

ON PEER-TO-PEER MULTIMEDIA CONTENT ACCESS AND DISTRIBUTION

Zhu Liu¹, Heather Yu², Deepa Kundur³, and Madjid Merabti⁴

¹AT&T Labs – Research, Middletown, NJ, USA

²Panasonic Princeton Research Laboratory, Princeton, NJ, USA

³Electrical & Computer Engineering Department, Texas A&M University, College Station, TX, USA

⁴School of Computing & Mathematical Sciences, Liverpool John Moores University, UK

ABSTRACT

This paper provides a brief overview of recent progress of peer-to-peer (P2P) technologies for multimedia applications. We provide an overview of the technical challenges, creative solutions and results related to P2P file sharing and content search, and P2P media streaming.

1. INTRODUCTION

Since Napster made its debut in 1999, peer-to-peer (P2P) content access and distribution has been gaining increased popularity. Many other P2P systems, which span relatively simple file sharing networks, more complex multimedia content search entities and sophisticated media streaming systems, have been subsequently deployed.

P2P computer networks are popular due to their inherent scalability and flexibility, which facilitates a broad spectrum of innovative multimedia applications. Such networks rely on the power of participant nodes of the network (called peers) for communications and computation. This is in contrast to the traditional client-server models for multimedia communications that employ dedicated networking infrastructure. Applications of P2P multimedia (P2P MM) include non-centralized file sharing and emerging content distribution and multimedia sensor networks. The paradigm shift to less hierarchical and more localized multimedia networking requires that many basic broadband communication challenges be revisited.

In this paper we look at current P2P multimedia systems and technologies, providing a brief tour of two focus areas in the subsequent sections. The next section targets file sharing and multimedia content search while Section 3 concentrates on media streaming.

2. P2P FILE SHARING AND CONTENT SEARCH

Due to the intrinsic scalability, robustness and decentralized structure of P2P systems, they are an ideal communication and computation paradigm for emerging collaborative multimedia applications. For example, P2P MM systems have been broadly adopted to share storage responsibilities,

bandwidth, and CPU cycles for a broad range of applications. Two recent expositions of the field provide interesting details [1, 2]. In this section, we focus on the file sharing aspects of P2P MM systems and review the recent progress of content searching within a P2P network.

2.1. Architecture of P2P File Sharing Systems

The first generation P2P network, represented by the well-known Napster music file sharing system, maintains a set of central servers to facilitate interaction (such as file exchange) between distributed peers. The central servers provide content indexing and search services, while the content itself is exchanged amongst the peers directly. Although the central servers are able to search content efficiently and accurately, the system is not scalable, and it has a single point of failure at the server. For this reason, most later P2P systems assume a partially or a purely decentralized architecture.

In partially decentralized systems (also called hybrid P2P systems in the literature), there exists a hierarchy of *leaf nodes* and *supernodes* as a tradeoff to meet the requirement of scalability and search efficiency. The role of each node is assigned dynamically depending on its capability and recent activity. Kazaa is one of the most successful P2P systems in the hybrid category. A group of leaf nodes are associated with a supernode, which maintains the leaf nodes' content indexes and serves as a proxy for their search requests. The supernodes are inter-connected, forming a more compact higher layer network for the sake of search efficiency. For Kazaa there is no logical correlation between nodes and content within the network, so locating content must employ a brute force approach, such as a flooding-based methodology and/or a random walk technique to query for content. Specifically, when a supernode receives a query from one node, it first checks its local content index. If the query cannot be satisfied, the supernode forwards the query to other linked supernodes.

The networks that most reflect the distributed and collaborative spirit of P2P networking are purely decentralized systems, where every node takes the same role in the network. All nodes serve as both servers and clients

simultaneously. Purely decentralized systems can be further categorized into *structured* and *unstructured* P2P networks based on whether a certain network overlay structure is preserved when nodes and content are added to the system. In unstructured P2P networks, nodes participate in a random manner, and the content location is independent of the nodes. While flooding or similar methods are the only means to find the content of interest, searching is more flexible as it is based on titles and keywords.

Structured P2P systems form nodes into a content-wise meaningful graph, such that query results are guaranteed in a limited number of hops. Chord is one such system. In Chord, nodes and content items are assigned *keys* (generated via consistent hashing of content title or metadata) that are uniformly distributed in an identifier space. Nodes are arranged in an identifier circle based on their key values; a content item with key k is appointed to the first node whose key is equal to or follows k . Each node maintains a list of pointers to its successors and predecessors, as well as a finger table to speed up the routing procedure. With this infrastructure, content can be efficiently located based on its key. The disadvantage is that content or keyword based search is not immediately supported because content search relies on the unique hash key that is created based on content title or its metadata. Besides Chord, other popular structured P2P systems include Content Addressable Network (CAN), Tapestry, Pastry, Kademia, Viceroy and others [1].

2.2. Content-Based Search in P2P Network

Most of the systems introduced in the previous section only provide limited content search capabilities, for example, search based on document title, keywords or descriptive text. More flexible and feature-rich search mechanisms are in high demand in emerging P2P MM networking applications. Content-based search strategies represent one class of promising solutions.

Shen and his colleagues [3] presented a P2P system that supports semantic-based content search. They proposed a general and extensible framework based on the concept of a hierarchical summary structure. There are three levels of summaries in the devised framework. A single document is summarized at the unit level, summaries of all content owned by a peer are combined at the peer level, and an overall content summary of a group of peers is maintained at the super level. Accordingly, indexes at these levels are created to help route the queries efficiently and effectively.

Lu and Callan explored content-based resource selection and document retrieval algorithms in hybrid P2P networks [4]. In their approach, the leaf node determines the retrieval results for certain queries using a probabilistic information retrieval algorithm, and the directory node (supernode) builds a unified content model for all of its leaf nodes and a set of neighboring directory nodes. The content model is used for routing query messages.

Recently, researchers have also applied content-based search technologies to multimedia content, including music, image, and video. Gao *et al.* investigated content-based music information retrieval in P2P networks [5]. In addition to the manually annotated metadata, such as artist and album names, the system can also search music based on automatically extracted acoustic features, including tempo, beat strength, and degree of harmonic change. Instead of a broadcast-style content discovery method, the authors proposed a rendezvous point based registration and query scheme to ensure search efficiency.

In [6], King *et al.* reported the DISCOVER (DISTRIBUTED Content-based Visual Information Retrieval) system, which allows the users to retrieve images using content-based features including color, texture, and shape. Within the system, peers sharing similar images (for example, sunset or tree) are grouped together based on their image feature similarity. The authors also proposed a content-based query routing strategy, called Firework Query Model. It is similar to Gnutella's flooding method, but takes advantage of the peer clustering structure to reduce network traffic and increase search performance.

3. P2P MEDIA STREAMING

Media streaming systems are distinct from file-sharing systems and are harder to deploy due to the real-time playback requirement at the receiver end. The fundamental application value of using P2P MM networks for media streaming lies in its potential to overcome the limitations of the client-server model where the streaming server is a potential bottleneck and represents a single point of failure. Cost reduction and improving scalability and reliability are just some of the many goals of P2P media streaming.

3.1 Popular P2P Media Streaming Models

A number of P2P media streaming schemes and systems can be found in the literature today. Some focus on small-scale P2P networks and some can support large-scale networking. In general, these P2P systems can be classified using several taxonomies. From the architecture point of view, analogous to file sharing, there are purely decentralized and partially decentralized models. With a purely decentralized architecture, all peer and content management operations are distributed. If *distribution trees* are employed for media streaming, construction and operation of the distribution trees are also distributed. In a partially decentralized system, servers host content and serve it to clients. Peers are used to reduce the servers' load in various ways. Based on the streaming media distribution protocol, the systems can be grouped into tree-based and non-tree-based classes. Constructing and maintaining an efficient distribution tree among the peers is a key challenge for tree-based systems. In some systems, multiple trees are built to improve system

fairness, scalability and/or reliability. Gossip-based protocols are popular in treeless P2P systems. They achieve operation decentralization via gossip, i.e. delivering data and/or control messages to peers randomly. Alternatively, we can classify existing P2P media streaming systems into receiver-driven versus sender-driven systems based on their control models. In a receiver-driven system, the receiver coordinates the peers, handles stream delivery from multiple peers, and performs load balancing. In contrast, the sender deals with peer coordination and stream distribution in a sender-driven system.

3.2 P2P media streaming solutions

CoopNet [7] addresses the server overload problem via client (i.e., peer) cooperation in content distribution. A server functions as the root and the manager of a distribution tree. In the live streaming mode, a client as a member of the distribution tree receives a live stream from its parent and further streams it out to its children. In the on-demand mode, a client (peer) caches recently viewed on-demand content streams. When the server faces overloading, it redirects new peers to other peers with the content in their caches. Each peer may have only a portion of the content stream. Hence the supplicant may need to contact multiple peers to get the complete content stream. To cope with the network dynamics, CoopNet introduces multiple description coding. A different description of the media signal is transmitted through a different distribution tree. Hence network dynamics and peer failure will only reduce the number of descriptions delivered to the receiver. Similarly in SplitStream [8], the streaming content is split into n stripes each sent using a separate multicast tree. The challenge is in creating a forest of trees such that an interior node in one tree is a leaf node in all the remaining trees. This is met by each tree having a groupID that differs from all other groupIDs (of the other trees) in the most significant digit. With the existence of the server, the task of locating content is simplified compared to completely decentralized P2P architectures. To control the end-to-end delay from the source to the receiver, Tran [9] proposed ZIGZAG. A set of rules are enforced to guarantee the height and node degree of the multicast tree that is built upon a hierarchy of bounded size clusters of peers. Further, by separating the administrative tree from the content distribution tree, ZIGZAG can recover from failure quickly and regionally and hence improves system robustness.

To improve the system scalability, many fully decentralized P2P media streaming systems have been proposed. PeerStreaming [10], for example, is a representative receiver-driven decentralized P2P system. It recognizes that peers are doing a favor for the sender and the receiver. They may go on or offline during a session and may only want to hold a portion of the media stream. In PeerStreaming, the receiver drives the P2P streaming process, connects to peers

that come online, redirects requests dropped by offline peers, and balances load among peers.

PROMISE [11], another treeless fully distributed P2P media streaming system that builds upon CollectCast, is deployed on top of P2P substrates such as Pastry. It does not use a multicast distribution tree for peer content access and distribution. Instead, a peer looks up peers with the media stream via the underlying P2P substrate first. A set of active peers called an active sender set is selected among the candidate peers using a topology-aware selection algorithm. The requester then receives the media stream via the parallel connections to all peers in the active sender set. The sending rate and data set of each peer sender are assigned by the receiver. To assure system robustness, PROMISE exploits standby peers in addition to the active peers. Peers monitor the bit rate of the incoming stream. Once failure is detected, the topology-aware selection algorithm will be used to find a new active peer from the standby peer set.

CoolStreaming [12], a data-centric design of a streaming overlay, also does not employ any tree structure but uses the data availability to lead the flow direction to deal with the high dynamics of nodes. Every node periodically exchanges data availability information with a set of peers, and retrieves unavailable data from one or more peers, or supplies available data to peers. There are no prescribed roles like parent/child. Instead, membership is managed with a gossiping protocol where a message is sent to a set of randomly selected peers from a node and then to other randomly selected peers from this set and so on, until the message is spread to all. Thus operation decentralization is achieved. Stream delivery is also decentralized by pulling data from multiple peers using selection and low-overhead scheduling algorithms. The gossiping protocol is again used in node failure notification. It achieves robustness and resilience via periodical partnership and availability updates. Noticeably, in many of the existing systems, the quality of a streaming session relies on the resources contributed by individual peers. Habib and Chuang [13] studied the impact of non-cooperative peers versus cooperative peers on streaming media quality and found a random peer selection scheme may result in highly variable streaming quality, especially in large-scale P2P streaming systems. To solve this problem, the authors propose a different approach to peer selection called a rank-based incentive mechanism that achieves cooperation through service differentiation. Contributors are rewarded with flexibility and choice in peer selection, resulting in high quality streaming sessions whilst free-riders are given limited options in peer selection.

The aforementioned systems are just some representatives among many existing ones in the literature. In summary, P2P provides a nifty way to deal with some of the challenges that traditional client-server based streaming systems face. Ideally it could provide high QoS at low cost. To realize a scalable and performance-guaranteed streaming session with efficient peer organization and/or selection algorithms, proficient content transmission strategies,

adequate fault tolerance capability, proper means to adapt to the network dynamics, and a cost effective and easy to implement P2P system, however, there are many challenges to face.

4. CONCLUSIONS

Multimedia content distribution is a prominent application area of P2P networks. As P2P MM technologies are still evolving, there are numerous open research problems. The basic challenges in building a successful and efficient peer-to-peer multimedia system include:

- Cost, peer organization and system efficiency in content access and distribution.
- Reliability/robustness, system's adaptation capability to the changing network and other environmental conditions.
- Scalability, system's stability and ability to deal with peer dynamics (i.e. peer joining and leaving), to maintain system performance attributes independent of the number of nodes or documents in the network, and to handle the heterogeneity of platform, device, user, and content.
- Security, system's ability to guarantee secure media access and distribution over the P2P network. P2P MM distribution architectures present added technical challenges due to their open and autonomous nature. For example, how to make sure the system is resilient to malicious attacks, renewable and able to recover from attacks; ensuring data integrity and authenticity, peer privacy and confidentiality; minimizing added complexity and end-to-end delays caused by security needs in a real-time multimedia application; and how to build a cost effective security system that does not degrade system QoS. Similarly, Digital Rights Management (DRM) remains an important challenge. Whilst P2P networks are designed to facilitate wide and efficient access to data, there remains a need to retain an element of control over content distribution. Both security and rights management relate closely to work on trust in P2P networks [14].
- P2P MM paradigms are also emerging in multimedia sensor network design [15,16]. Here, the extreme limitations on cost and portability of the peer device restrict the computation, communication, and memory resources making content sharing, streaming and security even more challenging. Many of the P2P routing paradigms are envisioned for MM sensor networks with the ultimate goal of reducing energy consumption and maximizing scalability.

5. REFERENCES

- [1] E.K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, "A Survey and Comparison of Peer-to-Peer Overlay Network Schemes," *IEEE Communications Surveys & Tutorials*, pp. 72-93, Second Quarter 2005.
- [2] S.A. Theotokis and D. Spinellis, "A Survey of Peer-to-Peer Content Distribution Technologies," *ACM Computing Surveys*, Vol. 36, No. 4, pp. 335-371, Dec. 2004.
- [3] H.T. Shen, Y. Shu, and B. Yu, "Efficient Semantic-Based Content Search in P2P Network," *IEEE Trans. on Knowledge and Data Engineering*, Vol. 16, No. 7, pp. 813-826, July 2004.
- [4] J. Lu and J. Callan, "Content-Based Retrieval in Hybrid Peer-to-Peer Networks," *Proceedings of ACM CIKM'03*, New Orleans, LA, pp. 199-206, Nov., 2003.
- [5] J. Gao, G. Tzanetakis, and P. Steenkiste, "Content-based Retrieval of Music In Scalable Peer-to-peer Networks," *Proceedings of Int.Conf. on Multimedia and Expo 2003*, pp. 309-312, Baltimore MD, July 2003.
- [6] I. King, C.H. Ng, and K.C. Sia, "Distributed Content-Based Visual Information Retrieval System on Peer-to-Peer Networks," *ACM Trans. On Information Systems*, Vol. 22, No. 3, pp. 477-501, July, 2004.
- [7] V. N. Padmanabhan, He. J. Wang , P. A. Chou, K. Sripanidkulchai, "Distributing streaming media content using cooperative networking," *Proceedings of the 12th international workshop on Network and operating systems support for digital audio and video*, Miami, Florida, USA, May 12-14, 2002.
- [8] M. Castro, P.Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh. "SplitStream: High-bandwidth multicast in cooperative environments," *Proceedings of SOSP'03*, Bolton Landing, NY, Oct. 2003.
- [9] D. Tran, K.Hua, and T. Do, "ZIGZAG: An Efficient Peer-to-Peer Scheme for Media Streaming," *Proceedings of IEEE INFOCOM'03*, San Francisco, CA, USA, April 2003.
- [10] J. Li, "PeerStreaming: A Practical Receiver-Driven Peer-to-Peer Media Streaming System," MSR-TR-2004-101, September 2004
- [11] M. Hefeeda, A. Habib, B. Botev, D. Xu, and B. Bhargava, "PROMISE: Peer-to-Peer Media Streaming Using CollectCast," *Proceedings of the 11th ACM international conference on Multimedia*, Berkeley, CA, USA, November, 2003.
- [12] X Zhang, J Liu, B Li, TSP Yum, "CoolStreaming/DONet: A data-driven overlay network for efficient live media streaming," *Proceedings of IEEE INFOCOM*, March, 2005.
- [13] A. Habib and J. Chuang, "Incentive mechanism for peer-to-peer media streaming," *Proceedings of the 12th IEEE International Workshop on Quality of Service*, Montreal, Canada, 2004.
- [14] S. Marti and H. Garcia-Molina, "Taxonomy of trust: Categorizing P2P reputation systems," *Computer Networks*, Vol. 50, No. 3, pp. 472-484, March 2006.
- [15] U. Okorafor and D. Kundur, "Efficient Routing Protocols for a Free Space Optical Sensor Network," *Proc. IEEE Int. Conference on Mobile Ad Hoc and Sensor Systems*, Washington, D.C., pp. 251-258, November 2005.
- [16] W. Luh and D. Kundur, "Distributed Privacy for Visual Sensor Networks via Markov Shares," *Proc. 2nd IEEE Workshop on Dependability and Security in Sensor Networks and Systems*, Columbia, Maryland, to appear April 2006.