

An Empirical Study on the Interactions between ALTO-assisted P2P Overlays and ISP Networks

Chaojiang 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—The recently proposed Application Layer Traffic Optimization (ALTO) framework has opened up a new dimension for Internet traffic management that is complementary to the traditional application-agnostic traffic engineering (AATE) solutions currently employed by ISPs. In this paper, we investigate how ALTO-assisted Peer-to-Peer (P2P) traffic management functions interact with the underlying AATE operations, given that there may exist different application-layer policies in the P2P overlay. By considering specific P2P peer selection behaviors on top of a traffic-engineered ISP network, we conduct a performance analysis on how the application and network-layer respective performance is influenced by different policies at the P2P side. Our empirical study offers significant insight for the future design and analysis of cross-layer network engineering approaches that involve multiple autonomous optimization entities with both consistent and non-consistent policies.

I. INTRODUCTION

Traffic Engineering (TE) techniques have been extensively investigated for achieving cost-efficient resource utilization in ISP networks. The basic TE strategy is to obtain optimized traffic routing and forwarding configurations in order to achieve desired network performance targets. Optimized traffic delivery paths are computed based on the mapping of the static long-term traffic demand (i.e. the traffic matrix - TM) onto the underlying network topology. To do this, the network operator needs to accurately predict the overall traffic demand between all ingress and egress routers, for instance based on Service Level Agreements (SLAs) established with customers and also from long-term traffic measurements. Nevertheless, with the increasing popularity of peer-to-peer (P2P) applications in recent years, the prevalent traffic patterns in operational ISP networks have become increasingly difficult to capture. The key challenge for operators in dealing with traffic flows incurred by P2P overlays is the *uncertainly in determining communication endpoints*, which are under the control of end users. More specifically, given the *ad hoc* peer selection and content swarming behaviors based on (instantaneous) content chunk availability, peer reputation, and uplink capacity etc., the overall P2P traffic demand has become much more difficult for ISPs to forecast [1]. As such, solely relying on a static traffic matrix in order to perform traditional *application-*

agnostic traffic engineering (AATE) does not seem to be practically effective. This is especially true when P2P-based content flows dominate the overall network traffic, as is the case in some ISP networks [2, 3].

More recently, proposals have appeared suggesting cooperation between applications and the underlying ISP network in order to achieve “win-win” solutions. In particular, the Application Layer Traffic Optimization (ALTO) framework is currently being investigated by the IETF [4, 5]. According to this approach, a dedicated ALTO server maintained by the ISP is responsible for providing necessary network information to the P2P overlay for supporting ISP-friendly peer selection. ALTO-based traffic optimization is no longer completely ISP-centric, as is the case with traditional AATE. Instead, traffic optimization can be “indirectly” enforced at the application layer, rather than by the ISP through manipulating traffic delivery paths in the network. Such paradigms open a new dimension in traffic optimization that has been traditionally performed only through routing and forwarding within the network. In this case, efficient network-aware (host-level) *endpoint selection at the network edge* can make a key contribution.

In this paper we conduct an empirical study in order to answer the following important questions concerning how to “position” the role of the new ALTO-based approaches in operational ISP networks for traffic optimization:

- What exactly is the behavioral interaction between ALTO and AATE if both are simultaneously but independently applied to operational ISP networks?
- How can different P2P overlay policies and objectives influence the overall network performance under joint traffic control by ALTO and AATE?
- What is the impact of the proportion of P2P flows in the overall Internet traffic on the network and application performance when ALTO and AATE interact with each other in a *cooperative* or *semi-cooperative* manner?

It is important to note that, with the coexistence of ALTO and AATE, traffic patterns are influenced by the decisions of two independent *autonomous entities*: the application provider (the P2P overlay) who is responsible for

performing ALTO¹, and the underlying ISP who is responsible for performing AATE. In comparison with the AATE objective which tries solely to optimize use of the overall network resources at the ISP side, the definition of objective(s) for the P2P overlay is more complex. In this paper we consider the following three distinct scenarios: (1) the P2P overlay is *fully-cooperative* with consistent objectives as AATE, and does not take into account any application-layer requirements; (2) the P2P overlay is *semi-cooperative*, meaning that peer selection decisions are based on both network performance and application-layer requirements such as end-to-end delay between peers; and (3) the P2P overlay is *non-cooperative*, in which case peers are selfishly selected only according to application layer requirements.

II. RELATED WORK

The key idea of the ALTO framework is to disseminate necessary information about the underlying network up to the application overlays, typically peer-to-peer systems, for assisting intelligent endpoint (peer) selections towards a win-win situation between the application and the network. Towards this end, a dedicated ALTO server is needed that is responsible for passing network information up to the decision-making engine on peer selections at the P2P side. Currently ALTO is being standardized in the IETF, both on the problem formulation [4] and protocol specification [5]. This type of collaboration between P2P systems and ISPs has also been proposed in the literature, such as the Oracle Service [3] and P4P [6]. Specifically, it is proposed that some network information can be shared with the P2P overlay for selecting the best partner peers at the application level so as to reduce P2P traffic within the network while improving application performance.

In the literature, a number of works [7, 8, 9, 10] 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 [7, 8]. 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 considering routing in the overlay. The second category of works [9, 10] 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, whereas in the CDNs of [9, 10] content servers are statically provisioned in the network for providing content delivery services. Secondly, we consider *symmetric* content exchange patterns: in P2P overlays content is provided by each peer to the other; 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.

III. ALTO AND AATE INTERACTIONS

We model the ISP (with AATE) and the P2P overlay (with fully-, semi-, or non-cooperative behaviors) as two autonomous and rational players who play *best-reply dynamics*: one player chooses the best response based on the other’s decisions in the previous round. Specifically, AATE aims to optimize the overall network performance by adjusting routing and/or forwarding decisions of customer traffic (including both P2P flows and non-P2P background traffic without differentiation) at the network layer. The changed routing/forwarding behaviors made by AATE for P2P traffic is then taken as input by the P2P overlay to take further actions. In the fully-cooperative case, the P2P overlay, assisted by the ALTO server and protocol, aims to exploit opportunities (by reselecting partner peers) for *further* improving the network performance which could not be achieved through routing or forwarding optimization by AATE in the previous round. In the semi-cooperative case, the P2P overlay considers the best trade-off between application-layer requirements and network performance according to its own strategy. Finally in the non-cooperative case, the P2P overlay selfishly performs peer reselections according only to application-layer objectives, especially for those peers whose performance has been affected by the previous AATE action due to route changes. As a result, the overall traffic performance can be affected by the adverse impact from the non-cooperative P2P overlay.

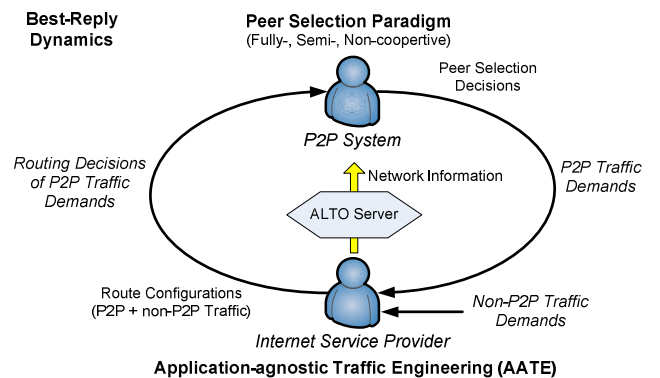


Figure 1: Iterative interaction between AATE and P2P overlay

The peer reselection action taken by the P2P overlay may further influence the overall traffic distribution within the network, possibly requiring further AATE operations from the ISP point of view. This is typically the case for semi- and

¹ More explicitly, here the concept of “ALTO” refers to the application layer traffic optimisation *enforcement* at the P2P side, rather than the actual ALTO server or protocol at the ISP side.

non-cooperative cases where the objectives of the two players are inconsistent or even conflicting with each other. As a result, multiple rounds of bargaining interactions between the P2P overlay and AATE can be conceived. Throughout such an iterative process, the P2P overlay and AATE adjust their own decisions according to each other's input from the previous round. We follow the generic assumption made in [9] that each player (i.e. the (P2P) content service overlay and TE) runs its optimization until the other has finished its operations in each round. Figure 1 shows a generic interaction process between the ISP and the P2P overlay, which can be modeled as best-reply dynamics (which will be introduced in Section IV.E).

IV. NETWORK AND P2P POLICY MODELING

A. Network & P2P Overlay Models

We 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 $a \in A$ is associated with 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 non-P2P background traffic is determined by AATE without distinction. Following the 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 that if multiple interconnected neighboring peers are clustered within the same PoP, then the associated bandwidth consumed by their local peering connections is ignored. Since the overall network traffic consists of both non-P2P background traffic and P2P traffic, let $t_{i,j}^{p2p}$ and $t_{i,j}^{np2p}$ denote respectively the P2P traffic demand and the non-P2P background traffic demand from PoP node i to j . Let $t_{i,j}$ be the total traffic demand, i.e. $t_{i,j} = t_{i,j}^{p2p} + t_{i,j}^{np2p}$. The parameter $f_{i,j}^a$ ($0 \leq f_{i,j}^a \leq 1$) is defined to be the traffic splitting ratio, computed by the AATE optimization engine, on each inter-PoP link $a \in A$ for the overall traffic demand from PoP node i to j (i.e. $t_{i,j}$). Specifically, AATE changes the routing/forwarding matrix $M = \{f_{i,j}^a\}$ for the entire network traffic, including both P2P and non-P2P flows to improve network performance through (re-) optimizing traffic splitting ratios across multiple paths.

Now we specify the P2P overlay model in which each peer is associated with one of the PoP nodes in the network topology. Without loss of generality, 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 P2P sessions, they are treated separately. Let V denote a set of active peers physically attached to network G . Each client peer in a P2P session z needs to connect to a set of partner peers from all available candidates in session z (denoted by $V(z)$), and download content from them at certain transmission rates. In this case the actual partner set for a

specific client peer u (denoted by $V_u(z)$) is a subset of all the available peers in session z , i.e., $V_u(z) \subseteq V(z) \subseteq V$.

B. Non-cooperative P2P Overlay Policy

A non-cooperative P2P overlay aims to exclusively optimize the performance experienced by end users, for example reducing end-to-end delay between interconnected peers. According to [11], end-to-end delay, which is of interest to P2P applications, can be often determined by the physical network distance (propagation delay) between PoP nodes in well over-provisioned networks. As far as each pair of interconnected peers is concerned, if the P2P flow is split across multiple physical paths, then the actual end-to-end delay between them is the one associated with the longest delay path. Let d_a be the delay of physical link $a \in A$, and the delay between a PoP node pair be the sum of the delay associated with each link constituting the longest path between them. The selfish objective by the non-cooperative P2P overlay can be formulated as:

$$\min \sum_z \sum_u \sum_{v \in V_u(z)} \sum_{a \in A} d_a Y_{u,v}^a \quad (1)$$

where $Y_{u,v}^a$ is the binary mapping coefficient, equal to 1 if link a constitutes the longest path from PoP node u to v , and equal to 0 otherwise. As such, it can be seen that the ultimate goal of the non-cooperative P2P is to minimize the overall end-to-end delay for individual peers at the application layer, without taking into account any network requirements or conditions.

C. Semi- and Full-cooperative P2P Overlay Policies

The ALTO proposal enables cooperative P2P approaches. As such, network conditions can be taken into account in order to assist intelligent peer selection at the application level. In this paper we consider two distinct scenarios in the cooperation context: semi- and fully-cooperative P2P overlays. Specifically, the semi-cooperative P2P overlay aims at multiple objectives concerning both network resource optimization and application-layer P2P performance when performing peer selections. The fully-cooperative P2P overlay follows exactly the same objective as defined by AATE for optimizing network resources, but without considering application layer requirements. Although it is as yet unknown whether P2P applications will be willing to adopt the fully-cooperative policy in practice, this scenario can be used as a reference point for evaluating the performance of the other cases.

Let us first look at the semi-cooperative scenario on the P2P side. The objective is to reduce the overall delay between the interconnected peers, which is consistent with Eq. 1 in the non-cooperative case. On the ISP side, we use the metric $p_{u,v}$ to denote the interface variable for optimizing network resources (known as P4P *p-distance interface* in [6]) that is computed by ALTO server based on the underlying network information, and then provided to P2P overlay for assisting ISP-friendly peer selections. With $p_{u,v}$ and a weighting coefficient β ($\beta > 0$) indicating relative weighting of two potentially conflicting objectives, the P2P side's objective function shown in Eq. 1 is extended for the semi-cooperative case:

$$\min \sum_z \sum_u \sum_{v \in V_u(z)} ((\sum_{a \in A} d_a Y_{u,v}^a) + \beta p_{u,v}) \quad (2)$$

Now we specify how $p_{u,v}$ is used here in order to include the AATE objective at the ISP side. In this semi-cooperative model we consider how to minimize the overall M/M/1 based network cost [8]. Such optimization is computed by the ALTO sever with the necessary knowledge of the network status. For traffic demand, let bt_a and pt_a be the overall non-P2P background traffic and P2P traffic over the physical link a respectively, which can be expressed as:

$$bt_a = \sum_{i \in N} \sum_{j \in N \setminus \{i\}} t_{i,j}^{\text{np2p}} f_{i,j}^a \quad (3)$$

$$pt_a = \sum_{i \in N} \sum_{j \in N \setminus \{i\}} t_{i,j}^{\text{p2p}} f_{i,j}^a \quad (4)$$

Following the same fashion as the M/M/1 network cost function defined in [8], we consider $l_a = \frac{1}{C_a - (bt_a + pt_a)}$ as the cost by transmitting a single packet on each physical link a . The objective function on the network side is then defined below and is same as the objective of AATE in section IV.D (minimizing overall network cost).

$$\min \sum_{a \in A} (bt_a + pt_a) l_a \quad (5)$$

$$\text{subject to } (bt_a + pt_a) < C_a \quad (5.1)$$

According to [6], we can decompose the problem (Eq. 5) separately by introducing a dual variable p_a for constraint (5.1) to construct a Lagrange dual function, so that Eq. 5 can be turned into:

$$\sum_{a \in A} p_a (bt_a - C_a) + \min \sum_{a \in A} (p_a + l_a) pt_a \quad (6)$$

$$\text{subject to } \sum_{a \in A} p_a C_a = 1, p_a \geq 0 \quad (6.1)$$

With equations (6) and (6.1), we can follow the same technique applied in [6] to update p_a and exchange $p_{u,v} = \sum_{a \in A} p_a Y_{u,v}^a$ between P2P overlay and ALTO server for both sides' optimization.

Finally, we specify the modeling of the fully-cooperative P2P overlay scenario that follows the same objective as AATE, i.e. to minimize network cost, but without considering any application-layer P2P performance. An interface variable $p_{u,v}$ also needs to be provided from the ALTO server to the P2P overlay for peer selection operations. However the calculation of $p_{u,v}$ is solely based on the network requirement reflected by the AATE objective. Since this P2P overlay policy does not take into account any application-layer requirement such as end-to-end delay, the objective function for the fully-cooperative P2P policy can be simply expressed as:

$$\min \sum_z \sum_u \sum_{v \in V_u(z)} p_{u,v} \quad (7)$$

It worth mentioning that the fully-cooperative P2P overlay theoretically produces the ideal situation the interaction, since P2P overlay and AATE have consistent objectives but perform the optimization at different layers (application level and physical level). Although it might not be always a practical scenario due to distinct operational objectives at the two sides in reality, we can certainly use it

as a baseline to evaluate the other two scenarios (non-cooperative and semi-cooperative P2P overlays) in our simulation.

D. Application-agnostic Traffic Engineering

AATE operations are applied in order to optimize the overall network performance, such as minimizing overall network cost. Thanks to network monitoring techniques, an ISP typically has full knowledge of network status and the estimated overall traffic volume between individual PoP pairs. Again, we emphasize that AATE aims to optimize the overall network performance rather than any specific type of traffic.

As we mentioned previously, we consider consistent objectives/cost functions between AATE and the fully-cooperative P2P policy at the network side. As far as AATE is concerned, we still use the M/M/1 network cost function, but represent it in a different manner for describing the ISP-centric AATE operations. A generic AATE function, $K(\cdot)$ is a non-negative, convex, and increasing function of the link load (M/M/1). Given the traffic demands between individual PoP node pairs, the overall traffic load T_a on the physical link $a \in A$ is the sum of all bandwidth demands of flow over this link, including both P2P traffic and non-P2P traffic, i.e. $T_a = bt_a + pt_a$. From the AATE point of view, T_a is also equal to $\sum_{i \in N} \sum_{j \in N \setminus \{i\}} t_{i,j} f_{i,j}^a$. The objective of AATE is to compute an

optimized value of the splitting ratio $f_{i,j}^a$ between each source-destination PoP node pair i and j on each link a in order to minimize the overall network cost. We formally model the AATE optimization function as follows:

$$\min \text{AATE} = \sum_{a \in A} k(T_a, C_a) = l_a T_a \quad (8)$$

$$\text{s.t. } \sum_{a:d(a)=z} t_{i,j} f_{i,j}^a - \sum_{a:s(a)=z} t_{i,j} f_{i,j}^a = \begin{cases} t_{i,j}, & z=j \\ -t_{i,j}, & z=i \\ 0, & \text{otherwise} \end{cases} \quad (8.1)$$

subject to where $\forall z, i, j \in N$, and $s(a)$ and $d(a)$ are the head node and tail node of link a respectively. With this convex increasing function, we formulate the AATE objective as a convex optimization when it has linear constraints. Global optimality can be reached through classical algorithms such as feasible direction methods [12]. It should be also noticed that it is required that $T_a < C_a$ in order to avoid a singularity point at $T_a = C_a$, as in [7, 8, 9].

E. AATE and P2P Overlay Interaction Analysis

In this section we model the 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 AATE and the P2P overlay take turns to optimize their own objectives in this interaction.

The strategy space that is applied by AATE can be described as a set of feasible traffic splitting ratio configurations between each PoP node pair $\langle i, j \rangle$. This can be expressed as:

$$S_{TE} = \langle \dots, \{f_{i,j}^a\}, \dots \rangle \quad (9)$$

On the other hand, the strategy space of the P2P overlay is a set of partner peers $V_u(z)$ of every single client peer u that are selected from all available peers candidates $V(z)$ in each session z . The non-cooperative P2P overlay selects the best partner peers to minimize the end-to-end delay among interconnected peers for each client in the session. Conversely, both network performance and user's experience are jointly taken into account (with weighting) by the semi-cooperative P2P paradigm. Finally, a fully-cooperative P2P overlay selects peer based on network engineering requirements only, while sacrificing its own P2P performance at the application layer.

$$S_{P2P} = \langle \dots, V_u(z), \dots \rangle \quad (10)$$

In the beginning, based on the current traffic distribution AATE takes the first turn to optimize the traffic splitting ratio on every physical link in the network in order to minimize overall network cost, without any difference between P2P traffic and non-P2P traffic. Since the actual delivery paths of P2P traffic are changed by this AATE operation, the P2P performance (e.g. end-to-end delay between existing interconnected peers) may be affected. In particular, if some P2P traffic is shifted to a longer path, the corresponding peers may experience higher end-to-end delay after such a change. In the non-cooperative case, in order to maintain the original quality of experience, the P2P overlay will reactively re-select some alternative partner peers (if they exist) within individual sessions in order to regain the original application performance that was affected by the path selection reconfigurations. Such a reaction at the P2P side *may* also occur in the semi-cooperative scenario, depending on the weighting coefficient between network and application objectives. Due to this reshuffling of peer connections following the previous round of AATE operation, the overall network traffic pattern may be changed significantly. As such, AATE may need to further readjust traffic splitting ratios in order to maintain its own objective in the next round. This is not the case for the fully-cooperative P2P overlay case where convergence can be reached since both parties follow exactly the same optimization objective (see section IV.C). On the other hand, the non-cooperative P2P (or semi-cooperative P2P) and AATE take turns to optimize their own operational objectives according to previous decision of the other player, which can be modeled as:

$$\begin{aligned} \{f_{i,j}^a\}_{(k+1)} &= \arg \min_{(k)} TE(V_u(z)) \\ V_u(z)_{(k+1)} &= \arg \min_{(k)} P2P(\{f_{i,j}^a\}) \end{aligned} \quad (11)$$

V. EXPERIMENTAL SETUP AND TEST SCENARIOS

A. Simulation Setup

We use the real PoP-level GEANT network topology [13] and traffic traces in our simulation experiments. The GEANT network consists of 23 nodes and 74 unidirectional

inter-PoP links according to data published in 2006 [11]. According to [11], the IGP link weight setting in GEANT is based on end-to-end latency, and hence customer traffic is effectively routed on the lowest delay paths. In addition, the actual GEANT traffic traces are obtained through the public dataset in [14].

The P2P traffic used in our experiments is synthetically generated according to the measured pattern of today's popular P2P applications [1]. We consider 20 concurrent P2P sessions, with each session attracting up to 1200 peers (specific numbers depends on the P2P traffic proportions under evaluation, see section V.B). The overall distribution of these peers in each PoP node is determined according to the actual population of each city (PoP), in which case larger PoP nodes have more peers assigned. Without loss of generality, our experiments include both popular and less popular sessions in the P2P overlay. For each requesting peer, there is a primary peer who feeds on average three times content as other auxiliary ones; this is also according to the measurement results from [1].

B. Simulation Scenarios

In our study, we consider the following three distinct traffic ingredient patterns: the overall P2P traffic demand is set to be *low*, *medium* or *high* proportion of the overall network traffic volumes, i.e. the P2P traffic accounts for 20% (low), 50% (medium), or 80% (high) of the 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% [1, 2]. It can be inferred that AATE becomes the main driving force for optimizing the overall traffic load in the scenario of low P2P proportion, while ALTO's impact on the overall network performance increases with a higher P2P traffic proportion. In addition, we set the value of weighting coefficient ($\beta=0.5$) in the semi-cooperative P2P to interact with AATE in our simulation. We model the three distinct P2P overlay behaviors (non-cooperative, semi-cooperative and fully-cooperative) to interact with AATE for 100 rounds in our evaluation. AATE initiates the interaction processes, so that AATE takes a turn at every odd round and the P2P overlay takes a turn at every even round. Since the fully-cooperative P2P and AATE have consistent objectives in optimizing the overall network cost, such a scenario is obviously able to achieve the best results (out of the three scenarios) at the network side, and hence it is used as the reference one in order to evaluate the network performance of the other two, including both overall network cost and maximum link utilization (MLU).

$$\frac{\text{Non/Semi-cooperative P2P}(t)}{\text{Fully-cooperative P2P}(t)}, (1 \leq t \leq 100) \quad (12)$$

In addition to the network performance we also investigate P2P side performance, such as end-to-end delay that can be reflected by the IGP distance [11]. In this case, the non-cooperative P2P overlay, whose objective is solely to improve user quality of experience (reducing delay), can be regarded as the reference point to evaluate the application-


oriented performance of the semi- and fully-cooperative P2P behaviors.

$$\frac{\text{Fully/Semi-cooperative P2P}(t)}{\text{Non-cooperative P2P}(t)}, (1 \leq t \leq 100) \quad (13)$$


VI. PERFORMANCE EVALUATION

A. Performance Analysis


Figure 2 shows the overall network cost performance of the non- and semi-cooperative P2P interacting with AATE in three different P2P traffic proportions of the overall network traffic (across 100 rounds). These are the relative ratios against the fully-cooperative scenario (see Eq. 12) whose overall network cost converges after the 2nd iteration according to our results, thanks to the completely consistent objectives between the two entities. More specifically, following the first round AATE operation, the fully-cooperative P2P overlay continues improving the network performance, without leaving any space for the next round of AATE optimization for any further enhancement. It is clear from Figure 2 that the overall network performance achieved by the semi-cooperative scenario is constantly better than the non-cooperative one, due to the more ISP-friendly objective. On the other hand, we can clearly observe the oscillatory behaviors of the non- and semi-cooperative cases in Figure 2. The reason is that the interaction between AATE and the P2P overlay is processed in an interleaved manner where AATE first obtains the best network cost solution, but the outcome of the P2P selfish behavior then leads to deteriorated network performance due to the inconsistent objectives with AATE. In response to the affected network performance caused by P2P overlay, AATE needs to re-compute the splitting ratios with the aim to regain the original performance. It is particularly worth mentioning that the *oscillation degree* (defined as the relative ratio between the maximum and minimum values across the 100 rounds) of the non-cooperative P2P becomes higher as the P2P traffic proportion increases (12%, 23% and 30% in the low, medium and high cases respectively). The reason for this situation is as follows. With the increase of the P2P traffic proportion, more peers have the opportunity to perform selfish peer (re-) selections based on their own end-to-end delay objectives. Such a significant traffic pattern change results in a larger optimization space for AATE in the next round that aims to regain the original optimized performance. We can clearly observe that the high-degree oscillation remains across the entire 100 rounds due to the dynamics of the high P2P traffic proportion influenced by both the P2P overlay and AATE with conflicting policies. However the oscillation degree of the semi-cooperative P2P stays lower (10%, 8% and 9% in the low, medium and high cases respectively) than the non-cooperative P2P scenario, and does not become higher as the P2P traffic proportion increases in Figure 2. This phenomenon is also due to the more “close” optimization objective in the semi-cooperative case in comparison to the non-cooperative one.



(a) Low P2P traffic proportion scenario



(b) Medium P2P traffic proportion scenario




(c) High P2P traffic proportion scenario


Figure 2: Network cost performance of the non- and semi-cooperative P2P vs. AATE

We now look at the relative performance gap of each specific policy across different proportions of P2P traffic. First of all, we can see that the overall network cost becomes higher for the non-cooperative scenario with the increase of P2P traffic proportion (on average 18%, 64% and 104% against the fully-cooperative scenario for the low, medium and high cases across the 100 rounds). Again, this outcome is expected: following a high number of non-cooperative peer reselection operations as it becomes more difficult for AATE to re-obtain the network performance achieved previously. This is in contrast to the semi-cooperative scenario which takes into account network objectives when performing peer reselections. As shown in Figure 2, the overall network performance does not change that dramatically across the three P2P traffic proportion cases (on average 12%, 34% and 29% on the low, medium and high cases respectively). In fact the performance gap between the semi- and non-cooperative scenarios becomes more significant with the increase of the P2P traffic proportion. This result indicates that, as far as network performance is concerned, the semi-cooperative policy is able to exhibit good optimization


capability in those ISP networks where P2P flows significantly dominate the overall traffic.



(a) Low P2P traffic proportion scenario



(b) Medium P2P traffic proportion scenario




(c) High P2P traffic proportion scenario

Figure 3: MLU performance of the non- and semi-cooperative P2P vs. AATE


In addition to the overall network performance, we also show in Figure 3 the maximum link utilization (MLU) dynamics. This is because MLU is another common metric that is often used to evaluate network performance. Similar oscillation performance curves can be observed for the same reason we mentioned before. Nevertheless, it should be noted that although both the network cost and MLU are popular metrics for evaluating network utilization, they are not always consistent with each other. The degree of MLU oscillation for the non-cooperative scenario does not strictly follow that of the network cost. A possible explanation is that the MLU is not as “sensitive” as the overall network cost, as the utilization of the most loaded link may remain the same despite the traffic dynamics elsewhere in the network that still influence the overall network cost.

We now investigate the end-to-end delay (a typical user-oriented metric) for P2P sessions, as shown in Figure 4. The same evaluation methodology is adopted as to the one used


in evaluating network-oriented performance. Nevertheless it is also worth mentioning that we use the delay of the non-cooperative P2P case (Eq. 13) as the reference point, given it is expected to achieve the best delay performance due to the selfish behavior at the P2P side. On the other hand, it is also important to note that, like the network cost, the actual end-to-end delay for the fully-cooperative scenario converges after the 2nd round as indicated before, while the reference curve from the non-cooperative case has oscillations. This is because the optimized delay achieved by the selfish peer selection can be significantly impacted by the follow-up AATE operation that aims at network resource optimization. It is not difficult to infer that the next round of peer selections pulls back to a re-optimized delay. Once again, we observe that the degree of oscillation in medium and high proportion P2P traffic is higher than that in low P2P traffic.



(a) Low P2P traffic proportion scenario



(b) Medium P2P traffic proportion scenario



(c) High P2P traffic proportion scenario

Figure 4: Delay performance of the semi- and fully-cooperative P2P vs. AATE

Last but not least, an interesting finding from Figure 4(c) is that the “worst-case” (i.e. peak) delay of the fully-cooperative scenario is not significantly worse than that of the semi-cooperative one in case of *high* P2P traffic. This

observation can be conflicting with the intuitive view that the fully-cooperative policy biases too much to the network optimization objective at the expense of P2P performance deterioration. According to our results, in those ISP networks that are heavily loaded with P2P traffic the fully-cooperative policy is preferable in order to result in improved network performance without impacting much on P2P application performance.

B. Key Observations

Based on our evaluations on the three different P2P overlay policies vs. application-agnostic TE in low, medium, and high P2P traffic proportion cases, we summarize the following key observations.

First of all, while it is intuitively desirable to perform simultaneously joint traffic optimization across application- and network-layers, i.e. ALTO and AATE as complementary to each other, different policies adopted at the P2P overlay may result in dramatically different performance at both the network and application side. Due to the potentially conflicting objectives by the two independent entities, both the *absolute performance* (e.g. the actual values of network cost, MLU and delay) and their *oscillation dynamics* can be affected. As such, any future joint optimization paradigms controlled by multiple parties/layers should consider both performance improvements as well as the overall system stability, as pointed out here. Second, the performance at both sides can be also significantly influenced by the *traffic type ingredients*, for instance the proportion of P2P traffic that can be directly manipulated by the ALTO optimizer at the overlay side. Given the dramatically varying proportions of P2P traffic in different ISP networks, extra complexity can be introduced in cooperative traffic optimization involving multiple autonomous parties across the entire Internet. Finally, it is easy to see that some trade-off between conflicting policies (e.g. the semi-cooperative scenario) are to be most likely to achieve desired trade-off between application and network-level performance, particularly in the scenario where P2P flows dominate the overall network traffic. Another interesting observation is that, although the fully-cooperative policy does not achieve as good application-level (i.e. delay) performance as the semi-cooperative one, the performance gap is not very big in case of high proportion of P2P traffic. Given the fact that the fully-cooperative policy has high capability in optimizing network performance, such a policy could be considered by the P2P overlay at the expense of moderate (i.e. less than 10%) worst-case performance deterioration.

VII. CONCLUSIONS

Thanks to the recently proposed ALTO framework, ISP-friendly peer selection approaches by P2P applications have become possible by taking into account underlying network performance. However, such solutions are not sufficient for optimizing the entire Internet traffic given the existence of different types of flows. With the dramatically varying proportion of P2P traffic across ISP networks, one important

issue that needs to be investigated is how ALTO-assisted P2P traffic management may interact with traditional application-agnostic traffic engineering (AATE), given that both are necessary and complementary to each other. In this paper we studied three distinct behavioral policies: non- (also denotes the situation where ALTO is not available), semi- and fully-cooperative that can be possibly adopted by the P2P overlay over the AATE operations performed by the ISP. Such interactions were modeled using the best-reply dynamics, according to the other entity's behavior in the last round. Our evaluation experiments based on the GEANT network data have yielded interesting results that can potentially shed light on future joint traffic management performed by multiple autonomous entities with their own policies. In particular, we have identified the following two major factors that need to be specifically considered when deploying such approaches: (1) the *degree of consistency* in the optimization objectives adopted by autonomous parties where some trade-off may be necessary, and (2) the specific proportion of P2P traffic that can be controlled by the application-layer optimizer. According to our results, these two factors may significantly impact the overall performance at both the network and the P2P application overlay.

VIII. ACKNOWLEDGEMENT

This work was partially funded by the EU FP7 UniverSelf project (257513) and EU FP7 COMET Project (248784).

REFERENCES

- [1] 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
- [2] A. Banerjee et al., "Dude, Where's My Peer?", Proc. IEEE Globecom 2006
- [3] V. Aggarwal et al., "Can ISPs and P2P Users Cooperate for Improved Performance?", ACM CCR, Vol. 37, No. 3, pp. 29-40, July, 2007
- [4] J. Seedorf et al, Application-Layer Traffic Optimization (ALTO) Problem Statement: <http://datatracker.ietf.org/doc/rfc5693/>
- [5] R. Alimi et al, "ALTO Protocol", IETF draft, draft-ietf-alto-protocol-05.txt, July 2010,
- [6] H. Xie et al., "P4P: Provider Portal for Applications", Proc. ACM SIGCOMM 2008
- [7] L. Qiu et al., "On Selfish Routing in Internet-Like Environments" Proc. ACM SIGCOMM 2003
- [8] Y. Liu et al., "On the Interaction Between Overlay Routing and Traffic Engineering", Proc. IEEE INFOCOM 2005
- [9] W. Jiang et al., "Cooperative Content Distribution and Traffic Engineering in an ISP Network", Proc. ACM SIGMETRICS 2009
- [10] D. DiPalantino et al., "Traffic Engineering vs. Content Distribution: A Game Theoretic Perspective", Proc. IEEE INFOCOM 2009
- [11] S. Uhlig et al., "Providing Public Intradomain Traffic Matrices to the Research Community", ACM CCR, Vol. 36, No. 1, pp. 83-86, Jan., 2006.
- [12] D.Bertsekas et al., "Data Networks (chapter 5: Routing in Data Networks)", Prentice-Hall Inc., 2nd edition 1992
- [13] The GEANT Network Topology, <http://www.geant.net>
- [14] The GEANT Traffic Traces: available at <http://totem.info.ucl.ac.be/>