Multicast vs. P2P

Anudhar Emani

anudhar@physics.rutgers.edu

Abstract

IP multicast offers scalable point-to-multipoint delivery necessary for using group communication applications on the Internet. However, the IP multicast service has seen slow commercial deployment by ISPs and carriers due to issues in technology as well as commercial dynamics. During the same time content distribution using peer-to-peer technologies have gained immense popularity. Once P2P is used mainly for illegal content distribution but today legal content providers like BBC, Redhat were using P2P for minimizing operational expenses. Today even live content streaming is possible through P2P. In this paper we see the Issues in commercial deployment of Multicast and Limitations of P2P and why Multicast may have chance for wide spread deployment in the feature.

1. Introduction:

Due to the explosion in internet usage over the last decade and rapid convergence of information, entertainment and communication industry today internet is one of the most viable option for content delivery. The content providers want to reach as many costumers as possible using internet. Various models are proposed and deployed to address content distribution.

One of the earliest of such proposals was multicast. Even though multicast addressing was as a separate address class from the beginning. It was only in late 80s that multicast extensions to the unicast routing mechanisms across datagram-based inter-networks were introduced, marking the beginning of IP multicast. As Internet was basically designed with one-to-one communication as prime objective Multicast requires some modifications to the internet infrastructure currently deployed.

Content distribution technologies like Peer-to-Peer(P2P) and Content Distribution Networks(CDN) gained popularity as they are effective even using one-to-one (TCP/IP) communication. Currently P2P is widely used for content distribution whereas Multicast is still to gain popularity.

2. IP Multicast Overview

Multicast is a content distribution service model where only one data stream from the source serves an unlimited number of receivers. Unicast distribution requires the content source to send separate data streams per requesting receiver whereas multicast allows the network to build and maintain a distribution tree from the source to all interested receivers. Multicast replicates the data stream in the routers per outgoing interface on behalf of receiver requests.

Networks using IP Multicast deliver source content to multiple users (hosts or receivers) that are interested in the data stream. A multicast channel refers to the combination of a content source IP address and the IP Multicast group address to which the content is being broadcasted. Unlike unicast/broadcast addresses, multicast groups do not have any physical or geographic boundaries, and receivers interested in joining can be located anywhere on a network or the Internet as long as a multicast-enabled path has been established.

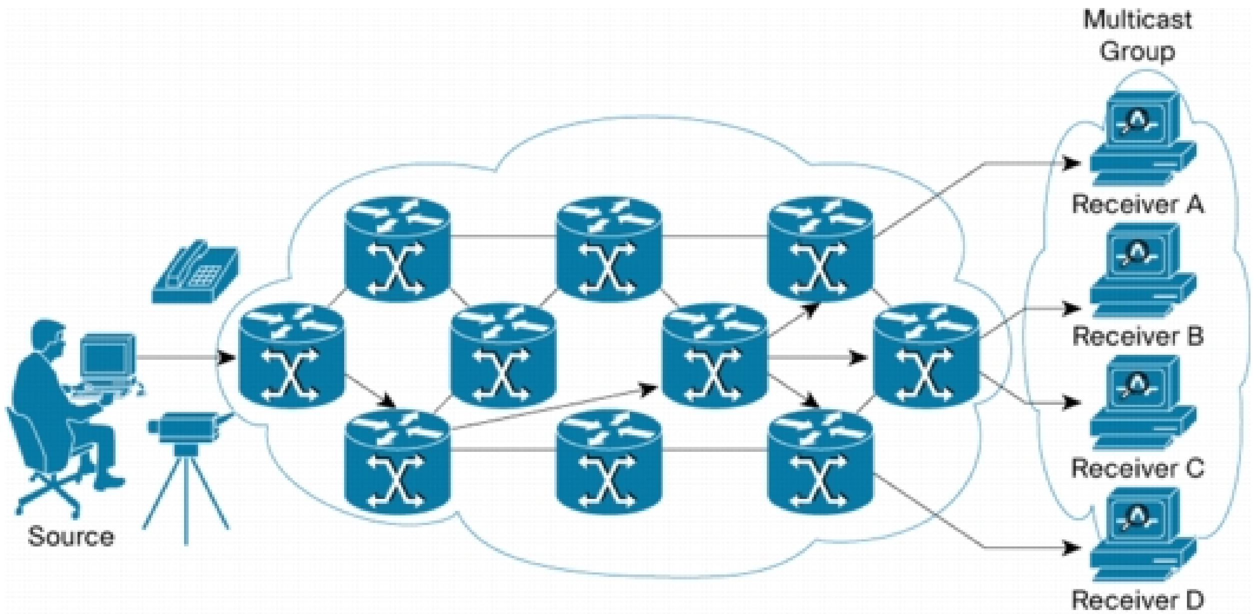

(from Cisco Systems[5])

IP multicasting relies on two mechanisms: a group management protocol(GMP) to establish and maintain multicast groups, and multicast routing protocols to route packets efficiently.

To receive a particular multicast data stream, hosts must join a multicast "group" by sending an Internet Group Management Protocol (IGMP) message to their local multicast router.

Once the receivers join a particular IP Multicast group, a multicast distribution tree is constructed for that group. The protocol most widely used for this is Protocol Independent Multicast or PIM. It sets up multicast distribution trees such that data packets from senders to a multicast group reach all receivers which have "joined" the group.

Advantages of Multicast:

We can see from above diagram that as the number of receivers increases bandwidth sharing also increases, unlike in unicast where each new receiver requires a new data stream. Robert Chalmers et al.[1] gave an estimate of bandwidth effeciency as a function of number of recievers of the multicast stream as shown in figure below.

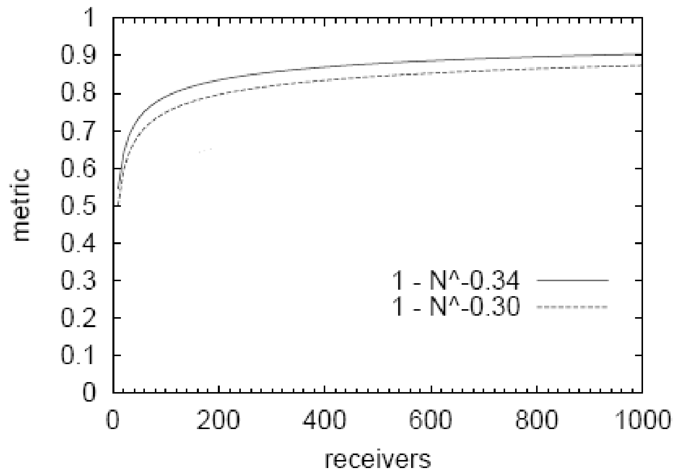

Efficiency estimate shown over 1,000 receivers for a range of efficiency factors.

Efficiency of zero means multicast and unicast use the same bandwidth and as it approaches one all receivers share a single path.

So in many applications like distribution of Live Video /Audio streaming, etc., can achive significant savings with Multicast. Multicast can be used along with "Digital Fountains"[ ] for delivering software updates etc.,

3. Issues with Multicast Deployment :

Since its introduction, IP multicast has seen slow commercial deployment in the Internet. Although it has been available through the experimental Mbone for a number of years, it is just beginning to see commercial support from carriers. There were both technical and commercial issues for such a slow growth.

Technical Issues:

The original technology was not suitable for adoption throughout the Internet. Basic parts had to be reengineered on the basis of experience. This took from 90's to early 2001

The original multicast network MBone, used a simple routing protocol called DVMRP (Distance Vector Multicast Routing Protocol). DVMRP is Dense Mode which floods multicast traffic to all receivers, until they ask it to stop. As there were only isolated sub-networks that wanted to deal with DVMRP, the old MBone used tunnels to get multicast traffic between DVMRP sub-networks. i.e., the multicast traffic was hidden and sent between subnetworks via unicast.

This mechanism was simple, but required manual administration and absolutely could not scale to the entire Internet. Also DVMRP creates its own routing table, and that requires substantial routing traffic which grows combinatorially with the size of the MBone. Thus multicast was perceived to be not scalable.

Commercial Issues:

Multicast basically is a very good solution for addressing content delivery but various commercial aspects have prevented its wide acceptance by ISPs as well as customers. Let us look at few of these important issues.

The original service model was designed without a clear understanding of commercial requirements or a robust implementation strategy. There is an overhead associated with providing an acceptable level of service using multicast. Many service providers have found alternative ways, such as Peer-to-Peer and CDN's, to serve information without multicast protocols. It can only work if all of the infrastructure players agree to deploy multicast. The users care very little whether data comes as a multicast packet or unicast packet. Since there is "bandwidth glut" of unused capacity that was built up during the late 1990s dotcom boom the impetus for implementing Multicast was very little until recently .

In the initial days of internet there was very little demand for Multicast as the bandwidths are smaller and also there are few applications which require multicast. So most of the internet has grown without support for Multicast, except in small islands like corporate networks etc., So now Multicast is in all or nothing proposition.

As current internet based on unicast has evolved over time and well tested, very few ISPs want to invest in Multicast unless it is easy to install and should impart little inconvenience to ISPs and their users. Also, before multicast can be even considered viable, some guarantee to providers must be made that their multicast address will be unique, at least for the duration of a particular session.

Finally, multicast users should be able to expect the same reliability as unicast users. The multicast architecture arguably requires more setup and administration than unicast. But even after implementing multicast, some carriers have reported that the current multicast architecture is unstable.

There is pattern in which the current internet infrastructure is upgraded. The newest and most expensive routers with advanced technology were deployed at the backbone and the old ones are gradually moved towards the edges. This gives the optimal returns for their investments on routers, as edge routers generally need not handle as much traffic as the ones in the backbones. But implementing Multicast involves that each and every router on the path supports relavent protocols. Since many old routers do not support Multicast they need to upgrade all of them, which requires huge investments.

Most routers deployed near the backbone are optimized for increased throughput and hence less intelligent. So implementing Multicast, which requires support for IGMP and corresponding routing protocols, congestion control mechanisms may place significant overhead on the routers. They may not be able to handle multicast without considerable performance issues.

Also of great concern to ISPs are inter-domain dependencies. Much of the multicast infrastructure requires trust among ISPs. It is entirely possible for the central core of a spanning tree to exist outside of one particular ISP's domain. Since ISPs generally have little control over the administration of the tree they also have no incentive to handle traffic for multicast groups for which they have no senders or receivers.

Even though it is not certain whether the cost of deploying multicast will lead to huge savings in feature[3], moreover as it involves close co-ordination among various ISPs to provide the required service, it requires revenue sharing among them. But there is no simple way to divide transmission costs between them. A study done on cost analysis of multicast systems [4] found that there exists no mechanism until now that is both efficient and budget balanced among all involved parties. So until such a mechanism is established many ISPs may be reluctant to provide Multicast service over their networks.

4. Multicast - Recent Trends

Starting about 1997, the building blocks for a multicast-enabled Internet were put into place. An efficient modern multicast routing protocol, Protocol Independent Multicast – Sparse Mode (PIM-SM), was deployed. The mechanisms for multicast peering were established, using an extension to BGP called Multiprotocol BGP (MBGP), and peering became routine.

The service model was split into:

1) a many-to-many part (e.g., for videoconferencing):

Any-Source Multicast (ASM), and

2) a one-to-many (or "broadcast") part:

Source-Specific Multicast (SSM).

By 2001, these had completely replaced the old MBone. The availability of better versions of these protocols is the primary driver behind multicast scalability.

Multicast was being deployed in Islands like Security Exchanges(NASDAQ, NYSE, AMSE, HKSE, etc..),Securities Trading Enterprises,Enterprises,Service Providers,MVPNs, DSL, Cable.Global multicast connectivity Still growing, but at a much slower-rate. It still needs to overcome the all-or-nothing barrier.

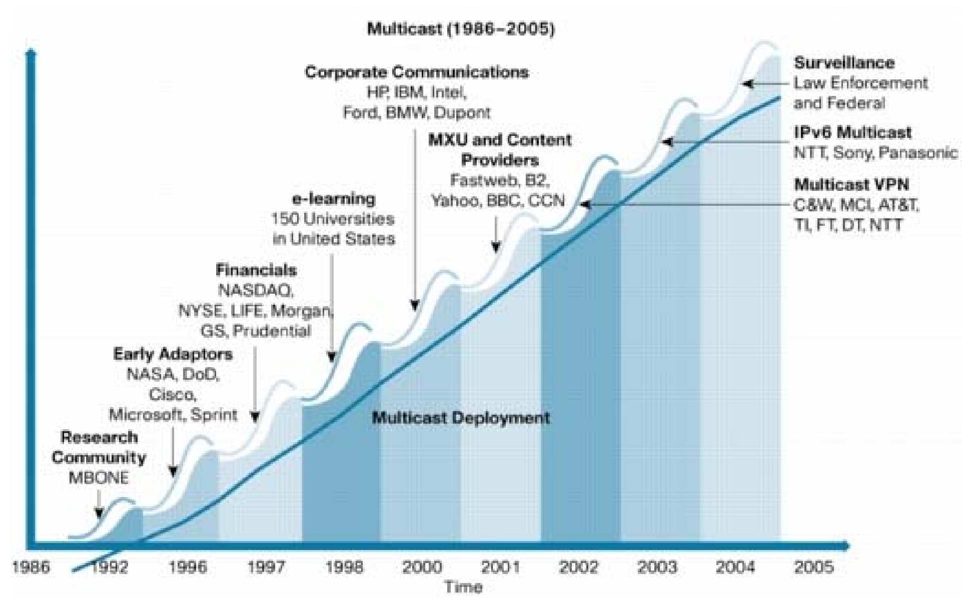

(from Cisco Systems[5])

5. Peer- to-Peer[6]:

As Multicast is not yet possible for content distribution over different domains other technologies based on existing internet infrastructure have gained immense popularity. Some of such technologies are Peer-to-Peer(P2P) and CDN's. P2P is the most widely used for content distribution currently. P2P networks are overlay networks depending on diverse connectivity between participants in a network and the cumulative bandwidth of participants rather than conventional centralized resources where a relatively low number of servers provide the core value to a service or application. P2P networks are typically used for connecting nodes via largely ad hoc connections. Such networks are useful for many purposes such as sharing content files containing audio, video, data, and real-time data, such as telephony traffic, live video streams, are also passed using P2P technology.

Advantages :

The various advantages in using P2P for content distribution over Unicast or Multicast are listed below:

Scalability: P2P services are inherently scalable. As each node(user) shares part of the load, more nodes mean not only more demand but also more capacity. By contrast, if a service runs on a central host, more users will eventually mean that more resources need to be added at the host.

Reliability: Due to wide scale distribution of data gives the potential for very high reliability because data is spread over many machines that are completely independent, and they're unlikely to fail at the same time. This has a lot of economic implications because you can achieve higher reliability with a lower cost machine.

Hardware Economics: Users pay about the same price per unit of computing and storage as a central provider would pay. But the users' machines are already bought and paid for, and they're mostly sitting idle. The incremental cost of assigning work to one of these machines is nearly zero. But in a centrally controlled system, new machines must be bought, and reserved for use in providing the service.

Disadvantages:

P2P is scalable only if bandwidth availability at the ISP backbone is scaled accordingly. As there is lot of duplication of data, demand on bandwidth increases accordingly. Various P2P technologies were proposed using similar protocols like Multicast in-order to reduce overall bandwidth consumption. Some of such protocols are MULTI+ protocol described in [7], or PeerCast [8], or Nemo [9]. But most of them require duplication of packets on physical links and hence suffer performance penalties and also incur larger end to end delay than IP Multicast. [10].

Dynamics underlying the success of p2p:

P2P gained most of its popularity due to illegal content distributed by various P2P networks like kazza, Bit-Torrents, eDonkey etc.,. Around the world in 2006, an estimated five billion songs, were swapped on peer-to-peer websites, while 509 million were purchased online. There is also no need to change all routers in the world or make an agreement with every single ISP in the world to make multicast work.

Several Content providers are using legal P2P networks to reduce their distribution cost to zero. The most popular cases are the BBC(iPlayer), Redhat, Joost. Even live video streaming is being explored my many players like Microsoft(Livestation)etc., Even VOIP application like skype based on P2P gained popularity quickly. The content provider need not buy extra bandwidth from ISP. However, they are basically pushing the distribution cost to the ISP.

6. Limitations of P2P:

Due to inherent nature of P2P the content distributers have limited control over data distributed using P2P unlike using unicast or multicast. This may further encourage piracy.

P2P protocols allow end hosts to pretty much fill their access bandwidth constantly. Total access bandwidth could exceed total backbone bandwidth by one or two orders of magnitude. In other words, the Internet is could be oversubscribed.

Over the last few years P2P traffic is growing at rapid rate due to broadband penetration. There are various studies indicating that P2P traffic ranges from 60% to 80% of the total internet traffic. ipoque has analyzed the Internet traffic in five regions of the world between August and September 2007. From the figure below we see that P2P traffic is about 74% of Total internet traffic in Germany and is at similar levels in other regions also.

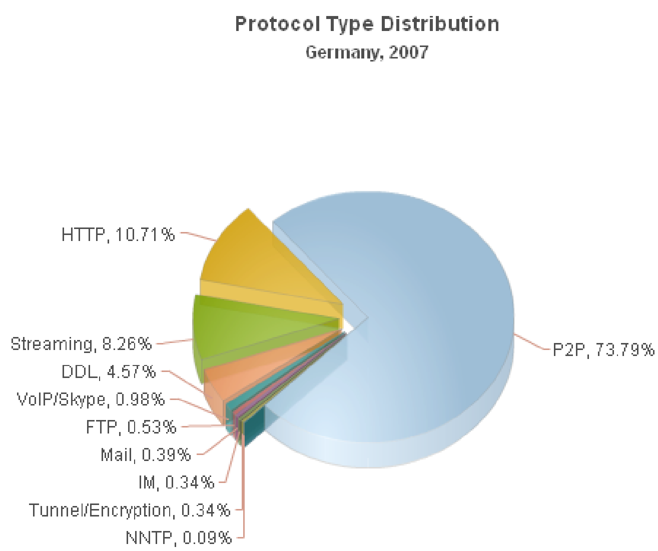

Since most of the local ISP providers are charging users based on bandwidth, they get lot of value out of P2P but nobody got paid extra for it. Also P2P topologies are not aware of the ISPs economics. The Internet has been built on the premise of core to edge distribution. But P2P assumes the internet to be flat.This makes P2P traffic to be routed over expensive transit links

So ISPs are actively exploring ways to limit P2P traffic. Some are using filters for de- prioritizing P2P traffic. This limits the availability of bandwidth for content distribution. Some ISPs have started charging based on volume of internet traffic rather than bandwidth alone. This motivates the nodes to restrict upload traffic, which result in serious degradence in performance of P2P networks. Some ISPs are exploring alternatives instead of limiting P2P traffic like caching the most frequently requested P2P files. But it is expensive to maintain and also it leads to legal problems.

So ISPs want legal content providers to pay for bandwidth, or they may limit its use. One way or another, legit content providers have to pay for distribution and current legal P2P models do not fit well with ISPs business models and one way to keep everybody happy is if ISPs deploy their own legal P2P technology.

7. Future of Multicast and P2P

There is no doubt that P2P is here to stay. Traffic levels of P2P are growing year by year. Acceptance of P2P technology by content providers like Microsoft, BBC, etc., is pushing many other leading players towards P2P for legal content distribution.

The existing Internet based on one-to-one communication model does not work for mass distribution. Media companies see P2P as the next "big wave" in distribution, allowing them to reach a huge audience with reduced price products at little/no cost of distribution beyond setup. End users love P2P as it gives them access to the media they want, when they want it and at high speed (depending on popularity).

The "bandwidth glut" built during dotcom boom is virtually eliminated as the new services are gobbling up huge amounts of bandwidth, to the point where in fact, we are speeding towards a "bandwidth crisis".
There were reports predicting that Internet traffic may grow 50 to 100 fold by 2015[11]. To support such a growth rate it requires huge investments in building bandwidth and related infrastructure. Since P2P traffic is about 60% to 80% of the total traffic one way are the other this is not going to be a cheap solution for content distribution.

But by using Multicast one can reduce bandwidth requirements drastically as more and more receivers are hooked up to the internet. If Multicast is available widely, even P2P can benefit a lot by building overlay networks using Multicast rather than unicast that is currently deployed. This leads to huge savings in terms of bandwidth. So P2P and Multicast complement each other very well.

8. Conclusion:

We have seen various aspects of content distribution using Multicast and Peer-to-Peer models. We also looked at underlying dynamics leading to success of P2P and relatively slow growth of Multicast. But this may change in near feature due to inherent advantages of using Multicast. Multicast may reach a threshold at which it may become commercially viable and advantageous, over Peer to Peer technologies, to deploy Multicast in large scale.

References:

[1]Modeling the Branching Characteristics and Efficiency Gains in Global Multicast Trees,Robert C. Chalmers, Kevin C. Almeroth, IEEE INFOCOM 2001.

[2] Agarwal, D. and Lawrence, O. “Using Multicast in the Global Communications Infrastructure for Group Communications”. 21st Seismic Research Symposium.

[3]Diot, C. et al. “Deployment Issues for the IP Multicast Service and Architecture”. 2000.

[4] Feigenbaum, J. et al. “Hardness Results for Multicast Cost Sharing”.

[5]https://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/prod_white_paper0900aecd804d5fe6_ns610_Networking_Solutions_White_Paper.html

[6] http://en.wikipedia.org/wiki/Peer-to-peer

[7] Luis Garc´es-Erice, Ernst W. Biersack and Pascal A. Felber, *MULTI+: Building Topology-AwareOverlay Multicast Trees*, Institut Eurecom, Sophia Antipolis, France, April 2004.

[8] Jianjun Zhang, Ling Liu, Calton Pu, Mostafa Ammar, *Reliable Peer-to-Peer End System Multicasting through Replication*, Fourth International Conference on Peer-to-Peer Computing (P2P'04), August 2004.

[9]. Stefan Birrer, Fabián E. Bustamante, *Resilient Peer-to-Peer Multicast from the Ground Up*, Network Computing and Applications, Third IEEE International Symposium on (NCA'04), September 2004.

[10]. Yang-hua Chu, Sanjay Rao, Hui Zhang and Srinivasan Seshan, *End System Multicast*, http://www.cs.cmu.edu/~sanjay/EndSystem/endsystem.html

[11]http://www.discovery.org/a/4444