On the feasibility of Commercial, Legal P2P Content Distribution*

Pablo Rodriguez
Microsoft Research
Cambridge, UK
pablo@microsoft.com

See-Mong Tan
Microsoft Corporation
Redmond, WA. USA
seemongt@microsoft.com

Christos Gkantsidis
Microsoft Research
Cambridge. UK
chrisgk@microsoft.com

ABSTRACT

A spirited panel was recently held at the 10th International Web Caching and Content Distribution workshop on the future of P2P in content distribution [1]. After more than ten years of content distribution research and technology efforts, P2P is emerging as an alternative solution to solve the mass distribution of large digital content. However, using P2P for commercial content distribution faces a number of serious challenges. Issues such as content protection, impact on ISPs, security, end-to-end connectivity, and business models need careful consideration before P2P can be used as an efficient tool by content providers. In this paper, we summarize the issues brought up in discussion, and delve deeper into the feasibility of commercial, legal P2P content distribution solutions.

Categories and Subject Descriptors

C.2.2 [**Computer-Communication Networks**]: Miscellaneous; C.2.4 [**Distributed Systems**]: Distributed Applications

General Terms

Design

Keywords

P2P, Content Distribution

1. INTRODUCTION

Internet Content Distribution first commenced with Web caching, cache cooperation and caching infrastructures. Then Akamai turned caching into a service for content providers, and Content Distribution Networks (CDNs) became one of the most important advances in Internet technologies over the last years [2].

In recent years, a new paradigm for Content Distribution has emerged based on a fully distributed architecture where commodity PCs are used to form P2P cooperative networks and share their resources (bandwidth, storage, CPU). By capitalizing on the bandwidth of end-nodes, P2P architectures offer great potential for providing a cost-effective distribution of bandwidth-intensive content to thousands of simultaneous users both Internet-wide and in private networks. To date, P2P has been widely associated with sharing of pirated content. In this paper we discuss the feasibility of using P2P systems for commercial content distribution.

*The views in this paper represent those of the authors and those expressed during the WCW [1], but do not necessarily reflect the position of Microsoft Corporation.

1.1 Content is King

Users are relying more and more on the Internet to obtain digital content (e.g. movies, music, games, etc.).[1] Consequently, traditional retailer distribution models (e.g. Blockbuster, Netflix, etc) may soon be complemented with online 'e-tailer' services that provide high quality digital content on-demand ([19]), changing the way people watch television and buy digital content.

Despite the large amount of potential content that can be made available for download, content providers have been quite reluctant to doing so for two primary reasons: a) security and content protection, and b) distribution costs.

Piracy and content security can be tackled to a great extent by protecting the content using DRM techniques (Digital Right Management).[2] Distribution costs, on the other hand, require efficient delivery mechanisms.

1.2 The need for FileCasting

Soon after the Internet was created, IP multicast promised to enable efficient point-to-multipoint transmissions. However, lack of proper billing mechanisms, infrastructure changes, and lack of coordination among ISPs slowed down the IP multicast deployment to the point that it is not clear whether multicast at the network level will ever happen [7].

Still, the cost of distributing large content files to many users is non-negligible. For instance, the cost of distributing a HD movie can be close to $2 at current bandwidth prices. Since the current retail prices of DVDs range from $12 to $18 in the US market, and the user gets a physically boxed product, the authors expect that the final price for a virtual download to remain close to that price point, even if it is in high definition. In addition to the egress charges, content providers need to provision for infrastructure support (e.g. server farms, load balancers, etc).

P2P networks, instead, implement multicast support from the network edge, removing most of the load form the content source. Consequently, P2P has emerged as a serious alternative to providing point-to-multipoint transfers without requiring much infrastructure support.

2. THE P2P PHENOMENON

[1]The RIAA reported that CD and other physical format sales decreased in 2005 from 345.4 million units to 326.1 million units. During the same period, there were 148.7 million digital downloads, a major increase from 58.6 million during the same period in the previous year [17].

[2]While DRM systems have been breached multiple times, modern on-line revocation and renewability mechanisms serve to mitigate the ill effects of any single crack and pose a significant barrier for users trying to break content protection.

P2P applications already account for roughly 60% of Internet's traffic [3]. For the most part, research efforts in P2P have focused on enabling large scale distributed search[4][5]. However, recently, a new trend is emerging where P2P systems are considered as an alternative solution to enabling large scale commercial content distribution (e.g. [10]).

For instance, systems such as BitTorrent have been used for the distribution of Linux RedHat distributions and are now being integrated in several mainstream browsers. Another example is the BBC's announcement on the use of a P2P network to distribute one week worth of TV programs in the UK ([8]). These examples indicate that publishers are eager to use P2P networks as a tool for large scale content distribution.

2.1 Benefits

P2P systems have a number of good properties that make them very attractive. These include:

- Scalability
- Low server requirements
- Fast deliveries from nearby nodes
- Cost effectiveness
- Enables Push Video-on-Demand

One interesting characteristic of P2P networks is their adaptability, e.g. they dynamically absorb load as demand rises. As an example, P2P provides an ideal distribution platform for the growing number of independent artists seeking a cost-effective channel for delivery (e.g. [9]).

However, there are a number of issues that must be tackled before P2P can become a widely used tool for commercial content distribution. In the next section, we detail some of them.

3. TOP 10 ISSUES WITH LEGAL P2P

Next we describe some of the most important issues with P2P and discuss some ideas on how they can be tackled.

1. End2End connectivity
Network Address Translation (NAT) enables multiple hosts on a private network to access the Internet using a single public IP address. However, NATs often prevent peer hosts from establishing communications with each other, in essence because NATs break the end-to-end principle and prevents the Internet from reaching its full potential. The problem seems to be getting worse with the rapid deployment of home wireless routers. For example, data collected by the MSN Messenger Customer Improvement Program indicate that the percentage of NATed nodes in the United Kingdom and in the US, as of April 2005, was more than 50% of the broadband users.

By preventing peers from communicating with each other, the efficiency of a P2P distribution system is seriously impacted. Nodes behind NATs cannot behave as servers for other peers, thus, reducing their ability to participate in the distribution process.

However, P2P file distribution does not pose such stringent requirements compared to other interactive P2P applications (e.g. Instant Messaging or VoIP) where all participating nodes need to be bidirectional reachable. In essence, because communications can be initiated from either end and connections can be used to transfer data both ways, the problem is easier to tackle. In fact, in a system where peers have the same capacities, roughly, it is enough for 50% of the population to be reachable to obtain a fully connected system.

Current P2P systems rely mostly on manual port configuration. However, techniques for automatic communicating through NAT boxes are becoming stream line, in essence, by layering another address on top of IPv4 (for instance, STUN uses the SIP URI, and Teredo uses IPv6). Guha and Francis also report encouraging practical results in [13]. Having a reliable architecture that enables direct communications among peers is a critical component in the success of these systems.

2. Corruption/Pollution/Privacy
In traditional client-server architectures, receivers often only need to trust a small set of servers deployed by the content provider (or a related entity such as a CDN). In P2P, on the other hand, any computer is a potential server and thus, building trust relationships becomes much harder. Instead, P2P networks need to assume that some nodes may be malicious and may introduce corrupted content. P2P networks need to identify corrupted information on-the-fly and be robust in the face of such attackers.

Another related issue with P2P networks is privacy. As users connect to other peers to download content, they can expose important personal information, such as their IP address, their geographical location, their viewing preferences, etc. By revealing such information, targeted marketing campaigns, spam, or massive security attacks (e.g. worms) can be launched. Thus, mechanisms must be in place to ensure that user's identity is kept private as much as possible.

3. Promoting distribution of legal content
Commercial P2P networks should make an effort to foster the distribution of legal content. Following several rulings against P2P file sharing systems (e.g. [11]), it is important for commercial P2P distribution systems to publicize and encourage their legal uses and applications. In this regard, applications such as those recently enabled by the BBC for distributing TV shows, the World of Warcraft ingame patching system, or the distribution of the independent film Star Wreck, can help improve perception over P2P networks.

4. ISP's impact. Traffic engineering
P2P networks impact ISPs in two different ways: a) increasing load and changing traffic patterns, and b) generating large amount of upload traffic.

P2P has made possible the exchange of very large files among Internet users. Such large file downloads are increasing the amount of traffic on ISPs. More importantly, they are changing the ISP's traffic patterns. Traditional On and Off traffic periods are being replaced with constant traffic levels across the day. These new traffic patterns create problems for ISPs, which re-use the same channel to provide access to many subscribers. With an increasing amount of constant traffic, ISPs cannot multiplex the same number of users in the same link and are forced to provision more resources or reduce the number of users served on the same link (e.g. average contention ratios are now close to 20:1 compared to 50:1, which were common earlier on).

This situation, however, it is not specific to P2P applications. A similar behaviour would occur if users were offered the possibility to download popular movies or other large files at a very low cost from a well-provisioned server-farm.

One problem that is a direct consequence of P2P systems is the increase of upstream traffic. With P2P, content first flows upstream from a user's machine and then downstream into another user's machine, roughly generating double the amount of traffic through the ISP's network compared to traditional client/server architecture. The amount of resources needed to cope with this extra upstream traffic are relatively small compared to the bandwidth needed to cope with a constant increasing demand for large files. Still, it affects the ISP in a critical way since ISP architectures assume that content flows from the core of the Internet to the edges, rather than among the edges. In this sense, P2P applications are increasing the pressure on ISPs to re-engineer their networks to better handle upstream traffic.

Increased upload traffic may prove especially important for those ISPs that depend on other carriers to deliver last-mile traffic (e.g., European Tier-2 ISPs). In fact, for such ISPs, utilizing their transit egress capacity might result in a more cost-effective solution rather than re-routing traffic internally.

5. Pricing schemes

In traditional client-server architectures, the Content Provider pays for all the distribution costs. With P2P, the distribution cost is pushed to the ISPs. End-systems become servers of content but ISPs do not see a corresponding increase in revenue since users are often charged flat rates.[3] One possible option for ISPs is to block, limit, or shape P2P traffic. However, this is an unpopular solution that often takes customers away. Another option is to implement volume-based billing and charge users for the amount of data downloaded and uploaded. But volume-based billing systems are hard to implement and support, and require additional infrastructure that permits users to monitor their usage. Competition among ISPs will also put more pressure to keep popular flat based billing schemes.

The most likely scenario is that more access to digital content will drive more broadband consumers and ISPs will implement a tiered pricing scheme where high tier pricing schemes will allow "all-you-can eat" download/upload schemes, while lower tier pricing schemes will only allow less traffic consuming applications. Irrespective of the pricing scheme finally used, ISPs will need to handle P2P traffic efficiently to service the highest number of subscribers without having to perform major network infrastructure upgrades.

6. Integration with caches

In order to minimize the impact of P2P traffic into ISPs, P2P networks need to be cache-aware. In doing so, ISPs can have better control on how they engineer their network. P2P caches ensure that traffic flows only downstream, thus, preventing the extra upstream requirements. Consequently, caches fit naturally with the traditional asymmetric architecture of today's ISPs, where the downstream channel is more heavily provisioned relative to the upstream, driven by the assumption that customers download more than what they upload. In addition, caching also provides significant savings since it reduces the amount of traffic over the ISP's transit link (hit ratio for current P2P systems is close to 50% ([6], [3]), and further reduces end-user's download times.

In an ideal world where caches would be deployed at every ISP and caches would cooperate, the need for P2P content distribution would be diminished. However, many ISPs may decide not to deploy caches because of the incurred costs, and cache cooperation among different ISPs is hard to achieve. Instead, P2P networks can be used to scale caching systems without requiring additional infrastructure investments and to provide delivery support in the areas where no caches exist. To a large extent, P2P and caching can be seen as complementary to each other. Consequently, several efforts to integrate P2P networks with standard Web proxies or specialized P2P caches are already in place (e.g. [12], [3]).

7. Performance Guarantees and Manageability

Current P2P networks do not offer the same quality of service assurances than those provided by infrastructure-based solutions. The primary reason is that peers are quite volatile and can have very fluctuating rates; indeed, it is often the case that users wait long periods of time before finishing their download. While infrastructure based solutions can be well monitored, dimensioned, and provisioned, this is not the case in current P2P systems. Consequently, it is difficult to assure a given performance level, and traditional billing schemes based on strict service level agreements may not be adequate for such P2P systems.

Content providers using a P2P network require, however, feedback about the performance they are providing to end-users. Similarly, P2P administrators must have good tools for measuring and monitoring performance, troubleshooting errors, or assisting customers in solving their connectivity problems.[4] Unfortunately, such tools do not exist today since current file-sharing applications rely on end-users monitoring and troubleshooting their own problems.

8. Incentive mechanisms for cooperation

One critical element in P2P networks is that they rely on receivers willing to participate in the distribution process. Current P2P applications enforce different levels of cooperation where nodes need to contribute to the system roughly as much data as they want to download from it. However, other incentives such as economic savings, access to premium content or community-based alternatives should be explored.

Motivating users is easier if the ISPs continue the current flat price model since users can only benefit by contributing their resources; however, if the ISPs adopt more elaborate traffic schemes for charging consumers based on their upload traffic volumes, then the customers will need to make complicated decisions of whether to contribute to the P2P network or not.

Another alternative is to include P2P technologies in home set-up boxes. In such case, P2P could be used by ISPs or cable providers to enable multicast support from the edge of the network, without requiring core infrastructure changes or explicit user cooperation. This could enable interesting push-VoD scenarios and similar applications.

9. Locality awareness

One major problem of current P2P networks is the fact that content is retrieved from random nodes in the network. In doing so, P2P applications often fetch content from nodes

[3] The ISPs will increase their revenues by subscribing new broadband customers, however, the extra revenues are small compared to the anticipated revenues of selling new services (i.e. ISPs becoming themselves content distributors).

[4] How can the administrator check that the user is connected or not in the P2P network? How can the administrator debug the slow rate experienced by a user?

in other ISPs, thus, consuming capacity at the local ISP's transit link. In fact, with current P2P systems, ISP's transit link is often used as much as with a client/server distribution model where many copies of the same file are distributed unnecessarily [6].

To mitigate the impact of P2P applications into ISPs, P2P applications can leverage the existence of copies of the same content in the local ISP. Consequently, a locality-aware P2P application can ensure that a large portion of the traffic does not to travel outside of the ISP's boundaries, thus, mimicking the performance of a cache system. In fact, initial reports suggest that a locality-aware P2P system only requires about 1.5 times the downlink link capacity compared to that of a cache [6].

Implementing a locality aware P2P system can be done with simple modifications to current peer discovery protocols or by deploying request redirectors at ISPs [14].

10. Play as you download? Streaming?

Even though P2P systems have been very successful in distributing large files, they fall short in providing a streaming media experience to the end user. Users often need to wait for the full file to be downloaded before it can be played, thus, adding significant delay before the content can be consumed. Some P2P networks have been used to provide live streaming broadcasts ([15], [16]), however, little attention has been paid at enabling Video-on-Demand services where each receiver can start playing at any point in time. An interesting question remains to whether P2P systems can be used to provide Near Video On Demand services where clients are able to start watching almost immediately, without impacting the experience of other users already watching the same file.

4. P2P AND CDNS

With the advent of P2P networks as a cost effective solution to providing distribution of large files, the question arises about the role of more traditional Content Distribution Networks. CDNs have been designed to transmit Web content with small files where quick responsiveness is key. P2P networks with fluctuating resources, can not provide the same levels of service level agreements and tight bounds on latency as those provided by a CDN. Moreover, CDNs provide the ability to publish dynamically generated content, by replicating databases and service intelligence close to the network edge. Finally, CDNs can act as robust infrastructure systems that intercept and absorb DoS attacks, preventing major Internet sites from coming down. Consequently, CDNs stand as a reliable and robust alternative to P2P networks for those applications that require infrastructure support against DoS attacks or fast interactive responses.

5. SUMMARY/CONCLUSIONS

Although P2P technology has been widely associated with the distribution of pirated content and has been subject to a barrage of attacks (e.g. DoS, spoofing and content pollution), as discussed in this paper, there are ways to decrease the risks associated with distributing content using P2P technology.

Strong economic incentives will remain for the content industry to continue refining its legal assault on illegal P2P activity. It is still unclear whether such legal attacks will significantly curtail P2P use. Recent surveys have shown that the global growth in P2P use has not slowed [18] despite significant high profiles wins for the content industry, such as the Grokster case before the US Supreme Court.

Government regulation may also affect the evolution of P2P technologies in the marketplace. Some possible mandates include requiring the checking of content against a registered official database of copyrighted material, or requiring watermarking or other filtering technologies in all P2P technologies.

The transition to successful commercial e-tailing of video content will happen once a reasonable and attractive alternative to pollution prone, legally risky illegal P2P occurs. The technical challenges discussed in this paper will need to be solved and combined with the appropriate business model. At the appropriate price point, ease of use and combined with compelling content, P2P distribution can serve as the engine to efficiently deliver digital downloads in the future.

6. ACKNOWLEDGEMENTS

We would like to thank Paul Francis, Andrew Parker, and Ernst Biersack for participating in the WCW panel discussion and greatly contributing to the ideas in this paper.

7. REFERENCES

[1] 10th Web Caching and Content Distribution Workshop, `http://2005.iwcw.org/`, Sophia Antipolis, Sep 2005.
[2] `http://www.akamai.com/`
[3] `http://www.cachelogic.com/`
[4] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, Chord: A scalable peer-to-peer lookup service for internet applications, *ACM Sigcomm 2001*
[5] A. I. T. Rowstron and P. Druschel, Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems, *IFIP/ACM 2001*
[6] T. Karagiannis, P. Rodriguez, and D. Papagiannaki, "Should ISPs fear Peer-Assisted Content Distribution?", *ACM USENIX IMC, Berkeley, 2005*
[7] M. Ammar, "Why Johnny Can't Multicast", *NOSSDAV Keynote, June 2003*
[8] BBC iMP, `http://www.bbc.co.uk/imp/`
[9] Open Media Network, `http://www.omn.org/`
[10] Christos Gkantsidis and Pablo Rodriguez, "Avalanche: Peer-Assisted Content Distribution", `http://www.research.microsoft.com/~pablo/avalanche.aspx`
[11] Gorkster Case, `http://www.eff.org/IP/P2P/MGM_v_Grokster/`
[12] `http://forum.emule-project.net/index.php?showtopic=14768`
[13] S. Guha and P. Francis "Characterization and Measurement of TCP Traversal through NATs and Firewalls", *ACM USENIX IMC*, Berkeley, 2005
[14] `http://www.sandvine.com/`
[15] `http://www.cs.cmu.edu/~esm/`
[16] `http://www.coolstreaming.it/`
[17] Recording Industry Association of America, 2005 `http://www.digitalmusicnews.com/data/100305first_half_riaa.pdf`
[18] Big Champagne, "P2P Usage Volumes", August 2005. `http://www.bigchampagne.com/`
[19] Can Netflix's Reed Hastings succeed in the battle to deliver movies online, "The Economist", *Jul 7th 2005*