

Multicast vs. P2P for content distribution

Abstract

Many different service architectures, ranging from centralized client-server to fully distributed are available in today's world for Content Distribution on the Internet. The IP Multicast seemed to be a wonderful technology with potential benefits for one-to-many or many-to-many content distribution applications. However, even after more than two decades of its invention, only very few commercial applications of multicast are in use. On the other hand, the recent wide-spread use of peer-to-peer (P2P) applications such as SETI, Napster, eDonkey, BitTorrent and Gnutella indicate that there are many potential benefits to fully distributed peer-to-peer systems. In this position paper, we briefly describe service model and architecture of both Multicast and P2P networks and put forth the reasons for P2P being a better technology than Multicast for Content Distribution.

1. Introduction

The number of Internet users is increasing enormously along with the processing power and bandwidth of computers. These users are scattered all over the world. The amount of content available on the web is also increasing rapidly as well. Examples of content distribution includes, but not limited to, bulk data transfer, streaming continuous media, shared data applications, data feeds, web cache updating, interactive gaming, etc. The content distribution industry exploded in the late 1990s and still continues to grow, especially in the distribution of audio and video content. The content providers need to distribute their respective content efficiently to users. This requires delivery of packets from one or more senders to a group of receivers. The traditional Unicast (one to one) models of content delivery are not really efficient. Consider the example of transferring a software upgrade from the software vendor to list of all users who needs this upgrade. Using Unicast model, it must transfer a

separate copy to each user independently, thus leading to more bandwidth consumption, which is a metered and limited resource. To overcome this and provide a better network and bandwidth utilization, a new model was proposed called as IP multicast which sets up a shared delivery tree spanning all the users. Packets sent down this tree are replicated at branch points so that no more than one copy of each packet traverses each link. Multicast thus greatly reduces the transmission costs involved in reaching large user populations.

Even though IP Multicast was proposed in early 80s, it has seen slow commercial deployment by ISPs and carriers due to several issues which we will be discussing later. Another alternative for content distribution evolved in the form of Peer-to-peer (P2P) technology. P2P is a type of transient internet network that allows a group of computer users with the same networking program to connect with each other and directly access files from one another's hard drives. It is a compelling content distribution paradigm because all the content is transferred directly between ordinary peers without passing through third-party servers. The goal of this position paper is to show P2P as a better technology than IP multicast for content distribution.

We begin in section 2 with a brief overview of Multicast and P2P networks. In section 3, we summarize the counter claims of Multicast and refute the same. In Section 4, we provide the reasons for supporting P2P over Multicast for content distribution. Finally, in section 5, we give our conclusions.

2. Overview

In this section, we give a brief overview of multicast model and P2P networks.

2.1 Multicast Model

IP multicast is a protocol for transmitting IP datagrams from one source to many

destinations in a LAN or WAN. With IP multicast, applications send one copy of the information to a group address. The information reaches all the receivers who are part of that group as shown in fig 1. Multicast technology addresses packets to a group of receivers rather than to a single receiver thus utilizing lesser bandwidth. IP multicast routing protocols provide efficient delivery of datagrams from one source to any number of destinations throughout a large, heterogeneous network such as the Internet. A client wishing to receive this multicast stream would have to inform the router that it wishes to participate. When accepted, the router will replicate the traffic to the client and all others in that 'host group'. Hosts can join a multicast group by sending IGMP (Internet Group Management Protocol) messages. They use this to notify a router on their local network of their desire to receive packets sent to a certain multicast group. Another protocol is clearly required to coordinate the multicast routers (including the attached routers) throughout the Internet, so that multicast packets are routed to their final destination. This functionality is accomplished by the network-layer multicast routing algorithms, such as PIM, DVMRP, and MOSPF. In March 1992, the Mbone carried its first event with 20 sites worldwide received multicast audio streams from a meeting of the Internet Engineering Task Force (IETF) in San Diego.

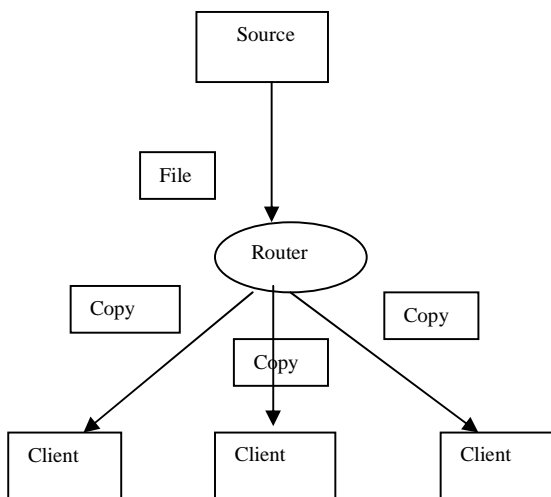


Fig 1: Multicast Network

2.2 Peer-to-peer (P2P) Networks

P2P is a type of network communication that allows file and data sharing between two peers (end users) without going through a server as shown in fig 2. In other words, it's a type of network in which each workstation has equivalent capabilities and responsibilities. This differs from client/server architectures, in which some computers are dedicated to serving the others. In P2P network, peers are computer systems which are connected to each other via the Internet. P2P networks are classified in two ways, one by what they can be used for (file sharing, content delivery, media streaming, etc) and other according to their degree of centralization. Different P2P networks according to the latter classification are pure P2P, centralized P2P, hybrid P2P, structured P2P, unstructured P2P, etc. The important aspect in all P2P networks is that all clients provide resources, including bandwidth, storage space, and computing power. Therefore the total capacity of the system increases as more nodes (peers) arrives and demand on the system increases.

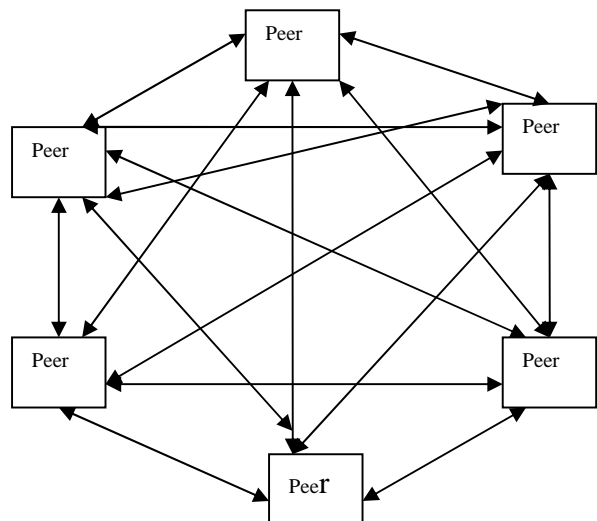


Fig 2: Peer-to-Peer Network

Pure P2P network is the one where peers acts as equals, merging the roles of clients and server with no central router or server

for managing the network. Some examples of pure P2P are Gnutella and Freenet which are used for file sharing [14]. Centralized P2P network involves using a large server to provide the directory service. Whenever a user launches the P2P file-sharing application, the application contacts the directory server which keeps track of IP address of a peer and its objects which are available for sharing. The file transfer between peers is decentralized, but the process of locating content is highly centralized. Napster was the first commercial company to deploy a wide-scale, P2P using centralized directory. In Decentralized P2P network, the responsibility of maintaining the content-location directory is given to a peer, called as group leader, with in a group of peers. To join a network, a peer first contacts a bootstrapping node which responds with the IP address of one of the group leaders, and the peer then establishes an edge with that group leader. The advantages of this design are that sizes of the database are relatively small and the directory information is now distributed over ordinary peers. KaZaA, a popular P2P application, was one of examples of decentralized P2P network Both structured and unstructured P2P networks uses Query Flooding technique in which a peer joins the overlay network, by contacting the bootstrapping node which gives the IP address of one or more existing neighboring peers in the overlay network. In an Unstructured P2P network (Gnutella, FastTrack, etc), if a peer wants to find a desired piece of data in the network, the query has to be flooded through the network to find as many peers as possible that share the data. Structured P2P (CAN, etc) uses distributed hash table (DHT) to assign the ownership of each file to a particular peer which ensures effective routing for a desired file, even if the file is extremely rare.

3. Counter Arguments

Supporters of Multicast technology believe that world wide deployment of Multicast is possible with existing model and architectural design. Below we examine their arguments.

3.1 Cost-effective

Ideally, Multicast is supposed to be cost effective as evident from the Multicast model which is a service that reduces the bandwidth required to transport data to multiple recipients. This bandwidth reduction acts as great incentive for corporate networks and ISPs to employ Multicast technology. However, to date, multicast is very costly in terms of deployment, installation and management. The cost of network service can be defined as sum of network related costs (which include sum of router state, processing, signaling and interdomain routing scalability) and management costs. Management costs include ease of deployment and maintenance in terms of human resources and infrastructure [2]. It makes sense for corporate and ISPs for adapting Multicast technology only if the deployment and management costs offset the bandwidth savings associated with it. Multicast has high initial cost (higher than Unicast) but adding new receivers should be cheaper when compared to adding the same in Unicast. This might be true if a new receiver is on the same domain but not if it is on the different domain or autonomous system (AS). Being on the different domain or AS increases management and network costs which would exceed the benefits of efficient multicast routing

3.2 Reliable

Reliable multicast protocols are available which provide communications services similar to TCP [3]. Multicast Dissemination protocol (MDP) [4] and the Reliable Multicast Transport Protocol (RMTP) [5] are some of the existing reliable multicast protocols. However, each of these protocols is designed in such a way that they serve only specific group of users and applications. This is because multicast applications which range from distributed databases to multicast ftp have different requirements and different reliable message delivery guarantees [3]. As a result there is no standard reliable multicast protocol available that can be thought as analogous to TCP. Only if such a protocol gets

developed, then only multicast applications can achieve true transparency.

Also the current Multicast service model does not provide group management which includes receiver authorization, transmission authorization and group creation. The four important distinct mechanisms that are missing or not addressed in current IP Multicast service and architecture are Authentication, Authorization, Encryption and Data Integrity. Source authentication and data integration are possible through IPsec services but not receiver authentication. As the current multicast address space is unregulated, it is quite possible that members of two sessions receive each other's data. Overall, the lack of access control in Multicast model may lead to a serious problems and number of possible attacks [2]. These include but not limited to flooding attacks, where high bit-rate useless data can be transmitted on the same multicast group causing congestion and packet loss, session collision, which arises if the two sessions use the same Multicast address which can interleave their data.

3.3 Successful Deployment

Supporters of multicast claim that there are few commercial applications, using multicast network, that are deployed successfully. One such example is Mbone which stands for Multicast Backbone being used for shared communication such as audio/video conferences. It marked the first widespread use of multicast in the Internet [6] and is still used by several researchers to develop group communication protocols and applications [7]. However, limits in performance, security and scalability, hold it back from becoming a truly enterprise-class multicast service [8]. Also few intranets have been developed successfully using multicast model but these are private networks where system or network administrator have complete control over multicast addresses and authorizations. The problems encountered in global multicast deployment don't always apply to private networks. Therefore, very few commercial applications of multicast are in use even after twenty eight years of its introduction. There is not

much multicast support from commercial routers yet.

4. Support for P2P

In this section, we provide reasons why P2P is better for content distribution applications rather than multicast.

4.1 No Common Infrastructure

Peer-to-peer networks require no common infrastructure. Anyone can enable P2P application by just installing the client software. By doing so, they connect to the members of the community and exchange files among them [11]. Therefore, joining a P2P network for any client (peer) requires no cost. For example to participate or connect to BitTorrent, a popular P2P file sharing protocol, all one needs to do is download one of the BitTorrent clients such as Azureus, Bitcomet, uTorrent, etc. In contrast, Multicasting requires common infrastructure. The network administrators must configure the routers so that they can support multicast and all the clients in different autonomous systems must conform to a common multicasting protocol. Achieving this seems to be difficult for the reasons mentioned in section 3.2.

4.2 Scalability

Scalability of a network indicates its ability to gracefully handle growing amounts of work. As the concept of Peer-to-peer and its usage is increasing tremendously, there is a need for the P2P networks to be highly scalable, which they are. Many algorithms have been developed to address this issue [9]. Consider a first generation peer-to-peer network Gnutella, where scalability can be achieved by simply making extensions to the basic behavior. For example, caching the query results in Gnutella nodes seems to be an excellent, downward compatible, upgrade to the algorithm [9].

On the other hand, Multicast faces many challenges in scalability. The scalability of multicast protocols is primarily determined by the usage of network resources. The routers enabled with multicast have to maintain information not only about the

individual networks but also about the multicast groups (typically an entry in the router) as well as the state of each interface. As the number of multicast groups increases the number of entries in the router also increases. Also, the memory usage is multiplied as the group size becomes large [10]. For example, MBone faces a lot of problems while scaling its existing structure to large groups as the members grow with regards to the resource utilization [8, 10].

4.3 Reliability

Reliability is a measure of the continuous delivery of correct service. Achieving reliability in P2P is very easy as peers communicate using a direct, reliable TCP connection whereas achieving reliability in multicasting is very difficult for reasons mentioned in section 3.2

4.4 Robustness

A system is said to be robust if either it can recover quickly or holds up well under exceptional circumstances. Because of its distributed nature, P2P increases robustness in case of failures as it replicates the content over multiple peers. As a result of that, peers can find data without depending on centralized index server. The main advantage with these applications is that there is no single point of failure. In contrast, Multicast which is a client-server model has a single point of failure. Achieving robustness in multicast is quite complex which requires lot of effort as mentioned in [15].

4.5 End-user incentive

The motivation behind P2P is to use the cumulative resources like bandwidth, storage space, and computing power of all the clients. As client has the benefit of gaining resources, they have an incentive to build and use new P2P applications. Also, as the number of clients increases, the total capacity of the system also increases. As a result of that, the concept of P2P and applications using P2P are increasing tremendously and the scenario is better than

it was eighteen years ago. The list of applications and their uses along with the protocols used is given in [12].

In multicasting, the end user has no benefit and cares less whether his data comes as a multicast or unicast packet. In other words, the end user has no incentive. Thus, it is the service providers who must drive the multicast development and needs to save the bandwidth. However, for reasons mentioned in section 3.1, service providers will not embrace multicasting until bandwidth savings offset deployment and management costs. Therefore, the scenario for multicasting is not better even after twenty eight years of its introduction [13].

5. Conclusion

P2P network scores over Multicast network with respect to scalability, availability and performance. P2P has become popular choice for content distribution as it has the ability to function, scale and self-organize in the presence of large group of nodes and networks and computer failures. Same can't be said about IP Multicast which still needs to address the marketing demands of corporate networks, ISPs and its customers. These require a service and a protocol which is not only easy to deploy, control and maintain but also highly scalable.

References

- 1) <http://www.multicast-isp-list.com/>
- 2) Christophe Diot. Et al, "Deployment Issues for the IP Multicast Service and Architecture".2000
- 3) Deborah A. Agarwal. Et al, " Using Multicast in the Global Communications Infrastructure for Group Communication".21st Seismic Research Symposium.

- 4) J.Macker and W. Dang, " The Multicast Dissemination Protocol version 1(mdpv 1) Framework", Technical white paper
- 5) J.Lin and S.Paul," RMTP: A Reliable Multicast Transport Protocol", IEEE INFOCOM, 1996
- 6) Sherlia Shi," A Brief History of Internet Multicast". 2002
- 7) A Brief History of Internet Multicast <http://www.arl.wustl.edu/~sherlia/thesis/chap1/node3.html>
- 8) The Multicast Backbone Network <http://www.iee.org.hk/iee/files/48.pdf>
- 9) Scalability of Peer-to-Peer Systems, HUT T-110.551 Seminar on Internetworking
- 10) Multicast Scalability <http://www.cse.iitk.ac.in/research/mtech1997/9711105/node7.html>
- 11) What is peer-to-peer networking? <http://www.solyrich.com/how-p2p-works.asp>
- 12) Peer-to-Peer Wikipedia <http://en.wikipedia.org/wiki/Peer-to-peer>
- 13) Multicast as a Viable Commercial Enterprise.
- 14) Computer Networking- A Top-Down Approach Featuring the Internet by James F. Kurose, Keith W. Ross
- 15) Stephen Deering. Et al," The PIM Architecture for Wide-Area Multicast Routing"