

Software Engineering 4C03 Winter 2005

Topic: An overview of the BitTorrent Protocol

By: Peter Antal – 9906390

Last Revision Date: 04/05/05

1. Introduction

Bittorrent (BT) is a peer-to-peer (P2P) file transfer protocol, used for sharing large files over the Internet. According to a study conducted by CacheLogic, 35% of all Internet traffic was carried by the BitTorrent protocol according to a report in November 2004.ⁱ The first such P2P protocol was Napster, which linked together “peer” machines sharing MP3 music from all over the internet, with central control by a server which kept track of files on the peers.ⁱⁱ The Napster system quickly gained notoriety for its use as an explicitly illegal mechanism for pirating music and raised the ire of the music industry in a short time.ⁱⁱⁱ In response to the hazards of litigation present for the handlers of a centralized file list, a second generation of P2P protocols emerged with decentralized control mechanisms.^{iv}

Based on lessons learned from past P2P protocols, Bram Cohen introduced the BitTorrent Protocol in 2003.^v However, unlike previous P2P protocols, BitTorrent is the first file sharing protocol which has also gained acceptance in corporate network file distribution schemes.^{vi} A notable example is the with the gaming company, Blizzard Software which employs the protocol to distribute patches and beta software to users.^{vii}

2. Benefits of the BitTorrent Protocol

The BitTorrent protocol helps to improve on previous generations of peer-to-peer software with a number of key innovations, which shall be discussed in the following subsections.

2.1. *High Speed Downloads*

Simply put, BitTorrent is quite popular as it parallelizes the process of file transfers. Unlike typical network behavior, in which bottlenecks cause performance to plummet; in BitTorrent, increases in demand can actually improve the performance of a computer network.^{viii} This stable scaling behavior makes the protocol admirably suitable for distributing large, popular content on the internet, for servers of limited capacity.^{ix} The protocol achieves this by breaking up large files (typically on the order of hundreds of megabytes) into uniform blocks of considerably smaller size, such as 256 kilobytes; by dividing up the blocks, the source components can be dynamically requested from multiple source machines, which do not need to possess the complete file to be able to upload blocks, meaning downloaders can quickly start contributing as uploaders for other users.

2.2. *Simplicity, Ease of Use*

Unlike previous P2P protocols which required the use of heavyweight, full featured client applications with elaborate menus and interaction schemes, the BitTorrent protocol operates on a web based distribution paradigm built over the HTTP protocol. Users simply download one of a few client implementations, which to a user act as an extension of their internet browser. By opening files of type .torrent, the protocol client will allow the user to save the file to a location of their choice on the hard drive in the same fashion as typical WWW file transfers. Until the user closes the transfer window, BitTorrent will continue sharing a file to other peers.

2.3. *Effective use of limited upload bandwidth*

A traditional problem in P2P systems is that home broadband connections have quite limited upload bandwidth available. To combat this the BT protocol employs a few mechanisms to mitigate for this problem. Once a machine has received a few initial blocks, it can then utilize its upload bandwidth to pass those blocks to other machines. Another aspect is that transfers are limited to 4 open sessions, and the number of files hosted is constrained to the few the user has made available; by focusing on a small number of activities, the control overhead over the network is reduced, allowing for more data to be transferred. By reducing the number of uploads per machine, and having considerably more machines uploading content, the time required by a client for downloading a large file is considerably less than traditional means.

2.4. *Overcoming the Free-Rider Problem*

In typical P2P systems, the performance of the overall system declines when users act in their own self-interest by attempting to obtain service while not contributing anything to the system. This is an old problem, as seen from the field of economics, and BitTorrent was designed with an approach of local optimization by allocating the most upload bandwidth to peers which are supplying the greatest amount of content.^x As a result of the design of the current BitTorrent Protocol, free-riders can only get 20% of the total system bandwidth.^{xi}

2.5. *High Content Integrity*

Some P2P systems (predominantly post-Napster 2nd generation ones), such as Kazaa which employed highly decentralized models, suffer from poor content integrity. Kazaa, which was a dominant standard prior to BitTorrent had cases of up to 70% of corrupted or falsified content, which would "pollute" the network.^{xii} The user frustration and gross waste of bandwidth arising from this state of affairs was recognized to be a critical concern, and so BitTorrent employs two mechanisms to combat it. First, to be accepted on a file tracking server, typically new content must be accepted by moderators who manually inspect it. To reduce the overload for moderators, regular suppliers of content can be promoted to unmoderated submitters after gaining enough trust.^{xiii} Once the original content is considered trusted, downloaded chunks are evaluated against a validation checksum present in the .torrent file supplied from the server.^{xiv}

3. Media Industry responses to the Protocol

Although originally designed for the legal distribution of large files such as open source software and movie trailers, currently, the top users of BitTorrent are none other than the many television and movie pirates present on the internet. Adoption of the protocol by users continues to grow. However, such technology has commercial potential if it can be harnessed; according to David Poltrack, Executive Vice President at CBS, "In our research with consumers, content-on-demand is the killer app. They like the idea of paying only for what they watch."^{xv} Based on estimates in the TV industry, the networks have 10 years to respond, before illegal downloading of television content becomes widespread.^{xvi} Currently the television industry seems to be interested in the potential of this protocol, as their revenues are derived from advertising which can still be employed in internet based variations, rather than consumer supported content sales. Based on reports from January 26, 2005 almost 10% of traffic on the Internet 2 academic network was carried on the BitTorrent protocol.^{xvii} A recent report announced that due to delays by broadcasters in airing new content, TV program pirating in Australia is rampant, accounting for 15.6% of all torrent traffic.^{xviii} One may observe that the "wait and see" attitude of the TV industry may be with a view of allowing this medium develop further, before introducing ventures in it.

The movie industry, which derives revenues from the sale of content to consumers is currently not at all fond of this technology. Like the music industry was affected by the introduction of Napster, the movie industry has to contend with the fact that user driven content exchange is incompatible with their current business model. Given the need for a timely response, the Motion Picture Association of America is responding to piracy employing BitTorrent and other technologies by announcing a campaign to slow the the adoption of the technology.^{xix} In December, 2004 the MPAA filed over 200 lawsuits against individuals operating servers managing the flow of pirated content on the internet. As an initial legal tactic to prevent the spread adoption of this technology, the lawsuits are targeting the cataloguing servers, which "provide users with lists of movies, songs and television shows that can be swapped with other users".^{xx} Like MP3 and similar audio technologies, which have only started to transform the music industry after over a decade of existence, BitTorrent and similar platforms will likely require quite a few years to gain an accepted role within the movie industry.

4. The future of BitTorrent

The widespread dissemination of this technology continues to spread like wildfire. Beyond future growth, the reader may be wondering what the future holds for this protocol. Some major considerations are as follows:

4.1. RSS and BitTorrent Integration

BitTorrent provides a high performance mechanism, particularly for heavily trafficked data, making an eminent match to the RSS protocol for disseminating new content.^{xxi} RSS is an XML based language originally developed by Netscape for managing headlines on news portals.^{xxii} It has grown to serve as a popular protocol for managing the syndication of discrete media content. With client applications, RSS allows users to subscribe to feeds of content which can be automatically downloaded and delivered to the

user. The most obvious applications of a combined BT with RSS are as a distribution mechanism for rich media such as television and radio shows. This can allow new media providers to provide high quality feeds, at far lower bandwidth costs than have traditionally been thought possible for such content.^{xxiii} For some on the internet, this technology is seen as the next logical step in content delivery, and may prove to be an effective way for small scale new-media firms to enter the market. For the TV industry, this may be the “critical app” to be able to offer content on-demand.

4.2. *Ethical and Legal Issues*

Unlike past P2P protocols, BitTorrent has not been designed exclusively for illicit uses. This has allowed the protocol to gain greater acceptance among corporate users, however the adoption of BT for applications for which the law is not clear is definitely a major issue at the present. As the existing media firms encounter growing revenue losses from this technology, litigation against operators of servers and heavy downloaders can be expected in the coming years. With this litigation, legal precedents in various jurisdictions can be expected, and the bounds of what is permissible with aspects such as fair use will become clearer. In the meantime, the reader is advised to practice caution and to educate themselves of the legality for their use of the protocol.

5. Conclusions

News of increased adoption of this revolutionary protocol continues to emerge, and this seems to be a rich area for continued research in computer networking, with significant potential benefits to society. Only time will tell if this protocol will be able to survive the challenges it will inevitably face in the legal system. If BitTorrent and similar protocols do find a legal place in the internet, it will certainly help pave the way for distributing rich media on the internet, and provide the impetus for new innovations.

Endnotes

- i Pasick, Adam. "LIVEWIRE – File-sharing network thrives beneath the radar" from <http://in.tech.yahoo.com/041103/137/2ho4i.html> (Accessed March 12, 2005).
- ii "File Sharing" http://en.wikipedia.org/wiki/File_sharing (Accessed March 12, 2005).
- iii "Cachelogic: Understanding Peer-to-peer history" <http://www.cachelogic.com/p2p/p2phistory.php#> (Accessed March 26, 2005).
- iv Ibid.
- v "BitTorrent Version History" http://www.bittorrent.com/bittorrent_versions.html (Accessed March 26, 2005).
- vi "BitTorrent" <http://en.wikipedia.org/wiki/BitTorrent> (Accessed March 12, 2005).
- vii Bangeman, Eric. "BitTorrent moves further into the Mainstream" <http://arstechnica.com/news.ars/post/20040317-3535.html> (Accessed March 27, 2005).
- viii "BitTorrent" <http://en.wikipedia.org/wiki/BitTorrent> (Accessed March 12, 2005).
- ix "BitTorrent" <http://en.wikipedia.org/wiki/BitTorrent> (Accessed March 12, 2005).
- x Cohen, Bram. "Incentives Build Robustness in BitTorrent" <http://www.bittorrent.com/bittorrentecon.pdf> (Accessed March 16, 2005).
- xi Qui, Dongyu and Srikant, R. "Modelling and performance analysis of BitTorrent-like peer-to-peer networks" <http://comm.csl.uiuc.edu/~srikant/Papers/sigcomm04.pdf> (Accessed March 25, 2005).
- xii Pouwelse, Johan. "The BitTorrent P2P file Sharing System" http://www.theregister.co.uk/2004/12/18/bittorrent_measurements_analysis/ (Accessed March 25, 2005).
- xiii Ibid.
- xiv "BitTorrent – Protocol" <http://www.bittorrent.com/protocol.html> (Accessed March 26, 2005).
- xv Thompson, Clive. "The BitTorrent Effect" <http://www.wired.com/wired/archive/13.01/bittorrent.html?pg=3> (Accessed March 30, 2005).
- xvi Ibid.
- xvii Bangeman, Eric. "BitTorrent moves further into the Mainstream" <http://arstechnica.com/news.ars/post/20040317-3535.html> (Accessed March 27, 2005).
- xviii "Impatient TV viewers turn to BitTorrent" http://news.com.com/Impatient+TV+viewers+turn+to+BitTorrent/2100-1026_3-5653362.html (Accessed April 5, 2005).
- xix Crawford, Krysten. "Hollywood turns up the heat on illegal downloading" <http://money.cnn.com/2004/12/14/news/fortune500/piracy/> (Accessed March 30, 2005).
- xx Ibid.
- xxi Grumet, Andrew. "Experimenting with BitTorrent and RSS 2.0" <http://blogs.law.harvard.edu/tech/bitTorrent> (Accessed March 30, 2005).
- xxii Pilgrim, Mark. "What is RSS" <http://www.xml.com/pub/a/2002/12/18/dive-into-xml.html> (Accessed March 30, 2005).
- xxiii Havens, Andrew. "BitTorrent: The 'long tail' and legal marketing" http://legalmarketing.typepad.com/blog/2005/01/bittorrent_the_.html (Accessed March 30, 2005).