An Improved Cooperative Model of P2P and ISP

Ying Liu[1,3], Shenglin Zhang[1,3], and Hongying Liu[2]

[1] Institute for Network Sciences and Cyberspace
Tsinghua University, Beijing 100084, China
[2] School of Mathematics and Systems Science
Beihang University, Beijing 100191, China
[3] Tsinghua National Laboratory for Information Science and Technology
Tsinghua University, Beijing 100084, China
liuying@cernet.edu.cn, slzhangsd@gmail.com,liuhongying@buaa.edu.cn

Abstract. P4P (Provider Portal for Applications) is a model aiming to incorporate P2P with ISP and improve the performance of both the ISP and the P2P applications. In this study, we analyze the relationship between the link traffic and the P-distance, which is the core interface of P4P, and illustrate the disadvantage of P4P in dealing with network topology with bottleneck links. Further, with *link utility function* as the optimization objective, we propose an improved model–*Improved-P4P*, making the traffic produced by P2P applications more homogeneous, which can reduce the peak link utilization and protect bottleneck links, and then improve both the network efficiency and the P2P performance. We have built a simulation platform based on BitTorrent and conducted extensive simulations. These simulations demonstrate that *Improved-P4P* achieves a lower cost for ISPs and a better performance for P2P applications than native P2P. Moreover, compared with P4P, *Improved-P4P* reduces traffic on bottleneck links without compromising on the performance of the P2P applications. We believe that relieving of pressure on bottleneck links hold great significance especially in extreme settings. *Improved-P4P* performs steadily in different swarm sizes, proving that it is scalable and easy to deploy.

Keywords: Subgradient Method, P2P, Traffic Engineering, Cooperative Model, Dual Function

1 Introduction

1.1 Background

The P2P (peer-to-peer) concept has fundamental advantages over the traditional C/S (Client/Server) model and the fixed infrastructure content distribution networks because of its excellent robustness and scalability, and plays an important role in modern networks. Some researches have found that more than 50% of the network traffic is introduced by P2P [1, 2], and the massive traffic generated by

P2P brings significant challenges in traffic engineering for ISP (Internet Service Provider)[3].

Because P2P applications are ignorant of the underlying network topology, most P2P applications apply application-level routing that is based only on the overlay network metrics [5]. Moreover, some P2P applications select the source of their downloading randomly, which may lead a P2P user in New York City to download from a user in Los Angeles, while this kind of data is available in New York City or in Washington DC. This kind of long-distance downloading may decrease the network efficiency and the performance of the P2P applications concurrently. P2P applications can avoid this by selecting neighbors with a lower delay or less router hops, but purely selecting neighbors with a low delay or less router hops may cause a Comcast user to select an AT&T user as its neighbor. This cross-ISP neighboring will generate unnecessary interdomain traffic, thereby significantly increasing the operational costs of ISPs. In conclusion, current P2P applications have the following problems:

(1) A P2P system may cause the dispersion of network traffic and make the traffic to unnecessarily flow through multiple intra-domain links. By conducting practical tests, [6] found that every bit of P2P traffic in Verizon needs 5.5 hops when passing through 1000 miles on average, and this average number of hops can be reduced to 0.89 without compromising the P2P application's performance.

(2) A P2P system may generate massive inter-domain traffic or cause massive traffic that when produced by multiple ISPs pass through a specific network [9]. In [8], Karagiannis studied the BitTorrent performance in a college network. He found that this low-efficiency inter-domain traffic may cause significant financial losses for ISPs. Even in the case of the top-level ISP (tier-1 ISPs who do not pay other ISPs), the inter-domain traffic caused by a P2P system can cause the traffic between these tier-1 ISPs to lose balance and then violate the P2P protocol.

However, a one-sided strategy of P2P and ISP is not ideal. As a result, a cooperative model of P2P and ISP should be built, making both sides exchange information and control the P2P traffic cooperatively, and thereby improving the network efficiency and the P2P performance simultaneously.

1.2 Related Works

[4] studied the advantages and disadvantages of P2P, and proposed a middle server called "*oracle*", by which ISPs can provide a neighbor selection policy for P2P users. After a peer sends its list of potential neighbors to "*oracle*", "*oracle*" will sort all the possible neighbors according to certain criteria, such as the nearest principle and the link bandwidth. The sorted neighbor list will guide peers to select neighbors and improve the P2P performance. At the same time, the ISPs can effectively manage massive the P2P traffic with this mechanism, assuring that the traffic does not pass across them and that it is led to the right path. With the abovementioned mechanism, ISPs can provide a better network service for their users.

[7] developed a cooperative model of CP (" *Content Provider*") and ISP, which aimed at co-optimization by sharing control between the ISP and the CP. This study analyzed the optimality of this model by using the game theory and compared this model with the traditional model by means of simulations. The simulation results showed the advantages of the sharing model under circumstances with different congestion levels, and pointed that under some conditions, if the complete network information is shared without any co-optimization control, the co-benefits of the CP and the ISP may be lower than before the sharing of the complete network information.

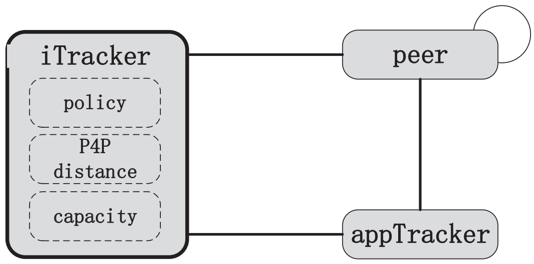

Fig. 1. The interactive between iTracker and P2P

2 Improved cooperative model and theoretical analysis

2.1 P4P

In the P4P model, each ISP maintains an iTracker for its network, and the iTracker has multiple interfaces for ISPs to communicate with P2P applications with respect to the following: (a) static network policy, (b) P-distance mirroring network policy and status, and (c) network capacity. The main interfaces are shown in Figure 1, where the P-distance interface is the core of P4P. Through the P-distance, an ISP can communicate to the P2P applications the current cost of its intradomain and inter-domain links. The P-distance reflects a network's preference and status with respect to the application's cost. The main algorithm of P-distance is the min-max link utilization with a distributed algorithm as follows:

$$\min_{\forall k: t^k \in T^k} \max_{e \in E} \big(b_e + \sum_k \sum_i \sum_j t_{ij}^k I_e(i,j)\big)/c_e, \tag{1}$$

where $b_e$is the background traffic (i.e. traffic that P4P can't control), c_e is the capacity of link e, $I_e(i,j)$ is the indicator link e on the route from PID$-i$ to PID$-j$, and T^k denotes the set of all feasible traffic solutions on the basis of the demand and the property of the P2P applications session k. $t^k = (t_{ij}^k)$, where t_{ij}^k denotes the traffic from PID$-i$ to PID$-j$ in the P2P session k.

The above mentioned PID (opaque ID) is a virtual network point. There are many types of PIDs, one of which is an aggregation point, i.e. a PID represents a set of points. In fact, a PID can also represent PoP, or a set of points with the same congestion state. In this study, each PID represents an aggregation point.

In particular, T^k is made up of t^k that satisfy the constraint condition as follows,

$$\sum_{j:j\neq i} t_{ij}^k \leq u_i^k, \forall i, \tag{2a}$$

$$\sum_{j:j\neq i} t_{ji}^k \leq d_i^k, \forall i, \tag{2b}$$

$$t_{ij}^k \geq 0, \forall i \neq j, \tag{2c}$$

$$t_{ij}^k \geq \underline{\rho}_{ij}^k \sum_{j'\neq i} t_{ij'}^k, \forall i, \quad j \neq i, \tag{2d}$$

$$\sum_i \sum_{j\neq i} t_{ij}^k \geq \beta * OPT, \tag{2e}$$

where u_i^k denotes the aggregation uploading capacity from PID$-i$ to other PIDs in session k, and d_i^k denotes the aggregation downloading capacity from other PIDs to PID$-i$ in session k. $\underline{\rho}_{ij}^k$ is the lower bound on the percentage of traffic from PID$-i$ to PID$-j$ among all the traffic from PID$-i$ to other PIDs. Note that $0 < \underline{\rho}_{ij}^k < 1$ and $\sum_{j\neq i} \underline{\rho}_{ij}^k < 1, \forall i$. β is the efficient factor that can be configured particularly to P2P applications in engineering. The OPT in (2e) is the lower bound of P2P applications performance. Because the cooperation of the ISP and the P2P application aims at improving the performance of both sides, the cooperation should not compromise the performance of P2P applications. Hence, in general, OPT can be set as the optimal value in the independent optimization of P2P applications . Typically, it can be set as follows:

$$OPT = \mathbf{maximize}_{t^k \in T^k} \quad \sum_i \sum_{j\neq i} t_{ij}^k, \tag{3}$$

i.e. P2P aims at matching downloading and uploading.

Suppose that $t_e^k = \sum_i \sum_j t_{ij}^k I_e(i,j)$, i.e. the total traffic produced by P2P in link e, then (1) equals to:

$$\mathbf{minimize}_{\alpha, t^k \in T^k, \forall k} \qquad \alpha \tag{4a}$$

$$\mathbf{subject \quad to} \qquad b_e + \sum_k t_e^k \leq \alpha c_e, \forall e \in E, \tag{4b}$$

The Lagrange dual function of (4a) is as follows:

$$D(p) = \min_{\alpha, \forall t^k \in T^k, k} \sum_e p_e(b_e + \sum_k t_e^k) + (\sum_e p_e c_e - 1)\alpha.$$

To make $D(p)$ finite, the coefficient of α should be zero. i.e.

$$\sum_e p_e c_e = 1.$$

Then

$$D(p) = \sum_e p_e b_e + \sum_k \min_{t^k \in T^k} \sum_e p_e t_e^k. \tag{5}$$

Its dual problem is

$$\textbf{maximize}_{p \geq 0} \quad D(p) \quad \text{subject to } \sum_e p_e c_e = 1. \tag{6}$$

This dual problem can be resolved into independent sub-problems on different sessions of applications with a distributed algorithm,

$$\textbf{minimize}_{t^k \in T^k} \qquad \sum_i \sum_{j \neq i} p_{ij} t_{ij}^k, \tag{7}$$

The aforementioned solution is the interactive optimization algorithm between ISP and a P2P application, i.e. the P2P application solves the subproblem (7) independently and delivers the optimal result $\bar{t}^k$ to iTracker, after which iTracker solves the main problem (6) to update p_e.

Assumption A In the following analysis, we suppose that there exists $\tilde{t}^k \in \tilde{T}^k$ that makes $b_e + \sum_k \tilde{t}_e^k < c_e, \forall e \in E$, i.e. there exists feasible flow solution $\tilde{t}^k$ that makes the restraint on the link capacity strictly feasible.

2.2 Properties of P-distance in P4P

Theorem 1. *Suppose that $\{\tilde{t}_e^k\}$ is the solution to (4), and $\{\tilde{p}_e\}$ is the solution to (6). Then there exists at least one link e whose link utilization is maximal and its corresponding $\tilde{p}_e > 0$. The $\tilde{p}_e$ whose corresponding links utilization doesn't achieve maximum is 0.*

Proof: (4) is an instance of convex programming, and according to assumption A, we know that the Slater constraint specification is true; hence, the strong dual theory is true. As a result, the solution of (4) and of its dual problem (6) satisfy the following:

$$b_e + \sum_k \tilde{t}_e^k \leq \tilde{\alpha} c_e, \forall e \in E \tag{8a}$$

$$1 - \sum_e \tilde{p}_e c_e = 0 \tag{8b}$$

$$\tilde{p}_e \geq 0, \forall e \in E \tag{8c}$$

$$\tilde{p}_e (b_e + \sum_k \tilde{t}_e^k - \tilde{\alpha} c_e) = 0, \forall e \in E \tag{8d}$$

where (8a) is the original feasible condition, (8b) and (8c) are the dual feasible conditions, and (8d) is the complementarity condition. Further, because of the optimization of (4)

$$\tilde{\alpha} = \max_{e \in E}(b_e + \sum_k \tilde{t}_e^k)/c_e,$$

where $\tilde{t}_e^k$ is the solution of (3). Based on (8d), we know that all $\tilde{p}_e$s whose corresponding links do not achieve the maximum utilization equal to 0, and $\tilde{p}_e$s whose corresponding links achieve the maximum are equal to or greater than 0. Moreover, by (8b), we know that there exists at least one $\tilde{p}_e$ that doesn't equal to 0. ■

As we see in the above theorems, all link prices of links that do not achieve the most congested state are equal to 0, i.e., for a flow, the link price of each non-most-congested link is equal to that of another such link. This property of the multiplier makes MLU invalid when dealing with networks with bottleneck links.

In Figure 2, suppose that the capacity of each link is 1, the traffic demand between node 1 and node 3 is 1, and the traffic demand between node 3 and node 4 is 0.9. If we set MLU as the optimization objective of the ISP, the traffic on link (1, 3) will be 0.9 and the traffic on link (1, 2) will be 0.1. Now, there are two bottleneck links in the network, i.e. link (1, 3) and link (3, 4), which lead to the situation that of the two links between node 1 and node 3, one is very congested and the other is free.

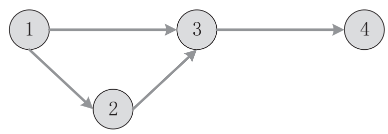

Fig. 2. An Example of Network with Bottleneck Links

2.3 Improved Cooperation Algorithm

We introduce a *link utility function* as the ISP optimization objective and verify that the new objective can better utilize the network resources by carrying out a theoretical analysis and experimental simulations. In this section, we propose a cooperative algorithm of ISP and P2P with a *link utility function* as the ISP optimization objective, and analyze some issues with this objective.

We follow the methods introduced in congestion control[11], and make the *link utility function* the ISP optimization objective. Consider that

$$\textbf{maximize}_{\{s_e\}, t^k \in T^k, \forall k} \quad \sum_e v_\beta(s_e) \tag{9a}$$

$$\textbf{subject to} \quad s_e \leq c_e - \sum_k t_e^k - b_e, \forall e \in E, \tag{9b}$$

where s_e is the free link capacity of link e and $v_\beta(s_e)$ is an increasing concave function. In this paper, we adopt the form as stated in [11].

$$v_\beta(s_e) = \begin{cases} \log(s_e), & \beta = 1 \\ \frac{s_e^{1-\beta}}{1-\beta}, & \beta \neq 1. \end{cases}$$

The Lagrange dual function of (9) is as follows:

$$\begin{aligned} D(p) = & \max_{\{s_e\}, t^k \in T^k, \forall k} \sum_e \Bigg(v_\beta(s_e) \\ & - \sum_e p_e \Bigg(s_e - c_e + \sum_k t_e^k + b_e \Bigg) \Bigg) \\ = & \max_{s_e} \sum_e \big(v_\beta(s_e) - p_e s_e \big) \\ & + \sum_k \min_{t^k \in T^k} \sum_e p_e t_e^k + \sum_e p_e \big(c_e - b_e \big). \end{aligned} \tag{10}$$

The dual problem of (9) is as follows:

$$\min_{p \geq 0} \; D(p). \tag{11}$$

Because $D(p)$ is not differentiable and (11) cannot be solved with the gradient method directly, we solve the problem by using the subgradient method. We can obtain the subgradient of $D(p)$ from [10],

$$\zeta_e = c_e - b_e - \tilde{s}_e - \sum_k \tilde{t}_e^k, \forall e \in E,$$

where $\tilde{s}_e, \{\tilde{t}_e^k\}$ is the solution of

$$\mathbf{maximize}_{c_e - b_e \geq s_e > 0} \big(v_\beta(s_e) - p_e s_e \big), \forall e \in E \tag{12}$$

and

$$\mathbf{minimize}_{t^k \in T^k} \quad \sum_e p_e t_e^k, \forall k. \tag{13}$$

On the basis of the subgradient projection method, p_e can be updated as follows,

$$p_e(\tau + 1) = \begin{cases} p_e(\tau) - \mu(\tau)\zeta_e(\tau), & p_e(\tau) > \mu(\tau)\zeta_e(\tau) \\ 0, & p_e(\tau) \leq \mu(\tau)\zeta_e(\tau) \end{cases}$$

where ζ_e is the subgradient and $\mu(\tau)$ is the step parameter. Theoretically, the step parameter $\mu(\tau)$ is of vital importance to the convergence of this algorithm. However, practically, owing to the continuous evolving of the network and the P2P applications, we can set the step parameter as a constant value.

After solving (11) by using the subgradient method, we obtain the distributed algorithm for solving (9), which is the interactive optimization algorithm of the

ISP and the P2P application. In other words, P2P applications and ISPs can solve the subproblem (13) and (12) independently at first and then transfer the optimal solution $\tilde{t}^k$ and $\tilde{s}_e$ to iTracker. In the next moment, iTracker will update p_e by solving (11).

With respect to $v_\beta(s_e) = \log(s_e)$, we can express the explicit solution of (12) and the engineering significance of its dual variable as follows:

Theorem 2. *When $\beta = 1$, the subproblem (12) has the explicit optimization*

$$\tilde{s}_e = \begin{cases} \frac{1}{p_e}, & p_e \geq \frac{1}{c_e - b_e}, \\ c_e - b_e, & 0 \leq p_e < \frac{1}{c_e - b_e}. \end{cases}$$

Theorem 3. *Suppose that $\{\tilde{t}_e^k\}$ is the solution of (9) and $\{\tilde{p}_e\}$ is the solution of the dual problem (11). Then $\forall e \in E$, and when $\tilde{s}_e = c_e - b_e$,*

$$\tilde{p}_e \in [0, \frac{1}{c_e - b_e});$$

when $\tilde{s}_e < c_e - b_e$,

$$\tilde{p}_e = \frac{1}{c_e - \sum_k \tilde{t}_e^k - b_e}.$$

As we see from theorem 3, to any link e, the larger its free link capacity is, the lower its link price ($\tilde{p}_e$); the smaller its free link capacity is, the higher its link price ($\tilde{p}_e$). This makes the ISPs to control traffic on the non-most-congested links more efficiently when dealing with networks with bottleneck links.

3 Simulation methodology

We have built a discrete-event package for simulation. We have followed the method de-scribed in [12] and performed the simulations by implementing the native BitTorrent protocol. We have also calculated the traffic on every link in order to estimate the link utilization. Our simulation includes keeping statistics of the traffic on bottleneck links, interdomain links, and intradomain links in P2P, P4P, and *Improved-P4P*. Further, we have studied the benefits of P2P applications in P4P and *Improved-P4P* by varying the swarm size.

3.1 Assumptions

- We have ignored the propagation delay because the propagation delay relates only to small control packets. We believe that this simplification has very little impact on the conclusion because of the following:
  1. The downloading time depends on the transmission time of the packets.
  2. In practice, the pipelined processing mechanism of BitTorrent reduces most of the propagation delay of the packets.
- We have followed the method proposed in [13] and assumed that all TCP sessions share the link's capacity equally in the stable state.
- After finishing downloading, all peers leave the network immediately.

3.2 Detailed Parameters

1. Bandwidth between PIDs: 100MBps (bidirection).
2. Size of block: 256 KB.
3. Number of seeds: 1
4. Upstream access link capacity of a seed: 5 MBps.
5. Upstream access link capacity of peers: uniform distribution between 550 KBps to 1000 KBps.
6. Downstream access link capacity of peers: twice the upstream capacity.

Table 1. network topology

Network	Region	Aggregation level	♯Nodes	♯Links
Abilene	US	PoP	11	28

We use PoP-level topologies of Abilene.Table 1 the detailed parameters of A-bilene. Notice that the capacity of each link in Abilene is 10 Gbps on both directions.

3.3 Neighbor Selection Policy

P4P and *Improved-P4P* improve the neighbor selection policy with the interaction of the ISP and the P2P application. For a peer,

1. Neighbor selection within PID: appTracker select a certain number of neighbors within peer i's PID; the ratio of this number to the number of peer i's neighbors should be limited below a certain percentage.
2. Neighbor selection between PIDs: Suppose that the link price between PID i and PID j is p_{ij}. For each $i \neq j$, if $p_{ij} \neq 0$, $w_{ij} = \frac{1}{p_{ij}}$, else we can set w_{ij} to be a very large value.

$$W_{ij} = \frac{w_{ij}}{\sum_{i \neq j} w_{ij}} \quad (14)$$

For peers in PID i, after finishing neighbor selection within PID, they will select a certain percentage of neighbors from other PIDs according to W_{ij}. For robustness, concave function $f(x)$ can be introduced to enlarge the relative weight of W_{ij}. In this paper, we apply $f(x) = \sqrt[6]{x}$.

3.4 Performance Metrics

We consider the following performance metrics:

1. Completion time: It includes statistics of time that all peers need to finish downloading and the time that each single peer needs to finish downloading.
2. P2P bottleneck traffic: It is the total P2P traffic on a link that achieves the maximum link utilization.
3. Interdomain traffic: It is the totalP2P traffic on links between ASes; this metric is used only in interdomain settings.

4 Simulation results

We have detected bottleneck link traffic and traffic on links between PoPs in Abilene, and kept a record of the peers' completion time for different swarm sizes in different topologies. Further, we have obtained the statistics of traffic on links between ASes in interdomain settings.

In this section, we considers a case in which all peers share a 250-MB file, and discuss the performance of P2P, P4P, and *Improved-P4P* for different swarm sizes and network topologies within AS .

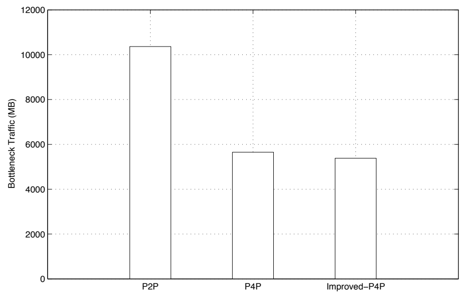

Fig. 3. Variance of Traffic (Swarm Size = 600)

Figure 3 shows that when the swarm size is 600, compared with the native Bit-Torrent that is based on P2P, BitTorrent integrated with *Improved-P4P* can reduce the bottleneck traffic substantially. Moreover, to some degree, the *Improved-P4P* can reduce the P4P's bottleneck. In particular, in the case of Abilene, the bottleneck traffic produced by the native BitTorrent is 1.93 times that produced by the BitTorrent integrated with *Improved-P4P*, and BitTorrent integrated with *Improved-P4P* can reduce the bottleneck traffic of BitTorrent integrated with P4P by 4.74%.

We can conclude from the above result that compared with P2P and P4P, *Improved-P4P* can considerably reduce the bottleneck traffic and better protect the bottleneck.

In conclusion, *Improved-P4P* can reduce bottleneck link utilization and the duration of high-level bottleneck traffic. Undoubtedly, *Improved-P4P* can relieve the pressure on bottleneck links.

We can observe from the result that compared with P2P and P4P, *Improved-P4P* makes the traffic on every link more even, which then reduces the traffic on the bottleneck links.

Figure 4 shows how the normalized completion time varies with the swarm size. The normalized completion time refers to the normalized value based on the maximum average downloading time of the peers in the native BitTorrent system.

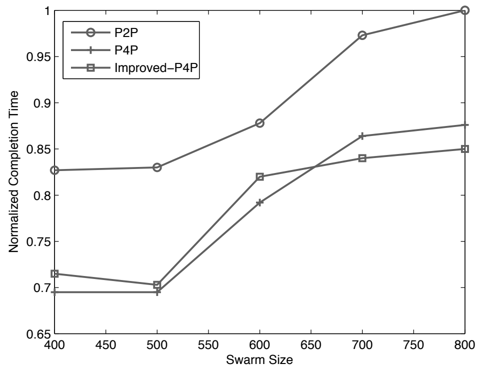

Fig. 4. Average Completion Time

In particular, in the case of Abilene, the BitTorrent integrated with *Improved-P4P* and that integrated with P4P has almost the same average completion time, and they all reduce the average completion time by approximately 20% compared with the native BitTorrent.

In other words, in intradomain settings, *Improved-P4P* can reduce the bottleneck link utilization and therefore relieve the pressure on networks with bottlenecks without degrading the performance of the P2P applications.

5 Conclusion

We analyzed the P-distance in a P4P cooperative model theoretically and showed the relationship between the P-distance and the link traffic. We concluded that for a network topology with bottlenecks, the application of MLU as the optimization objective of the P4P cooperative model does not achieve good performance. Therefore, we proposed an improved model, *Improved-P4P*, which introduced a *link utility function* as its objective. We analyzed the relationship between the link price and the traffic in this model, and demonstrated that *Improved-P4P* could make the traffic in a network more homogeneous. We carried out a considerable number of simulations on P2P, P4P, and *Improved-P4P* in order to verify that the *Improved-P4P* could implement the inter-active control of ISP and P2P with respect to the network traffic and that it could benefit ISPs with a reduction of the bottleneck traffic without compromising the performance of the P2P applications. Further, *Improved-P4P* performed stably for various swarm sizes, proving its feasibility, scalability, and effectiveness. In conclusion, *Improved-P4P* could solve the cooperation problem of P2P and ISP efficiently.

6 Acknowledgement

This research is supported by the National Natural Science Foundation of China (No.60873253), and National Key Basic Research Program ("973" Program) of China (No.2009CB320500 and 2009CB320501).

References

1. Karagiannis, T., Broido, A., Brownlee, A., Claffy, K.C., Faloutsos, M.: Is P2P dying or just hiding. In: Global Telecommunications Conference, pp. 1532-1538. IEEE Press, New York (2004)
2. Steinmetz, R., Wehrle, K.: Peer-to-Peer Systems and Applications, Springer Press, Berlin (2005).
3. Keralapura, R., Taft, N., Chuah, C., Iannaccone, G.: Can ISPs take the heat from overlay networks. In: Proceedings of HotNets III, pp. 8-15 (2004)
4. Aggarwal, V., Feldmann, A., Scheideler, C.: Can ISPs and P2P users cooperate for improved performance. In: ACM SIGCOMM Computer Communication Review, pp. 31-40. ACM Press, New York (2007)
5. Aggarwal, V., Bender, S., Feldmann, A., Wichmann, A.: Methodology for estimating network distances of gnutella neighbors. In: Proceedings of the Workshop on Algorithms and Protocols for Efficient Peer-to-Peer Applications, pp. 219-223 (2004)
6. Xie, H., Yang, Y.R., Krishnamurthy, A., Liu, Y.G.: P4P: Provider portal for applications. In: Proceedings of the ACM SIGCOMM 2008 conference on Data communication, pp. 351-362. ACM Press, New York (2008)
7. Jiang, W., Zhang, R., Rexford, J., Chiang, M.: Cooperative content distribution and traffic engineering in an ISP network. In: Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems, pp. 239-250. ACM Press, New York (2009)
8. Karagiannis, T., Rodriguez, P., Papagiannaki, K.: Should internet service providers fear peer-assisted content distribution. In: Proceedings of the 5th ACM SIGCOMM Conference on Internet Measurement, pp. 63-76. ACM Press, New York (2005)
9. Seetharaman, S., Ammar, M.: Characterizing and mitigating inter-domain policy violations in overlay routes. In: Proceedings of the 2006 14th IEEE International Conference, pp. 259-268. IEEE Press, New York (2006)
10. Palomar, D.P., Chiang, M.: Alternative distributed algorithms for network utility maximization: Framework and applications. IEEE Transactions on Automatic Control, pp. 52(12):2254-2269 (2007)
11. Mo, J., Walrand, J.: Fair end-to-end window-based congestion control. IEEE/ACM Transactions on Networking, pp. 8:556-567 (2000)
12. Bharambe, A.R., Herley, C., Padmanabhan, V.: Analyzing and improving a BitTorrent network's performance mechanisms. In: Proceedings of 25th IEEE International Conference on Computer Communications, pp. 1-12. IEEE Press, New York (2006)
13. Bindal, R., Cao, P., Chan, W., Medval, J., Suwala, G., Bates, T., Zhang, A.: Improving traffic locality in BitTorrent via biased neighbor selection. In: 26th IEEE International Conference on Distributed Computing Systems. IEEE Press, New York (2006)
14. PlanetLab,http://www.planet-lab.org/