Title	A heterogeneous peer-to-peer network testbed
Author(s)	Li, VOK; Cui, L; Liu, Q; Yang, GH; Zhao, Z; Leung, KC
Citation	The 1st International Conference on Ubiquitous and Future Networks (ICUFN 2009), Hong Kong, 7-9 June 2009. In Proceedings of the 1st ICUFN, 2009, p. 46-51
Issued Date	2009
URL	http://hdl.handle.net/10722/62059
Rights	International Conference on Ubiquitous and Future Networks. Copyright © IEEE.

A Heterogeneous Peer-to-Peer Network Testbed

Victor O.K. Li[1], Li Cui[2], Qiang Liu[2], Guang-Hua Yang[1], Ze Zhao[2], Ka-Cheong Leung[1]

[1]Department of Electrical and Electronic Engineering
The University of Hong Kong, Hong Kong, China
[2]Institute of Computing Technology
Chinese Academy of Sciences, Beijing, China

Abstract — **In this paper, we describe a heterogeneous peer-to-peer network testbed, which is developed as part of a joint research project to investigate novel resource discovery and content distribution protocols in a heterogeneous wired/wireless environment. We describe the testbed requirements, the testbed architecture, the multi-functional wireless node, and the software architecture. We also describe some of the proposed protocols to be developed and tested on the testbed.**

Keywords — *heterogeneous networks, peer-to-peer networks, testbed, resource discovery, content distribution.*

I. INTRODUCTION

The next generation network is a heterogeneous network of networks, with different component networks which may be wired and wireless, each with very different characteristics in terms of transmission speeds, errors, and interference tolerance. In addition, due to the broadcasting nature of a wireless channel, the mobile wireless host can send a single copy of data to several of its neighbouring nodes simultaneously. A static wired host, however, can transmit data to its neighbours only one at a time.

Peer-to-Peer (P2P) is a very popular communication paradigm for information sharing. Emerging P2P applications, such as BitTorrent (BT) [1], CoolStreaming [2], Gnutella [3], Kazaa Media Desktop (KaZaA) [4], Kazaar [5], Napster [6], and YouTube [7], have been the dominant applications in the Internet. These P2P applications can be classified into two different categories, namely, streaming and content transfer. A streaming application deals with the transfer of multimedia contents as a stream of packets received by a peer or user, whereas a content transfer application facilitates the reliable delivery of certain multimedia contents or files to a peer. These two types of applications impose very different quality of service (QoS) requirements. In this paper, we will focus on the latter.

To support content transfer applications in heterogeneous P2P networks, we need the following services. First, a resource discovery service is needed to search for and locate resources (including multimedia contents and other services supporting the applications). Second, a content distribution service is required to reliably transfer the requested content (located via the resource discovery service) to a set of peers.

A. Related Work on Resource Discovery

Existing approaches on resource discovery can be categorized into three different groups: centralized, decentralized, and flooding [8]. In the centralized approach, a peer queries the centralized server for the requested resource. It then requests the resource from the peers reported as having the resource.

This has been the original design of Napster [6]. The approach suffers from a single point of failure and performance bottleneck as all queries for resource discovery have to go through the only centralized server before any requested resource is located.

In the decentralized approach, directory information is maintained at several servers or peers. In KaZaA [4], each peer is either a supernode or an ordinary node. Each supernode acts as a mini-Napster hub so as to keep track of the addresses and the resources maintained by its descendants. A peer first sends a query to its supernode, which responds with matches or forwards the query to a subnet of supernodes repeatedly until the requested resource is located. This technique may fail when a supernode leaves the system as each ordinary node is generally associated with a single supernode. The dynamic join or leave of a peer is, however, quite common in the heterogeneous environment.

Distributed indexing service based on distributed hash tables has been proposed to facilitate the location of contents [9]. Peers and contents are mapped to a key space via a hash function. Thus, this method requires that peers and content indices are organized in a rigid structure according to their keys. This may not be practical in heterogeneous P2P networks since many peers may be mobile and join/leave the network unpredictably.

In the flooding approach, a peer broadcasts its query to all of its neighbours. If the queried peer has the requested resource, it sends back a response to the originating peer through the reverse path. Otherwise, it forwards the query to its neighbours and the process repeats. This technique may generate excessive query traffic in the network and thus is not scalable with the number of peers. In [3], limited scope query is applied to enhance scalability by limiting the number of neighbours repeating a broadcast and only to those within some given hops from the originating peer. However, these limitations may render the requested resource unlocated even if it is present in the system.

As far as we know, the existing approaches do not take various characteristics of the heterogeneous P2P networks into account for resource discovery. First, they do not provide any mechanisms to handle dynamic join/leave of peers which is common in heterogeneous P2P networks. Second, the placement of directory information and the construction of the directory overlay do not take into account of the heterogeneity of the P2P networks. For example, it may not be beneficial to place the directory information at a mobile peer. Third, there is no mechanism to perform traffic analysis and caching of the directory information of some popular resources so as to

optimize the network performance and the user response time for subsequent queries.

B. Related Work on Content Distribution

There have been a number of large-scale deployments of peer-to-peer content sharing systems in the Internet. BT [1] is perhaps the best known. A shared content is partitioned into multiple small chunks. Peers exchange information on which chunks they currently possess, and request missing chunks from others. A peer can maximize its download speed by requesting different chunks from different peers at the same time. A peer selection mechanism is employed to find to whom such a request should be made. In addition, a chunk selection technique is needed for a peer to decide which chunk(s) to request. Indeed, peer selection and chunk selection are two crucial issues that affect the global effectiveness of the content distribution process [10].

Existing approaches on peer selection [11, 12, 13, 14] rely on the estimates of the round-trip times between peer pairs and peer response times. They use some bandwidth probing mechanisms to optimize the selection of "best" peers. In [15], the "best" peer is selected by incorporating the path length and transmission power constraints in wireless networks. Since the estimates are generally updated periodically, performance suffers. Furthermore, requests are only forwarded to the "best" peer based on these estimates, which do not favor spreading out the requests across a set of peers, thereby causing further performance degradation due to load imbalance.

In [16], we proposed a generalized application-layer anycasting protocol, known as paracasting, to advocate concurrent access of a subset of replicated servers to cooperatively satisfy a client's request. Each participating server satisfies the request in part by transmitting a subset of the requested content to the client. An analytical model was developed to study how to effectively download a shared content from a set of replicated servers. However, the study assumes that all peers are homogeneous. In addition, it focuses on the selection of a subset of peers for content downloading, but does not consider the redistribution of any downloaded chunks to other peers.

In [17], three different replication strategies based on size, bandwidth, or display time of the multimedia clips in wireless peer-to-peer networks have been investigated. The goal is to determine the number of replicas needed for each clip so as to enable all home-to-home online (H2O) devices to support continuous display of a clip. However, only the string network topology has been considered in the analysis. Besides, the effect of channel interference among these H2O devices was ignored. Hence, these replication strategies are not directly applicable to heterogeneous P2P networks with conflict-prone wireless links.

The problem of broadcasting or multicasting a shared content as a single chunk in heterogeneous networks has been investigated in [18, 19]. It has been shown that the problem of finding an optimal broadcast schedule that minimizes the maximum completion time is NP-complete. In particular, the problem in [18] has been formulated in a network where peers have different processing times and the transmission times between different peer pairs also vary. It is shown that the fastest node first (FNF) algorithm is worse than the optimal by an unbounded factor. Heuristic algorithms, such as fastest edge first (FEF) and earliest completing edge first (ECEF), have also been developed. [19] studies the problem in a network in which the transmission times of a peer to all its neighboring peers are assumed to be the same. It proves that FNF gives a 1.5-approximation schedule. It is worth noting that [18, 19] only analyze the problem of distributing just one chunk, whereas a shared content is generally divided into multiple chunks in peer-to-peer content sharing systems. In addition, [18, 19] also assume that the transmission times are known a priori and transmission errors are ignored. These assumptions are rather unrealistic in heterogeneous P2P networks with error-prone wireless links.

BT employs the rarest element first (REF) algorithm, in which those chunks lacking in the most peers are downloaded first. REF is good at increasing the availability of different chunks in the network and is efficient at distributing all chunks from the original source to different peers across the network. However, the simulation results in [20] show that the performance of REF is far from optimal since REF does not take into account which peer(s) should have higher priorities as recipients based on their needs.

Instead of finding the "best" peer to fulfill a peer's request, concurrent access to a set of peers for satisfying requests, also known as parallel downloading, has been proposed [21, 22]. The general idea is to satisfy a request by the involvement of all available peers. A peer may download a certain chunk from each of the available peers, where the decisions of how much and which chunk to be downloaded from which peer are determined before the download begins [21]. Alternatively, a shared content may be partitioned into a large number of smaller equal-sized chunks and chunk-based requests are forwarded to all peers until all chunks are received [22]. However, as discussed in [22], issues such as the conditions under which it is beneficial to apply the proposed dynamic parallel access and the optimal/preferred chunk size are still unresolved. Furthermore, the methods described in [21, 22] do not consider any redistribution of any downloaded chunks to other peers.

In [20], we have investigated the collaborative file distribution problem in a homogeneous peer-to-peer network, where all peers have asymmetric upload and download bandwidths and the transmission times between any peers are the same. Several heuristic algorithms were developed to solve the problem. The performance of these algorithms was also studied. However, [20] focuses on wired peer-to-peer networks and assumes the same transmission time of a chunk on all logical links in the network. The scheduling process for data distribution is carried out in discrete cycles synchronously.

As far as we know, none of the existing work accounts for the broadcasting nature and the error characteristics of the wireless media. In addition, most existing work, such as the "rarest element first" algorithm in BT, are "greedy" algorithms

which attempt to optimize the objective in one step, but does not account for the effect of this short-term gain on subsequent steps.

C. Related Work on Testbed

There are many types of network testbeds. IBM has developed a wireless sensor network testbed with Bluetooth, WLAN and Zigbee modules to evaluate wireless mesh networking for range extension and reliability enhancement [23]. To access the diverse kinds of networks, multiple radio modules are equipped for each node, while for a single kind of network, only one radio module is deployed. [24-25] use wired cables to connect individual devices or equipments to construct a multi-radio platform. In [24], Intel x86-based notebook computers and FALINUX EX-X5 + EXPCMCIA boards were used to implement mesh routers. [25] combines two Soekris net4801 stations with a cross-over cable to form a single logical mesh point to implement dual-radio mesh relay points. Software defined radio (SDR) device is a flexible solution for wireless node. However, some SDR devices have to be driven by computers, thus confining their flexibility in deployment. Independently operated SDR devices are available but tend to be expensive for a testbed of reasonable size [26]. In this work, the proposed portable and flexible testbed with 30 multi-radio wireless nodes can form a highly integrated and cost-effective heterogeneous P2P testbed for various operations.

In this paper, we study peer-to-peer (P2P) applications in a heterogeneous wired/wireless environment, and describe the development of a testbed to evaluate two of the key P2P protocols, namely, the resource discovery and the content distribution protocols. In Section II, we shall describe our testbed. Section III describes novel protocols that will be developed and tested on the testbed, and we conclude in Section IV.

II. A HETEROGENEOUS PEER-TO-PEER NETWORK TESTBED

To study heterogeneous peer-to-peer networks, we design a testbed to evaluate and verify proposed protocols.

A.   Testbed Requirements

To facilitate the deployment and evaluation of proposed protocols, the testbed has to meet the following design criteria:

1. Heterogeneity

The testbed should be composed of both wired and wireless communicating components. The data rate in the wired portion is more than 100Mbps, but it may be much lower in the wireless portion.

2. Support resource discovery service

Resource discovery is a key function of P2P networks. The testbed must provide functions to (i) keep track of the join/leave information of peers dynamically, (ii) maintain the

The architectures of the proposed wireless nodes with different data rates are shown in Fig. 2 and Fig.3, respectively. Fig.2 shows the structure of a node with high data rate. The node uses two G2M5477 modules to allow simultaneous transmit-transmit, transmit-receive, or receive-receive operations with

directory information in the wired and wireless nodes, (iii) allow real-time traffic analysis, and (iv) facilitate caching of directory information at the nodes.

3. Support content distribution service under dynamic channel conditions

Content distribution service is another key function for P2P applications. Each node in the testbed is required to (i) be cognizant of the error characteristics of the wireless media to enable the application of opportunistic scheduling protocols, and (ii) exploit the broadcast nature of the wireless media to improve performance.

4. Support customized multi-interface protocol

The testbed needs to support multi-interface, multi-channel structure to enhance the flexibility of nodes to access different applications.

5. System service

System service software with a user-friendly interface is used to monitor and control the functions of the testbed, and to store the data and event log.

B. Testbed Architecture

The testbed consists of two major parts which are the heterogeneous peer-to-peer network and the monitoring interface, respectively. The system architecture of the testbed is shown in Fig. 1. The network part is constructed of multi-functional portable wireless nodes, wireless notebook computers, desktop computers, servers, wired and wireless gateways, switches, etc. The monitoring part connects to the network through the Ethernet. Users can operate the testbed in browser/server mode to set the parameters, control the network operation, and analyze the data. The database server stores the data and event log.

C. Multi-functional Portable Wireless Node

The multi-functional portable wireless node is an important component of the testbed. To match the high data rate of the wired network, the wireless node should support as high a data rate as possible. Since we plan to deploy a testbed of at least 30 nodes, we plan to buy off-the-shelf wireless modules. We have also considered a software-radio based testbed, but find the costs prohibitively high for a testbed of reasonable size. After considering such issues as data rates, availability, reliability, and costs, we decided on an embedded WiFi solution. We find that most WiFi modules available in the market integrate the radio frequency (RF) and the WiFi protocol with an antenna, and do not allow one to change the medium access control (MAC) protocol. Since we aim to develop new opportunistic scheduling protocols, we do need to access and modify the MAC protocol. After some investigation, we have adopted the G2C547 WiFi module from G2 Microsystems for our testbed. The G2C547 chip supports a 4Mbps throughput at the transport layer with TCP/IP.

open source firmware. Each G2M5477 module may be set to utilize Channels 1-11 and 14 at 2.4GHz. To minimize interference we use the non-overlapping Channels 1, 6, and 11, simultaneously. The CPU is XScale PXA270 with RAM, ROM, SD card port and Ethernet port. The RAM and ROM

are used for embedded Linux, and the SD card port may be used to provide large file storage. The Ethernet port is connected to the wired network and used to transfer test and control information. There are also interfaces designed to high data rate node to evaluate various application scenarios depending on different requirements. Fig. 4 is the photograph of a prototype node with medium speed. The high speed node is still under development.
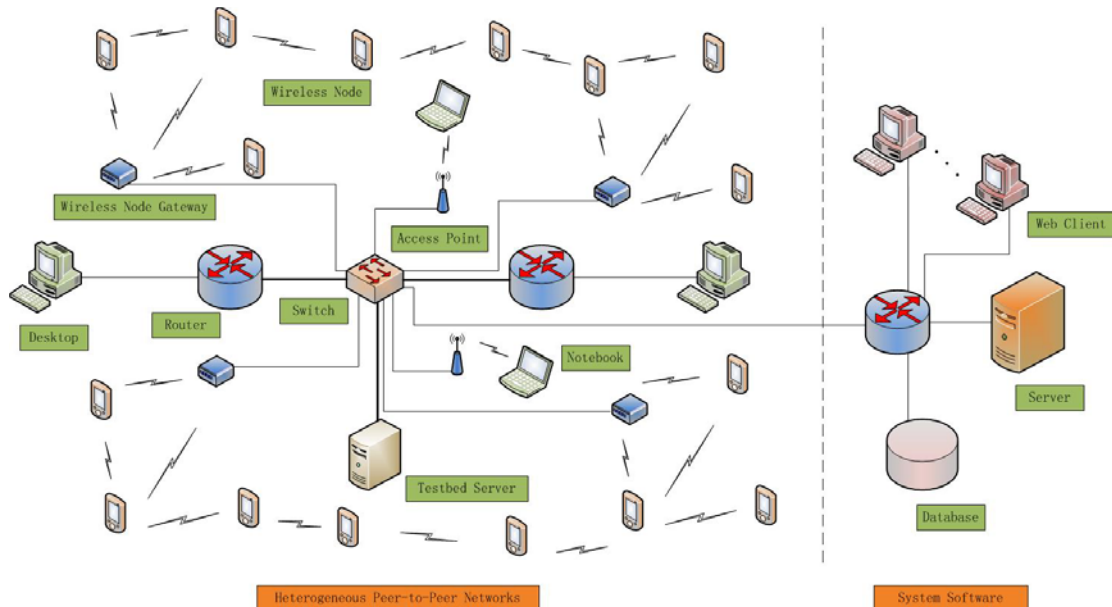

Fig. 1. Testbed architecture

connect with different types of sensors for various applications. Energy consumption gauge chips are employed to monitor the energy consumption level of the G2M5477 modules and of the whole node.

We have also designed another type of node with medium to low speed (see Fig. 3). The node is composed of MCU, data RF, monitoring RF, Flash ROM and external sensors. With CC2420 as the data RF, the node can provide 250Kbps at the physical layer. The WiFi module is used to connect to the testbed. This type of node can be used in conjunction with the
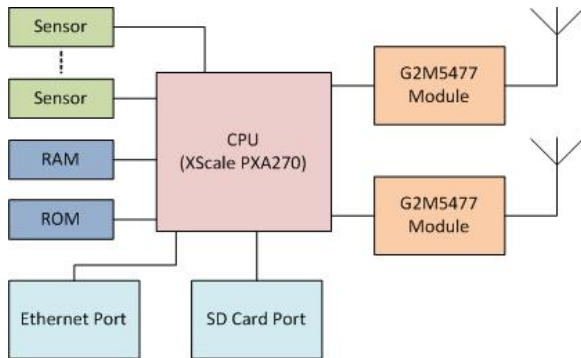

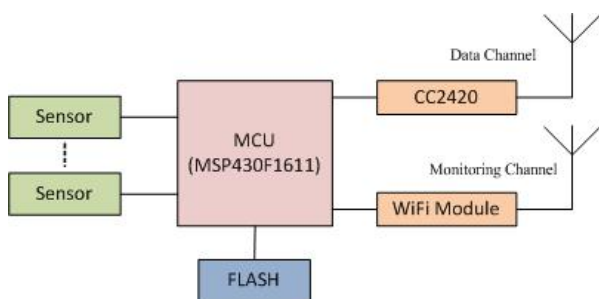Fig. 2. Multi-functional portable high speed wireless node

Fig. 3. Medium speed testbed node architecture
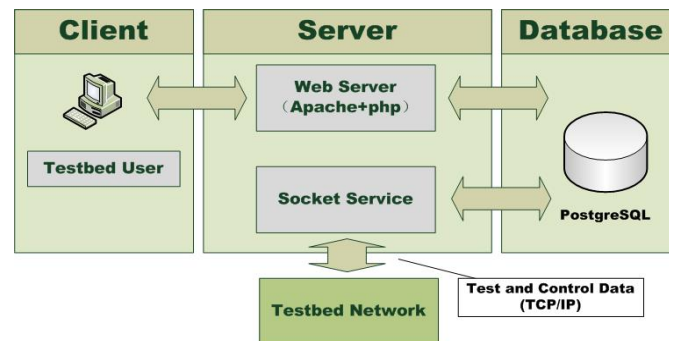

Fig. 4. Medium speed testbed node

Fig. 5. System software architecture

D. System Software

The monitoring software consists of two major components. One is Socket Service which deals with the events of the testbed such as the collection of the test data and the control of the testbed operation. The other is Web Server which is the interface to testbed users. Both components are based on the Linux platform. Web Server is constructed with Apache and PHP and the database server is constructed with PostgreSQL server. All of the platforms follow the GNU standard.

III. Proposed Research on the Testbed

We shall develop a resource discovery service and a content distribution service to support near-optimal resource discovery and collaborative multimedia content distribution in heterogeneous peer-to-peer (P2P) networks.

A. Design of Resource Discovery Service

Our new resource discovery service aims to satisfy the following objectives. First, the resource discovery service is robust to dynamic join/leave of peers which is common in heterogeneous P2P networks. Second, the placement of directory information and the construction of the directory overlay consider the heterogeneity of the P2P networks. Third, there is a mechanism to perform traffic analysis and caching of the directory information of some popular resources so as to optimize the network performance and the user response time for subsequent queries.

To improve the robustness of the resource discovery service, the directory will be replicated and maintained by a number of peers and/or network nodes as directory servers or trackers. When a tracker/directory server leaves the system, the directory information will still be available at other servers. In addition, the level of replication depends on the conditions (such as connectivity status and remaining energy in the case of mobile peers) of the peers/nodes in the networks.

Determining how to intelligently place and connect such trackers in heterogeneous P2P networks is still an open research problem. We propose to formulate an optimization problem to compute the optimal level of replication, the construction of the directory overlay, and the placement of directory in the heterogeneous P2P networks for resource discovery. Based on the computed optimal parameters, we will devise some directory management protocols to automate and optimize the placement of such trackers in heterogeneous P2P networks.

To improve the network and user performance, resource caching techniques will be investigated. The idea is to allow some popular multimedia contents and directory replicas to be cached in some local network nodes, which may or may not be peers. Alternatively, some network nodes may behave as peers so as to induce the P2P applications to distribute the multimedia contents and directory replicas to them. By doing so, the accesses to these contents and directory information may instead be entertained by these network nodes so as to reduce the traffic load and the user response time as well as to balance the network load. Some techniques for the content and traffic analysis will be studied to determine which of these resources will be cached, and where copies of the resources

should be maintained in the heterogeneous P2P networks. Resource caching protocols will be devised.

B. Design of Content Distribution Service

Our novel content distribution service aims to satisfy the following requirements. First, the content distribution service must go beyond the "greedy" approach of optimizing only the near-term gain. Second, the broadcasting nature and the error characteristics of the wireless media in heterogeneous P2P networks must be accounted for in the content distribution algorithms for performance enhancement.

In heterogeneous P2P networks, network nodes and peers are connected via wired or wireless links. Whenever a peer or node communicates to its neighbour via a wireless channel, its transmission can reach several neighbouring peers or nodes due to the broadcasting nature of the wireless media. Our content distribution protocols will exploit the broadcasting nature of the wireless channels.

Besides, a wireless link is generally error-prone. Opportunistic-based techniques should be employed to schedule on-going transmissions to achieve better network performance. We plan to incorporate opportunistic-based scheduling into our content distribution protocols to yield further performance improvement through the synergy of both medium access control (MAC)-layer frame scheduling and content distribution in heterogeneous P2P networks.

We shall formulate an optimization problem for collaborative content distribution. The objective is to minimize the completion time of delivering the whole multimedia content to a given set of peers, with respect to the energy constraints of the peers and the conditions of the wireless links in the network. Based on the computed optimal parameters, we will devise some content distribution protocols to automate and optimize the transmission schedule in heterogeneous P2P networks.

C. Evaluation of Our Proposed Solutions

We shall devise analytical and simulation models to evaluate and compare the performance of the proposed solutions for resource discovery and content distribution. Analytical and simulation results will be produced to validate the correctness of the proposed ideas and demonstrate the effectiveness of the proposed solutions. To demonstrate the feasibility of the implementation of our proposed solutions in real systems, we shall also deploy our proposed algorithms on the testbed. Experimental results will be obtained to validate the proposed ideas and models.

IV. Conclusion

In this paper, we have described a heterogeneous peer-to-peer network testbed, which is developed as part of a joint research project to investigate novel resource discovery and content distribution protocols in a heterogeneous wired/wireless environment. One of the key components of the testbed is a multi-functional multi-channel wireless node. To construct this node, and to allow us to test various opportunistic scheduling protocols, we have to find a WiFi module which allows us to access and modify the MAC protocol. We have also developed the hardware and software architectures of the

testbed. We are now in the process of developing resource discovery and content distribution protocols to be tested on the testbed.

REFERENCES

[1] D. Qiu and R. Srikant, "Modeling and Performance Analysis of BitTorrent-Like Peer-to-Peer Networks," ACM SIGCOMM Computer Communication Review, Vol. 34, No. 4, pp. 367-377, October 2004.

[2] X. Zhang, J. Liu, B. Li, and Y.-S. P. Yum, "CoolStreaming/DONet: A Data-Driven Overlay Network for Peer-to-Peer Live Media Streaming," Proceedings of IEEE INFOCOM 2005, Vol. 3, pp. 2102-2111, 13-17 March 2005.

[3] D. Hughes, G. Coulson, and J. Walkerdine, "Free Riding on Gnutella Revisited: The Bell Tolls?," IEEE Distributed Systems Online, Vol. 6, No. 6, June 2005.

[4] The official KaZaA website, http://www.kazaa.com.

[5] The official Kazaar website, http://www.kazaar.com.

[6] The official Napster website, http://www.napster.com.

[7] The official YouTube website, http://www.youtube.com.

[8] J. F. Kurose and K. W. Ross, "Computer Networking: A Top-Down Approach," Fourth Edition, Addison-Wesley, 2008.

[9] P. Trunfio, D. Talia, H. Papadakis, P. Fragopoulou, M. Mordacchini, M. Pennanen,, K. Popov, V. Vlassov, and S. Haridi, "Peer-to-Peer Resource Discovery in Grids: Models and Systems," Future Generation Computer Systems, Vol. 23, No. 7, pp. 864-878, August 2007.

[10] P. Felber and E. W. Biersack, "Cooperative Content Distribution: Scalability Through Self-Organization," Lecture Notes in Computer Science, Vol. 3460, pp. 343-357, May 2005.

[11] T. S. E. Ng, Y. H. Chu, S. G. Rao, K. Sripanidkulchai, and H. Zhang, "Measurement-Based Optimization Techniques for Bandwidth-Demanding Peer-to-Peer Systems," Proceedings of IEEE INFOCOM 2003, Vol. 3, pp. 2199-2209, 30 March – 3 April 2003.

[12] S. Ranjan, R. Karrer, and E. Knightly, "Wide Area Redirection of Dynamic Content by Internet Data Centers," Proceedings of IEEE INFOCOM 2004, Vol. 2, pp. 816-826, 7-11 March 2004.

[13] M. Sayal, Y. Breitbart, P. Scheuermann, and R. Vingralek, "Selection Algorithms for Replicated Web Servers," ACM SIGMETRICS Performance Evaluation Review, Vol. 26, No. 3, pp. 44-50, December 1998.

[14] E. W. Zegura, M. H. Ammar, Z. Fei, and S. Bhattacharjee, "Application-Layer Anycasting: A Server Selection Architecture and Use in a Replicated Web Service," IEEE/ACM Transactions on Networking, Vol. 8, No. 4, pp. 455-466, August 2000.

[15] A. K. H. Leung and Y.-K. Kwok, "An Efficient and Practical Greedy Algorithm for Server-Peer Selection in Wireless Peer-to-Peer File Sharing Networks," Lecture Notes in Computer Science, Vol. 3794, pp. 1016-1025, December 2005.

[16] K.-C. Leung and V. O. K. Li, "A Paracasting Model for Concurrent Access to Replicated Internet Content," IEEE Transactions on Multimedia, Vol. 8, No. 1, pp. 90-100, February 2006.

[17] S. Ghandeharizadeh and T. Helmi, "An Evaluation of Alternative Continuous Media Replication Techniques in Wireless Peer-to-Peer Networks," Proceedings of ACM MobiDE '03, pp. 77-84, 19 September 2003.

[18] P. B. Bhat, C. S. Raghavendra, and Viktor K. Prasanna, "Efficient Collective Communication in Distributed Heterogeneous Systems," Journal of Parallel and Distributed Computing, Vol. 63, No. 3, pp. 251-263, March 2003.

[19] S. Khuller and Y.-A. Kim, "On Broadcasting in Heterogeneous Networks," Proceedings of ACM-SIAM SODA '04, pp. 1011-1020, 11-14 January 2004.

[20] J. S. K. Chan, V. O.K. Li, and K.-S. Lui, "Performance Comparison of Scheduling Algorithms for Peer-to-Peer Collaborative File Distribution," IEEE Journal on Selected Areas in Communications, Vol. 25, No. 1, pp. 146-154, January 2007.

[21] C. Gkantsidis, M. Ammar, and E. Zegura, "On the Effect of Large-Scale Deployment of Parallel Downloading," Proceedings of IEEE WIAPP '03, pp. 79-89, 23-24 June 2003.

[22] P. Rodriguez and E. W. Biersack, "Dynamic Parallel Access to Replicated Content in the Internet," IEEE/ACM Transactions on Networking, Vol. 10, No. 4, pp. 455-465, August 2002.

[23] S. Furrer, W. Schott, H.L. Truong, B. Weiss, "The IBM Wireless Sensor Networking Testbed," Proceedings of IEEE TRIDENTCOM 2006, pp. 42-46, 1-3 March 2006.

[24] H. Song, B.C. Kim, J.Y. Lee, H.S. Lee, "IEEE 802.11-based Wireless Mesh Network Testbed," Proceedings of 16th IST Mobile and Wireless Communications Summit, pp.1-5, 1-5 July 2007

[25] V. Badarla, D. Malone, D.J. Leith, "Implementing TCP Flow-Level Fairness Using 802.11e in a Multi-Radio Mesh Testbed," IEEE Communications Letters, Vol. 12, No. 4, pp:262-264, April 2008.

[26] The Ettus Research LLC website, http://www.ettus.com/downloads/ettus_ds_usrp2_v2.pdf