

Northumbria Research Link

Citation: Zhang, Xu, Wang, Ning, Cao, Yue, Peng, Linyu and Meng, Haining (2018) A Stochastic Analytical Modelling Framework on ISP-P2P Collaborations in Multi-domain Environments. IEEE Systems Journal, 12 (3). pp. 2320-2331. ISSN 1932-8184

Published by: IEEE

URL: <http://dx.doi.org/10.1109/JSYST.2017.2725914>
<<http://dx.doi.org/10.1109/JSYST.2017.2725914>>

This version was downloaded from Northumbria Research Link:
<http://nrl.northumbria.ac.uk/id/eprint/31261/>

Northumbria University has developed Northumbria Research Link (NRL) to enable users to access the University's research output. Copyright © and moral rights for items on NRL are retained by the individual author(s) and/or other copyright owners. Single copies of full items can be reproduced, displayed or performed, and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided the authors, title and full bibliographic details are given, as well as a hyperlink and/or URL to the original metadata page. The content must not be changed in any way. Full items must not be sold commercially in any format or medium without formal permission of the copyright holder. The full policy is available online: <http://nrl.northumbria.ac.uk/policies.html>

This document may differ from the final, published version of the research and has been made available online in accordance with publisher policies. To read and/or cite from the published version of the research, please visit the publisher's website (a subscription may be required.)

A Stochastic Analytical Modelling Framework on ISP-P2P Collaborations in Multi-domain Environments

Xu Zhang, Ning Wang, *IEEE Member*, Yue Cao, *IEEE Member*, Linyu Peng and Haining Meng

Abstract—Cooperation between peer-to-peer (P2P) overlays and underlying networks has been proposed as an effective approach to improve the efficiency of both the applications and the underlying networks. However, fundamental characteristics with respect to ISP business relationships and inter-ISP routing information are not sufficiently investigated in the context of collaborative ISP-P2P paradigms in multi-domain environments. In this paper, we focus on such issues and develop an analytical modelling framework for analysing optimized inter-domain peer selection schemes concerning ISP policies, with the main purpose of mitigating cross-ISP traffic and enhancing service quality of end users. In addition, we introduce an advanced hybrid scheme for peer selections based on the proposed analytical theory framework, in accordance with practical network scenarios, wherein cooperative and non-cooperative behaviours coexisting. Numerical results show that the proposed scheme incorporating ISP policies is able to achieve desirable network efficiency as well as great service quality for P2P users. Our analytical modelling framework can be used as a guide for analysing and evaluating future network-aware P2P peer selection paradigms in general multi-domain scenarios.

Index Terms—P2P systems, peer selection algorithms, analytical modelling, Markov chain.

I. INTRODUCTION

THE Internet is composed of thousands of Internet Service Providers (ISPs) networks known as domains or autonomous systems (ASes). In general, there are two main types of relationships between neighbouring ISPs: 1) *customer-provider*, where one ISP pays the other to carry its traffic, and 2) *free peering*, where two ISPs freely exchange traffic between each other without specific payment, as roughly equal volumes of traffic flow between their networks. According to common inter-domain routing policies that are driven by business relationships between ASes, customer routes are preferred over routes via peering and provider routes [1]. Only local traffic, which is the traffic between the peering ISPs and their respective customer ISPs, can be exchanged over the peering links [2]. The overlay routing of peer-to-peer (P2P) traffic in the Internet, however, may indirectly

violate the business relationships of ISPs, as observed in [3]. In addition, carrying P2P traffic through application overlays may introduce unexpected extra economic cost to ISPs since P2P flows can travel across transit links to reach the demanded data, even if the content objects can be acquired in local ISPs. Such a scenario can impose unnecessary extra traffic volumes across domain boundaries which usually accounts for the Internet bottlenecks [4]. And the tussle is becoming even more challenging with the proliferation of P2P-based applications, e.g., P2P-style video streaming [5], and emerging social media applications such as Facebook/Pipe [6] and WebRTC [7].

Coordination between ISPs and P2P networks [4]-[10] has been proposed as an efficient approach to deal with the tussle of P2P overlay and ISP underlay interaction, by reducing substantial cross-ISP traffic while retaining desired P2P users Quality of Experiences (QoE). The basic idea is to provide a list of optimised peer candidates in proximity to each client peer by taking into account the context information provided by overlay-underlay collaboration. Such an approach is often known as the *locality-aware strategies* (or *non-cooperative strategy*). A collaboration entity located inside each ISP enables such peer selection ranking procedure, by collecting relevant context information of the local underlying network topology for locality-aware peer selection operations, e.g. the ALTO framework proposed in IETF [11]. However, under the traditional non-cooperative peering strategy, external peers located in remote domains are selected without distinguishing between inter-domain paths regarding the diversity in business relationships among ISPs. A few works have been proposed suggesting that ISP business relationships should be taken into consideration in order to encompass the economic benefits of ISPs [12]-[14], referred to as the *ISP policy-aware strategy* (or *cooperative strategy*).

It is worth mentioning that, while these approaches can effectively mitigate ISP costs among different inter-domain links, they are barely based on the hypothesis of an ideal all-cooperative environment, and there is still limited understanding of their performance on larger-scale collaborations across multiple autonomous ISP networks from an theoretical standpoint. On the one hand, it is difficult to practically enable such ideal cooperative behaviours for all ISPs, since some ISPs may not be willing to participate due to various reasons, such as different operational objectives, or simply privacy issues. In this case, a non-cooperative peering strategy is more applicable. On the other hand, even if there are incentives for such collaboration, potential risks can still exist that P2P traffic

Xu Zhang and Haining Meng are with the Xi'an University of Technology, Xi'an 710048, China, e-mail: {zhangxu,hnmeng}@xaut.edu.cn

Ning Wang is with the Institute for Communication Systems (IC-S), University of Surrey, Guildford, Surrey, GU2 7XH, UK, e-mail: n.wang@surrey.ac.uk

Yue Cao (corresponding author) is with the School of Computing, Northumbria University, Newcastle upon Tyne, NE2 1XE, UK, e-mail: yue.cao@northumbria.ac.uk

Linyu Peng is with the Department of Applied Mechanics and Aerospace Engineering, Waseda University, Shinjuku-ku, Tokyo, 169-8050, Japan, e-mail: L.Peng@aoni.waseda.jp

can be centralized over a small number of inter-ISP links under the ISP business promotion scenario, especially in the case of unlocalizable torrents. As such, costly transit P2P traffic can be alternatively redirected to free peering links in order to maintain ISPs profits.

While effective, the free agreement between ISPs can be challenged due to non-reciprocal benefits received regarding unbalanced traffic exchanged. In addition, congestions on critical inter-domain links can be incurred and in turn, users perceived service quality can be deteriorated as well [5] [15]. Concerning benefits and potential issues of ISP business relationships for peer selection across domains, we propose in this paper *a combination of non-cooperative and cooperative peering scheme* that can achieve the most efficiency for both network providers and P2P users. Given the increasing complexity of the Internet topology due to coexistence of cooperative and non-cooperative behaviours among different ISPs, peering and sibling operations, optimised peer selections in the inter-domain scenario have become more and more challenging for performance enhancement on both the service side and the network side.

In this paper, we aim to systematically address these aforementioned research issues by proposing a theoretical framework to provide comprehensive and accurate analysis on the following important research question: *How should peer selection procedure operate in a multi-domain scenario, with awareness of both ISP preferences and P2P users capacities diversity?* Our objective is to analytically quantify the P2P traffic optimization strategies across multiple autonomous domains, in order to help to understand the fundamental design criteria of collaborative ISP-P2P mechanisms in the research community.

The technical contributions from this work can be summarised as follows.

1) This paper develops a generic analytical modelling framework for optimized peer selection designs in a multi-domain network environments, based on which corresponding metrics are derived, regarding the network efficiency as well as utilities from both P2P systems and ISPs perspectives.

2) Based on the stochastic model, an advanced hybrid peer selection scheme is proposed in accordance with practical network scenarios. Numerical results show that the proposed strategy enables desirable network efficiency as well as great alleviation of P2P traffic load over critical inter-domain links.

3) Simple closed form boundary values are derived in order to guarantee desirable utilities for both P2P users and ISPs.

The rest of the paper is organized as follows. Related works are discussed in Section II. In Section III we develop the stochastic modelling framework and present our proposed solution. In section IV the most concerned performance attributes are derived. We then present extensive empirical results in Section V, followed by Section VI that concludes the paper.

II. RELATED WORK

The authors of [16] introduced the notion of peering localization in the context of BitTorrent. Based on payload packet traces and tracker-based logs, their simulation based results

showed that locality-aware solutions are able to significantly alleviate the induced cost at the ISPs, while providing an efficient performance for end users. The authors of [17] summarized specific interaction patterns between ISP and P2P systems, and concludes that both network operators and P2P applications could benefit from exchanging information with each other.

The schemes proposed in [4]-[10] are the representative paradigms to offer alternatives of locality-aware peer selection to improve ISPs efficiency and P2P systems performance by cooperation between the two layers. Simulation-based analysis is mainly used to verify that cooperation between ISPs and overlay systems can reduce cross-ISP traffic significantly while maintaining desired download experiences for users. This is typically achieved by introducing an entity coordinating between the P2P overlay and the underlying network, such as an oracle [4], [8] or by utilizing existing CDN [9] information. A solution is introduced in [10] to build an infrastructure-independent system to enable topology-aware BitTorrent Client, with an emphasis on downloading time and traffic reduction. However, most of these works mainly promote peer selections based on the measurement of AS-hops or latency, etc., without differentiating between individual domains regarding various business requirements. Relevant analytical works such as [18] mainly explore the impact of P2P traffic on the ISP business benefits, while [19] modelling the tussle between ISP and P2P systems. However, how peer selections can be evaluated in the scope of multi-domain has not yet been well examined.

The authors of [20] discussed the pitfalls for an ISP-friendly locality policy, and three main issues are discussed, with respect to the limited effect on the user side, the degradation on the P2P systems robustness, and conflicting interests between different ISPs. The first two issues are mostly addressed by works such as [21], which proposed a refined locality-aware peer selection to divide ASes into groups based on different swarms, in order to maintain fairness among P2P users in terms of balanced uploading and downloading capacities. Regarding the last issue, however, it is still not well understood.

P2P caches are deployed by many ISPs to reduce transit traffic through storing popular contents at local ISP [22]. However, analysis [23] shows that caching can lead to increased transit traffic in certain scenarios. In order to resolve this issue, [24] proposed a cache-to-cache scheme to enable collaboration between caches deployed by peering ISPs, which have shown the effectiveness of considering ISP business relationship into the concern of P2P traffic localization.

A few works proposed recently suggesting that peers in remote autonomous network systems should be ranked based on diversity ISP business requirements [12]-[14],[25]-[27]. While concerning BGP routing policies in peer selections, these works are mainly developed on the basis of a simple assumption, i.e., a fully cooperative scenario. However, congestion risks could potentially exist over critical inter-domain links under these proposals. On the other hand, ISPs are generally reluctant to collaborate due to privacy concerns. This would become a huge burden for the underlying network with P2P traffic constantly growing that accounts for a significant volume [3]

in the Internet.

This paper is an extension of our previous work [28], in which we show by analytical modelling that a hybrid peer selection can achieve enhanced performance for both ISPs and P2P users. However, in [28] a simple single-homed ISPs scenario was considered, while in practice asymmetrical ISP routing issues exist due to multi-homed ISPs, which are taken into account in this paper. Additionally, considering the importance and the necessity of maintaining the settlement-free relationship between peering ISPs, a boundary condition regarding the P2P traffic exchanged is analyzed. For comparison purpose, an enhanced locality-aware peer selection is considered and modelled. We also discuss in this paper the implementation issues of the proposed hybrid peer selection strategy.

III. A STOCHASTIC ANALYTICAL MODEL

Figure 1 illustrates a simple scenario consisting of multiple ISP networks. We assume that each ISP operates one single Autonomous System (AS) or domain, and hence we will use ISP network, domain and AS interchangeably in the rest of the paper. In Fig 1(a), ISP_3 is a transit provider of its two customers ISP_1 and ISP_2 . ISP_1 and ISP_2 establish a free peering link between each other to reduce transit fees through ISP_3 . Following the common definition [1], customer ISPs need to pay upstream ISPs for Internet access. In Fig 1(b), ISP_2 is multi-homed to transit service providers ISP_1 and ISP_3 at the same time. According to common practice in BGP routing policies, an ISP prefers the route learnt from its customer ISPs rather than that from its provider ISPs or peering ISPs if both are available.

However, ISPs policy with respect to business relationships can be violated by the prevalence of P2P applications. For instance in Fig 1(a), we assume that both user a and user b , subscribing to ISP_1 and ISP_2 , respectively, need the same content object available in ISP_3 . Due to the nature of P2P systems, the two users can provide each other the chunk of data fetched from ISP_3 directly instead of from the original source. From the P2P application layers view, the content object could traverse along the path ($ISP_3 - ISP_2 - ISP_1$ or $ISP_3 - ISP_1 - ISP_2$), which is conflicting with traditional ISPs polices that only local traffic can be exchanged via peering links. In addition, due to the agnostic of inter-ISP business relationships under the conventional locality-aware strategy, ISPs revenue can be suboptimal. Take Fig. 1(a) as an example again, from P2P users subscribing to ISP_1 viewpoint, candidate peers residing in ISP_2 or ISP_3 are treated equally based on the conventional locality-aware peer selection if measured by AS hops. And yet, costly transit fees would be incurred onto ISP_1 if peers in the provider domain ISP_3 are selected instead of from free peering ISP_2 .

Nevertheless, even if ISP business relationships are taken into account for peer selection in the inter-domain environment, issues may still exist especially in the case for multi-homed ISP networks.

Take Fig. 1(b) as an example, assuming a content object held by peers in customer ISP_2 (multi-homed to ISP_3) is




Fig. 1. Two Multi-domain Network Scenario

demand by users subscribing to ISP_1 . From the economic consideration of ISP_1 , it desires to obtain the object via the direct transit link from ISP_2 . However, the traffic flow generated from the ISP_2 does not necessarily follow the direct link from ISP_2 to ISP_1 . Instead, it may traverses through the alternative transit link (ISP_2 to ISP_3) and then through the peering link to reach ISP_1 . The rational is mainly due to the outbound route preferences of the multi-homed ISP_2 , e.g., cheaper transit cost via ISP_3 . Such a BGP path can be enforced by ISP_2 through setting of higher BGP local preference values. In this case, the content object effectively follows the path ($ISP_2 - ISP_3 - ISP_1$), rather than directly through the path ($ISP_2 - ISP_1$), which is “legal” under the traditional free peering relationship between ISP_1 and ISP_3 [2]. From this example, we can see that ISP_1 may suffer from potential revenue loss as the P2P traffic is injected via a peering link rather than a customer-to-provider one. In addition, the end-to-end delivery path is also suboptimal due to higher AS path length.

Therefore, it seems to be promising to take into account of ISP business relationships as well as inter-ISP routing preferences into peer selections. In practice, however, some ISPs would be reluctant to get involved into cooperation with other ISPs, such as stub ISPs with no customer domains, where they can only benefit from their free peering domains, as compared to costly transit traffic. On the other hand, peering ISPs may have incentives to tear up the free agreement due to non-reciprocal benefits received regarding unbalanced traffic exchanged [3][15]. Because of the risks, stub ISPs are more willingly to localize the P2P traffic within its own network if sufficient peers identified locally.

As elaborated previously, sub-optimality can still occur considering individual ISPs business objectives such as paying/gaining transition costs in carrying P2P traffic, let alone the additional consideration on the service performance concerned on the end user side. We thus propose a *hybrid peer selection mechanism*, which allows for *incorporating both non-cooperative and cooperative peering behaviours* to address the above tussles. With the hybrid peer selection procedure, users’ queries can be processed following the procedure as below.

(i) Queries are prioritized to be processed by utilizing pure cooperation-based (ISP-policy aware) peering strategies.

This certainly requires necessary information dissemination from the underlying ISP networks to the application-layer P2P. Given the availability of existing ISP-P2P collaboration paradigms, such as ALTO [11], such information can be certainly included, even though in this paper we will not

TABLE I
VARIABLE NOTATIONS

| Variable | Description |
|-------------|---|
| x | system state, $x \in \{0, L, P, O\}$ |
| λ_x | peer departure rate at state (x) |
| μ_x | object retrieval rate at state (x) |
| n_x | the number of peers participated in transferring objects in state (x) |
| E_x | mean uplink bandwidth of P2P users (Mbps) in state (x) |
| γ | the fraction of on-line peers |
| N | the number of peers sharing a common content object in a network |
| q | average request rate of a user in one request session (Mbps) |
| $P_i(t)$ | transient probability of the system process in state (i) at time t |
| c | fixed price that a user subscribing to the ISP has to pay |
| \bar{B} | average inter-domain bandwidth utilized by P2P traffic |
| C^d | transit cost for each unit of inter-domain bandwidth paid by a customer ISP |

specify how this will be practically realised.

(ii) If critical inter-domain links are highly congested, or the cooperative entities fail to operate their functionalities, queries are served under non-cooperative peer selections. Such strategy enables rough load balancing as well as desirable users service quality.

Intuitively, procedure (i) can be further extended to four processes: a) Queries are preferred to be served *inside* the local ISP network, followed by candidate peers b) from its *customer ISPs*, and peers from c) *peering ISPs*. And finally if still insufficient, d) queries are resolved at *provider ISPs* or even *further in the Internet which can be only reached via the provider ISP network*. Such a peer selection strategy is consistent with the current BGP routing policies driven by ISPs business relationships¹.

Noticing that the participation of P2P users is arbitrary and quite dynamic, peer behaviours (peer leaves/joins) in a session can be modelled perfectly well by stochastic methods [29]. In the following we develop a series of continuous-time Markov chain (CTMC) to analyse in a holistic way towards various peer selection strategies across multiple domains. Before that we make the following assumption regarding the time durations for peer selections.

Assumption 1: The mean time for peers participating in transferring desired objects is independently and exponentially distributed, with mean λ^{-1} , where λ is the rate of departure [29]. Similarly the mean time for peers accessing objects is independent and exponentially distributed, with mean μ^{-1} , where μ is the rate of object retrieval [29] (see Table I for notation).

Then the peer selection procedure can be described with a CTMC on the state space $\mathcal{L} = \{0, 1, 2, \dots\}$. Each state corresponds to one target object delivering environment x , $x \in \{0, L, P, O\}$, with respect to customer ISPs (L), peering ISPs (P) or provider ISPs (O), respectively. State (0) refers to the local domain. State transition between state (i), ($i \in \mathcal{L}$) and state (j), ($j \in \mathcal{L}$) is triggered with rate λ or μ , respectively, corresponding to either a content searching action or

a successful content retrieval. From an analytical point, the question regarding the peer selection across domains becomes: *In what circumstances can state transition be triggered and how?* In order to solve this question, we have the following definitions.

Definition 1: We denote by $P_i(t)$ the probability of the peering process in system state (i) at time t , and the transition rate from state (i) to state (j) is denoted as R_{ij} . Therefore, according to Markov process theories, we have the following differential equation [30].

$$\frac{dP_i(t)}{dt} = - \sum_{j \neq i} R_{ij} P_i(t) + \sum_{j \neq i} R_{ji} P_j(t) \quad (1)$$

Then according to Eq. (1), a set of differential equations can be acquired in a matrix format as $\frac{d\mathbf{P}(t)}{dt} = \mathbf{P}(t)\mathbf{Q}$, where vector $\mathbf{P}(t) = (\dots, P_i(t), \dots)$, $i \in \mathcal{L}$, and \mathbf{Q} is the generator matrix defined as $\mathbf{Q} = (R_{ij})$, $i, j \in \mathcal{L}$, $i \neq j$.

P2P connection sessions are heterogeneous due to different users capabilities (e.g., first-mile uplink bandwidths, hardware resources and on-line time, etc.). For simplicity, we use an expected uplink bandwidth of P2P users in the following to incorporate the impact of heterogeneous uplink on the peer selection.

Definition 2: We denote by z the number of categories regarding users uplink bandwidths, and then the expected uplink bandwidth for P2P users in state (x) ($x \in \{0, L, P, O\}$) can be described as below.

$$E_x[\eta] = \sum_{i=1}^z \sigma_i \eta_i \quad (2)$$

where σ_i is the percentage of the i th classification of users uplink bandwidth capacity. η_i specifies the i th classification of the uplink capacity.

Based on the above assumption and definitions, we can now refine the transition rate (λ and μ) as the following.

Definition 3: Let N denote the number of peers in a swarm, and γ denote the fraction of online peers. Given the average request rate of a user in a request session at rate q , the transition rate (λ and μ) can be refined as below.

$$\mu_x = \frac{n_x E_x[\eta] \gamma}{Nq} \quad (3)$$

where $x \in \{L, P, O\}$.

$$\lambda_x = \begin{cases} 1 & n_{x-1} I_{(\theta \neq 0)} = 0; \\ \frac{q}{n_{x-1} I_{(\theta \neq 0)} E_{x-1}[\eta] \gamma} & n_{x-1} I_{(\theta \neq 0)} > 0. \end{cases} \quad (4)$$

where $x \in \{L, P, O\}$.

wherein $n_x E_x[\eta] \gamma$ represents the capacity of the total available peers at state (x) [31][32]. And Nq refers to the total demand for the content object in the network. State ($x - 1$) refers to the system state before transiting to current state (x), and state ($L - 1$) refers to the local object delivering environment. Parameter θ corresponds to the asymmetrical ISP routing check, and $\theta \neq 0$ indicates a symmetrical case, and $\theta = 0$, otherwise. And thus $I_{(\theta \neq 0)}$ refers to the indicator function, which equals to 1 if the condition is met and 0, otherwise.

¹In this paper we do not focus on the competition of incoming requests for the uploading connection resources, which can follow Biased Unchoking (BU) paradigm [38][39].




Fig. 2. Continuous Time Markov Chain (CTMC) of the dynamics of random peer selection procedure in P2P systems

In special cases, there is few peers in a system state or the ISP routing is asymmetrical ($\theta = 0$), leading to $n_{x-1}I_{(\theta \neq 0)} = 0$ in Eq. (4). Then the transition rate follows $\lambda_x = 1$, implying a trigger of transition to the next system state.

In the following, we present analytical models for multiple native peering selection regimes and derive corresponding $P_i(t)$, $i \in \mathcal{L}$ based on Eq. (1). Following these, our proposed hybrid peering scheme is introduced.

A. Normal Randomized (Non-cooperative) Peer Selection Model

With conventional proximity-aware systems, queries are preferred to be served at local domain if there are sufficient source peers. Otherwise, candidates located at external domains are *randomly* selected. Then the dynamics of a random peer selection strategy is modelled in Fig 2 with a 4-state CTMC. Note that our study focuses on the peer selection associated with remote ISP networks, and thus initial state (0) involving only local peers is regarded as the starting point of the modelling.

Then we can derive a set of differential equations corresponding to the model according to Eq. (1) with generator matrix \mathbf{Q} :

$$\mathbf{Q} = \begin{pmatrix} -(\lambda_L + \lambda_P + \lambda_O) & \lambda_L & \lambda_P & \lambda_O \\ \mu_L & -\mu_L & 0 & 0 \\ \mu_P & 0 & -\mu_P & 0 \\ \mu_O & 0 & 0 & -\mu_O \end{pmatrix} \quad (5)$$

where λ_x follows $\lambda_L = \lambda_P = \lambda_O$, since external peers in adjacent domains are uniformly selected in a random mode. Solving Eq. (5) with the initial conditions $P_0(0) = 1$ and $P_1(0) = P_2(0) = P_3(0) = 0$, along with the boundary condition $\sum_{i=0}^3 P_i(t) = 1$ yields the probability that the process will be in state (i) at time t , $P_i(t)$, $i \in \{0, 1, 2, 3\}$.

B. Normal Locality-aware Peer Selection Model

Based on BGP routing policies, an ISP prefers the route learnt from its customer ISPs rather than from its provider ISPs or peering ISPs if both are available [2]. Thus we propose that external candidate peers are prioritized to be selected from its customer domains. Then the dynamics of an *enhanced locality-aware peer selection strategy* can be modelled as in Fig 3.

Then we have the following the generator matrix \mathbf{Q} :

$$\mathbf{Q} = \begin{pmatrix} -\lambda_L & \lambda_L & 0 & 0 \\ \mu_L & -(\mu_L + \lambda_P + \lambda_O) & \lambda_P & \lambda_O \\ \mu_P & 0 & -\mu_P & 0 \\ \mu_O & 0 & 0 & -\mu_O \end{pmatrix} \quad (6)$$




Fig. 3. Continuous Time Markov Chain (CTMC) of the dynamics of native locality awareness peer selection procedure

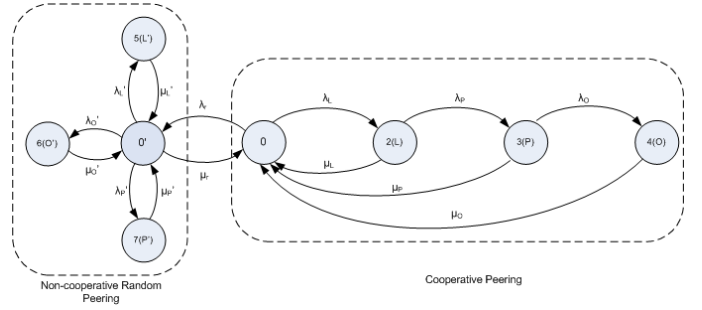


Fig. 4. Continuous Time Markov Chain (CTMC) of the hybrid peer selection procedure in a multi-domain scenario

where λ_P and λ_O follow $\lambda_P = \lambda_O$ for uniformly random choice between peering and provider domains. Solving Eq. (6) with the initial conditions $P_0(0) = 1$ and $P_1(0) = P_2(0) = P_3(0) = 0$, along with the boundary condition $\sum_{i=0}^3 P_i(t) = 1$ yields the probability that the process will be in state (i) at time t , $P_i(t)$, $i \in \{0, 1, 2, 3\}$.

C. Hybrid ISP Policy-aware Peer Selection Model

We develop a modelling framework as shown in Fig. 4 the dynamics of the proposed *hybrid peer selection procedure*, with concerns of both *cooperative* and *non-cooperative* behaviours. By adopting the stochastic method, we characterize the different strategy options for the peer selection as a number of stochastic states. Specifically, the non-cooperative procedure refers to the random approach, and the cooperative procedure corresponds to the ISP policy-aware approach.

The transition rate λ_r is defined to make mode transition decision from the ISP-policy based peering to the random strategy, in accordance with risks of potential congestion over critical inter-ISP links or failures on cooperative entities.

Then we can obtain the generator matrix \mathbf{Q} :

$$\mathbf{Q} = (\mathbf{A}_1 \quad \mathbf{A}_2) \quad (7)$$

$$\text{where } \mathbf{A}_1 = \begin{pmatrix} -(\lambda_L + \lambda_r) & \lambda_r & \lambda_L \\ \mu_r & -(\mu_r + \lambda_{L'} + \lambda_{O'} + \lambda_{P'}) & 0 \\ \mu_L & 0 & -(\lambda_P + \mu_L) \\ \mu_P & 0 & 0 \\ \mu_O & 0 & 0 \\ 0 & \mu_{L'} & 0 \\ 0 & \mu_{O'} & 0 \\ 0 & \mu_{P'} & 0 \end{pmatrix}$$

$$\mathbf{A}_2 = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \lambda_{L'} & \lambda_{O'} & \lambda_{P'} \\ \lambda_P & 0 & 0 & 0 & 0 \\ -(\lambda_O + \mu_P) & \lambda_O & 0 & 0 & 0 \\ 0 & -\mu_O & 0 & 0 & 0 \\ 0 & 0 & -\mu_{L'} & 0 & 0 \\ 0 & 0 & 0 & -\mu_{O'} & 0 \\ 0 & 0 & 0 & 0 & -\mu_{P'} \end{pmatrix}$$

With the initial conditions $P_0(0) = 1$, $P_i(0) = 0$, $i \in 1, \dots, 7$, along with the boundary condition $\sum_{i=0}^7 P_i(t) = 1$ yields the probability that the process will be in state (i) at time t , $P_i(t)$, $i \in \{0, \dots, 7\}$.

λ_x and μ_x are determined according to Eq. (3) and (4), with state (x) extended to $x \in \{L, P, O, L', P', O'\}$. Transition rate

λ_r and μ_r are denoted as $\lambda_r = \frac{k'}{kI(\nu \neq 0)} = \begin{cases} \frac{k'}{k} & \nu \neq 0 \\ 1 & \nu = 0 \end{cases}$,

$\mu_r = \frac{n_r E_r[\eta]}{Nq} \gamma$, respectively, and $\nu \neq 0$ represents a normal operation of the cooperative entity and $\nu = 0$, otherwise. Indicator function $I(\nu \neq 0)$ equals to 1 if the condition is met and 0 otherwise. Term $\frac{k'}{k}$ represents the congestion ratio constituted by P2P traffic over inter-ISP links, with k' number of inter-ISP links congested out of total k number of inter-ISP links utilized by P2P traffic. And term $n_r E_r[\eta] \gamma$ implies the total capacity of candidate peers under non-cooperative random peer selection.

IV. ATTRIBUTE MODELS IN PEER SELECTIONS

Following the presentation of the modelling in the previous section, next we derive a set of performance metrics, with concerns on ISP operational objectives and user service requirements when running their applications. Based on such comprehensive set of metrics, a novel peer selection technique will be introduced in section 4.4.

A. ISP Efficiency

ISP efficiency here refers to the ISPs capability to control P2P traffic in an optimised manner. As previously mentioned, we mainly consider the key objective of *reducing P2P content traffic across inter-domain transit links while maximising business revenues*. As such, a high ISP efficiency indicates a low P2P traffic volumes over transit links. From an analytical view, a desirable ISP efficiency is to maximize the probabilities of states $\{0, (L), (P), O', (L'), (P')\}$, which can be defined as the following.

$$e_{ISP} = 1 - [P_O(t) + P_{O'}(t)] \quad (8)$$

By adopting localization strategies, ISPs are able to reduce costly transit traffic, benefiting ISPs in terms of efficient network resources utilizations.

B. P2P User Efficiency

P2P user efficiency indicates the experiences for individual users to successfully download the desired content object. It is easy to derive that downloading experiences of users mainly depends on the aggregated retrieval rate of each system state. Thus according to Eq. (3), we denote the user efficiency as the following.

$$\mu = \frac{nE[\eta]}{Nq} \gamma + \sum_x \mu_x = \frac{nE[\eta]}{Nq} \gamma + \sum_x \frac{n_x E_x[\eta]}{Nq} \gamma \quad (9)$$

where $x \in \{L, P, O, L', P', O'\}$.

It is generally observed that the majority of P2P traffic traverses multiple inter-ISP links [33][34]. Thus there is a necessity to confine the P2P traffic within the bandwidth capacities of inter-ISP links, and Eq. (9) can be further refined as

$$\mu = \frac{nE[\eta] \gamma + k_c \bar{B}_u + m \bar{B}_p + k \bar{B}_d}{Nq} \quad (10)$$

where $k_c \bar{B}_u$ refers to the amount of traffic from customer ISPs, $m \bar{B}_p$ from peering ISPs and $k \bar{B}_d$ from provider ISPs, respectively.

Characterizing the accurate user efficiency is difficult, however, we can easily derive a lower bound of μ based on the following formular, so as to imply the minimal benefits of individual users towards the service perceived from the relevant network.

$$U_s = \lg(\alpha_s \mu + 1) - c \quad (11)$$

where α_s is a shape parameter related to a particular user. c is a fixed price that users pay ISPs for subscribing. Noticing that U_s is a strictly concave function in μ and as noted in [35], which is commonly used for describing elastic traffic and is also common used for performing distributed admission control. The \log function is chosen to model diminishing returns as μ increases. Thus a minimum μ can be derived by letting $U_s \geq 0$, which yields $\mu \geq \frac{e^c - 1}{\alpha_s}$.

C. ISP Economic Benefits

We next investigate the revenues generated by an ISP for carrying P2P traffic. In general, an ISP receives revenues from its subscribers (including customer domains and end users) and pay for the connection to its provider ISPs, which refers to the *transit cost*. Define the transit cost as C_i^d , which implies the cost for each unit of bandwidth paid to a provider ISP_i . And the cost is proportional to the mean allocated bandwidth \bar{B}_d . For simplicity we assume that there is an identical charge for both outbound and inbound traffic between a customer ISP and a provider ISP. Therefore, an ISPs profit can be expressed by

$$U_{ISP} = \left(\sum_{s=1}^n c I_{(U_s \geq 0)} + \bar{B}_d \sum_{z=1}^{k_c} C_z^d I_{(R_{c-z} \neq 0)} \right) - \bar{B}_d \sum_{i=1}^k C_i^d I_{(R_{o-i} \neq 0)} \quad (12)$$

where $I_{(\cdot)}$ is the indicator function, and equals to 1 if the condition in the bracket is met and 0 otherwise. Parameter n is the number of local users. R_{c-z} and R_{o-i} represent the traffic from a customer domain and from a provider domain, respectively. And thus $R_{o-i} = 0$ (or $R_{c-z} = 0$) means that there is no transit traffic incurred. $\sum_{s=1}^n c I_{(U_s \geq 0)}$ indicates the cost paid by users for subscribing. Term $\sum_{z=1}^{k_c} C_z^d I_{(R_{c-z} \neq 0)}$ refers to the revenues received from customer domains, and $\sum_{i=1}^k C_i^d I_{(R_{o-i} \neq 0)}$ indicates the paid transit fees.

Note that an ISP can be in deficit if $U_{ISP} < 0$, thus a minimum U_{ISP} needs to be maintained with $U_{ISP} \geq 0$. Together with the minimum users efficiency requirements by

letting $U_s \geq 0$ as analyzed previously, we thus have the following statement.

Theorem 1: *A lower bound on the number of candidate peers needs to be met, in order to guarantee desirable performance for both P2P systems and ISPs*

$$n_{min} = \max \left\{ \frac{\mu_{min} Nq - (k_c \overline{B_u} + m \overline{B_p} + k \overline{B_d})}{\gamma E[\eta]}, \frac{(k - k_c) \overline{B_d} C^d}{c} \right\}, \quad (13)$$

where $\mu_{min} = \frac{e^c - 1}{\alpha_s}$.

Proof: According to Eq. (11), $U_s \geq 0$ guarantees the minimal benefits for the user, thus $\mu \geq \frac{e^c - 1}{\alpha_s}$ can be derived accordingly, with $\mu_{min} = \frac{e^c - 1}{\alpha_s}$. Replace μ with μ_{min} in Eq. (10) and we thus have $n \geq \frac{\mu_{min} Nq - (k_c \overline{B_u} + m \overline{B_p} + k \overline{B_d})}{\gamma E[\eta]}$. According to Eq. (12) $U_{ISP} \geq 0$ can maintain profits for an ISP, we then can get $n \geq \frac{(k - k_c) \overline{B_d} C^d}{c}$ accordingly, given identical cost of transit links as C^d . Therefore, we obtain the theorem.

The above statement provides a benchmark on the number of local peers for efficient content delivering. Specifically, by comparing the number of local peers with updated n_{min} , the network provider is able to make decisions on *whether* to choose candidate peers from costly external domain. And an efficient peer selection strategy enables $n > n_{min}$ wherein n refers to the number of candidate peers located at local domains (including customer and peering domain).

Furthermore, a bound condition regarding transit links utilized needs to be maintained concerning the ISPs benefits, which can be derived as the following based on Theorem 1. As shown later in the numerical analysis, the number of transit links have more impact on ISPs side than on P2P users side, which further ensures the effectiveness of peering localization.

Corollary 1: *An upper bound on the number of transit links utilized exists, in order to maintain ISPs' profits.*

$$k_d - k_c \leq \frac{nc}{B_d C^d} \quad (14)$$

A bound on the number of transit links can be further refined with stub ISPs ($k_c = 0$), in order to maintain the utility for both ISPs and P2P systems.

$$k \in \left[\frac{\mu_{min} Nq - nE[\eta]\gamma}{B}, \frac{nc}{B_d C^d} \right] \quad (15)$$

Proof: $U_{ISP} > 0$ gives the condition for an ISP to be minimally benefited according to Eq. (12). Then $U_{ISP} \geq 0 \Rightarrow k_d - k_c \leq \frac{nc}{B_d C^d}$, given the number of available peers at local network. If $k_c = 0$, which indicates a stub ISP without customer ISPs subscribing to, an upper bound of transit links number can thus be obtained as $k = k_d \leq \frac{nc}{B_d C^d}$. According to Eq. (10), the lower bound of transit links can be derived as $k \geq \frac{\mu_{min} Nq - nE[\eta]\gamma}{B}$ by letting $U_s \geq 0$.

Theorem 1 together with Corollary 1 provides the main technique support for the proposed hybrid peer selection strategy, which answers the previous question that in what circumstances can remote peer selection be triggered. The peer selection procedure across domains depends on the benchmark (n_{min}) to follow an ISP-policy based peering process (or cooperative peering), and the link state regarding the utilization of

transit links can be monitored periodically by the cooperative entity during the process. Values beyond the range in Eq. (15) indicate a transition triggered from the cooperative to the non-cooperative peering scheme, details of which are further discussed in Section 4.4

Based on the ISP-policy based peer selection scheme (or cooperative peering selection), costly transit P2P traffic can be alternatively redirected to free peering links in order to maintain ISPs profits. While this can be effective, the free agreement between ISPs could be challenged due to non-reciprocal benefits received regarding unbalanced traffic exchanged. In order to maintain such ISP peering agreement, we have the following statement.

Corollary 2: *For two peering ISPs, e.g., ISP_A and ISP_B , a condition needs to be met so as to maintain a reciprocal peering connection in the scenario of an ISP-policy based peer selection.*

$$P_A - P_B \cong \frac{(n_o^B - n_o^A) E_o[\eta] \gamma}{\mu Nq} \quad (16)$$

where P_A refers to the probability if ISP_A selects peers from a transit ISP while ISP_B selects peers via the free peering link from ISP_A , and vice versa for P_B . $n_o^A E_o[\eta]$ indicates the total users capacity demand retrieved from ISP_A 's provider ISPs, similarly to $n_o^B E_o[\eta]$. Thus $\frac{n_o^A E_o[\eta] \gamma}{\mu Nq}$ refers to the percentage of traffic volumes retrieved from ISP_A 's transit provider.

The implication of Eq. (16) is quite straightforward, which indicates that if there is a potential risk that one peering ISP (e.g., ISP_B in Fig. 1) is served more from its peering ISP (ISP_A in Fig. 1), namely, $P_A > P_B$, more peers from its transit ISPs (n_o^B) should be considered alternatively to maintain the free peering agreement. The value of P_A (or P_B) can be derived from the model of Fig. 4 as P_3 for individual ISPs. Cooperative peer selection mode tends to increase the peering traffic between two peering ISPs, which can incur non-reciprocal benefits or even require additional upgrade on the port speeds at peering ports. Even if the upgrade can be cheaper compared to the costly transit traffic, the risk of tearing up such settlement-free relationship between two peering ISPs still exist, leading to the imperative introduction of Corollary 2. Such context information of the underlay could be monitored and collected by a trusty third party that is in charge of peer list ranking.

D. Enhanced peer selections in multi-domain environments

Without loss of generality, we assume the existence of a generic *collaborative entity* (CE) in each autonomous domain, which is responsible for collecting all necessary information from both the *P2P side* and the *network side* in order to compute optimised peer selections. How such entity can be specifically realised is outside the scope of this paper, as we mainly focus on the algorithm/theory side of the paradigm. In fact existing proposals like Oracle Service [4], P4P [8] and ALTO [11] can be all adapted for such purposes. On the one hand, individual P2P systems periodically (e.g., each time a swarm is established wherein all peers sharing the same popular object) transmit users states, including IP addresses

of peers, peer bandwidth capability distribution to the CE. On the other hand, necessary information of the underlying networks is also disseminated to the CE, including ISP business relationships and BGP routing information that can be potentially disseminated from the underlying BGP routing advertisements. Effectively, such information on the network side is rarely changed unless there are unexpected anomalies such as network failures. Available bandwidth on inter-ISP links can be mutually monitored by neighbouring ISPs in a periodical manner. In addition to the gathering of local information, CEs belonging to adjacent ISP networks also need to exchange network information with each other, as is indicated in the following.

There are two main steps involved in the peer selection strategy. Upon receiving all the necessary information input, the CE in ISP_i deduces the number of peers inside the local domain, n_{local} (based on their IP addresses), and compares it with the minimal satisfaction index n_{min} according to Theorem 1 that maintains the minimal benefits for users and the network. If $n_{local} > n_{min}$, IP addresses of these peers are added to the candidate peer list, and it is returned back to requesting users to start the content object transfer. Otherwise additional peers need to be identified outside the local ISP network.

Towards this end, the algorithm first sets ISP hop value $hop = 1$, and the CE then exchanges information (with transition rate at λ_x) with their counterparts belonging to neighbouring ISPs to identify sufficient candidate peers. The value of n_{min} is then updated and the condition $n_{local} > n_{min}$ is checked by the CE, in order to assure networks as well as users performance. Note the order for communicating with remote ISPs toward information exchange follows the ISP policy and inter-ISP routing strategy.

If there are still not sufficient peers located in neighbouring ISPs, then the algorithm sets ISP hop value to $hop = hop + 1$, and repeats the whole process. In this case, peer candidates will need to be identified from non-adjacent ISP networks. To support such feature, CEs need to communicate with each other even if they are not located in neighbouring domains. In general, CEs awareness about their counterparts in both neighbouring and remote domains can be achieved based on the dissemination of their location information through BGP route advertisement, as has been the case for other purposes like MPLS path computation elements (PCE) [36].

Algorithm Enhanced Peer Selection in Multi-domain.

- 1) n_i : the number of candidate peers for connection to peers inside ISP_i (initial value: 0)
- 2) n_{min} : the value deduced according to Theorem 1
- 3) Select peers in local domain
- 4) $n_i \leftarrow n_{local}$
- 5) **if** $n_{local} > n_{min}$ **then**
- 6) *CandidateList*[] $\leftarrow n_i$ IP addresses
- 7) *P2P systems select peers from CandidateList*[]
- 8) **else if** $n_{local} \leq n_{min}$ **then**
- 9) $ISP_{hop} \leftarrow ISP_{hop} + 1$
- 10) Select peers in customer domains
- 11) $n_i \leftarrow n_i + n_L$ and update n_{min}

- 12) **if** $n_{local} > n_{min}$ **then**
 - 13) *CandidateList*[] $\leftarrow n_i$ IP addresses
 - 14) *P2P systems select peers from CandidateList*[]
 - 15) **else if** $n_{local} \leq n_{min}$ **then**
 - 16) Select peers in peering domains
 - 17) $n_i \leftarrow n_i + n_P$ and update n_{min}
 - 18) **if** $n_{local} > n_{min}$ **then**
 - 19) *CandidateList*[] $\leftarrow n_i$ IP addresses
 - 20) *P2P systems select peers from CandidateList*[]
 - 21) **else if** $n_{local} \leq n_{min}$ **then**
 - 22) Select peers in provider domains
 - 23) $n_i \leftarrow n_i + n_O$ and update n_{min}
 - 24) **if** $n_{local} > n_{min}$ **then**
 - 25) *CandidateList*[] $\leftarrow n_i$ IP addresses
 - 26) *P2P systems select peers from CandidateList*[]
 - 27) **else if** $n_i \leq n_T$ **then**
 - 28) Go back to step 9
 - 29) **end if**
 - 30) **end if**
 - 31) **end if**
 - 32) **end if**
-

Finally, P2P systems select peers based on the information by quoting the CEs. Connections across domains are established with potential peers via candidate routes in the ranked order that obtained from the CE entities.

Note that in the peer selection process above, local available peers are considered with the highest priority. This promotes the peer connections at local area, such that users downloading experiences and inter-ISP traffic mitigation can be achieved.

For depiction simplicity, we can simply extend the peer selection scheme illustrated above to the hybrid peer selection scheme with minor corrections on the process between connections with external peers, with a further decision making after process 9 in the pseudo code. The decision can be made by the CE according to the number of inter-domain links utilized by P2P traffic (k). The CE computes the range of k and periodically compares the value of k with this range ($k \in [k_{min}, k_{max}]$) based on Corollary 1. If $k \leq k_{min}$, indicating that the P2P traffic is confined over limited inter-domain links that may constitute bottlenecks, the peer selection procedure should transit to the non-cooperative peer selection scheme in order to achieve simple load balancing over inter-domain links. Otherwise if $k > k_{max}$, implying an aggressive adoption of inter-domain links such that the economic benefits of the ISP cannot be maintained, then the peer selection procedure should be based on the ISP-policy based scheme.

E. Practicality Discussions

The purpose of the proposed hybrid peering scheme is to offer guidance on efficient peer selection across domains. Such a mechanism can be realised by a dedicated server managed by the ISP or by a trusty third party, as referred to the cooperative entity (CE) that is responsible to provide preferred information by collecting network and client information. The CE can be corresponding to the components developed by the ALTO Working Group as the ALTO server [11], or as an iTracker in

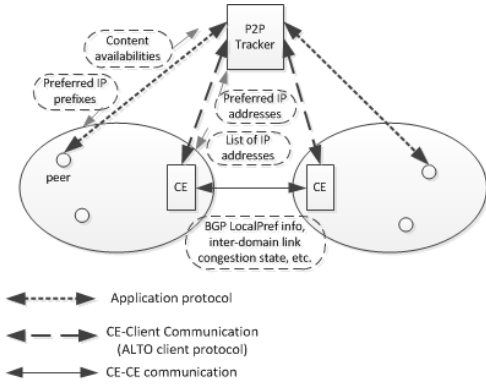


Fig. 5. Interaction between end peers, CEs and P2P trackers

P4P architecture [8]. The discovery of a CE can be enabled through DNS queries.

For a tracker-based P2P, the CE can be contacted by a P2P tracker directly to retrieve information about its preferences, such as preferred autonomous systems, IP ranges, etc. For a trackerless P2P that does not have central trackers but depends on mechanisms such as DHT, peers directly interact with CEs to obtain necessary information. In order to provide the optimized peer list, the CE is aware of both static and dynamic context information, collected from the network provider and application trackers (or peers). Static information, such as BGP preferences, does not change unless the AS-level topology is updated. If dynamic, such as the inter-domain bandwidth, the CE periodically collects such information and updates the peer list.

Specifically, in a tracker-based P2P system, such as BitTorrent, random peering is the default behaviour of the tracker that returns an arbitrary subset of active peers upon user request. Under a non-cooperative peering, the random peer list can be replaced by a list that is sorted in the ascending order by the CE based on AS hop count metric. In order to obtain the AS hop for individual peers, the tracker is assumed to be able to associate each peer (according to the IP address) to its ISP network, e.g., by using precomputed mapping information obtained from BGP tables, or by means of more sophisticated information as offered by P4P [8] in case ASes not equivalent to ISPs. Under a cooperative peering, the sorted list is based on BGP routing information to rank candidate peers rather than AS hops for remote candidate peers selection. In particular, the CE allows for the transition between cooperative and non-cooperative peering, by periodically monitoring potential congestions over critical inter-domain links.

Interactions between the CE and the P2P tracker, along with P2P peers, can be illustrated in Fig. 5. As discussed above, such interaction can be implemented by the protocol interactions between ALTO elements [37]. Specifically, the P2P tracker can be corresponding to the ALTO client and the CE can be implemented by the ALTO server. As shown in the figure, the CE has access to the P2P tracker to retrieve information of peers, e.g., a list of IP addresses. Such access can be realized by the ALTO client protocol. A preferred peer list can then be returned back to the tracker by the CE concerning preferences from the underlying network, e.g.,

BGP preferences, etc.


By referring to the Local Preferences (LocalPrefs) attribute recorded at a local BGP router [1], the CE is aware of which neighbouring domains are preferred by the local domain to set up route with. In this way, external peer selections across domains can follow ISP business concerns, such that the ISP economic benefits can be maintained. In order to enable efficient remote peer selections, the CE can also communicate with other CEs to exchange network information, regarding the symmetrical ISP routing check by comparing the LocalPref attribute values of two routes and inter-domain link congestion state, etc., which can be realized by the ALTO server discovery scheme through DNS queries. Such communications between CEs enables optimized peer selection in remote domains if insufficient peers identified locally.

The CE can be overloaded if frequent accesses by trackers (or peers). As such, we suggest that the tracker (or peer) accesses the CE in a periodic way in each swarm, and within that period a caching of the ranked list at the trackers (or peers) side can be retrieved as promoted by [11]. In this way, the query load on the CE can be alleviated.

V. NUMERICAL ANALYSIS

We specify the setting of experiment parameters below, and unless specified the values will keep fixed throughout the experiments. According to [31], the total number of concurrent users over the Internet sharing a popular content object can be assumed to be at the magnitude of 10,000.

- $N = 10^4$. The number of P2P users in the considered model which share a common content object.
- The bandwidth capacity of inter-domain (ISP network) links is $\bar{B} = 10$ Gbps.
- The value of γ , according to [4], the on-line fraction of peers can be in the range of 80-75%, so we set $\gamma = 0.8$ here, assuming that most of peers are on-line, and willing to share contents with each other
- We assume the average value of users uplink capacities can be calculated as $E_x[\eta] = 1$ Mbps. Since the upload bandwidth capacity of many users is much lower than the download bandwidth capacity based on ADSL environment, it is reasonable to assume the average users uplink capacity is around 1Mbps, compared to 10 Mbps of downlink capacity.
- We assume a unified cost for users subscribing to ISP is $\bar{c} = 1$ unit, the mean shape parameter is $\bar{\alpha}_s = 5$, thus the minimum rate efficiency can be obtained according to Eq. (11) as $\mu_{min} = \frac{e^c - 1}{\alpha_s} = 0.3$, which means that a number of users are minimally satisfied with their perceived services if at least one third of their queries can be retrieved successfully under the above assumption. The value of this parameter can be further tuned according to the sensitivity of the required object by users, such as on-line video sharing, wherein the value of μ_{min} can be set higher than the native file sharing (e.g., file downloading rather than real time video streaming, etc.) to meet users specific requirements.



(a) λ_x varying for fixed value of μ_x under ISP policy-based strategy (b) μ_x varying for fixed value of λ_x under ISP policy-based strategy (c) Under different peer selection strategies

Fig. 6. ISP efficiency evaluation

A. Evaluation of ISP Efficiency

1) *Impact of ISP policy on ISP efficiency:* We set $\mu_L = 0.7$, $\mu_P = 0.8$, $\mu_O = 0.9$ here corresponding to the fact that there are more peers holding the desired content in peering domains or in the public Internet (through provider domains) than those at local [20]. In Fig. 6(a) we plot the ISP efficiency under the ISP policy-based peering strategy (e_{ISP} , defined in Eq. (8)), with the value of λ_x varying for the first 20 time intervals (time intervals used here to show the dynamicity of the P2P system based on CTMC, and $t = 0$ is the time point at which the CTMC starts, and is in its initial state). As observed, as the transition rate of λ_x increases so as to maintain the desired user efficiency, the ISP efficiency tends to decrease. This result indicates a trade-off requirement between the ISPs and users efficiency regarding the number of transit links utilized, which can be achieved by referring to Corollary 1. As shown later the varying of inter-domain links can have bigger impact on the ISPs side rather than on the users side. Such observation further provides analytical verifications on the effectiveness of peering localization.

2) *Correlation between P2P user rate efficiency and ISP efficiency:* We vary μ_x to see how the retrieval rate influences ISP efficiency, with $\lambda_L = 0.9$, $\lambda_P = 0.8$, $\lambda_O = 0.7$ fixed. Relative high value of λ_x is set to show the lower bound of the performance. We plot in Fig. 6(b) the dynamicity of ISP efficiency e_{ISP} under the ISP policy-based peering strategy. As observed, the ISP efficiency increases with the augment of μ_x . We make the following observation based on Fig. 5 and 6 that the ISP policy-based peer selection scheme is able to benefit ISPs, since that a higher ISP efficiency indicates less traffic volume over transit links, which can in turn benefit users downloading experiences as well. However, as discussed previously, potential congestion risks over critical inter-domain links can be incurred based on the ISP-policy based peering scheme due to flash crowds when a large group of users begin to retrieve the content during a short period of time. We thus give analysis below on the ISP efficiency as well as the user efficiency based on the proposed hybrid peering scheme to show its efficiency of resolving these issues.

3) *Comparison of Different Peer Selection Strategies on Network Efficiency:* Fig. 6(c) compares the network efficiency under different peer selection strategies. We set $\mu_L = 0.7$, $\mu_P = 0.8$, $\mu_O = 0.9$ and $\lambda_L = 0.9$, $\lambda_P = 0.8$, $\lambda_O = 0.7$ here, corresponding to the fact that the number of peers varies

as the order of $n < n_L < n_p < n_O$, given the condition that $n_{x-1}I_{(\theta>0)} \neq 0$. As shown in the figure, peer selections with concerns of ISP policies can satisfyingly maintain desirable network efficiency and not surprisingly outperforms other peer selection schemes. The results are consistent with our analysis above that the ISP policy-aware peering scheme is able to offer a desirable performance for both underlying networks and P2P users, compared to native peer selections. Specifically, for $\lambda_L = 1$, indicating that $n_{x-1}I_{(\theta>0)} = 0$, the efficiency of the network is slightly reduced as shown in the figure, but still outperforms other native strategies substantially.

We also compare in Fig. 6(c) the ISP efficiency under coexistence of non-cooperative and cooperative peering scenarios (the hybrid peering scheme), with $\mu_L = 0.7$, $\mu_P = 0.8$, $\mu_O = 0.9$, $\mu_r = 0.5$ and $\lambda_L = \lambda'_L = 0.9$, $\lambda_P = \lambda'_P = 0.8$, $\lambda_O = \lambda'_O = 0.7$, $\lambda_r = 0.01$ fixed. Note that the hybrid peering scheme performs perfectly well in terms of network efficiency, similarly to the cooperative strategy (or ISP policy-based strategy) by mitigating a great amount of transit traffic as compared to the non-cooperative strategy as shown in the figure. The result implies that a desirable network efficiency can be maintained under the hybrid peering scheme, and at the same time, risks of high intensity of P2P traffic volume over critical inter-domain links can be mitigated effectively, such that congestions can be avoided accordingly.

Fig. 7 shows that the network efficiency under the hybrid peering strategy can be slightly decreased with the value of λ_r increasing, since non-cooperative is adopted with higher probability in this way that does not distinguish between different ISPs. But the results still outperform significantly those of non-cooperative or partially cooperative strategies.

We made the following observations based on Fig.6(c) and Fig. 7 that while a promising ISP policy-based peering strategy can achieve the most desirable network performance in terms of network efficiency, the hybrid peering scheme, however, is sufficient to retain such desirable performance as compared to the ISP policy-based peering scheme, and at the same time, congestions over critical inter-domain links can be eliminated.

4) *Comparison of Different Peer Selection Strategies on Reduction of Cross-ISP Traffic:* We also set here $\mu_L = 0.7$, $\mu_P = 0.8$, $\mu_O = 0.9$ and $\lambda_L = 0.9$, $\lambda_P = 0.8$, $\lambda_O = 0.7$. Fig. 8 illustrates the cross-domain traffic ratio under different peer selection strategies. As observed, the ISP policy-based localization strategy outperforms randomized strategies by reducing

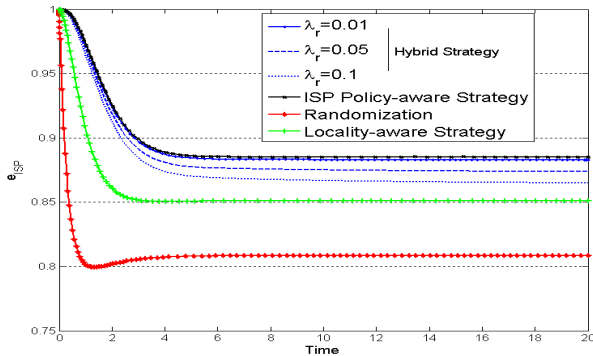


Fig. 7. ISP efficiency comparison with λ_r varying under different peer selection strategies

a substantial proportion of the inter-ISP traffic, a reduction 23.5% in the steady state. This result implies a promising capability of the ISP policy-based strategy to localize the P2P traffic that is similar to the effectiveness of locality-aware scheme. Specifically, for $\lambda_L = 1$ ($n_{x-1}I_{(\theta>0)} = 0$), the inter-ISP traffic mitigation ratio is slightly reduced but still outperforms the random strategy a lot. The results quantitatively prove the necessity of considering ISP policies in the peer selection procedure, in order to maintain desirable profits for networks and satisfied service quality for users at the same time.

As shown in Fig. 8, the hybrid peer selection scheme can maintain desirable performance as compared to the cooperative peering scheme in terms of great reduction of cross-ISP traffic. The results imply that the hybrid peering strategy is able to preserve the promising effectiveness of peering localization as compared to the ISP-policy based scheme. On the other hand, the slight decrement under hybrid peering scheme indicates that the P2P traffic volumes are not strictly confined within the limited number of inter-domain links, but rather can be distributed among multiple inter-domain links, which allows for a simple load balancing achievement.

Since the main purpose of introducing such hybrid peer selections proposition is to alleviate P2P traffic intensity over critical inter-ISP links, there is a potential increment of P2P traffic traversing via transit links. However, the increment is relatively trivial as shown in Fig. 8, only an increment of 3% for the value of $\lambda_r = 0.1$ as compared to the pure collaboration scenario. Then another concern may arise regarding the revenues loss for some lower tier ISPs due to potential adoption of transit links under non-cooperative strategy. As shown in the next section, the desirable benefits of individual ISPs can still be maintained if Corollary 1 and 2 can be guaranteed.

B. Evaluation of ISP Economic Benefits & User Efficiency

1) *Impact of Peer Selection Strategies on ISP Economic Benefits:* Fig. 9(a) compares the ISP economic benefits under the hybrid and cooperative peering scenario, with two more inter-domain links adopted under the hybrid strategy compared to the cooperative strategy. The linear display of the results shows the direct proportional relationship between the ISP economic benefits and the number of peers at local. As analyzed previously, since the hybrid peering strategy incorporates

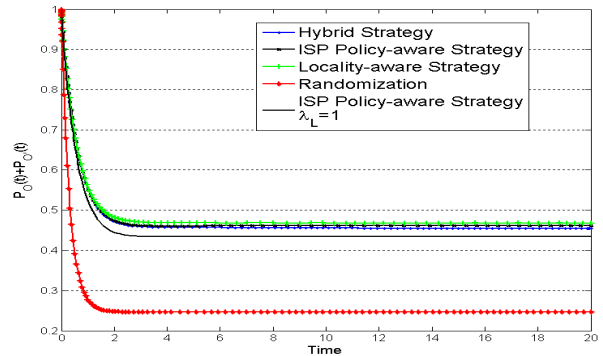



Fig. 8. Cross-ISP traffic mitigation ratio under different peer selection strategies

random selections among remote domains, it may lead to an increment involvement of the number of transit links to carry the P2P traffic. Thus the profits for an ISP can be damaged and also can be hardly predicible due to a probability concern of λ_r . As such, for some ISPs, e.g., stub ISPs without customer domains subscribing to, the profits can be affected much more than those of lower tier of ISPs, such as tier-1 ISPs. Therefore, as shown in the figure, revenues generated by an ISP can experience decrement under hybrid peer selection scenario compared to cooperation-based peering scenario. However, if specific requirements, such as Corollary 1 and 2, can be maintained in the hybrid scenario, the profit of an ISP can be guaranteed.

2) *Impact of the Number of Inter-domain Links:* We evaluate the P2P users efficiency in Fig. 9(b) under different peer selection strategies with $k = 4$. The linear display of the results shows the direct proportional relationship between the user efficiency and the number of peers at local. The results depicted in the figure indicate that users perceived service quality can be enhanced greatly under hybrid peering strategy, which is consistent with the previous analysis with an increment of $\frac{n_r E_r[\eta]}{Nq} \gamma$. A concern may arise that while localization of P2P traffic can enhance the network efficiency in terms of great reduction of costly transit traffic, the service quality perceived by P2P users could encounter degradation since fewer inter-domain links can be adopted. However, we argue that the number of inter-domain links can have bigger impact on ISP economic gains rather than on the users rate efficiency as shown in Fig. 9(c). In particular, the impact of the increment of inter-domain links can increase the transit traffic by 33.4% given the number of peers at local of around 100, while the gains for P2P users in terms of user efficiency are relatively small, with an improvement of only 8.4%. Thus the concern, that limited number of inter-domain links due to peering localization can have great impact on users service quality, can not necessarily be the case according to our analysis. The results further prove the effectiveness of localization promoted peer selections.

VI. CONCLUSION

In this paper, we develop a comprehensive analytical modelling framework for multi-domain peer selection, based on which closed forms of ISP efficiency, economic benefits



(a) ISP economic benefits under different peer selection strategies (b) User efficiency under different peer selection strategies (c) A comparison of benefits of networks & P2P users with the number of transit links varying

Fig. 9. ISP & P2P benefits evaluation

and user efficiency are derived. Based on this modelling framework, we propose an advanced hybrid peer selection scheme, taking into account of coexisting both cooperative and non-cooperative ISP behaviours in practice. The theoretical framework facilitates systematic analysis on different aspects of P2P system behaviours. In particular, we have derived bound requirements for ISPs to target in order to achieve desirable utilities for both ISPs and P2P systems. The numerical results show that the proposed mechanism is able to achieve significant performance gains for both P2P systems and ISPs. Specifically, risks of potential congestions over critical inter-ISP links could be greatly alleviated, and possible failure operations of cooperation strategies could be avoided as well.

REFERENCES

- [1] Caesar, M. and Rexford, J., *BGP Routing Policies in ISP Networks*. IEEE Network, vol. 19, no. 6, 2005.
- [2] W. B. Norton, *Internet service provider and peering*. in NANOG.
- [3] John S. Otto et al., *On Blind Mice and the Elephant: Understanding the Network Impact of a Large Distributed System*. in Proc. of ACM SIGCOMM 2011.
- [4] V. Aggarwal et al., *Can ISPs and P2P systems co-operate for improved performance?* in Proc. of ACM SIGCOMM Computer Communications Review (CCR), vol. 37, no. 3, 2007.
- [5] Cisco: *Cisco visual networking index: Forecast and methodology, 2015-2020, White Paper*, Jun. 2016.
- [6] The Independent News: *P2P file-transfer comes to Facebook with Pipe*, Jun. 2013.
- [7] WebRTC Standard: <https://webrtc.org/>
- [8] H. Xie et al., *P4P: Provider portal for applications*. in Proc. of ACM SIGCOMM, 2008.
- [9] D. Choffnes et al., *Taming the torrent: A practical approach to reducing cross-ISP traffic in P2P systems*. in Proc. of ACM SIGCOMM, 2008.
- [10] S. Ren, et al., *TopBT: A topology-aware and infrastructure-independent BitTorrent client*. in Proc. of IEEE INFOCOM, 2010.
- [11] J. Seedorf et al., *Application-Layer Traffic Optimization (ALTO) Problem Statement*. IETF RFC 5693, Oct. 2009.
- [12] Z. Dulinski, *Optimal Choice of Peers based on BGP Information*. in Proc. of IEEE International Conference on Communications (ICC), 2010.
- [13] Z. Dulinski, *Cost-driven Peer Rating Algorithm*. in Proc. of IEEE International Conference on Communications (ICC), 2011.
- [14] Peter Racz et al., *BGP-Based Locality Promotion for P2P Applications*. in Proc. IEEE 19th International Conference on Computer Communications and Networks (ICCCN), 2010.
- [15] R. C. Rumin et al., *Deep Diving into BitTorrent Locality*. in Proc. of IEEE INFOCOM, 2011.
- [16] T. Karagiannis et al., *Should Internet service providers fear peer-assisted content distribution?* in Proc. ACM IMC, 2005.
- [17] Gyorgy Dan et al., *Interaction Patterns between P2P Content Distribution Systems and ISPs*. in Proc. of IEEE Communication Magazine, vol. 49, no. 5, 2011.
- [18] Jessie H. Wang et al., *Modeling the Peering and Routing Tussle between ISPs and P2P Applications*. IWQoS, 2006.
- [19] Michele Garetto et al., *A Modeling Framework to Understand the Tussle between ISPs and P2P File Sharing Users*. Performance Evaluation, Elsevier, vol. 64, no. 9, 2007.
- [20] M. Piatek et al., *Pitfalls for ISP-friendly P2P design*. in Proc. of ACM HotNets, 2009.
- [21] Frank Lehrieder et al., *Mitigating Unfairness in Locality-aware Peer-to-Peer Networks*. International Journal of Network Management (IJNM), Special Issue on Economic Traffic Management, vol. 21, No. 1, 2011.
- [22] M. Hefeeda et al., *Traffic Modeling and Proportional Partial Caching for P2P Systems*. IEEE/ACM Trans. Netw., vol. 16, no. 6, 2008.
- [23] Frank Lehrieder et al., *The Impact of Caching on BitTorrent-like Peer-to-Peer Systems*. in Proc. of IEEE International Conference on Peer-to-Peer Computing (P2P), 2010.
- [24] Gyorgy Dan et al., *Cache-to-Cache: Could ISPs Cooperate to Decrease Peer-to-Peer Content Distribution Costs?* in Proc. of IEEE Trans. on Parallel and Distributed Systems (TPDS), vol. 22, no. 9, 2011.
- [25] Jie Dai et al., *On the Efficiency of Collaborative Caching in ISP-aware P2P Networks*. in Proc. of IEEE INFOCOM, 2011.
- [26] Eleni Agiatzidou et al., *Collaboration between ISPs for Efficient Overlay Traffic Management*. IFIP NETWORKING, 2011.
- [27] Marco Papa Manzillo et al., *CLOSER: A Collaborative Locality-aware Overlay Service*. IEEE Tran. on Parallel and Distributed Systems (TPDS), vol. 23, no. 6, 2011.
- [28] Xu Zhang et al., *A Hybrid Peer Selection Scheme for Enhanced Network and Application Performances*. in Proc. of IEEE Consumer Communications and Networking Conference (CCNC), 2013.
- [29] S. Ramabhadran et al., *Analysis of long-running replicated systems*. in Proc. of IEEE INFOCOM, 2006.
- [30] Boudewijn R. Haverkort et al., *Performance of Computer Communication Systems: A Model-Based Approach*. Wiley, 1998.
- [31] Tobias HoBfeld et al., *Measurement of BitTorrent Swarms and their AS Topologies*. Tech. Rep.464, University of Wrzburg, Jan. 2010.
- [32] Tobias Hobfeld et al., *Characterization of BitTorrent Swarms and Their Distribution in the Internet*. Computer Networks, Elsevier, vol. 55, no. 5, 2011.
- [33] H. J. Kolbe et al., *Monitoring the Impact of P2P Users on a Broadband Operators Network*. in Proc. of IFIP/IEEE IM 2009.
- [34] Kathy Brown et al., *Global Internet Report*. 2016.
- [35] F. Kelly, *Charging and Rate Control for Elastic Traffic*. European Trans. Telecomm., vol. 8, 1997.
- [36] A. Farrel *A Path Computation Element based Architecture*. IETF RFC 4655, August 2006.
- [37] M. Stiernerling et al., *ALTO Deployment Considerations*. Internet-Draft, draft-ietf-alto-deployments-16, Jul. 2016.
- [38] Oechsner S et al., *Pushing the performance of biased neighbour selection through biased unchoking*. In Proc. of IEEE International Conference on Peer-to-Peer Computing (P2P), 2009.
- [39] Ryan S. Peterson et al., *A Content Propagation Metric for Efficient Content Distribution*. in Proc. of SIGCOMM, 2011.