

A New Approach for Achieving Traffic-Exchange Localization in P2P-based Content Distribution

Chaojiong Wang, Ning Wang, Michael Howarth

University of Surrey

Guildford, United Kingdom

{C.Wang, N.Wang, M.Howarth}@surrey.ac.uk

Abstract — Due to the fact that P2P applications have dominantly accounted for the entire Internet traffic, how to efficiently manage P2P traffic has become increasingly important. It has been recently proposed that the underlying network information can be shared between ISPs and P2P service providers in order to achieve efficient resource utilization, with the locality-based peer selection being a specific example. Based on such collaboration, we propose a proportional traffic-exchange localization scheme for making efficient use of network resources. Our approach employs locality information in order to regulate the volume of traffic exchange between peers according to their physical distance between peers. The key objective of our approach is to further reduce both intra- and inter-autonomous system (AS) traffic compared with basic locality-based peer selection solutions. Our simulation-based results have shown that this approach is not only able to reduce a significant of inter-AS P2P traffic, but also to balance the network utilization in comparison to existing approaches.

Keywords- P2P Applications, Traffic Localization, Application Layer Traffic Optimization (ALTO)

I. INTRODUCTION

The Peer-to-Peer (P2P) technology, as a scalable solution for content distribution, has been in the past few years developed to offer various large-scale network services such as file sharing. In contrast to the traditional Client/Server (C/S) model where the server acts as a single point for content distribution, P2P takes advantage of every participant's resources (CPU, storage disk, and bandwidth etc.) to provide higher efficiency but lower cost services. On the other hand, with their increasing popularity in the Internet, including those for both file sharing and real-time multimedia-based content delivery, P2P applications have nowadays contributed to huge amount of network traffic that may potentially starve other non-P2P applications, due to their greedy and uncontrollable behaviours [4]. Such situation is not desired as far as Internet Service Providers (ISPs) are concerned, as this inevitably increases their operational cost on traffic engineering (TE). Although TE is traditionally considered by ISPs at an aggregate level, various types of network applications have been involved to influence the overall traffic pattern in the Internet [7]. As a result, how to efficiently manage P2P traffic at the application layer as complementary to traditional network layer traffic engineering, has been attracting numerous attentions from the research community.

Today most of the P2P applications running on top of the Internet apply peer selections that are agnostic to the underlying network information, e.g. the physical distance between individual peers. Without taking into account this information, a large amount of bandwidth resources are wasted

for carrying long-haul P2P traffic, especially across multiple Autonomous Systems (ASes) in the Internet. Currently the common practice is for ISPs to simply pose transmission limit for P2P traffic, which inevitably leads to a conflict between network providers (network efficiency) and P2P service providers (application performance). More specifically, TE considers reducing the network cost while P2P service providers concentrate on improving the service of quality (e.g. minimizing the delay) at the application layer. Nevertheless, it has recently been found that desired chunks of content can often be found at peers in nearby regions in the Internet due to the same language or similar culture. Therefore, Aggarwal et al. proposed a generic *Oracle service* [4] that takes advantage of relevant network layer information when performing locality-based peer selection. This approach allows ISP and P2P service provider to establish a collaborative relationship in gracefully provisioning P2P services across the Internet. A specific realization of such an oracle function based on BitTorrent [2] is to select peers according to the DNS redirection information gathered by content distribution servers [5]. Through the above proposals, the locality-based peer selection scheme has been proved for its feasibility and efficiency when ISPs cooperate with P2P service providers. On the other hand, in addition to the peer selection schemes, optimization on P2P traffic-exchange pattern among peers is another potential paradigm that may contribute to better utilisation of network resources. To better illustrate this idea, we first show some statistics from a SopCast [1] based measurement we have conducted. Figure 1 indicates the average traffic volume exchanged between individual local peer and its partners in the global Internet. We can see in Figure 1 (a) that 30% data of total local peer obtained from each partner who has 16~20 router-level hop-counts distance, which is 13% more than the partner that is only 10~15 hop counts away. The same problem also happened between 15~20 hop counts partners and 20~25 hop counts partners. This effectively means that peers do not tend to exchange more traffic with nearby partners than remote ones. Furthermore, in the Figure 1 (b) the largest proportion of data is exchanged with a peer having 4 AS hop counts, rather than a peer with 3 AS hop counts we expected. Notes that local peer located on campus, then it needs to take first ten hop counts (two AS hop counts) to traverse campus and education network. Based on such observations, we hence conclude even if peers are selected with locality awareness, a random traffic exchange pattern may still potentially lead to sub-optimal bandwidth consumptions. For instance a significant amount of P2P traffic may unnecessarily traverse across multiple network boundaries. In contrast, we believe that higher volume of traffic exchange with nearby peers as compared with remote ones can further improve the network resource utilizations.

In this paper, we propose a traffic-exchange localization algorithm that can be applied on top of the existing locality-based peer selection schemes in order to achieve further optimization of network resource utilization, especially to reduce bandwidth consumptions across inter-AS links. Similar to [4], this approach also requires that the ISPs provide necessary network information to the P2P service provider. In general, partners are selected based on conventional locality-based scheme [4]. We propose that traffic-exchange volume among peers should also take into account their physical distance (e.g. in terms of router-level hop counts). More specifically, peers with long distance may exchange less traffic than those that are physically closer to each other. This strategy can be easily extended to the scenario of inter-AS P2P traffic exchange for reducing bandwidth consumptions on inter-AS links. Given the fact that bandwidth resources on inter-AS links are more scarce, we thus proposed a penalty-based mechanism that aims to avoid incurring cross-ISP P2P traffic. Of course, this paradigm needs more intelligent implementation of the P2P content swarming mechanism at the application layer which is able to adapt the traffic exchange rate between partners to their physical locations in the Internet.

According to our simulation results, there is 43% decrease of inter-AS traffic even if the penalty on using inter-AS links is regarded as same as intra-AS links, while over 80% decrease is achieved when a high penalty (20 times as intra-AS link) is put on inter-AS links, as compared with the basic locality-based peer selection scheme. Both overall network cost and intra-AS bandwidth consumption reduce 10% respectively. Furthermore, the maximum link utilization (MLU) of inter-AS links in our approach reduces about 53% in comparison with the basic locality-based peer selection scheme without traffic-exchange localization.

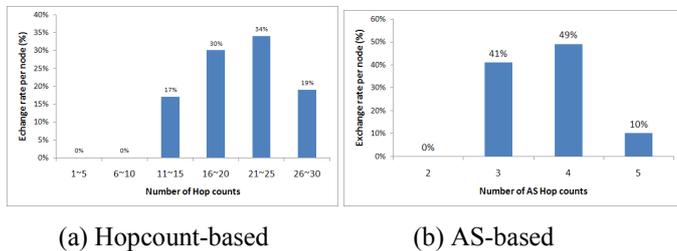


Figure 1: Traffic pattern of Sopcast

II. RELATED WORK

Based on the recent proposals [4, 5, 6], there are three types of peer selection schemes that have been proposed or developed. In order to achieve our key objective – to reduce resource consumption by P2P traffic, especially across inter-AS links, we propose that the new traffic-exchange localization scheme that is able to jointly work with a locality based peer selection paradigm (details in section III).

Random peer selection scheme: Most existing P2P applications employ this type of scheme to arbitrarily select peers without taking into account geographical distribution of peers in the Internet. Generally, this scheme does not have the capability to select optimized partners for enhancing the underlying network resource utilization. More specifically,

many remote peers are selected as partners for content swarming, rather than local ones even if they have obtained the content required. Due to large amount of unnecessary P2P traffic, non-P2P traffic may encounter starving problems which may lead to higher delay and packet loss.

Locality-based peer selection scheme: As an improvement to random peer selection scheme, peers are selected in terms of locality information such as hop count in order to reduce the ISP operational cost and improve the service of P2P applications. For instance, Aggarwal et al. [4] proposed to establish a collaborative relationship between network providers and P2P service providers by *oracles service* whose task is to provide network information from network provider to P2P service provider in order to optimize peer selections. By applying such locality-based peer selection approach, peers are selected in terms of locality information in order to reduce inter-AS traffic. The authors also suggested that locality information can be at either AS, Point of Presence (PoP) or router level. According to the common practice of operational network design, inter-AS links are usually considered as scarce resources. Therefore, if most partners are selected having shorter distance, especially those which are able to avoid crossing the multi-AS boundaries, the overall utilization of network resources can be improved. In [5], Choffnes et al. implemented extended BitTorrent based on the DNS redirection information to select peers. Through both simulation-based and real-application experiments, these proposals have been demonstrated that the locality-based optimization can efficiently make use of network resources.

Multi-objective peer selection scheme: Another new solution is a revolutionary P2P portal architecture called P4P [6]. The authors proposed the p4p-distance for helping to select peers optimally. This p4p-distance interface can be defined in various ways such as locality information (in terms of OSPF link weight), network status (congestion information), or the combination of them. Specifically, this solution aims to reduce the maximum link utilization (MLU) and balance the network load, while improving the performance of P2P applications.

III. TRAFFIC-EXCHANGE LOCALIZATION

This paper introduces an intelligent traffic-exchange localization approach that is to further reduce unnecessary P2P traffic on both intra-AS links and inter-AS links. It can be adapted on top of any of the peer selection schemes mentioned in section II. That is, our approach can be regarded as orthogonal to the existing peer selection schemes. More specifically, our approach can take the advantage of network information from network providers for traffic exchange localization, which is similar to the oracle service [4] for locality-based peer selection operations. The key objective of our approach is to further reduce inter-AS traffic, we thus chose to deploy our traffic-exchange localization approach on the locality-based peer selection scheme. Upon original locality-based peer selection mechanism, our approach takes advantage of the same locality information to achieve traffic-exchange localization without any other additional complications. In this section, we will describe our approach from two special cases – intra-AS and inter-AS to a general case.

A. Intra-AS Traffic-Exchange Localization

We first consider the basic intra-AS scenario where all peers are located in the same AS. Router-level hop-count is used as the metric to determine the distance between two peers. We now model the problem of traffic-exchange localization in single AS case. The physical network topology of a single AS is modelled as unidirectional graph $G_r = (V_r, E_r)$, where V_r is a set of nodes, and E_r is the set of links. In our case, a node represents a physical router in the network topology. For simplicity we ignore the scenario that one single router is associated with multiple peers, as the content exchange between them does not consume bandwidth resource on any network link $l \in E_r$. Let D_{ij} denote the traffic demand is transmitted from node j to node i , and P_{ij} is a path from node i to node j , which consists of some physical links $l (l \in E_r)$. $D_{ij}=0$ means node j is not selected by node i as its partner for content request. The overall bandwidth consumption on intra-domain links can be formulated as:

$$\text{Minimise } T_r = \sum_{i \in V_r} \sum_{j \in V_r} \sum_{l \in E_r} Y_{ij}^l * D_{ij}, \quad i \neq j \quad (1)$$

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

According to our proposed traffic-exchange localization scheme, the basic strategy is to determine D_{ij} for each peering partner pair i and j according to their physical distance in order to minimize T_r .

B. Inter-AS Traffic-Exchange Localization

It can be easily inferred that the traffic exchange localization approach can be also extended to inter-AS scenario where participating peers are distributed across multiple ASes. In this case it is essential to reduce the overall bandwidth consumption by P2P traffic across inter-AS links. The AS/domain-level topology is modeled as unidirectional graph $G = (V_d, E_d)$, where V_d is a set of ASes, and E_d is the set of inter-AS links. AP_{ij} represents the path from AS i to AS j ($i, j \in V_d$) which contains one or multiple inter-AS links $l (l \in E_d)$. To minimize the overall bandwidth consumption across all inter-AS links, the objective is:

$$\text{Minimise } T_d = \sum_{i \in V_d} \sum_{j \in V_d} \sum_{l \in E_d} AY_{ij}^l * D_{ij}, \quad i \neq j \quad (2)$$

$$\text{Where } AY_{ij}^l = \begin{cases} 1 & \text{if } l \in AP_{ij} \\ 0 & \text{otherwise} \end{cases}$$

Similar to the intra-AS scenario, the basic strategy is to determine D_{ij} for each peering partner pair and according to their physical AS-level distance in order to minimize T_d .

C. A General Traffic-Exchange Localization Case

Based on above two distinct cases, it is not difficult to integrate them into a general problem which is to determine the traffic exchange volume between peers both within one single AS and across multiple domains. The latter case can be imagined as the situation that some certain peers are not able to obtain sufficient content from existing local partners belonging to the same AS, so they need to further request content from remote ones in foreign ASes.

D. The Proposed Traffic-Exchange Localization Algorithm

In this section, we present our proposed traffic-exchange localization algorithm. As we have mentioned, the basic idea is that the traffic volume exchanged between individual peers is adaptive to their physical distance, and ideally in inverse proportion. Of course, this situation cannot be completely achieved due to some practical constraints such as content availability. Nevertheless, we argue that individual peers should gear towards this optimized situation once they have identified sufficient partners that are able to provide the required content. As we have mentioned previously, our proposed traffic exchange localization can be jointly applied with locality based peer selection paradigms. Now we briefly describe how this can be achieved. When a new peer joins the session, the application tracker first searches whether there are sufficient partners holding desired content in the local AS. If this is the case, the new peer directly contacts them to request content. Otherwise, all available local partners are selected, and in addition a number of remote partners are also needed. These remote partners are selected according to their AS-level distance to the requesting peer. More specifically, peers located in nearby ASes are preferred against remote ones.

Based on such locality-based peer selection scheme, we now consider how to apply the traffic-exchange localization scheme. As we have mentioned, the traffic-exchange pattern of the current P2P applications are random since data is arbitrarily pulled from partner nodes. Therefore space still exists for further reducing unnecessary P2P traffic on both intra-AS and inter-AS links. Without loss of generality, we introduce the physical distance H_{ij} between two peers i and j , where i and j can belong to either the same AS or different ASes. In general, the metric H_{ij} can be regarded as the combination of IH_{ij} and $AH_{ij} - IH_{ij}$ represents the number of intra-AS hop counts from i to node j , while AH_{ij} represents the AS-level distance. It should be noted that, inter-AS links are commonly believed to be more scarce resources compared with inter-AS ones. In order to put higher emphasis on the reduction of P2P traffic across inter-AS links, we propose to impose higher penalty on the usage of them. For instance we can conceive that the penalty on using one single inter-AS link is equivalent to that on using of n hops of intra-AS links. Towards this end, we introduce a weighting penalty n on the metric of inter-AS distance. More specifically $n=1$ means there is no differentiation between using of inter-AS link and intra-AS link, while $n>1$ means that a higher penalty is introduced on using inter-AS links. According to above description, a combined distance factor is formulated as:

$$H_{ij} = \sum_{i \in V_r} \sum_{j \in V_r} IH_{ij} + n * AH_{ij}, \quad n \geq 1, i \neq j \quad (3)$$

We use S_i to denote the overall content demand from a new peer i from all of its partners. In order to achieve traffic-exchange localization, the traffic volume exchanged between individual peer pair i and j is in the inverse proportion to their combined distance H_{ij} , which can be expressed as:

$$D_{ij} \propto \frac{1}{H_{ij}}, \quad H_{ij} \neq 0, \quad \sum_{j \in V_r} D_{ij} \geq S_i \quad (4)$$

Consequently, we can derive the proportion of traffic exchange rate from partner node j to node i .

$$R_j = (D_{ij} / \sum_{j \in V_r} D_{ij}) * S_i \quad (5)$$

IV. PERFORMANCE EVALUATION

In this section, we evaluate the performance of our proposed algorithm, locality-based peer selection & traffic-exchange localization (LS<), compared with other three schemes: random peer selection & traffic exchange (RS&RT), random peer selection & traffic exchange localization (RS<), and locality-based peer selection & random traffic exchange (LS&RT). These are effectively the four combinations of choices in locality-based peer selection and traffic exchange localization.

A. Metrics

In addition to the previously defined parameters, we introduce two metrics related to bandwidth capacity:

- C_r^l : Capacity of a physical intra-AS link l .
- C_d^l : Capacity of a physical inter-AS link l .

1. Overall bandwidth Consumption

The ultimate objective of our approach is to reduce bandwidth consumption by P2P traffic, especially on inter-AS links.

$$\text{Intra-AS traffic: } T_r = \sum_{i \in V_r} \sum_{j \in V_r} \sum_{l \in E_r} Y_{ij}^l * D_{ij}, \quad i \neq j \quad (6)$$

$$\text{Inter-AS traffic: } T_d = \sum_{i \in V_d} \sum_{j \in V_d} \sum_{l \in E_d} AY_{ij}^l * D_{ij}, \quad i \neq j \quad (7)$$

$$\text{Overall traffic: } T_a = T_r + T_d \quad (8)$$

2. Maximum Link Utilization (MLU)

Reducing the maximum link utilization is another important issue to be investigated, which aims to balance the overall network load. The lower the maximum link utilization is, the lower chance traffic congestion in the network will occur. We hence evaluate this objective for intra and inter-AS links respectively as:

$$\text{Link utilization (intra-AS): } f_r(l) = Y_{ij}^l * D_{ij} / C_r^l \quad (9)$$

$$\text{Link utilization (inter-AS): } f_d(l) = AY_{ij}^l * D_{ij} / C_d^l \quad (10)$$

$$\text{MLU on intra-AS links: } U_r = \max(f_r(l)), l \in E_r \quad (11)$$

$$\text{MLU on inter-AS links: } U_d = \max(f_d(l)), l \in E_d \quad (12)$$

3. 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 [8] to indicate the overall network cost on both intra- and inter-AS links:

$$\text{Intra-AS network cost: } \phi_r = \sum_{l \in E_r} \phi_l(U_r) \quad (13)$$

$$\text{Inter-AS network cost: } \phi_d = \sum_{l \in E_d} \phi_l(U_d) \quad (14)$$

$$\text{Overall network cost: } \phi_a = \phi_r + \phi_d \quad (15)$$

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

B. Experimental Setup

The network topology in our simulation is generated by the IGen topology generator [3]. Our simulation consists of a total of 5 inter-connected ASes, each having its own dedicated intra-AS topology. We assume 800 peers that are randomly associated with one of the nodes inside a specific AS. A set of border routers are randomly selected within each AS for inter-AS connection. As far as peer group dynamics is concerned, we consider a sequence of one-by-one node joins the session. We follow the observation that each new peer selects 20 partners for content request, which is similar to the real scenario of [1]. A fixed uploading capacity is set to 1000 Kbps, and a request download speed is set to 5000 Kbps in our measurements.

C. Simulator Scenarios

As we described in section II, our traffic-exchange localization approach can be adapted on top of any of peer selection schemes. In order to have a comprehensive comparison, we defined four schemes in two dimensions: random or localised peer selection schemes and random or localised traffic-exchange scheme, which are listed as follows:

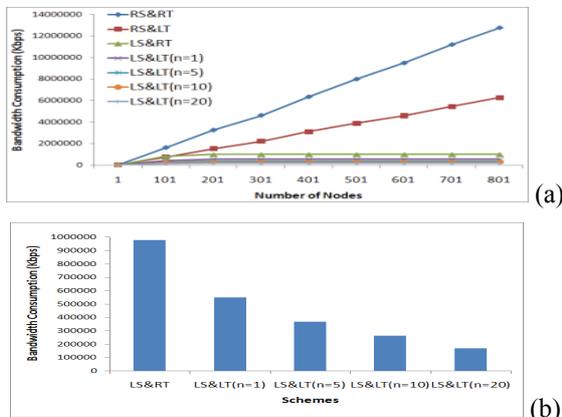
- (1) *Random Peer Selection & Traffic Exchange (RS&RT)*: new peer randomly selects partners holding desired contents and pulls random amount of data from its partners without taking into account distance information.
- (2) *Random Peer Selection & Traffic Exchange Localization (RS<)*: new peer randomly selects partners and exchanges more proportion of data with less hop-count partners.
- (3) *Locality-based Peer Selection & Random Traffic Exchange (LS&RT)*: new peer selects the partners having less hop-count, and arbitrarily exchanges data with its partners.
- (4) *Locality-based Peer Selection & Traffic Exchange (LS<)*: new peer also selects less hop-count partners, and exchange data in inverse to their distance.

By comparing the performance of (1) and (2), we can observe the improvement achieved by employing traffic-exchange localization algorithm based on a random peer

selection scheme. We also investigate the performance gap between locality-based traffic exchange scheme (2) and locality-based peer selection scheme (3). Finally, scheme (4) gives us a sense of the performance of our algorithm, as compared with schemes (1), (2), and (3). Furthermore, as we discussed, we put higher penalty on the usage of inter-AS links in order to further reduce redundant P2P traffic on inter-AS links, represented by the weighting parameter n in equation (4) above. For instance we set is the value of n to be 1, 5, 10, and 20 in our simulation scenarios.

D. Results

Figure 3 shows the overall bandwidth consumption on all inter-AS links connecting the five ASes. Figure 3 (a) illustrates the increase of bandwidth consumption with sequential join of individual peers, as we can see from the figure, the bandwidth consumed by scheme (1)-RS&RT and (2)-RS< is much higher than other schemes because both of them employ a random peer selection. On the other hand, scheme (2)-RS< only consumes around 50% of bandwidth resources on inter-AS links as compared with scheme (1)-RS&RT, thanks to our algorithm of traffic exchange localization. In order to clearly indicate the performance gap between scheme (3)-LS&RT, scheme (4)-LS< and the penalty cases, we show in Figure 3 (b) the final state when all 800+ peers have already joined the session. When there is no penalty on using inter-AS links against intra-AS ones, scheme (4) with locality-based peer selection scheme and traffic-exchange localization algorithm is able to reduce around 43% bandwidth consumption compared with plain locality-based peer selection in scheme (3). In addition, we can observe that the higher penalties that are put on using inter-AS links, the less bandwidth consumed on them. In the extreme case, we set penalty $n=20$, there is 83% reduction between this penalty case and scheme (3).

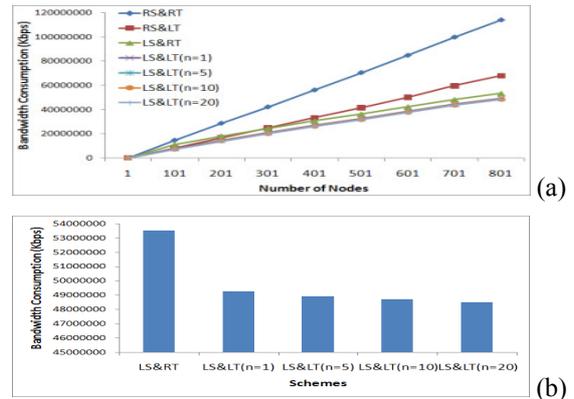


(a) Overall performance (b) Last traffic patterns

Figure 3: Overall inter-AS bandwidth consumption

Our traffic-exchange localization algorithm reduces not only inter-AS traffic but also intra-AS traffic. As Figure 4 (a) shows, the scheme (1)-RS&RT consumes the highest intra-AS bandwidth, and around 40% more bandwidth consumed compared with scheme (2)-RS<, while scheme (2) has 21% higher bandwidth consumption than scheme (3)-LS&RT according to Figure 4 (b). When we apply the locality-based peer selection scheme, the scheme (4)-LS< achieves 8%

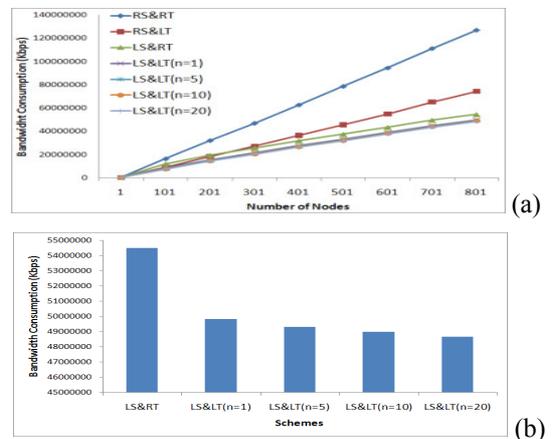
reduction of bandwidth consumption in comparison with scheme (3) in Figure 4 (b). Furthermore, in the penalty situations, we can clearly see that less significant improvement is achieved since penalties are put on inter-AS links for reducing inter-AS traffic.



(a) Overall performance (b) Last traffic patterns

Figure 4: Overall intra-AS bandwidth consumption

Let's consider the overall case where bandwidth consumption on both intra- and inter-AS links are included, as shown in Figure 5. Since our algorithm with locality-based peer selection is mainly to reduce inter-AS traffic that accounts for small proportion of total network traffic, thus the traffic pattern of overall network is similar with intra-AS's pattern. We can still see scheme (1)-RS&RT has the highest values, and scheme (2)-RS< stays at second position in Figure 5(a). In comparison with scheme (3)-LS&RT, scheme (4)-LS< still has 9% lower bandwidth consumption as shown in Figure 5(b). In addition, the penalty cases have also another improvements (11%) achieved compared with scheme (3).

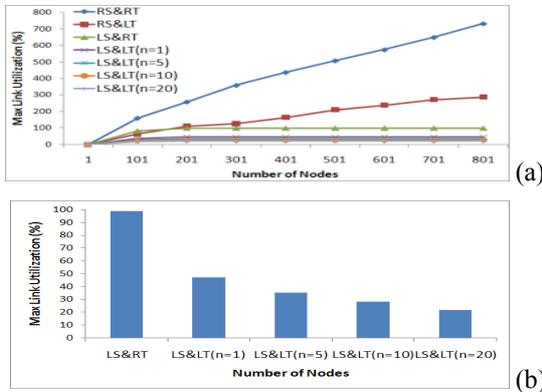


(a) Overall performance (b) Last traffic patterns

Figure 5: Overall network bandwidth consumption

Figure 6 illustrates the maximum link utilization performance across inter-AS links. We can see scheme (1)-RS&RT and scheme (2)-RS< have the highest MLU values among all four schemes, since peers are randomly selected which leads to less balanced traffic distribution. There is a significant gap between scheme (1) and other schemes since both plain locality-based peer selection scheme and plain

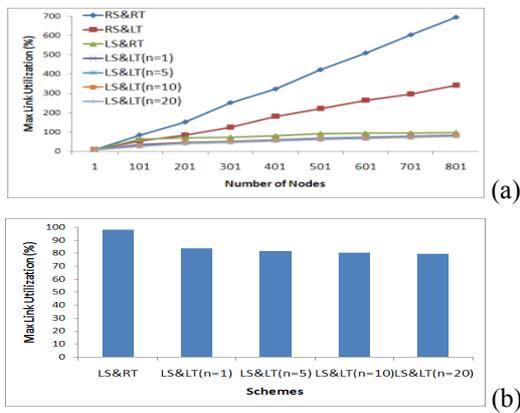
traffic-exchange localization approach are able to significantly reduce inter-AS traffic. Scheme (3)-LS&RT has a higher value of MLU than scheme (2) when there are around 100 active in the network. In the Figure 6 (b) it can be clearly observed that our scheme (4)-LS< reduces 53% MLU compared with scheme (3) because less traffic traverses via inter-AS links. In the penalty cases, for example, the $n=20$ case achieves 78% decreases of MLU compared with scheme (3) while 54% decreases in comparison with scheme (4).



(a) Overall performance (b) Last traffic patterns

Figure 6: Max link utilization on Inter-AS links

We also study the MLU performance on the intra-AS links. Compared with inter-AS case, in Figure 7 (a) scheme (1)-RS&RT and (2)-RS< have higher MLU values than the other two schemes. On the other hand, the MLU performance by scheme (3) is very close to that by scheme (4) according to the Figure. Figure 7 (b) shows that our scheme (4)-LS< reduces 15% MLU compared with scheme (3)-LS&RT, and we also find the improvement of penalty scenarios become trivial because the objective of making penalty is to reduce the inter-AS traffic, rather than intra-AS traffic.



(a) Overall performance (b) Last traffic patterns

Figure 7: Max link utilization on Intra-AS links

Figure 8 shows the network cost performance of three different aspects respectively – intra-AS, inter-AS, overall network cost with the four schemes. In order to clearly show the improvement of our algorithm achieved, we exclude the penalty cases. In Figure 8(a) and (b), we can see that network

cost increases as the number of peers increases, and the network cost on overall and intra-AS cases are very similar with each other. The scheme (4)-LS< incurs only around 50% of network cost by scheme (3)-LS&RT. On the other hand, Figure 8(c) shows that the network cost in inter-AS case always remains the lowest. This result indicates that our algorithm is able to significantly reduce the network cost on inter-AS links.

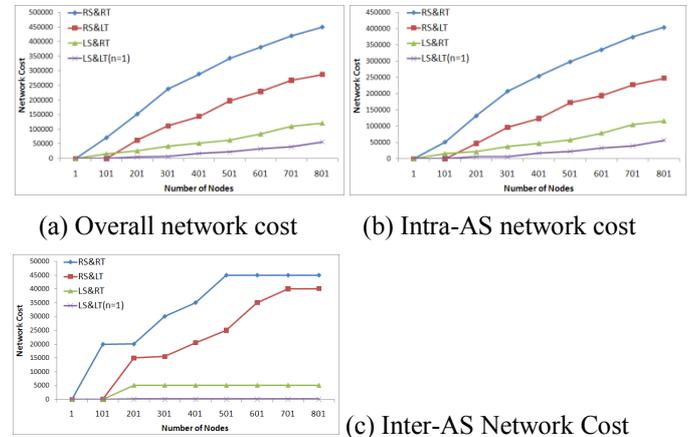


Figure 8: Overview of Network Cost

V. CONCLUSION

By using network information to efficiently select peers in P2P applications has emerged in the past few years. This approach is able to reduce cross-AS traffic and improve the performance of P2P applications. Despite plain locality-based peer selections, the traffic-exchange pattern among peers is still random, which leaves spaces for further reducing the cross-AS P2P traffic. In this paper, we introduce an efficient traffic-exchange localization algorithm that can jointly work with traditional locality-based peer selection paradigms for further performance improvement. According to our algorithm, each peer tries to obtain higher proportion of data from nearby partners than remote ones. According to our simulation results, the proposed traffic-exchange localization algorithm is able to significantly reduce inter-AS traffic and the maximum link utilization in comparison with other random traffic exchange approaches.

REFERENCES

- [1] Sopcast, <http://www.sopcast.com>
- [2] BitTorrent, <http://www.bittorrent.com>
- [3] IGen, <http://www.info.ucl.ac.be/~bqu/igen/>
- [4] V. Aggarwal et al., “Can ISPs and P2P Users Cooperate for Improved Performance?”, ACM CCR, Vol. 37, Issue 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] E. Marocco et al., “Application-Layer Traffic Optimization Problem Statement”, <http://www.ietf.org/proceedings/08jul/agenda/alto.html>
- [8] Fortz and M. Thorup, “Internet Traffic Engineering by Optimizing OSPF Weights”, Proc. IEEE INFOCOM 2000