Mobile File-Sharing over P2P Networks

Lu Yan
Åbo Akademi, Finland

INTRODUCTION

Peer-to-peer (P2P) computing is a networking and distributed computing paradigm which allows the sharing of computing resources and services by direct, symmetric interaction between computers. With the advance in mobile wireless communication technology and the increasing number of mobile users, peer-to-peer computing, in both academic research and industrial development, has recently begun to extend its scope to address problems relevant to mobile devices and wireless networks.

The mobile ad hoc network (MANET) and P2P systems share key characteristics including self-organization and decentralization, and both need to solve the same fundamental problem: connectivity. Although it seems natural and attractive to deploy P2P systems over MANET due to this common nature, the special characteristics of mobile environments and the diversity in wireless networks bring new challenges for research in P2P computing.

Currently, most P2P systems work on wired Internet, which depends on application layer connections among peers, forming an application layer overlay network. In MANET, overlay is also formed dynamically via connections among peers, but without requiring any wired infrastructure. So the major differences between P2P and MANET that concern us in this article are:

a. P2P is generally referred to the application layer, but MANET is generally referred to the network layer, which is a lower layer concerning network access issues. Thus, the immediate result of this layer partition reflects the difference of the packet transmission methods between P2P and MANET: the P2P overlay is a unicast network with virtual broadcast consisting of numerous single unicast packets, while the MANET overlay always performs physical broadcasting.
b. Peers in P2P overlay are usually referred to static node though no priori knowledge of arriving and departing is assumed, but peers in MANET are usually referred to mobile node since connections are usually constrained by physical factors like limited battery energy, bandwidth, computing power, and so forth.

BACKGROUND

Since both P2P and MANET are becoming popular only in recent years, the research on P2P systems over MANET is still in its early stage. The first documented system is Proem (Kortuem et al., 2001), which is a P2P platform for developing mobile P2P applications, but it seems to be a rough one, and only IEEE 802.11b in ad hoc mode is supported. 7DS (Papadopouli & Schulzrinne, 2001) is another primitive attempt to enable P2P resource sharing and information dissemination in mobile environments, but it is rather a P2P architecture proposal than a practical application. In a recent paper, Lindemann and Waldhorst (2002) proposed *passive distributed indexing* for such kinds of systems to improve the search efficiency of P2P systems over MANET, and in ORION (Klemm, Lindemann & Waldhorst, 2003), a broadcast-over-broadcast routing protocol was proposed. The above works were focused on either P2P architecture or routing schema design, but how efficient the approach is and what the performance experienced by users is—these are still in need of further investigation.

Previous work on performance study of P2P over MANET was mostly based on simulative approach, and no concrete analytical mode was introduced. Performance issues of these kinds of systems were first discussed in Goel, Singh, and Xu (2002), but it simply shows the experiment results and no further analysis was presented. There is a survey of such kinds of systems in Ding and Bhargava (2004), but no further conclusions were derived. A sophisticated experiment and discussion on P2P communication in MANET can be found in Hsieh and Sivakumar (2004). However, all above works fall into a practical experience report category, and no performance models are proposed.

There have been many routing protocols in P2P networks and MANET respectively. For instance, one can find a very substantial P2P routing scheme survey from HP Labs in Milojicic et al. (2002), and U.S. Navy Research publishes ongoing MANET routing schemes (MANET, n.d.); but all of the above schemes fall into two basic categories: broadcast-like and DHT-like. More specifically, most early P2P search algorithms, such as in Gnutella (www.gnutella.com), Freenet (freenet.sourceforge.net), and Kazaa (www.kazaa.com), are

Figure 1. Broadcast over broadcast

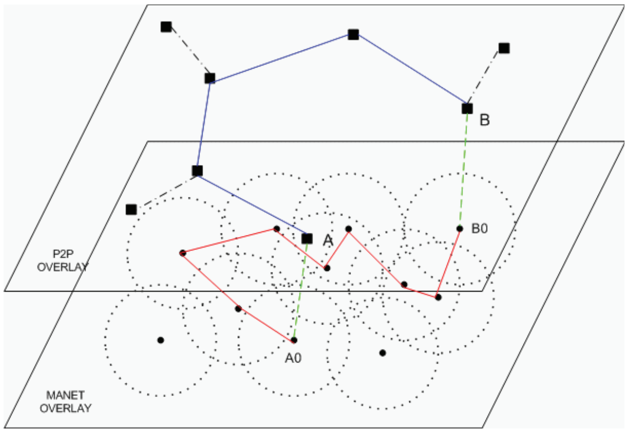

broadcast-like, and some recent P2P searching, like in eMule (www.emule-project.net) and BitTorrent (http://bittorrent.com/), employs more or less some feathers of DHT. On the MANET side, most on-demand routing protocols such as DSR (n.d.) and AODV (n.d.) are basically broadcast-like. Therefore, we here introduce different approaches to integrate these protocols in different ways according to categories.

BROADCAST OVER BROADCAST

The most straightforward approach is to employ a broadcast-like P2P routing protocol at the application layer over a broadcast-like MANET routing protocol at the network layer. Intuitively, in the above settings, every routing message broadcasting to the virtual neighbors at the application layer will result in a full broadcast to the corresponding physical neighbors at the network layer.

The scheme is illustrated in Figure 1 with a searching example: peer A in the P2P overlay is trying to search for a particular piece of information, which is actually available in peer B. Due to broadcast mechanism, the search request is transmitted to its neighbors, and recursively to all the members in the network, until a match is found or it times-out. Here we use the blue lines to represent the routing path at this application layer. Then we map this searching process into the MANET overlay, where node A0 is the corresponding mobile node to the peer A in the P2P overlay, and B0 is related to B in the same way. Since the MANET overlay also employs a broadcast-like routing protocol, the request from node A0 is flooded (broadcast) to directly connected neighbors, which in turn flood their neighbors and so on, until the request is answered or a maximum number of flooding steps occur. The route establishing lines in that network layer are highlighted in red, where we can find that there are few overlapping routes between these two layers, though they all employ broadcast-like protocols.

We have studied a typical broadcast-like P2P protocol, Gnutella (Clip2, 2001), in previous work (Yan & Sere, 2003). This is a pure P2P protocol, in which no advertisement of shared resources (e.g., directory or index server) occurs. Instead, each request from a peer is broadcast to directly connected peers, which themselves broadcast this request to their directly connected peers and so on, until the request is answered or a maximum number of broadcast steps occur. It is easy to see that this protocol requires a lot of network bandwidth, and it does not prove to be very scalable. The complexity of this routing algorithm is O(n) (Ripeanu, Foster, & Iamnitch, 2002; Chawathe, Ratnasamy, Breslau, & Shenker, 2003).

Generally, most on-demand MANET protocols, like DSR (Johnson & Maltz, 1996) and AODV (Perkins & Royer, 2000), are broadcast-like in nature (Kojima, Harada, & Fujise, 2001). Previously, one typical broadcast-like MANET protocol, AODV, was studied (Yan & Ni, 2004). In that protocol, each node maintains a routing table only for active destinations: when a node needs a route to a destinations, a path discovery procedure is started based on a RREQ (route request) packet; the packet will not collect a complete path (with all IDs of involved nodes) but only a hop count; when the packet reaches a node that has the destination in its routing table, or the destination itself, a RREP (route reply) packet is sent back to the source (through the path that has been set up by the RREQ packet), which will insert the destination in its routing table and will associate the neighbor from which the RREP was received as the preferred neighbor to that destination. Simply speaking, when a source node wants to send a packet to a destination, if it does not know a valid route, it initiates a route discovery process by flooding the RREQ packet through the network. AODV is a pure on-demand protocol, as only nodes along a path maintain routing information and exchange routing tables. The complexity of that routing algorithm is O(n) (Royer & Toh, 1999).

This approach is probably the easiest one to implement, but the drawback is also obvious: the routing path of the requesting message is not the shortest path between source and destination (e.g., the red line in Figure 1), because the virtual neighbors in the P2P overlay are not necessarily also the physical neighbors in the MANET overlay, and actually these nodes might be physically far away from each other. Therefore, the resulting routing algorithm complexity of this broadcast-over-broadcast scheme is unfortunately $O(n^2)$, though each layer's routing algorithm complexity is O(n) respectively.

It is not practical to deploy such a scheme for its serious scalability problem due to the double broadcast; and taking the energy consumption portion into consideration, which is somehow critical to mobile devices, the double broadcast will

Figure 2. DHT over broadcast

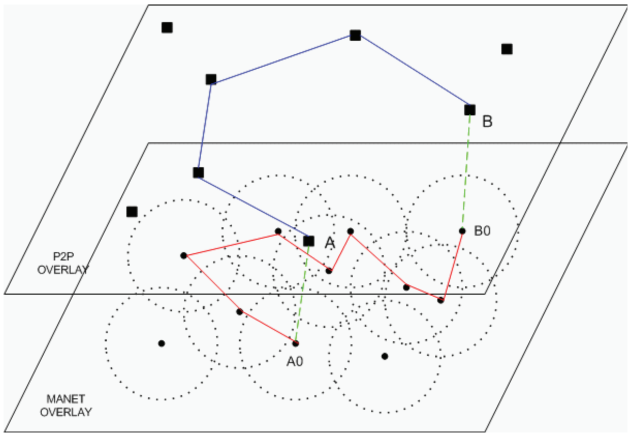

also cost a lot of energy consumption and make it infeasible in cellular wireless data networks.

DHT OVER BROADCAST

The scalability problem of broadcast-like protocols has long been observed, and many revisions and improvement schemas are proposed (Lv, Ratnasamy, & Shenker, 2002; Yang & Garcia-Molina, 2002; Chawathe et al., 2003). To overcome the scaling problems in broadcast-like protocols where data placement and overlay network construction are essentially random, a number of proposals are focused on structured overlay designs. The distributed hash table (DHT) (Stoica, Morris, Karger, Kaashoek, & Balakrishnan, 2001) and its varieties (Ratnasamy, Francis, Handley, Karp, & Schenker 2001; Rowstron & Druschel, 2001; Zhao et al., 2004) advocated by Microsoft Research seem to be promising routing algorithms for overlay networks. Therefore it is interesting to see the second approach: to employ a DHT-like P2P routing protocol at the application layer over a broadcast-like MANET routing protocol at the network layer.

The scheme is illustrated in Figure 2 with the same searching example. Compared to the previous approach, the difference lies in the P2P overlay: in a DHT-like protocol, files are associated to keys (e.g., produced by hashing the file name); each node in the system handles a portion of the hash space and is responsible for storing a certain range of keys. After a lookup for a certain key, the system returns the identity (e.g., the IP address) of the node storing the object with that key. The DHT functionality allows nodes to put and get files based on their key, and each node handles a portion of the hash space and is responsible for a certain key range. Therefore, routing is location-deterministic distributed lookup (e.g., the blue line in Figure 2).

DHT was first proposed in Plaxton, Rajaraman, and Richa (1997) without intention to address P2P routing problems. DHT soon proved to be a useful substrate for large distributed systems, and a number of projects are proposed to build Internet-scale facilities layered above DHTs; among them are Chord, CAN, Pastry, and Tapestry. All take a key as input and route a message to the node responsible for that key. Nodes have identifiers, taken from the same space as the keys. Each node maintains a routing table consisting of a small subset of nodes in the system. When a node receives a query for a key for which it is not responsible, the node routes the query to the hashed neighbor node towards resolving the query. In such a design, for a system with n nodes, each node has O(log n) neighbors, and the complexity of the DHT-like routing algorithm is O(log n) (Ratnasamy, Shenker, & Stoica, 2002).

Additional work is required to implement this approach, partly because DHT requires a periodical maintenance (i.e., it is just like an Internet-scale hash table or a large distributed database); since each node maintains a routing table (i.e., hashed keys) to its neighbors according to DHT algorithm, following a node join or leave, there is always a nearest key reassignment between nodes.

This DHT-over-broadcast approach is obviously better than the previous one, but it still does not solve the shortest path problem as in the broadcast-over-broadcast scheme. Though the P2P overlay algorithm complexity is optimized to O(log n), the mapped message routing in the MANET overlay is still in the broadcast fashion with complexity O(n); the resulting algorithm complexity of this approach is as high as O(n log n).

This approach still requires a lot of network bandwidth and hence does not prove to be very scalable, but it is efficient in limited communities such as a company network.

CROSS-LAYER BROADCAST

A further step of the broadcast-over-broadcast approach would be a cross-layer broadcast. Due to similarity of broadcast-like P2P and MANET protocols, the second broadcast could be skipped if the peers in the P2P overlay would be mapped directly into the MANET overlay, and the result of this approach would be the merge of application layer and network layer (i.e., the virtual neighbors in P2P overlay overlaps the physical neighbors in MANET overlay).

The scheme is illustrated in Figure 3, where the advantage of this cross-layer approach is obvious: the routing path of the requesting message is the shortest path between source and destination (e.g., the blue and red lines in Figure 3), because the virtual neighbors in the P2P overlay are de facto physical neighbors in the MANET overlay due to the merge of two layers. Thanks to the nature of broadcast, the algorithm complexity of this approach is O(n), making it

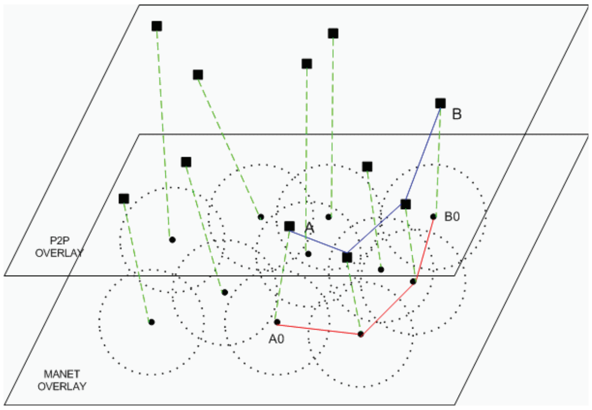

Figure 4. Cross-layer DHT

suitable for deployment in relatively large-scale networks, but still not feasible for Internet-scale networks.

CROSS-LAYER DHT

It is also possible to design a cross-layer DHT in Figure 4 with the similar inspiration, and the algorithm complexity would be optimized to O(log n) with the merit of DHT, which is advocated to be efficient even in Internet-scale networks. The difficulty in that approach is implementation: there is no off-the-shelf DHT-like MANET protocol as far as we know, though recently, some research projects like Ekta (Pucha, Das, & Hu, 2004) towards a DHT substrate in MANET are proposed.

CONCLUSION

In this article, we studied the peer-to-peer systems over mobile ad hoc networks with a comparison of different settings for the peer-to-peer overlay and underlying mobile ad hoc network. We show that the cross-layer approach performs better than separating the overlay from the access networks in Table 1. Our results would potentially provide useful guidelines for mobile operators, value-added service providers, and application developers to design and dimension mobile peer-to-peer systems.

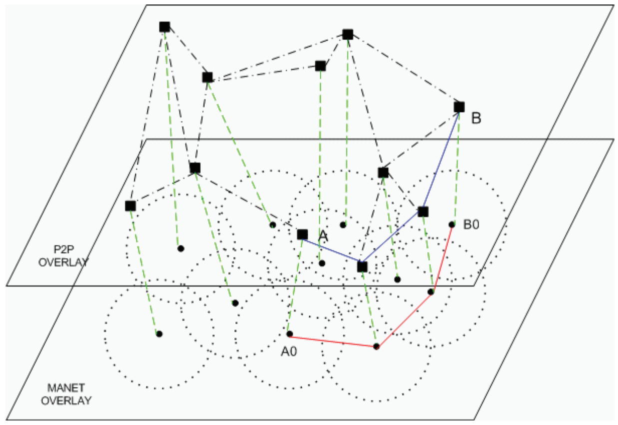

Figure 3. Cross-layer broadcast

REFERENCES

AODV. (n.d.). *AODV IETF draft v1.3.* Retrieved from http://www.ietf.org/internet-drafts/draft-ietf-manet-aodv-13.txt

Chawathe, Y., Ratnasamy, S., Breslau, L., & Shenker, S. (2003). Making Gnutella-like P2P systems scalable. *Proceedings of ACM SIGCOMM.*

Clip2. (2001). *The Gnutella protocol specification v0.4* (document revision 1.2). Retrieved from http://www9.limewire.com/developer /gnutella protocol 0.4.pdf

DSR. (n.d.). *DSR IETF draft v1.0.* Retrieved from http://www.ietf.org/internet-drafts/draft-ietf-manet-dsr-10.txt

Ding, G., & Bhargava, B. (2004). Peer-to-peer file-sharing over mobile ad hoc networks. *Proceedings of the 2nd IEEE Conference on Pervasive Computing and Communications Workshops.*

Goel, S. K., Singh, M., & Xu, D. (2002). Efficient peer-to-peer data dissemination in mobile ad-hoc networks. *Proceedings of the International Conference on Parallel Processing.*

Table 1. How efficient does a user try to find a specific piece of data?

	Efficiency	Scalability	Implementation
Broadcast over Broadcast	$O(n^2)$	n/a	Easy
DHT over Broadcast	O(n log n)	Bad	Medium
Cross-Layer Broadcast	O(n)	Medium	Difficult
Cross-Layer DHT	O(log n)	Good	n/a

Hsieh, H.Y., & Sivakumar, R. (2004). On using peer-to-peer communication in cellular wireless data networks. *IEEE Transactions on Mobile Computing, 3*(1).

Johnson, D. B., & Maltz, D. A. (1996). *Dynamic source routing in ad-hoc wireless networks. Mobile computing.* Kluwer.

Klemm, A., Lindemann, C., & Waldhorst, O. (2003). A special-purpose peer-to-peer file sharing system for mobile ad hoc networks. *Proceedings of the IEEE Vehicular Technology Conference.*

Kojima, F., Harada, H., & Fujise, M. (2001). A study on effective packet routing scheme for mobile communication network. *Proceedings of the 4th Symposium on Wireless Personal Multimedia Communications.*

Kortuem, G., Schneider, J., Preuitt, D., Thompson, T. G. C., Fickas, S., & Segall, Z. (2001). When peer-to-peer comes face-to-face: Collaborative peer-to-peer computing in mobile ad hoc networks. *Proceedings of the 1st International Conference on Peer-to-Peer Computing.*

Lindemann, C., & Waldhorst, O. (2002). A distributed search service for peer-to-peer file sharing in mobile applications. *Proceedings of the 2nd IEEE Conference on Peer-to-Peer Computing.*

Lv, Q., Ratnasamy, S., & Shenker, S. (2002). Can heterogeneity make Gnutella scalable? *Proceedings of the 1st International Workshop on Peer-to-Peer Systems.*

MANET. (n.d.). *MANET implementation survey.* Retrieved from http://protean.itd.nrl.navy.mil/manet/survey/survey.html

Milojicic, D. S., Kalogeraki, V., Lukose, R., Nagaraja, K., Pruyne, J., Richard, B., Rollins, S., & Xu, Z. (2002). *Peer-to-peer computing.* Technical Report HPL-2002-57, HP Labs.

Papadopouli, M., & Schulzrinne, H. (2001). A performance analysis of 7DS, a peer-to-peer data dissemination and prefetching tool for mobile users" *IEEE Sarnoff Symposium Digest.*

Perkins, C. E., & Royer, E. M. (2000). *The ad hoc on-demand distance vector protocol. Ad hoc networking.* Boston: Addison-Wesley.

Plaxton, C., Rajaraman, R., & Richa, A. (1997). Accessing nearby copies of replicated objects in a distributed environment. *Proceedings of ACM SPAA.*

Pucha, H., Das, S. M., & Hu, Y. C. (2004). Ekta: An efficient DHT substrate for distributed applications in mobile ad hoc networks. *Proceedings of the 6th IEEE Workshop on Mobile Computing Systems and Applications.*

Ratnasamy, S., Francis, P., Handley, M., Karp, R., & Schenker, S. (2001). A scalable content-addressable network. *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications.*

Ratnasamy, S., Shenker, S., & Stoica, I. (2002). Routing algorithms for DHTs: Some open questions. *Proceedings of the 1st International Workshop on Peer-to-Peer Systems.*

Ripeanu, M., Foster, I., & Iamnitch, A. (2002). Mapping the Gnutella network: Properties of large-scale peer-to-peer systems and implications for system design. *IEEE Internet Computing, 6*(1).

Rowstron, A., & Druschel, P. (2001). Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. *Proceedings of the IFIP/ACM International Conference on Distributed Systems Platforms.*

Royer, E. M., & Toh, C. K. (1999). *A review of current routing protocols for ad-hoc mobile wireless networks.* IEEE Personal Communications.

Stoica, I., Morris, R., Karger, D., Kaashoek, F., & Balakrishnan, H. (2001). Chord: A scalable peer-to-peer lookup service for Internet applications. *Proceedings of ACM SIGCOMM.*

Yan, L., & Ni, J. (2004). Building a formal framework for mobile ad hoc computing. *Proceedings of the International Conference on Computational Science.*

Yan, L., & Sere, K. (2003). Stepwise development of peer-to-peer systems. *Proceedings of the 6th International Workshop in Formal Methods.*

Yang, B., & Garcia-Molina, H. (2002). Improving search in peer-to-peer networks. *Proceedings of the International Conference on Distributed Systems.*

Zhao, B. Y., Huang, L., Stribling, J., Rhea, S. C., Joseph, A. D., & Kubiatowicz, J. (2004). Tapestry: A Resilient global-scale overlay for service deployment. *IEEE Journal on Selected Areas in Communications.*

KEY TERMS

AODV: Ad hoc On-demand Distance Vector routing.

DHT: Distributed hash table.

DSR: Dynamic source routing.

MANET: Mobile ad hoc network.

P2P: Peer-to-peer.