

On the Interactions between Non-Cooperative P2P Overlay and Traffic Engineering Behaviors

Chaojiong Wang, Ning Wang, Michael Howarth
University of Surrey
Guildford, United Kingdom
{C.Wang, N.Wang, M.Howarth}@surrey.ac.uk

George Pavlou
University College London
London, United Kingdom
G.Pavlou@ee.ucl.ac.uk

Abstract—Emerging Peer-to-Peer (P2P) technologies have enabled various types of content to be efficiently distributed over the Internet. Most P2P systems adopt selfish peer selection schemes in the application layer that in some sense optimize the user quality of experience. On the network side, traffic engineering (TE) is deployed by ISPs in order to achieve overall efficient network resource utilization. These TE operations are typically performed without distinguishing between P2P flows and other types of traffic. Due to inconsistent or even conflicting objectives from the perspectives of P2P overlay and network-level TE, the interactions between the two and their impact on the performance for each is likely to be non-optimal, and also has not yet been investigated in detail. In this paper we study such non-cooperative interactions by modeling best-reply dynamics, in which the P2P overlay and network-level TE optimize their own strategies based on the decision of the other player in the previous round. According to our simulations results based on data from the ABILENE network, P2P overlays exhibit strong resilience to adverse TE operations in maintaining end-to-end performance at the application layer. In addition, we show that network-level TE may suffer from performance deterioration caused by greedy peer (re-)selection behavior in reacting to previous TE adjustments.

I. INTRODUCTION

Today, P2P flows account for some 50%-70% of the overall Internet traffic, according to recent traffic measurements [1, 2]. Under such circumstances, network capacity for other types of services, such as conventional web-based applications, may be impacted due to resource competition with this large volume of P2P traffic. In the literature, traffic engineering (TE) techniques have been proposed for ISPs to optimize customer traffic in order to improve overall network performance, for example by performing load balancing and network cost reduction. It should be noted that, in general, TE solutions do not distinguish between P2P flows and conventional Internet traffic, which means that traffic optimization is performed in an aggregate fashion, regardless of specific types of flows. In P2P overlay networks, the current implementation of peer selection paradigms are often based on application-layer optimization for enhancing the quality of experience by end users. For instance, real-time multimedia P2P systems usually select partner peers that are associated with low delay in order to achieve fast playback at the user side. On the other hand, the objective of TE is to improve the overall performance at

the network side, instead of focusing on individual users. As such, there is an obvious misalignment between the TE objectives and the selfish P2P peer selection in the application layer. As for the two autonomous entities – the P2P overlay and the network-level TE – the decisions that are made by each one of them may influence the performance of the other. Such interactions may adversely impact the relevant performance on both sides due to conflicting behavior. For instance, TE may adjust the underlying routing decisions in order to re-optimize network performance, but such a change may also shift some P2P traffic to alternative paths with sub-optimal user-perceived QoS performance (e.g. higher end-to-end delay due to longer paths). As a result, the P2P overlay may react to such dynamics by re-selecting partners in each P2P session in order to regain the original performance at the application layer. Such a behavior will once again change the overall traffic condition so that the underlying TE mechanisms need to react accordingly once again. This adjustment of network configurations may further trigger re-selection of peers in the P2P overlay. In this paper we investigate the interaction between selfish peer selections and optimized routing configurations in non-cooperative environments.

In the literature, a number of papers [3, 4, 5, 6] have investigated the interaction between TE and overlay network operations. We can classify these works into two categories. The first category focuses on the interactions between network-layer routing configurations decided by TE and logical overlay routing on top [3, 4]. In this scenario, TE and the overlay respectively adjust their own *routing* strategies in turn, based on each other's decisions. Compared with this type of interaction, the key difference from our work is as follows: we focus on the P2P overlay side which only considers how to select the best partner peers (i.e. the other endpoint of individual P2P connection sessions), rather than changing the routing configuration in the overlay. The other category [5, 6] focuses on CDN (Content Distribution Network) –like paradigms, and considers the interaction between network-layer routing decisions made by TE and application-layer content server selections. Our work differs from this category in the following three features. Firstly, in P2P overlay networks, peers, as both content producers and consumers, have highly dynamic join/departure patterns, while in the CDNs of [5, 6] content servers are statically provisioned in the network for providing content delivery services. Secondly, we consider *symmetric* content exchange patterns: in P2P overlays a peer

not only requests data from, but also provides content to other peers; this differs from the previous studies in which a specific set of clients only download data from a number of dedicated content servers. Finally, in P2P overlays each peer needs to simultaneously fetch chunks of content from a set of partners, while in conventional CDNs a client typically requests content from one specific server at a time.

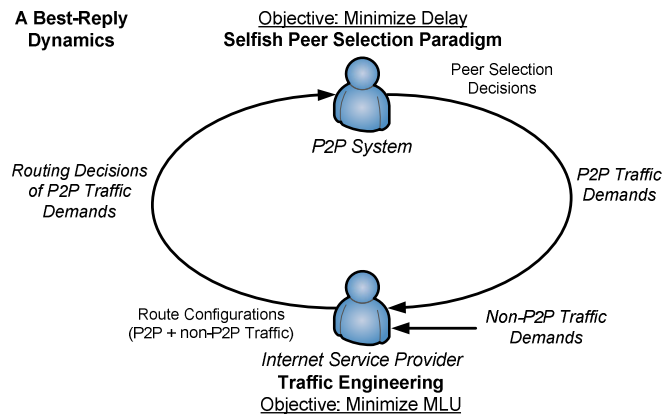


Figure 1: Dynamic interaction between TE and P2P overlay

In this paper we model TE and P2P overlay as two rational players respectively who play the *best-reply* dynamics [4]: one player chooses the best response based on the other player’s decisions in the previous round. As shown in Figure 1, TE aims to optimize the overall network performance (e.g. load balancing) through adjusting routing decisions for customer traffic (including both P2P and non-P2P background flows) in the network layer. The outcome of path selection changes by TE for the P2P traffic is then taken as input by the P2P overlay to re-select partner peers in order to regain the original application-level performance (e.g. minimize end-to-end delay between individual peering partners) in case of adverse impacts from TE. Such peer reselection further influences the overall traffic distribution in the network, requiring further TE operations, and so on. Note that TE and the P2P overlay may optimize their own objectives at different timescales, but here we follow the assumption made in [5] that each player runs its optimization until the other has finished its operations in each round. With such interactions, both TE and P2P overlay adjust their own decisions in turn according to each other’s previous behavioral changes. Based on this model, we investigate how well TE and P2P overlay may react to each other in such a non-cooperative environment. As far as TE is concerned, today’s solutions can be classified into IP-based TE and Multi-Protocol Label Switching (MPLS-based) TE. The latter is more flexible than the former, in the sense that it allows explicit routing and arbitrary splitting of traffic across multiple active label switched paths (LSPs), even at short time-scales. In this paper we specifically focus on the interaction between MPLS-based TE and P2P overlays, since MPLS-based TE is ideal for online traffic adjustment that is agile to short time-scale traffic dynamics. In our study, we aim to answer the following questions:

1) *Is there an equilibrium point in this interaction between TE and P2P overlay behaviors?*

2) *If the answer is yes, then does this equilibrium point converge to an improved point?*

3) *What is the potential impact on the performance of the P2P overlay and the network performance under such interactions?*

We believe that a good understanding of such interactions will offer significant insight into the future development of intelligent Internet P2P traffic management paradigms in dynamic environments.

II. TRAFFIC ENGINEERING & PEER SELECTION

In this section we first describe the modeling of the interactions between MPLS-based TE and selfish peer selection in P2P overlays. After this, a *best-reply* dynamics model between the two players is presented for analyzing behavioral interactions between them.

Let’s first consider a physical Point-of-Presence (PoP) network topology that is modeled as a unidirectional graph $G = (N, A)$, where N is a set of PoP nodes and A is the set of inter-PoP links. Each physical link $a \in A$ has a bandwidth capacity C_a . The tuple $\langle i, j \rangle$ is defined as a PoP node pair where $i, j \in N$ refer to a source and a destination PoP node respectively. In our model, each peer is associated with one of the PoP nodes in the PoP-level network topology. The routing of both P2P traffic and conventional background traffic is determined by TE, and the two traffic types are not distinguished. Let P_{ij} represent a set of explicit LSPs between PoP nodes i and j , with each LSP consisting of one or multiple inter-PoP links. In a common practice of ISP network design, bandwidth resources within a single PoP are usually highly over-provisioned, so we only focus on bandwidth resources on inter-PoP links in A . This means if multiple peering neighbors are clustered within the same PoP, then the associated bandwidth consumed by their local peering connections is ignored. According to common MPLS-based TE approaches, multiple LSPs are established between each PoP node pair in order to allow adaptive splitting of the overall traffic demand across them for achieving dynamic load balancing. Let t_{ij}^{p2p} and t_{ij}^{np2p} denote respectively the overall P2P traffic demand and the overall non-P2P background traffic demand from PoP node i to j . Let t_{ij} be the overall traffic demand $t_{ij} = t_{ij}^{p2p} + t_{ij}^{np2p}$, and x_{ij}^p ($0 \leq x_{ij}^p \leq 1$) be the traffic splitting ratio on each specific LSP $p \in P_{ij}$.

A. Traffic Engineering

TE operations are normally applied by ISPs in order to optimize the overall network performance, such as load balancing and network cost reduction. In our model, we consider the objective of TE to be minimizing the maximum link utilization (MLU), since this has been widely used as a TE performance metric in the literature. Once again, we emphasize that TE aims to optimize the *overall* network performance rather than any specific type of traffic, and in our model, TE does not differentiate between P2P traffic and non-P2P background traffic.

As previously mentioned, m ($m > 1$) disjoint LSPs are pre-established between each PoP node pair in our analysis. We introduce a binary mapping coefficient y_p^a to indicate the

relationship between LSP p and physical link a : Y_p^a equals 1 if physical link a is on physical LSP $p \in P_{ij}$, and 0 otherwise. The overall traffic load $T_a = \sum_{i \in N} \sum_{j \in N \setminus \{i\}} \sum_{p \in P_{ij}} x_{ij}^p * t_{ij} * Y_p^a$, ($\sum_{p \in P_{ij}} x_{ij}^p = 1$) on the physical link $a \in A$ is the sum of all demands of flow over this link, including both P2P traffic and non-P2P traffic. With a demand matrix $(t_{ij}, \forall i, j \in N)$, the objective of TE is to compute an optimized value x_{ij}^p (splitting ratio) across LSPs between each source-destination PoP node pairs in order to minimize the overall MLU across the network, i.e.:

$$\min_{\forall a \in A} \max \left(\frac{T_a}{C_a} \right) \quad (1)$$

$$\text{s.t.} \quad \sum_{a:d(a)=z} x_{ij}^p * t_{ij} * Y_p^a - \sum_{a:s(a)=z} x_{ij}^p * t_{ij} * Y_p^a = \begin{cases} t_{ij}, & z = j \\ -t_{ij}, & z = i \\ 0, & \text{otherwise} \end{cases}$$

where $\forall z, i, j \in N$, where $s(a)$ and $d(a)$ are source node and destination node of link a respectively.

An adaptive MPLS-based TE is considered, which can be operated at a short variable time-scale (e.g. at the scale of minutes). To allow for fast reactions to changing traffic conditions (e.g. due to frequent peer reselections in P2P overlay), we consider an online heuristic solutions without involving global optimizations each time (i.e. only local adjustments are applied). Similar to [6], we consider the generic TE strategy as follows: a set of physical links with high link load is first identified and TE tries to iteratively shift some traffic demand away from them in order to better balance the overall traffic distribution. More specifically, TE identifies the top k inter-PoP links with the highest utilization across the entire network, and then some traffic currently traversing those links needs to be shifted away from them. To achieve this, TE needs to re-configure the traffic splitting ratio for m LSPs between some node pairs $\langle i, j \rangle$ that involve those highly-loaded links. The shifting action is chosen so that it does not introduce any new hot spots with higher link load than the original loads before the adjustment. This traffic shifting through re-optimizing splitting ratios at individual source PoPs can be recursively performed until no further improvement is achieved.

B. Selfish Peer Selection in the P2P Overlay

The P2P overlay aims to optimize the performance experienced by end users, for example by reducing the end-to-end delay between individual peers. This is typically done in the application layer by localized partner peer selection

In the modeling of the P2P overlay behavior, each peer is associated with one of the PoP nodes in the network topology. We consider multiple simultaneous P2P sessions running over the network, with each session containing a distinct set of active peers sharing the same content. If one end-user participates in multiple sessions, it is treated as an independent peer in each of them. More specifically, let V denote a set of active peers physically attached to network G . Each client peer in P2P session t needs to connect to a set of partner peers from all available peers in the session (denoted by $V(t)$), and

download content from them at certain transmission rates. In this case the actual partner set for a specific client peer u (denoted by $V_u(t)$) is effectively a subset of all the available peers in session t , i.e., $V_u(t) \subseteq V(t)$. In addition, let D_a be the delay of physical link $a \in A$, and let the delay between a PoP node pair be the sum of the delay associated with each link constituting the LSP between that node pair. While intra-PoP delay is ignored, the formal objective in peer selection is to minimize the delay between each single client peer u and each of its selected partner peers i.e.:

$$\min_{\langle u, v \rangle, v \in V_u(t)} \left(\sum_{a \in A} D_a * Y_{u,v}^a \right) \quad (2)$$

where $Y_{u,v}^a$ is a binary mapping coefficient, equal to 1 if link a constitutes the LSP from PoP node u to v , and is equal to 0 otherwise.

With such a selfish peer selection paradigm with latency-localization, the P2P overlay may dynamically readjust the partnership connections for every client peer based on own measurements. As we mentioned previously, both non-P2P traffic and P2P traffic are shifted by TE without distinction so as to improve the overall network performance. For the P2P traffic, the end-to-end delay from some current partner peers to the client peer may become higher if the traffic is shifted from a shorter delay LSP to a longer one. In this case, in the following round, the P2P overlay has the opportunity to re-select some new partners elsewhere with lower delay. Such a change at the P2P side then changes the traffic condition input for TE, which may in turn makes further adjustment actions.

C. TE and P2P Overlay Interaction Analysis

We consider such interaction as *best-reply* dynamics where each of the two rational players decides its own best strategy in response to the change of behavior of the other player in the previous round. The MPLS-based TE and P2P overlay take turns to optimize their own objectives respectively in this interaction.

In our analysis, the strategy space that is applied by MPLS-based TE can be described as a set of feasible traffic splitting ratios across m distinct LSPs $\{x_{ij}^{P_1}, \dots, x_{ij}^{P_m}\}$ between each PoP node pair $\langle i, j \rangle$. This can be expressed as:

$$S_{TE} = \langle \dots, \{x_{ij}^{P_1}, \dots, x_{ij}^{P_m}\}, \dots \rangle \quad (3)$$

Traffic splitting in MPLS-TE is performed on a per-flow basis instead of per-packet [7], and therefore the P2P flow between each peering partner always follows one single path at any time.

On the other hand, the strategy space of P2P overlay is a set of partner peers $V_u(t)$ of every single client peer u that are selected from all available peers $V(t)$ in each session t . By selecting the best partner peers, the end-to-end delay among peers can be maintained as a minimum for each client peer in the session.

$$S_{P2P} = \langle \dots, V_u(t), \dots \rangle \quad (4)$$

Based on the above specifications, we now consider the case where TE optimizes the path selection decisions for both non-P2P traffic and P2P traffic without distinction. Since the

P2P traffic paths are changed by the TE operation, the delay performance from partner peer to client peer may change. As described above, if some P2P traffic is shifted from one LSP to another longer path, the corresponding peers may experience higher end-to-end delay after such a change. In order to attempt to maintain the original quality of service, the P2P overlay may then reactively select alternative partner peers within individual sessions, based for example on measured delay, in response to the changed MPLS paths by TE. Due to such reshuffling in peer connections in the P2P overlay following the previous TE operation, the overall traffic condition within the network changes again. It should be noted that, given the fact that P2P traffic dominates the overall Internet traffic today, behavioral changes in the P2P overlay may have a significant impact on the overall network traffic patterns. As such, MPLS-based TE may need to further readjust traffic splitting ratios in order to maintain its own objective. Generally, MPLS-based TE and the P2P overlay take turns to optimize their own operational performance, adjusting to the previous decision of the other player. This can be modeled as:

$$\begin{aligned} \{x_{ij}^{P_1}, \dots, x_{ij}^{P_m}\}_{(k+1)} &= \arg \min_{(k)} TE(V_u(t)) \\ V_u(t) &= \arg \min_{(k+1)} P2P(\{x_{ij}^{P_1}, \dots, x_{ij}^{P_m}\}_{(k)}) \end{aligned} \quad (5)$$

Using our model of TE and P2P overlay behavior, we investigate whether best-reply dynamics converge to an equilibrium point in this interaction between non-cooperative behaviors. If there is a converging equilibrium point, we may further investigate whether it is a globally optimal point such as a Pareto point [5], or a “biased” one in favor of either side. In addition to convergence, we also study the impact on the relevant performance of the P2P overlay and TE respectively. Through this analysis, a clear picture can be drawn on whether P2P overlay can synergistically interact with TE in current network environments, and useful guidance can be further derived on how to mutually improve the performance of both P2P overlay and TE paradigms in order to potentially achieve a ‘win-win’ situation between the two rational players.

III. PERFORMANCE ANALYSIS

In this Section we consider the performance of the TE and P2P overlay interactions. We first describe our simulation environment, including the network topology and the setup of the P2P overlay. The simulation scenarios and parameters are then presented. Following this, we present results for the performance metrics associated with the TE and P2P overlay.

A. Simulation setup

We use the real ABILENE network topology [8] at the Point-of-presence (PoP) level. The ABILENE network consists of 11 nodes and 28 unidirectional inter-PoP links. Each link has its actual link capacity and IGP link weight configuration. According to [9], in the ABILENE network the IGP link weight setting is based on end-to-end latency, and hence customer traffic is effectively routed on the lowest delay paths. In addition, the ABILENE traffic traces that are

measured through NetFlow are used (with scaling) as background traffic in our simulation.

The P2P traffic used in our experiments is synthetically generated based on the measured pattern of today’s popular real-time multimedia based P2P applications [10]. We consider 20 concurrent P2P channel sessions, with each channel attracting up to 1000 peers. Hence altogether we consider up to 20,000 peers that are distributed across the 11 PoP nodes in the ABILENE network. The overall distribution of these peers in each PoP node is determined according to the actual population of each city (PoP), where larger PoP nodes have more peers assigned. The channel session selected by each peer is randomly determined. Without loss of generality, there are both popular channels and unpopular channels on the P2P overlay side. In addition, we follow the observation that each client peer has on average 40 peering connections in order to satisfy the overall downloading rate requirement for stable playback (1Mbps, [10]). For each requesting peer, there is one top peer partner that provides on average three times as much content as the other (auxiliary) ones.

B. Simulation Scenarios

To make the analysis more comprehensive, we use three scenarios to analyze the interaction between MPLS-based TE and P2P overlay. Following [4, 5], we set the overall P2P traffic demand as a *low*, *medium* or *high* proportion of the overall network traffic volume, i.e. the P2P traffic accounts for 40% (*low*), 60% (*medium*), or 80% (*high*) of overall network traffic. Such configurations are reasonable, as it has been observed that the actual proportion of P2P traffic in the Internet varies significantly and it may peak at 80%.

C. Performance Analysis

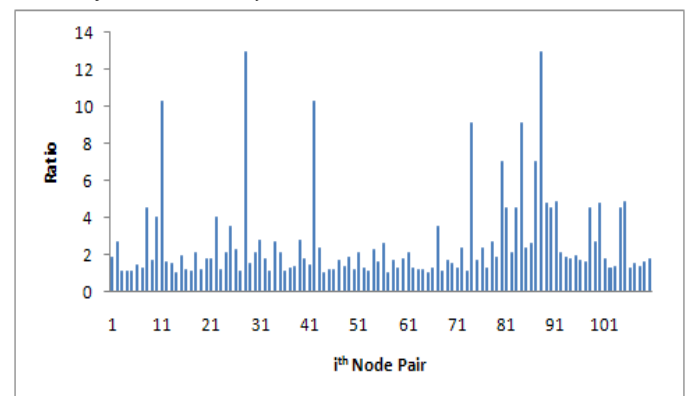


Figure 2: Ratio of delay of the longest path to the shortest path

As far as traffic splitting in MPLS-based TE is concerned, we first show in Figure 2 the end-to-end delay ratio across parallel LSPs, i.e. the ratio of the longest delay LSP to the shortest delay LSP between each pair of PoP nodes in the ABILENE network. This ratio effectively indicates the actual change of delay experienced by the peers whose flows are shifted by TE traffic splitting adjustments from the longest delay LSP to the shortest delay LSP. From Figure 2 we can see that the maximum delay ratio is 13:1, minimum ratio is 1:1, and the average is around 2.6:1. Such difference between the paths may easily result in selfish partner reselections by some

affected peers whose flows are shifted from the shorter path to the longer one.

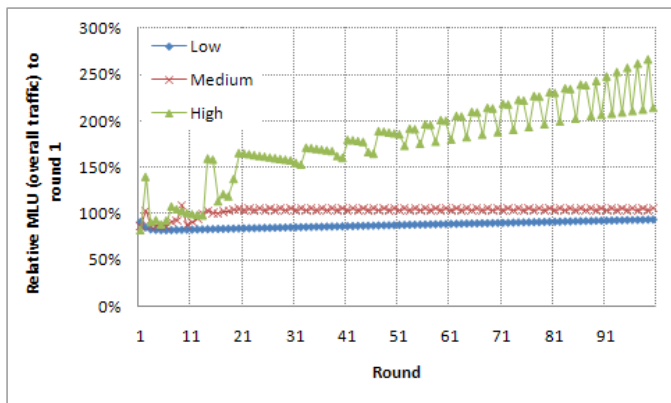


Figure 3: Relative change of MLU for overall traffic

Figure 3 indicates the pattern of the overall MLU performance change relative to the initial state (round 1) on per round basis (we consider 100 rounds in our experiments), i.e.

$$\frac{MLU(t)}{MLU(1)}, (1 < t \leq 100)$$

We can clearly see that different proportions of P2P traffic have yielded distinct performance curves. Specifically, the low scenario converges to an equilibrium point which has 8% decrease compared with the MLU value of in the initial state (100%). Similarly, in the medium scenario we can also observe the convergence towards an equilibrium point, but interestingly, the final converged performance has a 5% increase compared with the initial state. Based on the above results, we can see that even if a specific equilibrium exists, it is not always the case that the overall TE performance will converge to an *improved* performance. The reason is that, after the adjustment of TE, the P2P overlay may selfishly reselect new partner peers which may lead to significantly worse network performance compared with the situation before the TE operation, and the next round of TE operation might not be able to achieve further better performance than the previous round. By investigating the high scenario in Figure 3, we can clearly see some oscillation patterns on the MLU performance as the number of rounds increases, and more importantly there is no specific equilibrium and indeed the MLU performance becomes worse with time. As we mentioned above, P2P overlay selfishly re-selects the peering partners if the original ones experience higher delay following the adjustment of TE. It can be inferred that such selfish peer reselection behavior may have some significant (negative) impact on the TE performance in the non-cooperative environment, especially when P2P flows dominate network traffic.

In addition to the overall MLU performance, we also show how the interactions between TE and P2P overlay impact the performance of background non-P2P traffic. Figure 4 indicates the change of non-P2P traffic utilization compared to the initial state. We can see that in both low and medium scenarios, the background traffic condition is not significantly impacted by the interaction. However, if P2P flows substantially dominate the overall traffic (high scenario), the utilization of non-P2P

background traffic may become less stable, with some oscillation observed (Figure 4). The reason is that a large number of peer reselections performed by the P2P overlay cause TE to unilaterally adjust the traffic splitting ratio across multiple LSPs. Since TE optimizes P2P and non-P2P background traffic without distinction, the non-P2P traffic is impacted by the TE optimization in response to P2P reselection behavior.

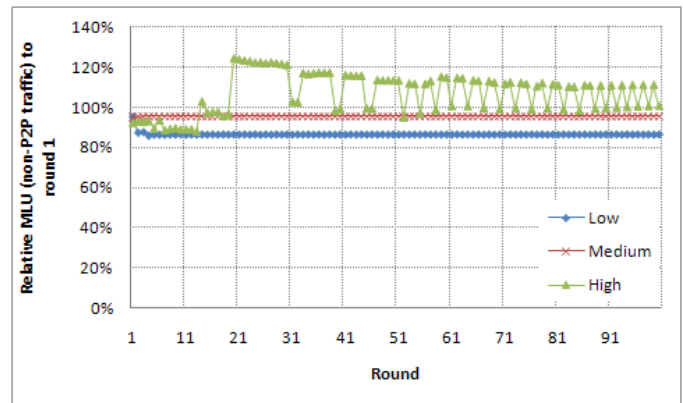


Figure 4: Relative change of MLU for non-P2P traffic




Figure 5: Relative change of delay

We now consider the P2P network delay performance. For the P2P overlay, we first show in Figure 5 the change of end-to-end delay performance between individual peers upon the completion of peer reselections after each round. We only consider the peer partners whose connections have been actually affected by the TE operation. We recall that the average delay ratio, i.e. the ratio of the longest LSP delay to the shortest LSP delay, between each PoP node pair is 2.6. In the Figure we can see that for all the three scenarios the end-to-end delay is not significantly negatively impacted by the TE adjustment for most of the period, and sometimes such performance can be even improved compared with the initial state. This is due to the P2P overlay's greedy reactions to the changed paths – alternative partners can be often identified without significant extra delay as compared to the initial state. In Figure 5, the low and medium scenarios finally converge to equilibrium points that have 10% and 7% increase in delay respectively compared with the initial state. The medium scenario has a regular oscillation pattern in the last 40 rounds, oscillating between a 5% increase and a 20% decrease compared to the initial state. This observation indicates that the

P2P overlay has generally high *resilience* capability in maintaining end-to-end performance assurance against the change of underlying path selections by TE operations, thanks to the selfish peer selection behaviors at the application level.




Figure 6: Relative change of partner peer churn ratio

Finally we investigate how the P2P connection stability performance is affected by the TE. We define the *peer churn ratio* metric to analyze the relevant performance at the P2P overlay side. As mentioned before, due to the experiences of higher delay following the adjustment of TE, a set of affected peers may need to re-select some of their partners to replace the original ones whose end-to-end delay performance becomes higher. We define a requesting peer that has *any* affected partners (as a result of TE operation) to be an in-churn peer. This metric for the whole system stability can be defined as:

$$\frac{\text{No. of requesting peers that need to find new partner peers}}{\text{Total No. of requesting peers}}$$

The reason for evaluating such a metric is that the transient time period during reselections may lead to a perceivable service disruption for real-time P2P services, and this is in particular the case if a large number of existing partners need to be replaced for the requesting peer. According to Figure 6, we can see that the high scenario has the lowest peer churn ratio among all the three scenarios. The low scenario has lower churn ratio than first step and can converge to an equilibrium point (60% of the initial state). We also find that medium scenario has significant oscillations after 56th round, but the peer churn ratio is a maximum of 40%.

D. General Observations

Based on these simulation results we now discuss the issues we raised in Section I. First of all, we can see that it is not always the case that an equilibrium point can be reached between the two rational players. Even if such equilibrium point exists, it is not always a Pareto point. According to the MLU performance patterns, we find that in the non-cooperative environment, TE does not seem to be able to drive the overall network performance to an improved situation when interacting with selfish P2P overlays. This is particularly the case when P2P flows dominate the overall network traffic. On the other hand, the P2P overlay exhibits high resilience capability and is able to avoid end-to-end delay degradation following TE adjustments. Nevertheless, a high proportion of

partner re-selection may lead to perceivable service disruption for real-time P2P services.

IV. CONCLUSIONS

In this paper we use *best-reply* dynamics to model the non-cooperative interactions between the P2P overlay and network-level TE behaviors, each of which has distinct optimization objectives – the P2P overlay aims to improve the experienced quality in terms of delay for individual peers, while TE aims to optimize the overall network resource utilization. The decision made by each has significant impact on the performance of the other when they optimize potentially conflicting objectives simultaneously. Through the analysis of such interactions based on our simulations, we show that in the non-cooperative network environment TE does not seem to be efficient in optimizing network performance when interacting with a selfish P2P overlay. On the other hand, the P2P overlay is generally resilient against potentially adverse TE operations in terms of end-to-end delay performance. However, with the P2P-based real-time streaming applications being more and more popular, a high proportion of peer partner re-selection in reaction to TE operations may introduce service disruptions, and this phenomenon needs to be further investigated. We intend to investigate relevant interactions in more detail in our future work.

V. ACKNOWLEDGEMENT

This work was partially funded by the EU FP7 COMET Project (248784).

REFERENCES

- [1] V. Aggarwal et al., “Can ISPs and P2P Users Cooperate for Improved Performance?”, ACM CCR, Vol. 37, No. 3, pp. 29-40, July, 2007
- [2] The IETF ALTO WG: <http://www.ietf.org/dyn/wg/charter/alto-charter.html>
- [3] L. Qiu et al., “On Selfish Routing in Internet-Like Environments” Proc. ACM SIGCOMM 2003
- [4] Y. Liu et al., “On the Interaction Between Overlay Routing and Traffic Engineering”, Proc. IEEE INFOCOM 2005
- [5] W. Jiang et al., “Cooperative Content Distribution and Traffic Engineering in an ISP Network”, Proc. ACM SIGMETRICS 2009
- [6] D. DiPalantino et al., “Traffic Engineering vs. Content Distribution: A Game Theoretic Perspective”, Proc. IEEE INFOCOM 2009
- [7] S. Kandula et al., “Walking the Tightrope – Responsive Yet Stable Traffic Engineering”, Proc. ACM SIGCOMM 2005
- [8] ABILENE network, <http://www.internet2.edu/network/>
- [9] S. Uhlig et al., “Providing Public Intradomain Traffic Matrices to the Research Community”, ACM CCR, Vol. 36, No. 1, pp. 83-86, Jan. 2006.
- [10] X. Hei et al., “A measurement study of a large-scale P2P IPTV system”, IEEE Trans. on Multimedia, Vol. 9, No. 8, pp. 1672-1687, Dec. 2007