.t. their bandwidth, so that their path cost can be fairly compared and their routing coordinated; the equivalence condition applies grouping all the inter-cone flows at each side.

[1]The research presented in this paper has been carried out within the INCAS S.JRA.7 activity of the FP7 Euro-NF Network of Excellence (Network of the Future - From Theory to Design). [a] was also partially funded by the I-GATE project of the Institut Telecom (Networks of the Future Lab), France.

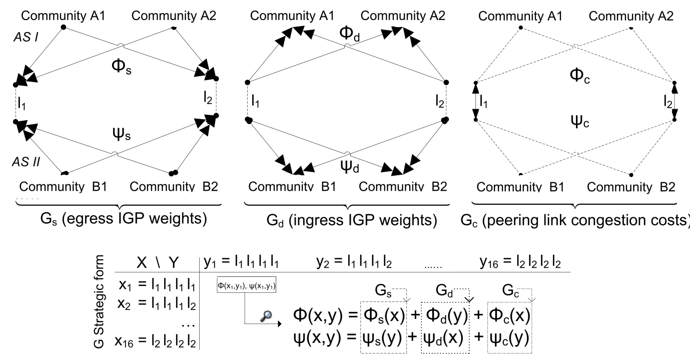

Fig. 1. Multi-pair 2-link ClubMED game composition example.

Practically, a destination cone can be identified by a BGP 'community id' tag in order to give to the BGP decision process the means to identify the scope of application of the ClubMED game. The game is to be built only at the ClubMED nodes connecting the destination cones; its 'solution' relies on a coordinated peering equilibrium policy indicating an 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. While the IGP path costs can be coded with little primitive extensions via a composite MED attribute in BGP announcements, in order to build the congestion game the bitrate of each inter-cone flow should be known by each ClubMED node (e.g., via Netflow).

Two examples of ClubMED game are given in Fig. 2 (with G_s and G_d only). Mathematically, it is a particular potential game, in which the Nash equilibria correspond to the minima of a potential function, and viceversa. It is possible to have a single equilibrium (as in the top of Fig. 2), or many (as in the bottom), depending on the IGP path cost settings. Moreover, the Nash equilibria may not be Pareto-efficient (as in the bottom).

Furthermore, the game setting is expected to change when the IGP weights are reconfigured (after IGP Weight Optimization, IGP-WO, operations). Their possible variations are to be precomputed and taken into account so as to select robust equilibria, which actually corresponds to select equilibria with a potential value below a specifically computed potential threshold, hence generating larger sets of Nash equilibria and Pareto-superior profiles. Finally, the usage of G_c could also allow reacting fastly to sudden peering link failures, independently to IGP-WO operations.

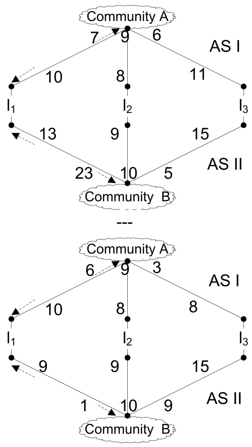

I\II	l_1	l_2	l_3
l_1	$(17,36)^6$	$(19,32)^2$	$(16,38)^8$
l_2	$(15,23)^4$	$\mathbf{(17,19)}^0$	$(14,25)^6$
l_3	$(18,18)^7$	$(20,14)^3$	$(17,20)^9$

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	$\boldsymbol{(14,18)}^0$	$\mathbf{(17,18)}^0$	$(11,24)^6$

Fig. 2. 3-link examples.

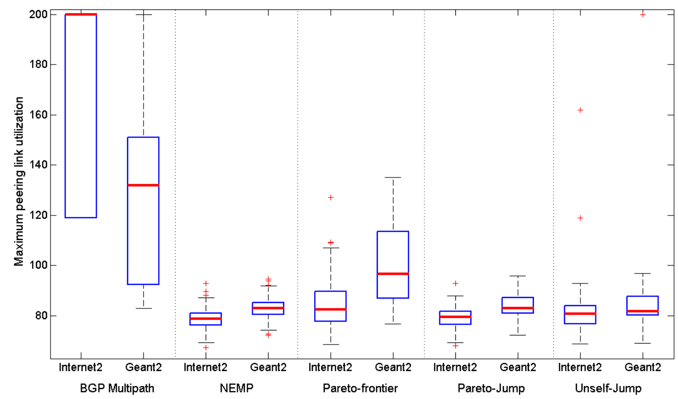

Fig. 3. Maximum peering link utilization boxplot statistics.

III. Peering Equilibrium MultiPath (PEMP)

There is thus a coordination problem to fine-select efficient and robust equilibria over broad sets. In the following, we present the Peering Equilibrium MultiPath (PEMP) routing policies that can be implemented upon the ClubMED game; they are characterized in detail in [2].

1) Nash Equilibrium MultiPath (NEMP): Assuming that ClubMED remains a fully non-cooperative framework, its implicit solution strategy to which to coordinate without any signalling message is: *play the equilibria of the Nash set, only the Pareto-superior ones if any*. In the bottom of Fig. 2, e.g., AS I may balance the load on l_2 and l_3, being aware that AS II may balance its load on l_1 and l_2.

2) Pareto-frontier: Given that the game Pareto-frontier may not contain equilibria, in a repeated ClubMED context, an explicit coordination strategy is: *play the profiles of the Pareto-frontier*, with the threat of a selfish choice otherwise. The ClubMED game would be repeated an indefinite number of times, indeed. From “folk-theorem”-like results [3], this strategy is an equilibrium of the repeated game and grants a maximum gain for the players in the long-run. Nevertheless, the unilateral trust for such a strategy could decrease whether in a short period of analysis the gains reveal to be unbalanced and in favor of a single peer. The reciprocal trust among peers can thus affect the reliability of such a Pareto coordination.

3) Unselfish-Jump: Another strategy is conceivable to guarantee balancedness in gains in the short term, and thus helping to keep a high level of reciprocal trust. 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: Instead, if the jump is constrained toward a Pareto-superior profile only (not necessarily in the Pareto-frontier), one can avoid unselfish sacrifices.

E.g., in the bottom 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 last two policies are not binding: it would be enough to associate the policy 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.

IV. Performance Evaluation

We evaluated the performance of the PEMP routing policies with realistic simulations, comparing them to each other and to BGP Multipath (the concurrent implemented alternative). We created a virtual interconnection scenario among the Geant2 and the Internet2 ASs, and built the game over many successive traffic matrix configurations.

To report some results, Fig. 3 reports the Boxplot statistics maximum link utilization as seen by each peer, with all the methods. All the PEMP strategies but the Pareto-frontier one never caused congestion on peering links (utilization above 100%). The enabling of the Multipath mode in BGP does not have a significant effect on the peering link congestion. With ClubMED, instead, the multipath routing choice is carefully guided toward efficient solutions. The NEMP, Pareto-Jump and Unselfish-Jump policies show the median, the upper and lower quartiles always above 85%, remembering that with full BGP Multipath one would have a $200/300 = 66,7\%$ utilization. The Pareto-frontier strategy does not guarantee, however, a congestion-free solution, with a median close to 100% utilization. The reason for this behavior are still the highly asymmetric cost profiles introduced by the Pareto-superiority condition in the solution.

V. Further work

We are currently working on a possible generalization of the game-theoretic model to several IGP scenarios (robust IGP routing, absence of IGP-WO, reaction against intra-AS link failures, etc). Moreover, the game theoretical modeling could be adapted to other contexts being studied for the future Internet such as, e.g., Locator/Identifier Separation Protocols.

References

[1] S. Secci et al., “ClubMED: Coordinated Multi-Exit Discriminator Strategies for Peering Carriers”, in *Proc. of NGI 2009*. Best Paper Award.
[2] S. Secci et al., “PEMP: Peering Equilibrium MultiPath routing”, in *Proc. of GLOBECOM 2009*.
[3] R.B. Myerson, *Game Theory: Analysis of Conflict*, Harvard Univ. Press.