Bittorrent traffic optimization in Wireless Mesh Networks with ALTO service

Francesco Paolo D'Elia, Giovanni Di Stasi, Stefano Avallone, Roberto Canonico
Università di Napoli Federico II, Dipartimento di Informatica e Sistemistica,
Via Claudio 21, 80125, Napoli - Italy

***Abstract*—Wireless Mesh Networks (WMNs) have recently gained wide attention as a way to provide Internet connectivity in hard-to-reach areas as well as in big cities. The main problem of such networks is their limited bandwidth availability compared to wired networks. Hence, efficient usage of available communication resources is of crucial importance for WMNs. Peer-to-peer applications, these days, produce a considerable amount of traffic over the Internet. Such applications build their own overlay network by not taking into account the topology and the status of the underlying infrastructure, so causing an inefficient utilization of physical communication resources. This problem is particularly severe for Wireless Mesh Networks. In this paper we propose a cross-layer solution for P2P Bittorrent traffic optimization in WMNs, based on the ALTO (Application Layer Traffic Optimization) Service recently defined by the IETF. Our solution is able to significantly optimize Bittorrent traffic and it does not require any modification to the Bittorrent application. The tests we performed show a reduction of about 27% of download time and a reduction of about 66% of traffic crossing the WMN borders.**

***Keywords*-wireless mesh network, p2p traffic optimization**

I. Introduction

Wireless Mesh Networks (WMNs) are gaining more and more attention as a viable and economic solution for providing Internet access in rural areas as well as in big cities, thanks to their reduced deployment costs. WMNs are comprised of a set of *wireless routers* that provide wireless access to the clients, thus relaying the internal traffic towards the wired Internet, through a limited number of *wireless gateways*. The main limitation of this kind of networks is their scarce bandwidth, mainly due to the limited resources of the shared wireless medium and interference among transmissions. Therefore, for WMNs it is important to optimize wireless spectrum resource usage. Interference may be limited in multi-radio WMNs by properly assigning channels to wireless router interfaces, in combination with the use of 802.11a/g orthogonal channels.

However, WMN resource management schemes should also take into account traffic patterns of today's most popular applications. P2P applications, in particular, produce between 50% and 85% of the overall Internet traffic [1]. In this paper, we focus on the Bittorrent tracker-based file sharing application [2]. Trackers are software entities that play a central role in the creation of the peer-to-peer Bittorrent overlay. A peer that wishes to download a file distributed by Bittorrent, in fact, obtains from the tracker a list of randomly sorted peer nodes to be contacted. The Bittorrent P2P traffic pattern, then, is a direct consequence of the tracker's algorithm used to sort the peers. It has been proved that a random peer selection may lead to several inefficiencies from an ISP point of view. In fact, this means that a node may choose to connect to a peer that is many hops away in the underlying network. In particular, the peer-to-peer traffic may cross several ASes (Autonomous Systems) and ISPs (Internet Service Providers), drastically increasing the provider management costs.

To enforce traffic locality in P2P networks, several solutions have been proposed [3]. For example, peers may infer the characteristics of an end-to-end overlay link by directly probing the network [4], or by exploiting information retrieved from Content Delivery Networks [5], or by querying a decentralized network coordinate system [6]. Alternatively, ISPs may implement specific services that can be used to optimize the overlay creation process [7], [8]. Along this latter line, the ALTO IETF Working Group has defined a service that is aimed at providing P2P applications with information useful to make better-than-random neighbor peer selection [9].

Exploiting a cross-layer approach, in this paper we introduce a strategy for P2P Bittorrent traffic optimization in WMNs. Our strategy is founded on the exchange of network layer information between the WMN routing agents and the Bittorent tracker through an ALTO-based architecture. In particular, we implemented an ALTO service which is able to order the peers to be returned by the tracker in terms of increasing distance from the asking peer. The ALTO service takes advantage of network layer information provided by a number of WMN agents co-located with the WMNs which have peers participating in the swarm. We believe the architecture proposed has several advantages, compared to existing solutions: i) it does not require any modifications to the Bittorrent clients, as commonly required by other optimizations strategies; ii) it exploits, without additional efforts in terms of control messages, the information already gathered on wireless links by the routing daemons; iii) such network layer information is finely tailored to the WMN characteristics, as a pure wireless metric is employed; iv) the architecture is easily deployable in current WMNs as it just requires the deployment of an *WMN agent* by the WMN providers; v) being ALTO based, it is easily extendable, through the definition of new types of agents, *Network* *agents* in the following, each of which is responsible of a different region, e.g., WiMax network; vi) it is able to optimize at a global level, as the information gathered by the Network Agents is centrally managed by the ALTO service.

The rest of the paper is organized as follows. In section II we describe the basic Bittorrent concepts and related terminology. In section III we describe the ALTO Service architecture as it has been proposed by the ALTO Working Group of the IETF. In section IV we describe our main contribution: an ALTO-compliant architecture specifically tailored for Wireless Mesh Networks. In section V we illustrate an experimental evaluation of the effectiveness of our system to control Bittorrent traffic in a WMN testbed connected to the Internet. In section VI we compare our approach with other recent proposals aimed at improving peer-to-peer traffic locality. In section VII we draw our conclusions.

II. Tracker-based Bittorrent Protocol

Bittorrent is a peer-to-peer application specifically designed to efficiently distribute large amounts of data over the Internet [2]. Bittorrent typical use is for file sharing. In particular, Bittorrent is useful when large files need to be distributed to large communities of users, e.g. when a new distribution of a popular Operating System is released. The basic idea behind Bittorrent is that a peer, while downloading a large file, may contribute to the spreading of the content by making it available to other peers. In this way, the whole community of peers may benefit from a greater aggregate upload bandwidth. The set of all the nodes (*peers*) sharing the same resource is called a *swarm*.

Peers wanting to join a swarm, first need to obtain the metadata *torrent* file which contains the URL of the tracker entity associated to that specific swarm. The tracker will provide peers with the list of other peers to be contacted in order to join the swarm. Bittorrent clients contact peers on different TCP connections. Peers that own the complete shared resource are called *seeders*.

III. The ALTO Service architecture

The Application Layer Traffic Optimization Service is a recently established IETF Working Group which has the objective of defining ”an information-sharing service that enables applications to perform better-than-random peer selection” [9]. The idea behind the Application Layer Traffic Optimization is to provide P2P applications with knowledge, from the provider perspective, about the underlying network. This kind of information may be exploited by the P2P application, improving user perceived performance while at the same time optimizing resource usage. In a possible overlay network topology creation, the ALTO Service would be very helpful in the choice of the connection endpoints. The ALTO Service, directly controlled by the network manager, is also aware of the network resources availability, the critical bottlenecks, and the location of the peer nodes inside the network. An ALTO Server, then, provides oracle-like functionalities, supplying network information to applications. Typical information may be node connectivity properties, location of endpoints and costs for paths between network hosts.

Therefore, applications may no longer need to rely on measuring link performance metrics by themselves, and obtain information from the ALTO servers.

The ALTO protocol structure is composed of two main services:

- Server Information Service which provides the list of information that can be retrieved from the ALTO server (or the other ALTO Servers that refer to the one contacted), like cost metrics and possible operations.
- ALTO Information Services: several kinds of information may be provided by the ALTO server, each one encoding different kinds of network provider preferences.

The ALTO architecture defines entities called “Sources of topological information” whose function is to provide the ALTO server with information on specific network regions.

IV. An ALTO architecture for wireless mesh networks

The main contribution of this work is the definition of a strategy for the optimization of Bittorrent traffic and the definition of an ALTO-based architecture which acts accordingly.

In Section III we showed the service defined by the ALTO Working Group, the ALTO Service, that provides mechanisms to convey network’s preferences to applications. Accordingly with the ALTO specifications, we implemented an inter-domain ALTO server, to which are conveyed information from the perspective of different network regions. In particular, in order to provide the ALTO Server with information of the WMNs involved, we implemented a WMN Agent, an entity responsible for the peer ranking inside a WMN region. The WMN agent is deployed on a router inside the WMN, preferably on a mesh gateway. In this way it can easily have access to the routing information and link performances.

Our architecture, in general, defines the *Network agents*, software entities which manage and acquire information on specific network regions on behalf of the ALTO Server. Network agents act, as defined by ALTO specifications, as “Sources of topological information” for the ALTO Server.

The architecture we propose is depicted in figure 1. The process of constructing the peer list to be returned to the asking peer goes through three main steps. The first step, performed by an inter-domain ALTO server, is to group in network regions the peers participating to the swarm. A network region is defined as the set of peers which is

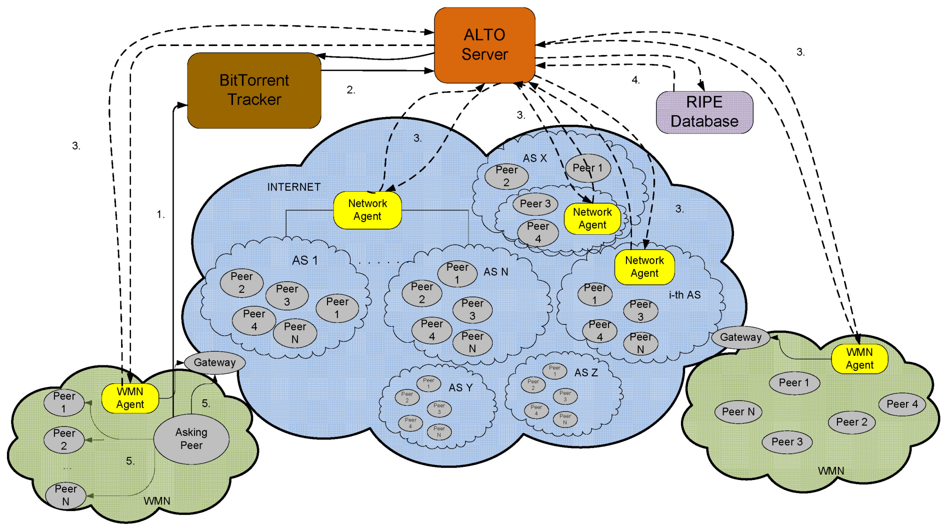

Figure 1. Architecture of the ALTO Service. **(1)** The peer node queries the tracker for a list of peers; **(2)** The tracker is configured to relay the request to the ALTO Server, providing the peer list to be ordered and the IP address of the asking Peer; **(3)** The ALTO Server verifies the Agents presence and delegates the nodes ranking to the appropriate Agents; depending on whether the Agent belongs to the same network region as the asking peer or not, the peer ranking is performed by taking into consideration the cost of the network path towards the asking peer or towards the gateway; **(4)** The nodes that are not covered by an Agent are ranked directly by the ALTO Server, exploiting the AS-region information obtained by queuing the RIPE server. to the Server: in the event that the requesting node belongs to a wireless network operated by the agent, the cost from the asking peer is considered; otherwise, the Agents take into consideration the cost of each peer from the gateway. Then, the ALTO Server collects the information from all the available Agents and it returns the complete peer list to the tracker, which will turn it over to the asking peer node; **(5)** Finally, the node establishes connections with the nodes on the ranked peer list.

managed by an Agent, e.g., a WMN managed by an WMN Agent. Peers not managed by any Agent are treated like groups with one element. The second step is to order the regions in terms of increasing *AS-distance* from the asking peer. The AS-distance is defined as the number of ASes (Autonomous Systems) to be traversed from the source to the destination considered, i.e. the asking peer and the region gateway in our case. The third step, accomplished with the help of the Network agents, is to order the peers inside every region. Each Network Agent receives from the ALTO Server the list of peers belonging to the region it controls and returns the same list ordered in terms of increasing cost either from the region gateway or from the asking peer, respectively if the asking peer is external or inside the Network Agent region.

In addition to order WMN peers depending on the wireless metric, the WMN agent also accomplishes the important function of allowing direct connections between WMN peers in presence of a NAT. In such a case, peers announce themselves to the tracker with the NAT public ip address and port. For this reason, while ranking the WMN peers, the WMN agent translates the NAT public ip address and port couples of peers which are located in the same WMN of the asking peer with their local counterparts. This allows the asking peer to establish direct local connections with peers in its WMN, without the need of traversing the NAT gateway or the need to implement a local peer discovery process.

A. Implementation details

We implemented the WMN agent as a plugin for *OLSRD*, a widely-used routing agent for wireless mesh networks.

The plugin, having access to all the routing information kept by OLSRD, i.e., nodes, edges and the cost of edges in the chosen metric, is able to calculate, by applying the Dijkstra algorithm, the shortest distances from the asking peer, or from the WMN gateway, to all the co-located peers and then sort them. The metric used for the cost of links is the ETX (Expected Transmission Count) [10]. The ETX metric minimizes the expected total number of packet transmissions (including retransmissions) required to successfully deliver a packet to the ultimate destination. It incorporates the effects of link loss ratios, the asymmetry in the loss ratios between the two directions of each link and the interference among the successive links of a path. Such metric, being a pure wireless metric, is able to significantly improve the overlay creation, as shown in the experimental section. Other wireless metric, however, are supported by OLSRD and can be easily used with the proposed architecture just by changing the routing daemon configuration.

V. Experimental evaluation

The objective of the experimental campaign is to provide a practical example of our architecture at work and to show the goodness of the optimization strategy we are proposing.

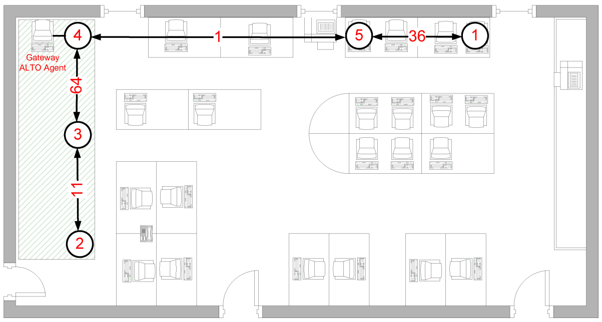

Figure 2. Wireless mesh network setup

We run on this purpose two different kinds of tests: the first one involving a single WMN in order to show the improvement in the overlay creation inside the WMN and the consequent reduction of the download time; the second one involving nodes from a global context, i.e. nodes from both the WILEE testbed and PlanetLab Europe, in order to show the improvements in terms of traffic locality, i.e. the improvement in the amount of traffic kept inside the same AS. Both the kinds of tests consisted in the execution of a series of complete file transfers between a seeder, i.e., the host with the file to spread, and a few leechers, i.e. the hosts that download the file.

The nodes used belong to WILEE (WIreLess Experimental), a wireless mesh testbed located at the University "Federico II" of Naples. It consists of SBCs (Single Board Computers) based on the Intel Network Processor IXP4225 which run as operating system OpenWRT, the well-known Linux distribution for embedded devices.

A. Single WMN experiment

For the first set of experiments we set up the WILEE testbed to form a chain topology of five-hop of length (see Figure 2). This setup, even if simple, is appropriate to show the improvement obtainable by carefully creating the overlay. A wrong overlay creation strategy, in fact, can lead a node to choose a neighbor, at the overlay level, which is many hops away from it, e.g. node 1 could choose node 5 as neighbor in the overlay, which is five hops away, thus laying the overlay link on all the physical links and wasting resources.

The setup of the testbed was the following. Node 4 was chosen as seeder of the experiment, i.e., it was the node with the file to spread. This node was also connected to the WMN Gateway and hosted the WMN Agent. The tracker was installed on a host located on an subnet adjacent to the WMN testbed. In order to take into account the limited size of the testbed, we set the number of nodes to be returned by the tracker to 2.

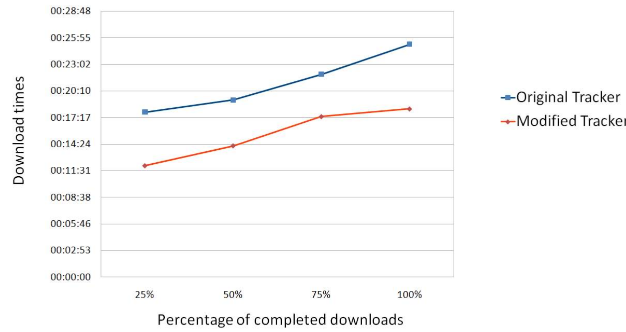

Figure 3. Download times with and without the ALTO Service

The experiments consisted in the complete download of a file owned by the seeder by all the nodes of the WMN testbed.

We employed for the experiments alternatively two different architectures: the first one featuring only a standard Bittorrent tracker, *Quash* [1]; the second one featuring the same tracker, modified by us in order to be ALTO compliant, an ALTO Server and the WMN Agent relative to the WILEE testbed. Both architectures featured *Transmission* [2], an open source Bittorrent client. We outline that no modification were needed at client side, but it was just needed to enable the tracker for the interaction with the ALTO server.

The results of ten experiments, five for each architecture, are reported in Fig. 3. In the x-axis is represented the percentage of the nodes that have completed the download, while on the y-axis are represented the complete download times. It can be seen that the time required for all the peers to complete the download in case our ALTO architecture is employed is significantly smaller than the time needed in case the standard tracker is employed. In particular, we have a reduction in the download times of about 27%. The results can be easily explained if we analyze the overlay created in the two cases. When the ALTO architecture was used, each node was downloading from its immediate neighbors, e.g. node 2 was downloading from nodes 3 and 4; at the contrary, when the standard Tracker was used, nodes were often downloading from nodes which were several hops away, e.g. node 2 was downloading from nodes 1 and 5.

B. Global-scale experiment

In order to test our architecture in real world conditions, we performed a second set of experiments in a scenario where, in addition to the WMN testbed, nodes belonging to the PlanetLab Europe testbed were involved (see Fig. 4).

The purpose of this experiment was to show the improvement, in terms of percentage of traffic locality, i.e.

[1] Quash - http://code.google.com/p/quash/
[2] Transmission http://www.transmissionbt.com/

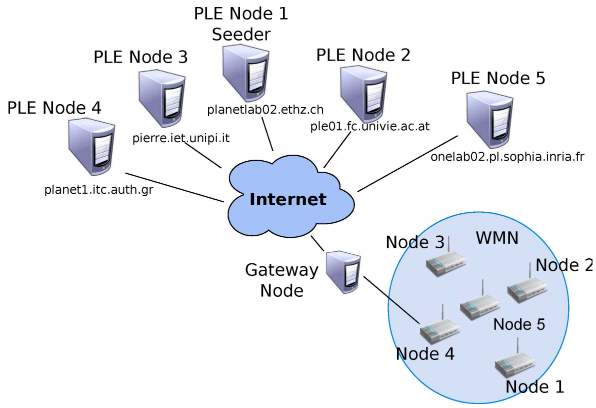

Figure 4. Global-scale experiment setup.

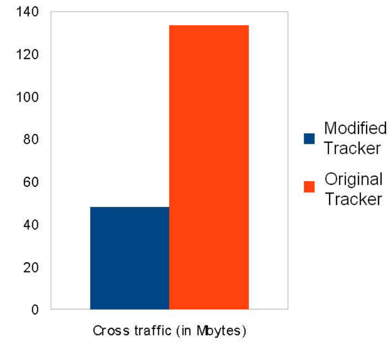

Figure 5. Amount of exchanged cross traffic with and without the ALTO Service

percentage of traffic kept in the same AS (Autonomous System), obtainable where our ALTO based architecture is employed. As previously stated, it is important for a p2p application to be able to keep its traffic as much as possible inside the same AS, so as to lower the cost for the ISPs.

We performed two sets of experiments in this new scenario, with the same architectures described in the previous section. In Fig. 5, we report the results, averaged over 10 different repetitions, which show that the amount of Bittorrent traffic flowing through the WMN Gateway, i.e., the amount of traffic exchanged between the nodes inside the WMN and the external nodes, was significantly lower when the ALTO Service was exploited. In particular, the results show a reduction of about 66% of the overall amount of traffic crossing the WMN borders. This means that the ALTO Service significantly helps the network provider in localizing the traffic inside the WMN (or, similarly, inside the same Autonomous System).

VI. Related work

Several paper have been recently published on the peer-to-peer traffic optimization problem. In [7], differences in overlay and underlay routing are analyzed, like topology creation, routing criteria and different dynamics. In particular, the overlay topology creation problem is addressed. Each node should pick the neighbor with smallest delay, possibly highest throughput, and in the same AS: in the paper the cross-layer information exchange is welcomed and an "oracle" service is studied. The "oracle" is a service supplied by the ISP to the overlay P2P users. The paper shows how explicit cross-layer information exchange can really improve overlay operations and at the same time improve underlay network usage. Despite the difficulties in oracle management, the results suggest that cross-layer information exchange is a valid starting point in solving the overlay-underlay frictions.

In [11] the authors evaluate how the choice of local peers improves the download time and the locality of connections. The locality is defined in terms of the ASs (Autonomous Systems) the peers belong to. We also take into account the AS a node belong, but we further optimize the choice in case of a WMN participates to the swarm, by ordering the peers inside the WMNs depending on their distance from the asking peer.

One of the first works considering ALTO as an architectural choice for traffic optimization in P2P systems is [12]. The work is more focused on the ALTO architecture as it is being designed, and on its relations with existing solutions. No attempt is made in order to leverage its features for specific optimization cases, like when a significant amount of peers is located inside a WMN.

Other efforts have been made in order to improve the behavior of P2P applications in the context of wireless networks. In [13], a complete new Bittorrent client is developed and the WMN is flooded with control messages in order for local peers, i.e. peers within the same WMN, to be able to find and rank themselves depending on their reciprocal distances, and then instantiate local connections. In [14], the authors propose a method for improving efficiency of Bittorrent using network information in the MANET environment, in particular the authors propose network usage information as a parameter for peer selection. In order to obtain what they want, the authors have to substantially modify the OLSR routing protocol and the P2P Bittorrent client. Our approach, which is ALTO based, does not require modifications to the client and does not require to flood the network with control messages, as the needed information is taken from the standard OLSR routing daemon. In [15], the authors propose a joint peer selection and utility maximization problem as a mixed integer nonlinear programming, and they introduce a penalty-based heuristic based on genetic algorithm for the solution of the MINLP. The model proposed by the authors is not realistic for Bittorrent, though, since each peer is allowed to download only from one seeder.

Moreover, finding the single and only source for each peer in the network with a genetic algorithm is infeasible for practical implementation in large scale p2p networks, due to the long calculation times. Finally, a non-standard MAC layer is used, nor auto-rate and variable transmitting powers are considered, and the WMN considered is very small and isolated, with no gateways or external peers. In [16] the authors present an adaptation of Chord for Wireless Mesh Networks which requires nodes position and exploits the 1-hop broadcast transmission peculiarity. Work [17] is probably the most similar to ours. The authors propose the interaction between the application layer Gnutella client and the routing agents, so as to enhance the speed of the overlay creation and improve the adherence of the overlay to the WMN physical topology. The main limitation of this work is that it just exploits the local view of the network given by the routing agents, and do not try to harmonize such view in the global context, i.e. overlay creation and traffic optimization does not extend beyond the borders of a single WMN. Our approach, which is based on a hierarchy of entities, i.e., a central ALTO Server and several Network Agents, has a global view of the networks involved, and therefore can optimize the overlay creation and management on a global scale.

VII. Conclusions

In this paper we addressed the P2P traffic management problem in Wireless Mesh Networks. We outlined how traffic flow control is a priority in WMNs, since wireless channel configuration is finely tailored on the expected traffic flows. A way to control P2P application traffic is to control the choice of the connection endpoints in a P2P swarm. We referred to tracker-based file sharing Peer-to-Peer applications and to the IETF Application Layer Traffic Optimization Working Group solutions. We exploited an ALTO compliant solution and integrated it with an WMN Agent, which is responsible for ranking the mesh nodes. The WMN Agent is usually deployed at wireless network access gateways and therefore it is able to directly evaluate link performance and to point out the best routing path between two nodes. We implemented the proposed ALTO architecture, i.e. an ALTO Server and a WMN Agent, and shown the improvements obtained for the end user application and the reduced costs for the network operator.

References

[1] T. Karagiannis, A. Broido, N. Brownlee, K. Claffy, and M. Faloutsos, "Is p2p dying or just hiding?" in *Proceedings of IEEE GLOBECOM '04*, vol. 3, november-december 2004, pp. 1532–1538.

[2] B. Cohen, "The bittorrent protocol specification," http://www.bittorrent.org/beps/bep_0003.html.

[3] V. K. Gurbani, V. Hilt, I. Rimac, M. Tomsu, and E. Marocco, "A survey of research on the application-layer traffic optimization problem and the need for layer cooperation," *IEEE Communications Magazine*, vol. 47, pp. 107–112, August 2009.

[4] C. Tang and P. McKinley, "On the cost-quality tradeoff in topology-aware overlay path probing," in *Proceedings of the 11th IEEE International Conference on Network Protocols, ICNP 2003*, nov. 2003, pp. 268–279.

[5] D. R. Choffnes and F. E. Bustamante, "Taming the torrent: a practical approach to reducing cross-ISP traffic in peer-to-peer systems," *ACM SIGCOMM Computer Communications Review*, vol. 38, no. 4, pp. 363–374, 2008.

[6] F. Dabek, R. Cox, F. Kaashoek, and R. Morris, "Vivaldi: a decentralized network coordinate system," *ACM SIGCOMM Computer Communications Review*, vol. 34, no. 4, pp. 15–26, 2004.

[7] V. Aggarwal, A. Feldmann, and C. Scheideler, "Can ISPs and P2P systems co-operate for improved performance?" *ACM SIGCOMM Computer Communications Review*, vol. 37, no. 3, pp. 29–40, July 2007.

[8] H. Xie, R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz, "P4P: Provider Portal for Applications," *ACM SIGCOMM Computer Communications Review*, vol. 38, no. 4, pp. 351–362, October 2008.

[9] J. Seedorf and E. Burger, "RFC5693 - Application-Layer Traffic Optimization (ALTO) Problem Statement," Oct 2009. [Online]. Available: http://www.faqs.org/rfcs/rfc5693.html

[10] D. Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," *Wireless Networks*, vol. 11, no. 4, pp. 419–434, 2005.

[11] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, and A. Zhang, "Improving traffic locality in BitTorrent via biased neighbor selection," in *Distributed Computing Systems, 2006. ICDCS 2006. 26th IEEE International Conference on*. IEEE, 2006, p. 66.

[12] J. Seedorf, S. Kiesel, and M. Stiemerling, "Traffic localization for P2P-applications: The ALTO approach," in *Peer-to-Peer Computing, 2009. P2P'09. IEEE Ninth International Conference on*. IEEE, 2009, pp. 171–177.

[13] M. Sbai, C. Barakat, J. Choi, A. Hamra, and T. Turletti, "Adapting BitTorrent to wireless ad hoc networks," *Ad-hoc, Mobile and Wireless Networks*, pp. 189–203, 2008.

[14] S. K. Ko and Y. H. Kim, "Network-aware p2p content sharing over manet," in *Mobile Data Management: Systems, Services and Middleware, 2009. MDM '09. Tenth International Conference on*, may 2009, pp. 661 –665.

[15] M. Zulhasnine, C. Huang, and A. Srinivasan, "Penalty function method for peer selection over wireless mesh network," in *Vehicular Technology Conference Fall (VTC 2010-Fall), 2010 IEEE 72nd*, sept. 2010, pp. 1 –5.

[16] S. Burresi, C. Canali, M. Renda, and P. Santi, "Meshchord: A location-aware, cross-layer specialization of chord for wireless mesh networks," mar. 2008, pp. 206 –212.

[17] M. Conti, E. Gregori, and G. Turi, "A cross-layer optimization of gnutella for mobile ad hoc networks," in *Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*. ACM, 2005, pp. 343–354.