

An Adaptive Peer Selection Scheme with Dynamic Network Condition Awareness

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

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

Abstract— Locality-based peer selection paradigms have been proposed recently based on cooperation between peer-to-peer (P2P) service providers, Internet Service Providers (ISPs) and end users in order to achieve efficient resource utilization by P2P traffic. Based on this cooperation between different stakeholders, we introduce a more advanced paradigm with adaptive peer selection that takes into account traffic dynamics in the operational network. Specifically, peers associated with low path utilization as measured by the ISP are selected in order to reduce the probability of network congestion. This approach not only improves real-time P2P service assurance but also optimizes the overall use of network resources. Our simulations based on the GEANT network topology and real traffic traces show that the proposed adaptive peer selection scheme achieves significant improvement in utilizing bandwidth resources as compared to static locality-based approaches.

1. INTRODUCTION

The application of Peer-to-Peer (P2P) technology has in the past few years expanded from simple file sharing to real-time multimedia-based services such as IP Telephony [1] and IP Television (IPTV) [2, 3]. In particular, P2P-based IPTV applications have emerged as a popular live content distribution service that has been enjoyed by millions of customers across the Internet. Compared with the traditional non-real-time P2P applications that still work well under the best-effort based Internet traffic delivery paradigm, P2P-based IPTV services demand stringent Quality of Service (QoS) requirements from the underlying network. Nowadays it is common practice to establish Service Level Agreements (SLAs) with the underlying ISP in order to achieve guaranteed QoS assurance. Nevertheless this has not been the case for P2P-based IPTV services which is simply rely on the availability of sufficient Internet bandwidth resources, and hence the corresponding QoS performance experienced by end users cannot be guaranteed. P2P applications have however always caused operational problems for ISPs due to their greedy and uncontrollable behavior in consuming Internet bandwidth. According to recent traffic measurements [5,6], P2P applications account for some 50%-70% of the overall Internet traffic, and hence how to efficiently manage P2P traffic in order to best utilize the underlying network resources has become an increasingly important research topic. Unfortunately, today's ISPs only adopt very simple approaches to treating P2P traffic, such as limiting the utilization of network resources by P2P applications. With this approach, non-real-time applications may continue to operate, as they do

not need stringent bandwidth requirements. In contrast, multimedia based IPTV services may suffer significantly from such a treatment policy, since bandwidth resources are vital to achieve QoS guarantees to individual end users.

More recently, new research topics have been identified on how to intelligently manage P2P services in the Internet. Proposals have been made towards the optimization of Internet P2P traffic through close collaboration between P2P service providers, who offer various P2P applications, and the underlying ISPs who own the physical network resources. The aim is to develop sophisticated peer selection schemes with *locality awareness* in order to conserve bandwidth resources consumed by P2P applications. Aggarwal et al. proposed an *Oracle service* [4] that allows ISPs and P2P service providers to establish a collaborative relationship in provisioning P2P services at large scale. The Oracle service is provided by the ISP and its function is to gather relevant network layer information, such as the physical distance between potential P2P end systems, for instance in terms of hop count at the router-level, the Point-of-Presence (PoP) level or even the Autonomous System (AS) level. Such information is released to the P2P application tracker maintained by the service provider, whose task is to provide localized partner selection instructions to the incoming new peers, by taking into account the network layer information. A specific implementation of such a function based on BitTorrent selects peers according to the DNS redirection information gathered by content distribution servers, as recently suggested by [5]. Xie et al. also proposed a revolutionary P2P portal architecture called P4P [6] to optimally use network resources for supporting generic P2P systems.

As we have mentioned, real-time P2P IPTV applications are very sensitive to bandwidth resource availability. On the other hand, a distinct observation is that traffic patterns in today's operational networks are highly dynamic even within a single day [7]. This effectively means that the quality of real-time P2P services can be further impacted by the unavailability of the bandwidth resources along the peering connection paths between individual peers, in addition to peering group churn. To solve this problem, we propose an advanced peer selection paradigm that takes into account not only the static network layer information such as the physical distance between peers, but also *dynamic* network conditions. This may help to optimally select peers with higher bandwidth availability in their connections, in which case higher P2P service assurance can be achieved. In addition, given that P2P flows dominate today's Internet traffic, such an intelligent peer selection

paradigm may also benefit the underlying network resource utilization, as individual peers tend to select their partners with least utilized paths. From this point of view, our proposed scheme can be regarded as a twofold solution that benefits both P2P services (improving service quality) and ISPs (load balancing from the application layer). Nowadays, many network operators periodically perform network measurement regarding the traffic volume and bandwidth utilization, possibly with a period of 5 to 15 minutes [7]. In our proposal, the underlying ISP may also periodically provide up-to-date network condition information at some level of abstraction to the P2P service layer base on recent measurement, for instance to notify the P2P application tracker to avoid selecting peers that would further increase the load of some close-to-congestion paths. More specifically, we propose to use both static and dynamic network layer information in order to optimally select peers while both localizing P2P traffic and dynamically balancing the load at the application layer. Towards this end, we present a prioritized algorithm that selects partners for new incoming peers by taking into account these objectives.

According to our simulation results based on the GEANT network topology [8] and its real traffic traces over a 24 hour period, the overall increase in maximum link utilization in our proposed approach for optimally accommodating P2P traffic is only 30% of that by the static locality-based peer selection scheme without taking into account dynamic network conditions, and 11% of that by random peer selections. Furthermore, the overall bandwidth consumption of our approach is very close to the existing locality-based scheme, which means the original traffic locality objective is not significantly compromised.

2. PEER SELECTION OPTIMIZATIONS

2.1 System Overview

Before presenting details of our proposed peer selection algorithm, we first illustrate its overall operation at the system level. The application tracker maintained by the P2P service provider needs to periodically gather information about underlying network conditions in order to obtain optimized peering decisions that are adaptive to dynamic network conditions. This is effectively achieved through the communication between the application tracker at the P2P service provider side and the network manager at the ISP side, as proposed in [4, 6]. Figure 1 illustrates how the network manager, the application tracker and end users (peers) interact with each other in order to achieve such a goal. According to the implementation of most P2P IPTV systems [2, 3], each new peer first needs to contact the application tracker and request a list of existing peers that have already been receiving the content of the channel. The application tracker then returns the list of candidate peers. Meanwhile, the network manager deployed by the ISP periodically measures the network conditions, for instance every 15 minutes in the case of the GEANT network [7]. The measured conditions are recorded in its local network repository, based on which abstracted network condition information is passed on to the application tracker in order for the latter to compute an optimized set of peering candidates. It is worth mentioning that, for privacy reasons, detailed information may not necessarily be released from the ISP to any third-party, such as the actual utilization of each physical network link. In this case, we believe some level of abstraction on the network conditions need to be provided to

the application tracker, for instance the condition of the virtual paths between node pairs at the geographical Point-of-Presence (PoP) level. In this case, the physical network infrastructure is still treated as a black box from the P2P service provider's point of view, but the end-to-end path condition between any logical PoP node pair is still available for peer selection.

When a new peer (the requestor) joins a P2P IPTV channel, it first contacts the application tracker in order to obtain the peer list that has already been in the channel. The application tracker thereafter returns a list of potential partners according to the optimized peer selection algorithm (details in section 2.3), based on the most recently obtained network condition information from the network manager. It is worth mentioning that the application tracker does not necessarily contact the network manager on the arrival of each new peer; instead, the peer selection decisions for all newly joined peers during each measurement/reporting interval are made according to the same conditions reported for that interval.

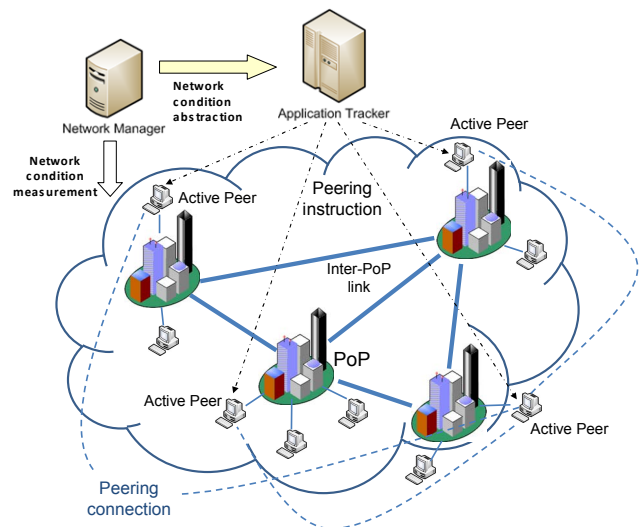


Figure 1. System overview

2.2 Objectives

We now formulate the problem of peer selection optimization with both traffic locality and load balancing. The network topology can be modeled as a unidirectional graph $G = (V, E)$, where V is a set of PoP nodes and E is the set of inter-PoP links. According to our modeling, each peer is associated with one of the PoP nodes in the physical network topology. According to the common practice of operational network design, bandwidth resources within a single PoP are usually highly over-provisioned, so we only focus on bandwidth resources on inter-PoP links in E . This means the bandwidth consumption is ignored if the peering neighbors belong to the same PoP. Let P_{ij} represent the physical path between PoP nodes i and j , consisting of one or more inter-PoP links. The bandwidth utilization of each physical link $l \in E$ is defined as u_l .

Our proposed objective of P2P traffic load balancing at the application layer is to minimize the maximum link utilization (MLU) through optimized partner selection for each joining peer given the recently measured network conditions. Towards this end, the peer selection algorithm takes into account the end-to-end path utilization reported by the ISP, which is effectively the utilization of the most loaded network link (i.e.

the bottleneck) along the considered path P_{ij} . The bottleneck of link utilization of the physical path between nodes i and j can be formulated as:

$$U_{ij} = \max(u_l), l \in P_{ij} \quad (1)$$

As we have mentioned previously, the application tracker does not need to know from the ISP the actual location of the bottleneck link, but instead only the end-to-end path conditions. As far as traffic locality is concerned, PoP level hop-count is the metric we take into account. According to [4], the strategy of traffic locality is to select physically nearby peers instead of remote ones, in order to conserve bandwidth. The physical distance between two PoP nodes i and j is represented by the number of hops between them, which is formulated as:

$$H_{ij} = \sum_{l \in E} Y_{ij}^l \quad (2)$$

where

$$Y_{ij}^l = \begin{cases} 1 & \text{if } l \in P_{ij} \\ 0 & \text{otherwise} \end{cases}$$

2.3 The Proposed U-H-based Peer Selection Algorithm

In this section we present an efficient Utilization-HopCount (U-H) based algorithm for peer selection in order to achieve the two objectives described in the last section. As we mentioned previously, the partner selection for all newly joining peers during the period between two adjacent network condition reports by the ISP is performed based on the most recently measured network performance. Figure 2 indicates a simplified repository maintained by the application tracker, called *PoP-level Path Condition Table (PPCT)* for recording abstracted network layer information reported from the ISP. The *PPCT* maintains both static and dynamic path information between each PoP node. The static information refers to the physical distance between each PoP H_{ij} (i.e. PoP-level hop counts), while the dynamic information indicates the bottleneck path utilization between each PoP node pair U_{ij} , as updated periodically according to the network condition reported by the ISP. Note that the traffic contributing to the path utilization is all components including both P2P traffic and other background traffic (http, ftp etc.); hence the path conditions measured by the ISP are the result of utilization by both types of traffic.

Source PoP	Destination PoP	Path Utilization	Distance
...
$i \in V$	$j \in V, j \neq i$	U_{ij}	H_{ij}
....

Figure 2 The *PPCT* structure

Once each new peer requests to join the channel, the application tracker will identify a list of existing peers that have already been in the channel as potential partners for the new peer by using the peer list maintained by the application tracker. First, all the active peers that are located in the same PoP are automatically selected as the partners of the new peer: this information can be directly obtained from the Oracle service as proposed in [4], where the geographical location of active peers is known, including their distribution in individual PoPs. In cases where the local peers cannot support the content availability for the new peer, the application tracker needs to select additional peers from remote PoPs. In this scenario, our proposed peer selection scheme will be used in order to achieve efficient use of network resources on inter-PoP links.

According to our proposed algorithm, it is desirable to select those remote active peers associated with PoP-level paths with low traffic load towards the PoP node where the new peer is attached. To achieve this goal, the application tracker needs to retrieve the *PPCT* (Fig. 2) and examine the path utilization U_{ij} for each remote PoP. Starting from the one with the least loaded path towards the PoP node attached with the new peer, the application tracker identifies further active peers as candidate partners. In case these additional peers are still unable to satisfy the content available requirement, the next PoP with least loaded path will be examined for further peer selection. In cases where two or more PoPs have the same path condition towards the local node attached with the new peer, the active peers attached to the PoP node with the least distance (in terms of PoP-level hop counts) will be selected as tie breaking. The algorithm iterates with new candidate partners being selected for the new peer until its content availability requirement is satisfied. Once finished, the application tracker compiles a list of peer candidates and instructs the new peer to contact them. A flow chart for the entire operation is shown in Figure 3.

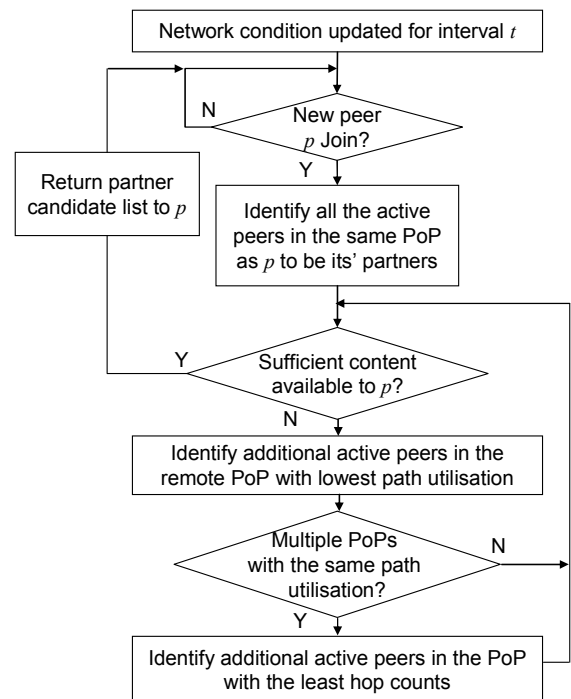


Figure 3. The U-H based peer selection algorithm

3. PERFORMANCE EVALUTATION

In this section we evaluate the performance of our algorithm compared with the other two schemes, namely *random peer selection* and *static locality-based peer selection* as proposed in [4]. In the random peer selection scheme, the application tracker arbitrarily chooses a list of active peers across the entire network as potential partners for the newly arrived peer. According to the locality-based peer selection scheme, peers are intelligently selected according to their physical distance to the newly joined peer, but end-to-end path conditions that can be influenced by overall network dynamics are not taken into account.

3.1 Metrics To Be Evaluated

We first present the definitions for the following parameters.

- C_l : the bandwidth capacity of inter-PoP link l ;
- B_l : the volume of background traffic (non-P2P traffic) that is projected on link l ,
- T_l : the volume of P2P traffic that is projected on link l .

- *Maximum Link Utilization*

The primary goal of our proposed algorithm is to balance the overall traffic in order to avoid traffic congestion that both impacts the real-time P2P service quality and leads to suboptimal utilization of the underlying network resources. Towards this end, the maximum link utilization (MLU) across the entire network is examined in our experiment, and is defined as:

$$\max(U_l) = \max((B_l + T_l)/C_l), \forall l \in E \quad (3)$$

- *Overall Bandwidth Consumption*

Conserving the overall network bandwidth consumption by P2P traffic is an important objective for the underlying ISPs. The U-H peer selection algorithm is proposed for this purpose. The metric of overall bandwidth consumption is defined as:

$$\sum_{l \in E} (B_l + T_l) \quad (4)$$

It should be noted that none of the three considered peer selection algorithms aims to control the non-P2P background traffic. Hence the actual objective of U-H based algorithm is to

$$\text{Minimize } \sum_{l \in E} T_l \quad (5)$$

- *Overall Network Cost*

The piece-wise linear cost function has been widely used for evaluating traffic engineering purposes. In this paper we use the cost function proposed in [9], that is:

$$\varphi = \sum_{l \in E} \varphi_l(U_l) \quad (6)$$

Where for all $l \in E$, $\varphi_l(0) = 0$ and

$$\varphi_l(x) = \begin{cases} 1 & \text{for } 0 < x \leq 1/3 \\ 3 & \text{for } 1/3 < x \leq 2/3 \\ 10 & \text{for } 2/3 < x \leq 9/10 \\ 70 & \text{for } 9/10 < x \leq 1 \\ 500 & \text{for } 1 < x \leq 11/10 \\ 5000 & \text{for } 11/10 < x \leq \infty \end{cases}$$

3.2 Experimental Setup

Our simulation experiment is based on the GEANT network [8] topology and its traffic traces across 24 hours. The GEANT network topology consists of 23 PoP nodes and 74 unidirectional inter-PoP links. According to [7], the GEANT traffic traces are measured every 15 minutes through NetFlow. In our simulation we took hourly samples of these traces during the period of one single day. Figure 4 shows the measured MLU performance in the GEANT network across these 24 hours (starting from 12:00 noon), which is effectively indicated with 24 traces. It can be clearly seen that the overall traffic volume is highly dynamic, for instance the minimum value of MLU during the period is around 35%, while the maximum value in the same day can reach as high as 85%. In our

simulation, we use the scaled volume of these traffic traces to emulate the non-P2P background traffic behavior.

The P2P traffic used in our experiment is synthetically generated according to the flow characteristics of today's popular P2P IPTV applications. We consider 6 IPTV channels, with each channel attracting up to 1200 peers. Hence altogether we consider 7000+ peers that are randomly distributed across the 23 PoP nodes in the GEANT network. More specifically, we consider a sequence of one-by-one peer joins during the 24-hour period, with each randomly assigned to one of the GEANT PoP nodes. The channel selected by each peer is also randomly determined. In addition, we use the observation that each new peer has around 80 peering connections in order to satisfy the content availability requirement for playback in a stable peering state. We also assume that the download rate is 800 Kbps which is the case for most of today's popular P2P IPTV applications [2, 3].

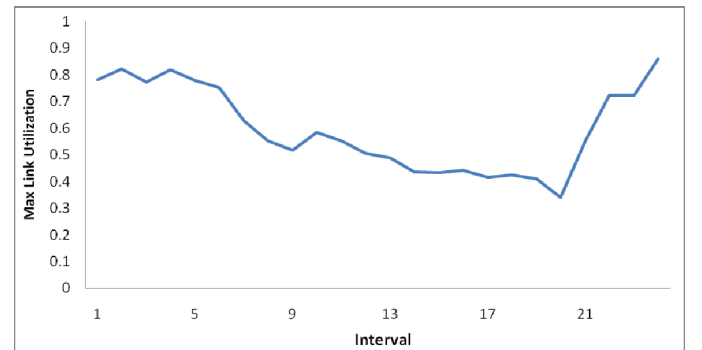


Figure 4. Background traffic dynamics in GEANT

3.3 Results

Figure 5 plots the maximum link utilization for the three schemes, namely random peer selection, locality-based peer selection and our proposed U-H based peer selection. We can clearly see from the figure that the MLU in the random peer selection scheme increases sharply as more peers join the groups. In comparison, the other two schemes have much lower MLU, thanks to their locality awareness, which tends to conserve the overall network bandwidth resources. On the other hand, our proposed U-H based scheme still outperforms substantially the plain locality based algorithm, as it further strives to intelligently select peers associated with least-loaded paths; this leads to much more balanced traffic distribution across the entire network. In effect, after accommodating all the peers at the end of the procedure, the MLU with the U-H based algorithm has only increased by 28.9%, in contrast to 95.1% for the locality based scheme and 257.8% for the random selection scheme. In other words, the increased MLU according to our proposed algorithm is only 30.4% and 11.2% of that by the static locality-based approach and pure random peer selection approach respectively. As we have mentioned, adaptive peer selection according to network conditions not only reduces the chance of traffic congestion that may severely disrupted ongoing P2P IPTV service quality, but also results in efficient resource utilization from the ISP's point of view. The performance improvement shown in the figure clearly indicates that our proposed adaptive peer selection scheme is a promising paradigm that will benefit both content providers and network operators.

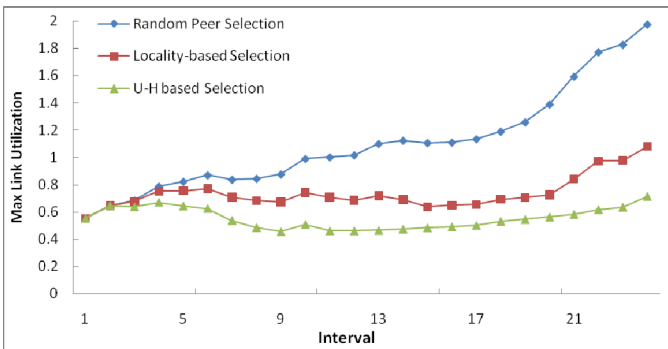


Figure 5. Maximum Link Utilization performance

In addition, we show in Figure 6 the overall bandwidth consumption for the three different schemes. The total bandwidth consumption for the random scheme increases significantly as the number of active peers increases. This means that the random scheme could generate much more P2P traffic across the entire network. On the other hand, the contribution of the locality-based scheme is to localize peering connections, and it thus saves around 40% of the overall bandwidth consumption. Our proposed U-H-based approach has only moderate increasing (9% on average) compared to the locality-based scheme. The reason is that we give higher priority to selecting peers associated with least utilized paths rather than those with shortest distance. Hence P2P flows may occasionally travel through some longer paths within the network, consuming more bandwidth resources.

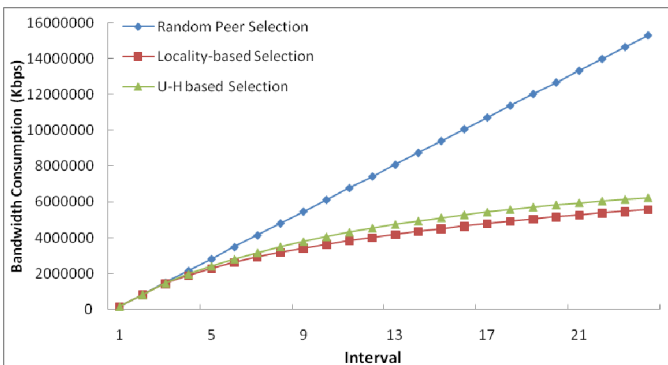


Figure 6. Overall bandwidth consumption

Figure 7 shows the overall network cost incurred by the three schemes. From Figure 7(a) we can clearly see that the network cost by the random scheme is much higher than the other two approaches, as the overall network resources are not optimized at all. In order to clearly show the performance gap between locality-based and U-H based schemes, we exclude the performance of random peer selection in Figure 7(b). As we can see, with the increase of active peer population, the gap between the two schemes becomes more significant. This is especially the case after 21th interval in the simulation when the network cost in the locality-based approach increases sharply. This occurs because the utilization of some links in the network rises significantly (0.9 or higher) and the cost given by equation (6), therefore rises dramatically.

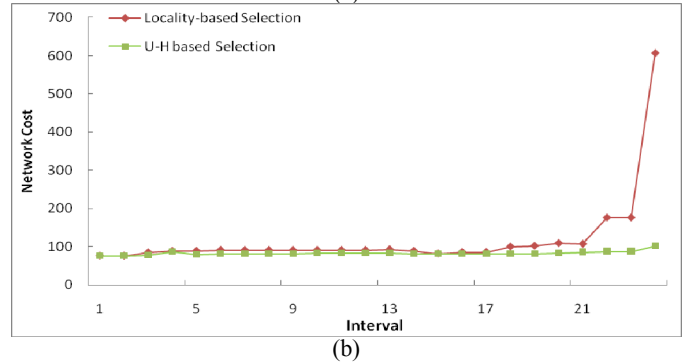
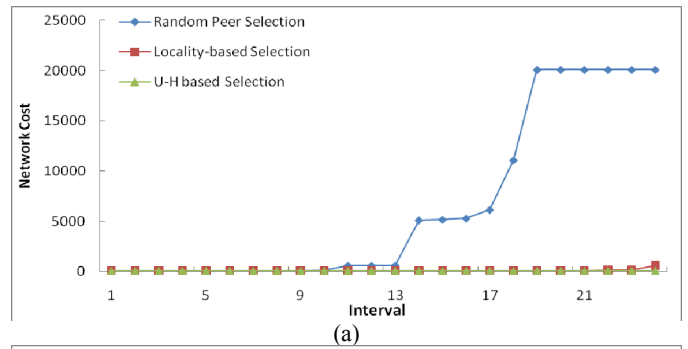


Figure 7. Network Cost performance

4. CONCLUSION

Proposals have been put forward for P2P service providers to collaborate with the underlying ISPs in order to achieve both higher service assurance and more efficient utilization of network resources. However, network dynamics such as traffic upsurges may significantly impact the service quality and operational network efficiency. In this paper, we introduce an intelligent peer selection paradigm based on the collaboration between P2P service providers and ISPs. It not only considers the static traffic locality requirement but also takes into account traffic dynamics that may significantly impact the P2P service and the network operational performance. Our simulation results based on the GEANT network topology and its traffic traces indicate that the proposed adaptive peer selection scheme achieves substantially higher network efficiency that benefits both P2P service assurance and network resource utilization performance.

5. ACKNOWLEDGEMENT

This paper was supported in part by the EU EMANICS Network of Excellence on the Management of Next Generation Networks, IST-026854.

REFERENCES

- [1] Skype, <http://www.skype.com>
- [2] PPLive, <http://www.pplive.com/en/index.html>
- [3] SopCast, <http://www.sopcast.com>
- [4] V. Aggarwal et al., "Can ISPs and P2P Users Cooperate for Improved Performance?", ACM CCR, Vol. 37 (3), 2007
- [5] D. Choffnes et al., "Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in P2P Systems", Proc. ACM SIGCOMM 2008
- [6] H. Xie et al., "P4P: Provider Portal for Applications", Proc. ACM SIGCOMM 2008
- [7] S. Uhlig et al., "Providing Public Intradomain Traffic Matrices to the Research Community", ACM SIGCOMM Computer Communications Review, 36 (1), January 2006
- [8] The GEANT Network topology, <http://www.geant.net>
- [9] Fortz and M. Thorup, "Internet Traffic Engineering by Optimizing OSPF Weights", Proc. IEEE INFOCOM 2000