On the Efficiency of Collaborative Caching in ISP-aware P2P Networks

Jie Dai[†], Bo Li[*†], Fangming Liu[†], Baochun Li[‡], Hai Jin[§]
[†]*Hong Kong University of Science and Technology*, [‡]*University of Toronto*
[§]*Huazhong University of Science and Technology*

***Abstract*—Collaborative ISP caching has been advocated to reduce the otherwise significant amount of costly inter-ISP traffic generated by peer-to-peer (P2P) applications. The fundamental design criteria employed by ISP cache servers are, however, not well understood, with respect to dynamic P2P traffic patterns, ISP peering policies and cache server capacity constraints. In particular, there is a lack of investigations on the design and analysis of resource allocation mechanisms with awareness of inter-ISP traffic and ISP policies in the context of collaborative ISP caching — which is our focus in this study. In this paper, by characterizing practical inter-ISP traffic patterns, we have developed a theoretical framework to analyze representative cache resource allocation schemes within the design space of collaborative caching, with a particular focus on minimizing costly inter-ISP traffic. The optimization framework incorporates both locality-aware and locality-unaware peer selection strategies and ISP peering agreements, in order to examine their respective effects on the design of ISP collaborative caching mechanisms. Our analyses not only help us understand the traffic characteristics of existing P2P systems in light of realistic elements, but also offer fundamental insights into designing collaborative ISP caching mechanisms.**

I. Introduction

The Internet has witnessed a remarkable increase in Peer-to-Peer (P2P) applications over the years, which has also posed significant pressure to Internet Service Providers (ISPs) with tremendous data volume traversing the Internet. As recently reported in [1], P2P applications generate 70% of the Internet traffic worldwide, which raises intolerable inter-ISP traffic costs to ISPs. It is not surprising that several ISPs have started to limit the bandwidth consumption of P2P applications by proactively detecting and throttling their data packets [2]. However, in practice, the perceived service quality of P2P users is of great importance for ISPs to maintain their customer bases, for which ISPs cannot proactively intercept P2P traffic over the long run. Thus, one critical challenge is to make P2P applications friendly towards ISPs.

The aforementioned tussle between two entities is originated from the disparity between the P2P overlay and the ISP underlay. In a P2P network where peers form an application-layer overlay, neighbors can be fairly distant from the viewpoint of the underlying Internet topology. Specifically, measurement studies on the BitTorrent system have shown that $50\% - 90\%$ of existing local pieces of data in active peers are downloaded externally [3]. To resolve this tussle, locality-aware peer selection mechanisms have been introduced in the literature, such as P4P [4] and TopBT [5]. Using proximity information provided by the ISP infrastructure or end-to-end distance estimation driven by peers, the proposed topology navigating algorithms guide peers to select nearby candidates when overlay connections are established. The locality of streaming or file sharing traffic can therefore be improved. However, the proximity-driven biased neighbor selection also raises doubts regarding the robustness of the P2P overlay [6]. The system performance becomes vulnerable due to the dynamic nature of P2P networks.

Our solution in this paper is to provide collaborative caching to achieve the traffic locality in P2P networks. Caching in the current Internet is mainly used for web traffic [7], but its benefits and applicability to P2P networks have also started to draw research attention [8]. Collaborative caching mechanisms represent an alternative way to unify the concerns on mitigating inter-ISP traffic costs and improving the service quality to P2P users. On the one hand, traffic redirected to cache servers deployed at the edges of ISPs can be controlled so that it does not incur transit traffic costs. This is the primary design objective of the new collaborative caching algorithms. On the other hand, the latency of each P2P packet can be correspondingly reduced, which also improves the experiences of end users, leading to a "win-win" situation for both the ISPs and users alike.

Unlike traditional web caching strategies, new collaborative caching schemes in P2P networks aimed at the mitigation of inter-ISP traffic are concerned with the following important characteristics: (i) *Inter-ISP traffic patterns*. To mitigate the traffic cost, cache servers have to understand fundamental characteristics of inter-ISP traffic. This requires the collaboration between P2P applications and ISPs to trace the dynamic patterns of P2P traffic; (ii) *Cache server resource allocation*. In a P2P content delivery network, especially on a video distribution platform, the cache hit ratio of a caching scheme is not the only factor contributing to the system performance. This is due to the limited uploading capacity provided by cache servers compared to the enormous requirements from miscellaneous end users. Therefore, resource allocation strategies need to be investigated, with respect to both storage and bandwidth constraints; (iii) *ISP peering agreements*. ISP peering relation has emerged as an important way to mitigate Internet transit traffic costs [9]. Since P2P service is typically built over the public Internet across multiple ISPs, it further requires the collaboration between individual ISPs and corresponding

[*]The research was support in part by a grant from RGC under the contracts 615608, and by a grant from Huawei Technologies Co. Ltd. under the contract HUAW18-15L0181011/PN.

cache servers to minimize the volume of transit traffic by exploiting ISP peering.

In response to those challenges and unexposed properties, we propose in this paper a unifying optimization framework to systematically analyze the efficiency of collaborative caching schemes in ISP-aware P2P networks. Taking the video distribution platform as an example, we first develop a theoretical model to characterize the inter-ISP traffic generated by practical P2P networks. We observe that both ISP scales and dynamic properties of P2P video sessions, such as the channel popularity, have explicit effects on the resulting traffic. By characterizing practical inter-ISP traffic patterns, we propose a representative resource allocation scheme with a particular focus on the mitigation of inter-ISP traffic. Our performance analysis shows that our framework achieves a significant reduction in traffic, with respect to both locality-aware and locality-unaware peer selection schemes. We further investigate the effects of ISP peering agreements on the performance of our optimization solution. It is not surprising to find that the peering relationships have positive effects on the mitigation of inter-ISP traffic. An evolved collaborative caching scheme tailored to ISP peering relationships has then been formulated.

The remainder of the paper is organized as follows. In Sec. II, we discuss our contribution in the context of related works. In Sec. III, we present the fundamentals of our inter-ISP traffic model, along with the optimization framework in mitigating inter-ISP traffic. Sec. IV discusses both the effects of ISP peering agreements and the corresponding variations in our proposed mechanisms. Sec. V proceeds to the system trace-driven analyses of our theoretical framework. Finally, we conclude the paper in Sec. VI.

II. Related Work

ISP-friendly design has recently received significant research attention in various studies, which can be generally categorized into three classes: peer-driven biased neighbor selection, ISP-driven biased neighbor selection and ISP caching. Bindal *et al.* [10] proposed a biased neighbor selection scheme to improve the traffic locality based on topology maps, AS mappings and other related metrics. Liu *et al.* [11] observed that PPLive [12] adopts a latency-based neighbor selection mechanism which implicitly mitigates costly inter-ISP traffic. A similar locality-aware neighbor selection framework is proposed by Ren *et al.* [5], in which TCP ping tools are used to estimate the proximity among peers in the system. These studies fall into the category of peer-driven biased neighbor selection.

P4P [4] is a representative framework towards ISP-driven biased neighbor selection. In P4P, ISPs provide an interface to advertise preferred paths to P2P applications. This allows them to explicitly and efficiently implement traffic control between applications and network providers. Our work in this paper differs from these related works in important ways. First, our work adopts ISP caching as the basic scheme rather than biased neighbor selection, which potentially impairs the robustness of P2P overlay [6]. Second, a collaborative caching scheme is transparent to end users while a locality-aware peer selection scheme explicitly interferes with P2P overlay construction. Finally, our work can build upon locality-aware P2P systems to further eliminate the inter-ISP traffic.

The caching strategy is traditionally applied to the web [7] for reducing user access latencies to web pages. With a remarkable increase of P2P traffic over the Internet backbone, several commercial cache products specifically designed for P2P traffic have appeared in the market, such as CacheLogic [13]. In the view of caching algorithms, Hefeeda *et al.* [14] have proposed a proportional partial caching scheme targeting on general P2P systems. However, these works mainly focus on independent server caching with the primary concern of improving the byte hit ratio. Neither the collaboration between ISPs nor cache server bandwidth constraints are included in the consideration. Our work not only provides a theoretic framework with respects to both storage and bandwidth utilization, but also incorporates concerns about ISP peering relation.

Dan [15] introduced a collaborative caching scheme with awareness of peering agreements among ISPs to reduce the volume of transit traffic crossing ISPs. However, this model simplifies the scenario as a rate allocation scheme among multiple cache servers, which lacks investigations into the fundamental properties of inter-ISP traffic and other practical constraints in real-world P2P applications. Compared to this work, our study first develops an inter-ISP traffic model based on practical P2P application properties, and then elaborates on the design of a new collaborative caching scheme with respect to server storage and bandwidth constraints, peer selection strategies and ISP peering agreements. This paper provides a general and unified framework with practical constraints in mind.

III. Inter-ISP Traffic Model and Caching Resource Allocation

In this section, we first develop an inter-ISP traffic model by taking P2P video streaming as a representative application. Both locality-unaware and locality-aware peer selection are taken into account. Based on this model, we then propose our optimization framework to address challenges in allocating resources on collaborative cache servers. Specifically, we explore two sets of server strategies to allocate storage and bandwidth resources on cache servers, with a focus on inter-ISP traffic mitigation. The collaboration between P2P applications and cache servers is essential for the implementation of our proposed mechanisms.

A. Inter-ISP Traffic Model

We first present our inter-ISP traffic model by considering a P2P video streaming system. Without loss of generality, suppose there are a total of N_i video channels in the system, which can be represented as $\mathcal{I} = \{1, 2, ..., N_i\}$. For the video distribution platform and any video channel $i \in \mathcal{I}$, relevant important properties are summarized as follows:

x: The number of concurrent users in the P2P video streaming system.

x_i: The number of concurrent users of video channel i in the system. The index of each channel i is assigned by following the reverse order of the corresponding channel popularity, ***i.e.***, the set of channels shall satisfy $x_1 \geq x_2 \geq \ldots \geq x_{N_i}$.

r_i: The streaming rate of video channel i.

f_i: The size of video channel i. We assume that each channel has the same streaming length, thus the size of video channel i only depends on its streaming rate.

d_{in}: The in-degree of individual peers, which defines the number of connections established for receiving data packets. To simplify our analysis, we assume that the value of peer out-degree equals the value of peer in-degree.

We are now ready to present properties of the set of existing ISPs $\mathcal{K} = \{1, 2, ..., N_k\}$, where N_k denotes the number of ISPs in which peers are currently viewing video channels. The index of each ISP k is also assigned by following the reverse order of the corresponding ISP popularity. Relevant notations are summarized as follows:

S_k: The storage capacity provided by deploying cache servers in ISP k.

U_k: The uploading bandwidth capacity provided by deploying cache servers in ISP k.

a_{ik}: The percentage of video channel i being currently stored in cache servers of ISP k. Since each cache server can store part of the available video, the value of a_{ik} shall satisfy $0 \leq a_{ik} \leq 1$. If the video channel i has been fully stored in cache servers of ISP k, we denote it as $a_{ik} = 1$.

u_{ik}: The uploading bandwidth assigned to video channel i by cache servers of ISP k.

x_{ik}: The number of concurrent users of channel i in ISP k.

We first attempt to capture the user population over multiple channels of the system, in order to drive the viewer population over ISPs. The Zipf-Mandelbrot distribution [14] is used to represent the probability of how frequently a P2P object will be accessed over the long term. This model can also be applied to formulate the user population over streaming channels. We define the probability that any online user is currently accessing to the video channel i as:

$$p_c(i) = \frac{1/H}{(i+q)^\alpha} \quad \text{where} \quad H = \sum_{i \in \mathcal{I}} \frac{1}{(i+q)^\alpha} \tag{1}$$

In this model, the higher the value of q, the flatter the tip of the distribution will be. We determine the value of parameters α and q in Sec. V by fitting the curve of the peak number of concurrent online users in a real-world system. This model also satisfies the property that channel indices are in a reverse order of the user population.

The second population model we elaborate here is the ISP user population. To illustrate the relation between the inter-ISP traffic model and ISP user population, a tunable parameter β is applied to set up different scenarios of ISP user populations. We define the probability that any online user is currently in ISP k as:

$$p_{isp}(k) = \frac{(N_k - k + 1)^\beta}{\sum_{k \in \mathcal{K}} (N_k - k + 1)^\beta} \tag{2}$$

By adjusting the value of β, we can produce different ISP distributions with particular properties. When $\beta = 0$, each ISP has the same amount of concurrent users. The higher the value of β, the more unbalanced the ISP user population will be. Note that the value of β is greater than or equal to 0.

Therefore, the formulation of x_i and x_{ik} can be derived from the channel popularity distribution and ISP user distribution that we have just introduced:

$$x_i = x \cdot p_c(i) = \frac{x/H}{(i+q)^\alpha} \tag{3}$$

$$x_{ik} = x_i \cdot p_{isp}(k) = \frac{x p_{isp}(k)/H}{(i+q)^\alpha} \tag{4}$$

1) Locality-Unaware Peer Selection: In a locality-unaware peer selection scheme, neighboring peers are evenly distributed among all candidate peers. Thus, the number of neighbors from different ISPs is decided mainly by the ISP user population as described above. We use the hyper-geometric distribution $H(x_i, x_{ik}, d_{in})$ to illustrate the probability that m neighbors are selected from the same Internet provider. We can subsequently obtain the result of the amount of traffic from external ISPs — also known as the inter-ISP traffic. Especially when $x_i < d_{in}$, P2P application streaming servers that supply insufficient bandwidth allocation will be regarded as external sources. Note that the responsibility of the streaming servers denoted here is not relevant to the responsibility of cache servers deployed by ISPs, which focus on intercepting costly inter-ISP traffic.

$$Pr(local = m) = \binom{x_{ik}}{m}\binom{x_i - x_{ik}}{d_{in} - m} \Big/ \binom{x_i}{d_{in}} \quad \text{where} \quad m \in [\max(0, d_{in} + x_{ik} - x_i), \min(d_{in}, x_{ik})] \tag{5}$$

The inter-ISP traffic rate generated by channel i in ISP k without involving deployed cache servers can then be formulated as:

$$\begin{aligned} T_{ik}^{n-c} &= x_{ik} \sum_m Pr(local = m)(d_{in} - m)\frac{r_i}{d_{in}} \\ &= \frac{x \cdot r_i \cdot p_{isp}(k) \cdot (1 - p_{isp}(k))}{H(i+q)^\alpha} \end{aligned} \tag{6}$$

Remark: Eq. (6) provides a formulation of the inter-ISP traffic rate generated by a single channel in a particular ISP k. The total amount of inter-ISP traffic rate for ISP k can be obtained by summing up results over all available video channels. We first note that the inter-ISP traffic caused by any particular channel positively correlates to the channel popularity as expected. Furthermore, we observe that the ISP user population has an interesting impact on the resulting traffic. On the one hand, if two ISPs have similar numbers of concurrent users as $p_{isp}(k) \approx p_{isp}(k')$, we have $T_{ik}^{n-c} \approx T_{ik'}^{n-c}$. On the other hand, if two ISPs have widely different scales, the resulting inter-ISP traffic is still very similar. For example, if two ISPs occupy 90% and 10% of the total online users respectively, any particular channel will generate the same amount of inter-ISP traffic to both. At first glance this is surprising, but with a little explanation it is not against intuition: under the locality-unaware peer selection scheme, neighbors of peers in large ISPs are mostly chosen locally, which generate comparatively less per-peer inter-ISP traffic. However, a large number of concurrent users results in a considerable volume of traffic. Relatively speaking, the opposite is true for small ISPs under the same circumstances.

2) Locality-Aware Peer Selection: For locality-aware peer selection [4] [5], peers give priority to nearby candidates when selecting neighbors. With particular concern to inter-ISP traffic in this work, the proximity is evaluated by the ISPs that peers belong to. The locality-aware peer selection strategy leads peers within the same ISP to form high-density clusters. Consequently, the inter-ISP traffic can be restricted. However, peers still maintain a number of connections to the external ISP, which allows them to have the visibility to the available blocks from the outside world [10]. We define the number of persistent external links as η.

- $x_{ik} \geq d_{in} - \eta$, peers viewing channel i in ISP k can sustain playback without generating inter-ISP traffic except for η persistent external links. The estimated inter-ISP traffic rate per peer can be represented as $\frac{\eta r_i}{d_{in}}$.
- $x_{ik} < d_{in} - \eta$, peers viewing channel i in ISP k cannot sustain playback without generating inter-ISP traffic since there are insufficient local candidate neighbors. The estimated inter-ISP traffic rate per peer can be represented as $\frac{(d_{in}-x_{ik})r_i}{d_{in}}$.

We summarize two different scenarios above, and obtain a formulation of the inter-ISP traffic rate generated by channel i in ISP k as:

$$\begin{aligned} T_{ik}^{n-c} &= \frac{\max(\eta, d_{in} - x_{ik})}{d_{in}} x_{ik} r_i \\ &= \frac{x \cdot r_i \cdot p_{isp}(k) \max(\eta, d_{in} - x_{ik})}{H(i+q)^{\alpha} d_{in}} \end{aligned} \tag{7}$$

Remark: We compare Eq. (7) to Eq. (6) to discuss the impact of peer selection strategies on inter-ISP traffic. Two different coefficients, $1 - p_{isp}(k)$ in the locality-unaware strategy and $\frac{\max(\eta, d_{in}-x_{ik})}{d_{in}}$ in the locality-aware strategy, make a difference in these scenarios. The latter coefficient indicates that the resulting traffic in the locality-aware strategy is relevant to practical system design. In a real-world P2P streaming system, the value of d_{in} is set to around 30 [16] and the value of η varies between 5 and 10. For a sufficiently large ISP such as $p_{isp}(k) = 80\%$, $1 - p_{isp}(k) \approx \frac{\max(\eta, d_{in}-x_{ik})}{d_{in}}$, in which the locality-aware/unaware peer selection strategies create a similar amount of the inter-ISP traffic. In addition, when a particular ISP is sufficiently small as $p_{isp}(k) \to 0$, both coefficients have their values close to 1. This demonstrates another scenario that peer selection strategies are not relevant to the inter-ISP traffic. However, the value of $1 - p_{isp}(k)$ is consistently larger than $\frac{\max(\eta, d_{in}-x_{ik})}{d_{in}}$ in most cases, which implies the advantage of locality-aware peer selection.

B. Caching Resource Allocation Mechanisms

In Sec. III-A, we have defined the inter-ISP traffic rate for both ISP-aware and ISP-unaware peer selection strategies. Given the constraint on the cache server storage capacity S_k and the uploading bandwidth capacity U_k, we present our optimization framework to analyze resource allocation mechanisms in collaborative caching. Without loss of generality, we assume that peers in any particular channel are evenly distributed along the channel. We then derive a formulation of the inter-ISP traffic rate for ISP k, taking into account the traffic reduction by deploying cache servers:

$$\begin{aligned} T_k^c &= \sum_{i \in \mathcal{I}} T_{ik}^{n-c} - \textit{traffic cached by ISP servers} \\ &= \sum_{i \in \mathcal{I}} (T_{ik}^{n-c} - \min(u_{ik}, a_{ik} T_{ik}^{n-c})) \end{aligned} \tag{8}$$

With particular focus on the inter-ISP traffic mitigation and relevant constraints, the cache resource allocation problem for ISP k is formulated as:

$$\begin{aligned} \text{Minimize} \quad & T_k^c = \sum_{i \in \mathcal{I}} (T_{ik}^{n-c} - \min(u_{ik}, a_{ik} T_{ik}^{n-c})) \\ \text{Subject to:} \quad & \sum_{i \in \mathcal{I}} a_{ik} f_i \leq S_k \\ & \sum_{i \in \mathcal{I}} u_{ik} \leq U_k \\ & u_{ik} \geq 0, 0 \leq a_{ik} \leq 1 \quad \forall i \in \mathcal{I} \end{aligned} \tag{9}$$

In a practical video streaming system, the uploading bandwidth assigned by cache servers for a particular channel does not exceed the rate of the potential inter-ISP traffic that can be intercepted. The optimization problem is then transformed to:

$$\begin{aligned} \text{Maximize} \quad & \sum_{i \in \mathcal{I}} u_{ik} \\ \text{Subject to:} \quad & \sum_{i \in \mathcal{I}} a_{ik} f_i \leq S_k \\ & \sum_{i \in \mathcal{I}} u_{ik} \leq U_k \\ & 0 \leq u_{ik} \leq a_{ik} T_{ik}^{n-c} \quad \forall i \in \mathcal{I} \\ & 0 \leq a_{ik} \leq 1 \quad \forall i \in \mathcal{I} \end{aligned} \tag{10}$$

Theorem 1: Given a limited server storage capacity S_k, bandwidth capacity U_k and relevant constraints in Problem (10), the upper bounds of inter-ISP traffic mitigation can be achieved by the optimal resource allocation are:

$$a_{ik}^* = \begin{cases} 1, & \text{for } i = 1, ..., z-1; \\ \frac{S_k - \sum_{i=1}^{z-1} f_i}{f_z}, & \text{for } i = z; \\ 0, & \text{for } i = z+1, ..., N_i, \end{cases} \tag{11}$$

where video channels are sorted as $\frac{T_{1k}^{n-c}}{f_1} \geq \frac{T_{2k}^{n-c}}{f_2} \geq ... \geq \frac{T_{N_i k}^{n-c}}{f_{N_i}}$, and $z = \min\{j : \sum_{i=1}^{j} f_i > S_k\}$.

If $\sum_{i \in \mathcal{I}} a_{ik}^* T_{ik}^{n-c} \leq U_k$, then:

$$u_{ik}^* = a_{ik}^* T_{ik}^{n-c} \quad \forall i \in \mathcal{I} \tag{12}$$

If $\sum_{i \in \mathcal{I}} a_{ik}^* T_{ik}^{n-c} > U_k$, then:

$$u_{ik}^* = \frac{a_{ik}^* T_{ik}^{n-c}}{\sum_{i \in \mathcal{I}} a_{ik}^* T_{ik}^{n-c}} U_k \quad \forall i \in \mathcal{I} \tag{13}$$

Proof: We observe from constraints of optimization problem (10) that $\sum_{i \in \mathcal{I}} u_{ik} \leq \sum_{i \in \mathcal{I}} a_{ik} T_{ik}^{n-c}$. Thus, there exists two upper bounds for the objective function such that $\sum_{i \in \mathcal{I}} u_{ik} \leq \min(\sum_{i \in \mathcal{I}} a_{ik} T_{ik}^{n-c}, U_k)$, where a_{ik} satisfy the constraints of problem (10).

We solve the following optimization problem with loose constraints compared to problem (10), which aims to drive the upper bound of $\sum_{i \in \mathcal{I}} a_{ik} T_{ik}^{n-c}$:

$$\begin{aligned} \text{Maximize} \quad & \sum_{i \in \mathcal{I}} a_{ik} T_{ik}^{n-c} \\ \text{Subject to:} \quad & \sum_{i \in \mathcal{I}} a_{ik} f_i \leq S_k \\ & 0 \leq a_{ik} \leq 1 \quad \forall i \in \mathcal{I} \end{aligned} \tag{14}$$

This problem is known as the *continuous knapsack problem* and the solution can be obtained as follows [17]: sort video channels in non-decreasing order with the index $\frac{T_{ik}^{n-c}}{f_i}$, then apply the greedy algorithm to allocate as much storage space as needed for video channels with higher priorities. The optimal results are denoted as a_{ik}^* in Eq. (11).

We can achieve the upper bound of $\sum_{i \in \mathcal{I}} u_{ik}$ as $\sum_{i \in \mathcal{I}} u_{ik} = \min(\sum_{i \in \mathcal{I}} a_{ik}^* T_{ik}^{n-c}, U_k)$, by using the caching bandwidth resource allocation scheme indicated in Eq. (12) and Eq. (13). ■

Remark: Theorem 1 represents the upper bound solution to our optimization framework with respect to the caching resource allocation. It conveys important design guidelines of collaborative caching mechanisms. The optimal strategy requires several system parameters of P2P video streaming systems such as the number of concurrent online users, the dynamic channel popularity, the file size and the streaming rate of channels. Meanwhile, it would be a problem for ISP cache servers to precisely identify the content requests of P2P packets without the help of P2P applications. Therefore, the ISP caching needs to collaborate with P2P applications. In the view of P2P users, the deployed cache servers, in return, benefit these users by reducing end-to-end latencies of P2P packets. Such mitigation of inter-ISP P2P traffic also prevents the risk of P2P traffic being throttled by ISPs. With both arguments, P2P applications and ISPs should closely collaborate with each other to implement the proposed collaborative caching mechanisms.

Algorithm 1 An Optimization-based Collaborative Caching Framework for Inter-ISP Traffic Mitigation

1) The P2P application actively transmits system states, including the population-based channel index i and concurrent users x, to ISP cache servers.
2) Compute the critical performance index $\frac{T_{ik}^{n-c}}{f_i}$, $\forall i \in \mathcal{I}$, allocate the storage and bandwidth as a_{ik}^* and u_{ik}^*.
3) Cache servers intercept video content requests to external ISPs. If the content is stored and the average uploading rate to the channel is smaller than u_{ik}^*, return the requested content.
4) Monitor the dynamics of P2P system states, adjust resource allocation according to Theorem 1.

IV. Improving Caching with ISP Peering Agreements

In Sec. III, we present our inter-ISP traffic optimization framework towards collaborative caching between ISPs and P2P applications. In this section, we further incorporate the real-world collaboration between ISPs in the form of ISP peering agreements into our collaborative model. We analyze how ISP peering affects the effectiveness of our traffic mitigation scheme, and investigate ways to improve collaborative caching with the presence of ISP peering agreements.

A. ISP Peering Agreements

An ISP peering agreement is an interconnection relationship whereby ISPs provide connectivity to each others' transit customers [9]. Traffic along the reciprocal peering connection is free of charge if the traffic balancing condition is satisfied. In this way, ISPs potentially alleviate the costly transit traffic which share a common objective with our proposed caching mechanisms. From the perspectives of collaborative caching, ISP peering leads to two positive outcomes: first, each peer connects with a larger group of traffic-free candidate neighbors; second, it is impractical to store all the P2P content in cache servers of a single ISP when storage resources are constrained, and ISP peering can be beneficial to mitigate such a scalability challenge. ISPs with peering agreements can strategically select P2P content to store and deliver as the cross-boundary traffic becomes free.

The ISP peering relation is a binary relation that satisfies the following properties: (1) Reflexive, each ISP can be regarded as self-peering since no transit traffic exists between peers of the same ISP; (2) Symmetric, ISPs reciprocally provide free connections. We use a symmetric Matrix $\mathbf{E}$ to illustrate the ISP peering relation as follows:

$$E_{ij} = E_{ji} = \begin{cases} 0, & \text{If ISP } i \text{ and } j \text{ have no peering,} \\ 1, & \text{If ISP } i \text{ and } j \text{ have peering.} \end{cases} \quad \forall i, j \in \mathcal{K} \tag{15}$$

B. Impact of ISP Peering

We first turn our attention to the scenario where cache servers do not deliver content to peers of peering ISPs. This implies that only peers in peering ISPs help to mitigate the inter-ISP traffic.

For locality-unaware peer selection, we have:

$$Pr(local = m) = \binom{x_{i(k-p)}}{m}\binom{x_i - x_{i(k-p)}}{d_{in} - m} / \binom{x_i}{d_{in}} \tag{16}$$
$$m \in [\max(0, d_{in} + x_{i(k-p)} - x_i), \min(d_{in}, x_{i(k-p)})]$$

Where $x_{i(k-p)}$ denotes the number of viewers in channel i that belong to ISP k or peering ISPs of ISP k:

$$x_{i(k-p)} = \sum_{k' \in \mathcal{K}} \frac{x E_{kk'} p_{isp}(k')}{H(i+q)^{\alpha}} \tag{17}$$

Correspondingly, the inter-ISP traffic rate caused by video channel i in ISP k can be formulated as:

$$T_{ik}^{n-c} = \frac{x r_i p_{isp}(k)(1 - \sum_{k' \in \mathcal{K}} E_{kk'} p_{isp}(k'))}{H(i+q)^{\alpha}} \tag{18}$$

For locality-aware peer selection, we have:

- $x_{i(k-p)} \geq d_{in} - \eta$, the estimated inter-ISP traffic rate per peer can be represented as $\frac{\eta r_i}{d_{in}}$.
- $x_{i(k-p)} < d_{in} - \eta$, the estimated inter-ISP traffic rate per peer can be represented as $\frac{(d_{in} - x_{i(k-p)}) r_i}{d_{in}}$.

Then, we formulate the inter-ISP traffic rate of channel i in ISP k as:

$$T_{ik}^{n-c} = \frac{x r_i p_{isp}(k) \max(\eta, d_{in} - x_{i(k-p)})}{H(i+q)^{\alpha} d_{in}} \tag{19}$$

An intuitive observation conveys the message that the inter-ISP traffic is reduced for both scenarios owing to the expansion of traffic-free neighbor candidates. However, the impact from different types of peering agreements still remains concealed. We analyze such impact in depth on our collaborative caching framework in Sec. V.

C. Improving Caching with ISP Peering

As a potentially large amount of P2P objects are shared by P2P networks, cache servers can hardly store the entire set of P2P content within a single ISP. This bottleneck will affect the efficient utilization of cache server bandwidth resources. However, the peering relation among ISPs provides an opportunity to combine strategies of distributed cache servers into a global cooperative caching framework. In this scenario, cache servers not only upload P2P content to peers of their own ISPs, but also contribute to the elimination of the inter-ISP traffic of peering ISPs by responding to content requests forwarded by them. This indicates a peering-based *full collaboration* where the challenge is to coordinate the storage and uploading bandwidth assignments of different ISPs.

We introduce the notation $u_{ik'}^{k}$ to denote the bandwidth assigned by ISP k' to ISP k for a particular channel i. We use $u_{ik'}^{k}$ to replace u_{ik} since cache servers are currently allowed to deliver content to users of peering ISPs. The constraint of $u_{ik'}^{k} \leq a_{ik'} T_{ik}^{n-c}$ in Problem (10) is still effective in the bandwidth assignment. In this scenario, the inter-ISP traffic rate generated by channel i in ISP k can be represented as:

$$T_{ik}^{c} = T_{ik}^{n-c} - \min(\bigcup_{k' \in \mathcal{K}} u_{ik'}^{k}, T_{ik}^{n-c}) \tag{20}$$

The operation $\bigcup_{k' \in \mathcal{K}} u_{ik'}^{k}$ has the combined concern of the storage and bandwidth assignment of distributed cache servers. Given an initial condition $a_{ik'} = 1$ for all $k' \in \mathcal{K}$, we have $\bigcup_{k' \in \mathcal{K}} u_{ik'}^{k} = \sum_{k' \in \mathcal{K}} u_{ik'}^{k}$ since any request to channel i can be served in the case of sufficient bandwidth provisioning. Consider that the available videos are mostly stored with integral copies on cache servers from proposed caching mechanisms, we could approximate it as $\bigcup_{k' \in \mathcal{K}} u_{ik'}^{k} \approx \sum_{k' \in \mathcal{K}} u_{ik'}^{k}$.

Now we can formulate the global optimization problem as:

$$\begin{aligned}
\text{Maximize} \quad & \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{I}} \min(\sum_{k' \in \mathcal{K}} u_{ik'}^{k}, T_{ik}^{n-c}) \\
\text{Subject to:} \quad & \sum_{i \in \mathcal{I}} a_{ik} f_i \leq S_k \quad \forall k \in \mathcal{K} \\
& \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{I}} u_{ik'}^{k} \leq U_{k'} \quad \forall k' \in \mathcal{K} \\
& u_{ik'}^{k} = E_{k'k} u_{ik'}^{k} \quad \forall i \in \mathcal{I}, k, k' \in \mathcal{K} \\
& 0 \leq u_{ik'}^{k} \leq a_{ik'} T_{ik}^{n-c} \quad \forall i \in \mathcal{I}, k, k' \in \mathcal{K} \\
& 0 \leq a_{ik} \leq 1 \quad \forall i \in \mathcal{I}, k \in \mathcal{K}
\end{aligned} \tag{21}$$

The solution to problem (21) offers a performance upper bound of collaborative caching mechanisms in mitigating the inter-ISP traffic. We evaluate its benefits under particular peering relations in Sec. V. However, this requires a centralized solution which is inappropriate for implementation in a practical system. With respect to peering agreements, resource allocation constraints and the rationale of using limited a_{ik} to serve maximum $T_{ik'}^{n-c}$, we propose a distributed collaborative caching scheme in Algorithm 2.

Algorithm 2 An ISP Collaboration-based Distributed Caching Framework for Inter-ISP Traffic Mitigation

1) Cache servers announce the surplus bandwidth capacity and the storage allocation a_{ik} to peering ISPs.
2) Upon receiving the surplus bandwidth information from ISP k', ISP k computes the critical performance index $\frac{a_{ik}}{T_{ik}^{n-c}}$, $\forall i \in \mathcal{I}$, and sorts the channels in descending order. For channel i in front of the queue, if $a_{ik'} = 1$ and channel i is not forwarding to servers of peering ISPs, bandwidth request for channel i is sent to k'.
3) Upon receiving the request from ISP k', ISP k allocates the bandwidth and confirms to k'.
4) Upon receiving the confirmation from ISP k', ISP k evicts contents being confirmed by k', reallocates storage and bandwidth to channels such that $a_{ik'} = 0$. It then broadcasts the surplus information to peering ISPs.

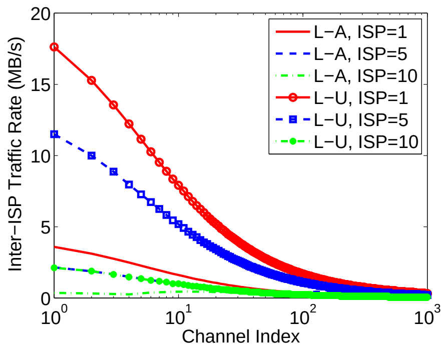

Fig. 1. The inter-ISP traffic rate of individual ISPs vs. the channel index, under different settings of peer selection strategies.

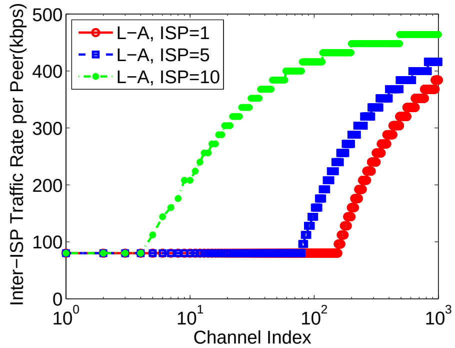

Fig. 2. The per peer inter-ISP traffic rate of individual ISPs vs. the channel index, under the locality-aware peer selection strategy.

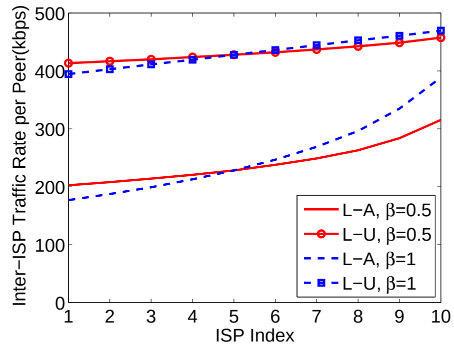

Fig. 3. The per peer inter-ISP traffic rate vs. ISP index, under different settings of ISP population parameter β and peer selection strategy.

V. Performance Evaluation

In this section, we carry out a series of numerical analyses using parameters driven from the real-world trace data to evaluate various perspectives of our collaborative caching mechanisms.

A. Trace-Driven Analyses

To make our performance evaluation more practical and comprehensive, we use a number of statistical results from measurement studies on a real-world P2P streaming system, UUSee [18]. In this video steaming system, the total number of video channels is 993. The channel with the sorted index of 100 has around 100 concurrent users at peak time. The number of concurrent users over the system is approximately 100000. We determine a set of parameters that approximately fit the curve of peak time users for video channels: the value of α in the Zipf-Mandelbrot model is assigned to 0.78 and the value of q is assigned to 4. We also set $d_{in} = 30$ and $\eta = 5$ based on our practical experience in P2P video systems.

B. Evaluation of Inter-ISP Traffic Pattern

As indicated in Sec. III-A, the inter-ISP traffic of P2P networks is relevant to both P2P content popularity and ISP popularity. The L-A (locality-aware) and L-U (locality-unaware) peer selection strategies also affect the resulting traffic. Here we quantitatively demonstrate how these factors contribute to the inter-ISP traffic pattern we summarized.

Fig. 1 depicts the inter-ISP traffic rate generated by video channels. There are in total $N_k = 10$ ISPs in the system. Results for ISPs with indices $1, 5, 10$ are shown in the figure, which represent popular, moderately popular, and unpopular ISPs, respectively. It can be observed that L-A significantly outperforms L-U by mitigating a substantial amount of the inter-ISP traffic, especially for those popular video channels involving a large user population. This is consistent with previous analytical studies on the benefits of locality-aware peer selection. Given a total number of users, an ISP with a larger user population always generates more inter-ISP traffic compared to that of an unpopular ISP, especially in the L-U scenario. This is because those larger ISPs can have a large number of concurrent users in channels, which generate a considerable amount of traffic in total. Such inefficiency is significantly improved when applying locality-aware peer selection in the system, as larger ISPs also provide better potential for local content sharing.

We then evaluate the per peer inter-ISP traffic rate of different channels in Fig. 2 while focusing on locality-aware peer selection scenario. The inter-ISP traffic rate is constantly around 80 kbps for the most popular channels in ISPs of different scales. The reason is that both ISPs have sufficient local candidate neighbors to provide for a popular channel. Therefore, η external links required by the P2P system lead to inter-ISP traffic. The per peer inter-ISP traffic rate increases remarkably as the channel popularity decreases. The rationale is that it is more difficult for peers in unpopular channels to find a sufficient number of partners in a single ISP to satisfy the streaming requirement. Hence, they have to seek more bandwidth resources from peers in other ISPs. In particular, the effect of locality-aware peer selection for mitigating inter-ISP traffic will be degraded when there is a relatively small number of participating peers. This quantitatively confirms the necessity for ISP collaboration mechanisms, without which there would be tremendous costs incurred by inter-ISP traffic.

Fig. 3 demonstrates how the balance of ISP user population affects the per peer inter-ISP traffic rate. We summarize the overall traffic rate of all existing channels. The unbalance of ISP user population increases when the value of β varies from 0.5 to 1. As can be observed from the figure, smaller ISPs are likely to generate a larger volume of inter-ISP traffic when $\beta = 1$, as peers in smaller ISPs can hardly find a sufficient number of local partners due to the unbalanced ISP user population. However, the increase of concurrent users in a larger ISP can reversely help decrease the inter-ISP traffic for its users. Another interesting observation from Fig. 3 shows that the inter-ISP traffic rate for the smallest ISP has a large variation as 30% increase in the L-A peer selection, which indicates that it is more sensitive to the balance of ISP user population. The rationale is that most channels of the smaller ISP are not "self-sufficient." Therefore, the variation of the inter-ISP traffic is proportional to the variation of available local candidates. Comparatively speaking, larger ISPs are slow with the variation of β due to a large number of channels already being "self-sufficient."

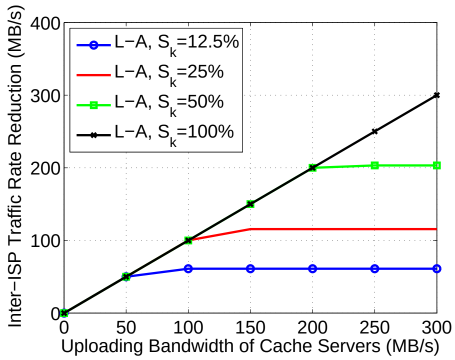

Fig. 4. The inter-ISP traffic rate reduction of ISP of index 5 vs. uploading bandwidth of cache servers, under different settings of storage capacity of cache servers.

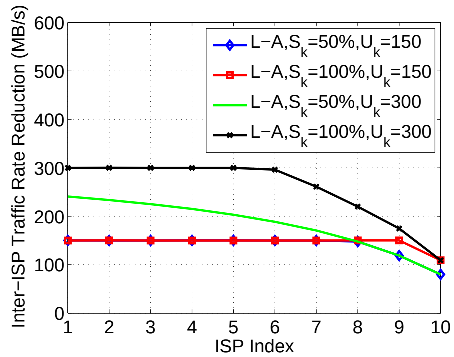

Fig. 5. The inter-ISP traffic rate reduction vs. ISP index, under different settings of uploading bandwidth and storage capacity of cache servers.

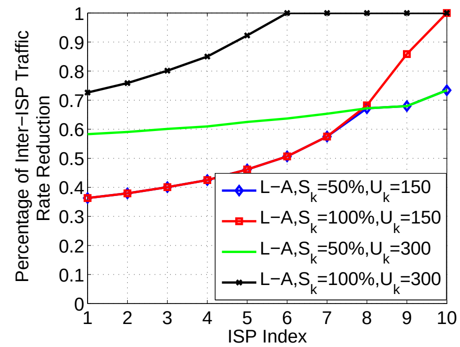

Fig. 6. The percentage of inter-ISP traffic rate reduction vs. ISP index, under different settings of uploading bandwidth and storage capacity of cache servers.

C. Evaluation of Collaborative Caching Mechanisms

The aforementioned analysis on the inter-ISP traffic pattern validates the need for our proposed collaborative caching mechanisms. Fig. 4 clearly demonstrates the potential benefits of collaborative caching for reducing the inter-ISP traffic. In particular, with more cache storage capacity provisioned, the uploading bandwidth contribution from cache servers can save more inter-ISP traffic. However, such traffic mitigation is very sensitive to the cache storage capacity. If the cache storage capacity is insufficient, with $S_k = 12.5\%, 25\%$ or 50% of the total storage requirement, the bandwidth capacity of cache servers cannot be effectively utilized. This gives us the motivation of finding a way to break the bottleneck brought by insufficient cache storage provisioning. Fortunately, the ISP peering-based full collaboration can resolve the challenge, which will be evaluated in Sec. V-D.

We further evaluate the case with all existing ISPs of different scales, as indicated in Fig. 5. When the storage capacity $S_k = 50\%$ and the bandwidth capacity $U_k = 150$, the reduction of inter-ISP traffic rate is consistently 150 for the 8 most popular ISPs. This indicates that the bandwidth capacity is the bottleneck of resource utilization, which can be proved by the scenario that raises the storage capacity S_k to 100%. However, when U_k is set to 300 while S_k remains 50%, the reduction of the traffic rate is gradually decreased along the curve. The reason behind this phenomenon is that the insufficient storage allocation restricts the bandwidth utilization. Meanwhile, the gradually falling channel popularity leads to this declining curve. The scenario becomes more interesting when we set $S_k = 100\%$ and $U_k = 300$. The consistent rate reduction still shows insufficient provisioning of the bandwidth resource for the 6 popular ISPs, while the curve declines starting from ISP=7. Although the main reason for the phenomenon is still the falling channel popularity, the inter-ISP traffic rate is already mitigated, shown by the flattened header of the curve. Fig. 6 further convinces our analysis by plotting the percentage of the inter-ISP traffic rate mitigation for different ISPs. Both Fig. 5 and Fig. 6 implies that a balanced allocation of storage and bandwidth capacity is critical in a collaborative caching design. From our theoretical caching framework, the balance point represents the scenario that $\sum_{i\in\mathcal{I}} a^*_{ik} T^{n-c}_{ik} = U_k$.

D. Evaluation of ISP Peering Agreements

To evaluate the impact of ISP peering agreements on the performance of our collaborative caching mechanisms, we study three representative peering scenarios with $N_k = 10$ as follows:

1) Scenario 1: $E(i, i+1) = 1, i \in \{1, 3, 5, 7, 9\}$;
2) Scenario 2: $E(i, i+N_k/2) = 1, i \in \{1, 2, 3, 4, 5\}$;
3) Scenario 3: $E(i, N_k+1-i) = 1, i \in \{1, 2, 3, 4, 5\}$.

Peering scenario 1 represents an extreme case since the ISP population becomes more unbalanced after peering. Peering scenario 3 denotes another extreme where the population becomes relatively balanced. Peering scenario 2 still holds the property of the original ISP population. Fig. 7 depicts the per peer inter-ISP traffic rate to exhibit the entire peering and non-peering design space. Clearly, ISP peering can help reduce costs due to the inter-ISP traffic, since peers can benefit from neighbors either in their own ISP network or other peering ISP networks. Meanwhile, the figure also demonstrates that our analysis can characterize the interesting ISP economics. More specifically, balanced peering results in a balanced inter-ISP traffic reduction across different ISPs, while unbalanced peering has a relatively larger deviation of traffic reduction. This suggests that a collaboration among all ISPs can benefit them in a fair manner, so as to encourage smaller ISPs to participate in such a peering relation. In contrast, larger ISPs may not prefer this, as their revenue brought by the inter-ISP traffic initialized by smaller ISPs could potentially decrease. This reveals a conflict of interest between larger and smaller ISPs in peering strategies when collaborative caching is applied.

Fig. 8 represents the correlation between different peering scenarios and the percentage of the inter-ISP traffic mitigation driven by our collaborative caching mechanisms. It further confirms that ISP peering can achieve more inter-ISP traffic reduction through the proposed framework. The five most unpopular ISPs have nearly 100% mitigation of the inter-ISP traffic with conventional settings as $U_k = 200$ and $S_k = 100\%$. More interestingly, for larger ISPs, the scenario of

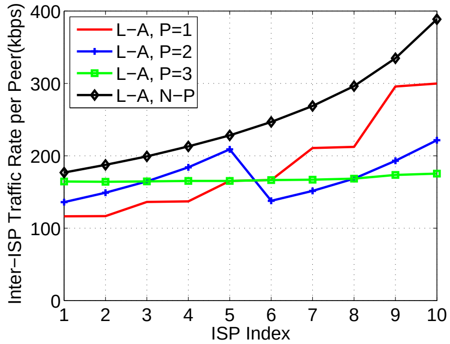

Fig. 7. The per peer inter-ISP traffic rate vs ISP index, under different settings of ISP peering scenarios.

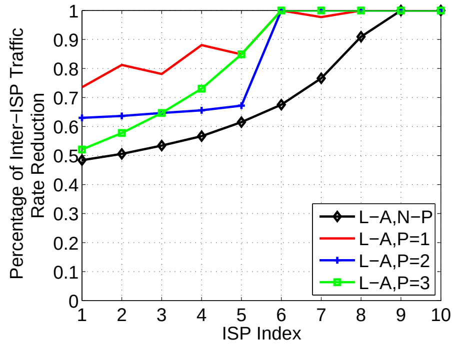

Fig. 8. The percentage of inter-ISP traffic rate reduction vs ISP index, under different settings of ISP peering scenarios.

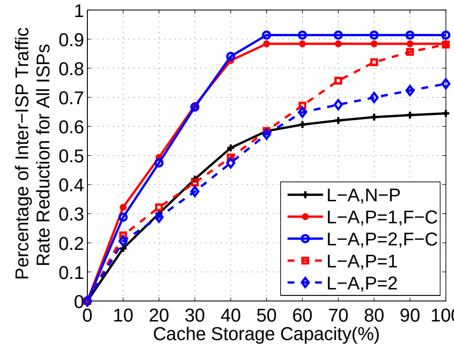

Fig. 9. The percentage of inter-ISP traffic rate reduction for all ISPs vs cache storage capacity, under different ISP peering scenario settings.

more unbalanced peering brings more benefit to them, whereas both peering scenarios would be useful as the scale of ISP decreases. This demonstrates that the proposed collaborative caching can resolve the conflict of interest between ISPs of different scales in peering strategies, by allocating adequate caching resources to smaller ISPs.

Combined with the above discussions on different peering and non-peering scenarios, Fig. 9 summarizes how collaborative caching mechanisms behave towards storage capacity constraints. The notation "F-C" denotes the full collaboration as discussed in Sec. IV-C. Not surprisingly, full collaboration outperforms the remaining strategies especially when the cache storage capacity is insufficient. Cache servers help deliver the content of peering ISPs, which efficiently utilize the bandwidth capacity when the storage is constrained. Another observation shows that, peering with collaboration between peers performs similarly to the case of non-peering when the storage capacity is insufficient. The underlying reason is that most channels stored in cache servers are popular channels in the case of insufficient storage provisioning. It makes no difference whether it is a peering or non-peering relation if the locality-aware peer selection is applied. Fig. 9 implies a substantial guideline to the practical collaborative caching design: the full collaboration between cache servers of peering ISPs is preferred to a lack of cache storage; however, native peering relations between ISPs are more desirable when the storage is relatively sufficient, as they provide satisfactory inter-ISP traffic mitigation with much less operational costs.

VI. Conclusion

The urgent requirement for mitigating an enormous amount of inter-ISP P2P traffic motivates us to investigate collaborative caching mechanisms in this paper. Based on the proposed inter-ISP traffic model, we have developed a cache resource allocation framework with respect to both resource constraints and ISP peering relations. The trace-driven analysis shows that inter-ISP P2P traffic can be efficiently removed with fine tuned system parameters and strategies.

Our analytical framework not only provides important properties of inter-ISP traffic, but also conveys important messages to caching system design. For example, we put forward guidelines for cache storage and bandwidth allocation within our design space. We also unveil how collaborative caching can be improved in the context of ISP peering. Other interesting topics can be explored in future work, including benefits to the user experience, which reflect additional advantages of our collaborative caching mechanisms.

References

[1] J. Erman, A. Mahanti, M. Arlitt, and C. Williamson, "Identifying and discriminating between web and peer-to-peer traffic in the network core," in *Proc. ACM WWW*, May 2007.

[2] "List of ISPs throttling P2P traffic," http://www.p2pon.com/guides/list-of-internet-service-providers-that-throttle-p2p-traffic/.

[3] T. Karagiannis, P. Rodriguez, and K. Papagiannaki, "Should Internet service providers fear peer-assisted content distribution?" in *Proc. ACM IMC*, Oct. 2005.

[4] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4P: provider portal for applications," in *Proc. ACM SIGCOMM*, Aug. 2008.

[5] S. Ren, E. Tan, T. Luo, S. Chen, L. Guo, and X. Z. M. Corporation, "TopBT: A topology-aware and infrastructure-independent BitTorrent client," in *Proc. IEEE INFOCOM*, Mar. 2010.

[6] M. Piatek, H. V. Madhyastha, J. P. John, A. Krishnamurthy, and T. Anderson, "Pitfalls for ISP-friendly P2P design," in *Proc. ACM HotNets*, Oct. 2009.

[7] P. Cao, J. Zhang, and K. Beach, "Active cache: caching dynamic contents on the web," in *Proc. IFIP Middleware*, Sep. 1998.

[8] G. Shen, Y. Wang, Y. Xiong, B. Y. Zhao, and Z. li Zhang, "HPTP: Relieving the tension between ISPs and P2P," in *Proc. USENIX IPTPS*, 2007.

[9] W. B. Norton, "Internet service providers and peering," in *Proc. NANOG*, Jun. 2000.

[10] R. Bindal, P. Cao, W. Chan, J. Medval, G. Suwala, T. Bates, and A. Zhang, "Improving traffic locality in bittorrent via biased neighbor selection," in *Proc. IEEE ICDCS*, Jul. 2006.

[11] Y. Liu, L. Guo, F. Li, and S. Chen, "A case study of traffic locality in Internet P2P live streaming systems," in *Proc. IEEE ICDCS*, Jun. 2009.

[12] "Pplive," http://www.pptv.com/.

[13] "Cachelogic," www.velocix.com/.

[14] M. Hefeeda and O. Saleh, "Traffic modeling and proportional partial caching for peer-to-peer systems," *IEEE/ACM Trans. Netw.*, vol. 16, no. 6, 2008.

[15] G. Dan, "Cooperative caching and relaying strategies for peer-to-peer content delivery," in *Proc. USENIX IPTPS*, Feb. 2008.

[16] C. Wu, B. Li, and S. Zhao, "Exploring large-scale peer-to-peer live streaming topologies," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 4, no. 3, 2008.

[17] V. Cerone and F. Croce, "A continuous knapsack problem formulation for the robustness analysis of a polytope of polynomials," in *Proc. American Control Conference*, 1995.

[18] Z. Liu, C. Wu, B. Li, and S. Zhao, "UUSee: Large-scale operational on-demand streaming with random network coding," in *Proc. IEEE INFOCOM*, Mar. 2010.