4

The Implications of Video Peer-to-Peer on Network Usage

Kevin Werbach

University of Pennsylvania

Rewritten by machine and new technology,
and now I understand the problems you can see.
Oh-a oh
I met your children
Oh-a oh
What did you tell them?
Video killed the radio star.[1]

Introduction

The tsunami is upon us. A rising tide of video Peer-to-Peer[2] activity is already beginning to affect data networks. And video Peer-to-Peer traffic will inexorably grow in the years ahead. Video Peer-to-Peer will expand beyond unauthorized sharing of commercial pre-recorded content, becoming a significant driver of broadband usage and potentially creating new revenue streams. Meanwhile, because of its sheer bulk and technical characteristics, video Peer-to-Peer traffic will place significant strains on broadband networks. Thus, video Peer-to-Peer will influence both the outputs and the inputs of the Internet of the future.

The network usage implications of video Peer-to-Peer are not widely appreciated. To date, most of the attention devoted to Peer-to-Peer has focused on the content of the files being transferred.[3] The unauthorized dissemination of copyrighted material through Peer-to-Peer systems has considerable implications for users, artists, network operators, technology developers, device manufacturers, investors, and content distributors. Yet

E.M. Noam and L.M. Pupillo (eds.), *Peer-to-Peer Video*, doi: 10.1007/978-0-387-76450-4_4,
© Springer Science + Business Media, LLC 2008

Peer-to-Peer, and especially video Peer-to-Peer, would have significant impacts even if none of the files involved were subject to intellectual property protection. And, though it is quite early in the development of the market, there are indications that video Peer-to-Peer will be used more actively than audio Peer-to-Peer for sharing content not subject to copyright limitations.

Focusing on commercial content provides only a partial view of the economics of video Peer-to-Peer. In the eyes of lawyers and entertainment industry executives, there is a vast difference between my home movies, which I freely make available to any prospective viewers, and the latest *Star Wars* prequel. One is a potential source of several hundred million dollars in revenue; the other is, at best, a reason to spend a few hundred dollars on a video camera. From a network engineering perspective, though, both are simply large amalgamations of data. And to me personally, both are valuable, though in different ways. Understanding likely usage patterns and network impacts is critical for any realistic assessment of the consequences of video Peer-to-Peer.

Bits and Bytes

All information transferred across digital networks such as the Internet is ultimately fungible. A bit is a bit. Thus, a large file requires more network capacity, typically expressed in terms of bandwidth, than a small file, regardless of what that file contains. However, data in motion is not equivalent to data at rest. In other words, there are differences in network impacts based on how files are used. A popular file consumes more network resources, because it is transferred more times, even though it takes up the same number of bytes on a hard drive as an unpopular file. This creates a rough parallel with the content-oriented perspective. It is quite likely that more viewers would download a pre-released print of the *Star Wars* movie than the Werbach family's summer vacation highlights.

On the other hand, what matters from a network perspective is the aggregate impact of all user activity. There are many more home movies shot each year than Hollywood feature films. And some non-commercial content achieves significant popularity – witness the spread of the Paris Hilton sex video, or the original installment of what became the animated show *South Park*, mailed out as a video Christmas card by a Hollywood executive. Before the Internet came along, distributing video content was beyond the means of individuals, and small commercial operations had limited reach. Now, when any content can become instantly available to a global audience of more, a bright-line distinction between commercial and non-commercial video is more difficult to draw. Each is likely to have significant economic and network capacity consequences as it becomes more feasible to distribute them across the Internet.

Technology adds another layer of complexity to the economic analysis. A television episode is subject to the same intellectual property law regardless of how it is delivered across the network. Yet content creates very different network loads depending on whether it is streamed, downloaded explicitly, or downloaded automatically in the background over a period of time. Video captured in real-time from a wireless device will create different choke demands than fixed video files such as recorded television programs. A video file pulled together from fragments stored on many different users' hard drives will have different network impacts than one delivered in a single piece from a central server on demand.

The network impacts of video Peer-to-Peer activity are significant for several reasons. Network capacity has a cost. If video Peer-to-Peer imposes costs on network operators, they will have incentives to limit it, especially if they do not receive a commensurate benefit from the activity. How the network infrastructure delivers Peer-to-Peer content also influences user behavior. To take one example, an individual may be inclined to use – or pay for – a Peer-to-Peer service that quickly and reliably delivers content, but not a service that is slow and prone to connection errors. Some applications, and some users, have more tolerance for poor quality than others.

Those tolerances are loosely connected to the type of content involved, though the relationship is synergistic. A Peer-to-Peer service used to swap Hollywood movies will have a different profile than one used to aggregate assorted content streams into a personalized television network, or one used to share running video diaries of life experiences with family and friends. What the network supports will influence which of these applications are more popular, and application popularity will influence network design. The legal and economic battles now being fought over Peer-to-Peer activity are exogenous factors which will have a significant bearing on this dynamic. Impediments to certain types of Peer-to-Peer activity will make some applications impractical or impossible, whether or not those applications are the actual target.

None of these issues is necessarily specific to video. However, because of video's unique characteristics, it will create different challenges than other media types. The rise of video Peer-to-Peer promises much greater network impacts than the largely music-dominated traffic heretofore. And video as a medium lends itself to certain applications and usage patterns that have heretofore not been widely adopted.

The remainder of this paper sketches the likely consequences of video Peer-to-Peer on networks and usage patterns. The sections focus on: reviewing of data on the level of video Peer-to-Peer activity, concluding that it already represents a substantial share of Internet traffic, and that share is likely to grow; Peer-to-Peer technology for transferring video, in particular the new class of swarming file-transfer software typified by BitTorrent; four major categories of video Peer-to-Peer applications that are likely to become significant in the coming years; evaluation of how the proliferation of these applications will affect different segments of the Internet; and consideration of how network operators and other service providers will respond to the capacity challenges and revenue opportunities of video Peer-to-Peer.

Rise of the Videonet

To a first approximation, video Peer-to-Peer file transfers *are* the Internet.

Comprehensive and reliable data on video Peer-to-Peer are just starting to become available. The distributed nature of Peer-to-Peer services, the short time that video has been an appreciable component of Peer-to-Peer activity, and the difficulty in segmenting among types of Peer-to-Peer traffic all make it hard to obtain accurate measurements.[4] Nonetheless, what we do know is striking: virtually everywhere measurements are done, Peer-to-Peer transfers of large files such as video are the single biggest component of network utilization.

Peer-to-Peer file-sharing in general represents a substantial proportion of Internet traffic. A study by network monitoring appliance vendor CacheLogic, released in July 2004, found that Peer-to-Peer is the largest consumer of data on Internet service provider (ISP) networks, and is still growing.[5] CacheLogic is one of a new breed of vendors whose equipment is capable of deep packet inspection at the application layer, allowing it to monitor and differentiate Peer-to-Peer applications more finely than previously possible.[6] In a single 30-day period, one CacheLogic appliance tracked Peer-to-Peer accesses from 3.5 million unique IP addresses. CacheLogic estimates the worldwide simultaneous Peer-to-Peer user base at 10 million – over 10% of all broadband accounts. And those users are sharing over 10,000,000,000 megabytes (10 petabytes) of data.[7]

Despite legal action by the record industry against individual users as well as distributors of Peer-to-Peer file-sharing software, Peer-to-Peer traffic continues to represent an increasing share of Internet traffic.[8] Peer-to-Peer file transfers represent an absolute majority of traffic on many networks, as high as 80% in some cases.[9] The numbers are especially high for broadband access networks.[10] Email and the World Wide Web, the "killer apps" of the Net, are small by comparison.

Video in particular is a significant and growing element of Peer-to-Peer traffic, especially since 2003. A study of BitTorrent, a Peer-to-Peer system popular for video, found between 237,500 and 576,500 daily BitTorrent transfers in progress between December 2003 and January 2004, of which roughly 100,000–150,000 per day were movies.[11] These numbers represent a small percentage of the Peer-to-Peer file-sharing user base, most of which is still engaged in trading music files.[12] However, video Peer-to-Peer transfers are rapidly eclipsing Peer-to-Peer distribution of music files in bandwidth usage. The CacheLogic study concluded that BitTorrent alone consumed more than one-third of all Internet bandwidth worldwide.[13]

The reason is simple: video files are enormous. As Table 4.1 shows, a feature motion picture, encoded using common compression mechanisms, may be a thousand times the size of a song, or even larger.[14]

Table 4.1 Relative sizes of different file types

File type	Approximate size (kilobytes)
Email message	5–100
Web page	25–500
Music audio file (MP3)	2,000–10,000
Music video	50,000–200,000
Feature film	500,000–4,000,000

Thus, video files have an impact on network usage that is grossly disproportionate to the number of users sending or receiving them. Ten thousand people viewing a 100-kilobyte (kB) Web page would move the same number of bits through the network as a single person downloading a 1-gigabyte (GB) movie. In fact, this comparison may actually understate the differences between the two traffic types, for reasons discussed below.[15]

The data bear out these predictions. One study of traffic on the popular KaZaA Peer-to-Peer file-sharing service found that while 91% of requests were for objects smaller than 10 megabytes (MB), a majority of the bytes transferred (65%) were from objects greater than 100 MB, primarily video files.[16] According to the Organization for Economic Cooperation and Development (OECD), video and other transfers made up a majority of Peer-to-Peer traffic in OECD countries in 2003, for the first time exceeding music.[17] Another study by economists Peter Lyman and Hal Varian concluded that video files represented 59% of total file-sharing traffic in 2003, compared to 33% for audio files.[18]

If these numbers sound shocking, it may be because the United States is something of a laggard in usage of video Peer-to-Peer. Video distribution on Peer-to-Peer networks is significantly more common in other parts of the world. One reason is that US Broadband penetration lags many other countries in Europe and Asia.[19] Countries such as Korea, where users have significantly more bandwidth available to them at significantly lower prices, unsurprisingly have higher rates of usage of Peer-to-Peer networks for video.[20]

Even today's relatively small level of video Peer-to-Peer activity is influencing overall network demand. In contrast to music, therefore, network operators may feel the economic impacts of video Peer-to-Peer distribution before commercial content owners do. Over the coming decade, video Peer-to-Peer usage will expand greatly, for two reasons: enhanced technology and new applications.

Video Peer-to-Peer Technology

Peer-to-Peer file sharing burst on the scene with the 1999 release of Napster, a software application written by college student Shawn Fanning. Fanning was an inexperienced programmer, and it showed. Though millions downloaded Napster and it quickly created an earthquake in the music industry, the application itself was unsophisticated. It wasn't even truly Peer-to-Peer. The files involved were transferred directly between users, but those transfers were coordinated through a central directory maintained by Napster.[21] That central coordination point was Napster's undoing on both a legal and practical level[22] when the company was sued by the record industry.

Peer-to-Peer technology has come a long way since Napster. The functions of today's Peer-to-Peer file-sharing applications may be similar, but the mechanisms differ in important ways. How Peer-to-Peer services operate has important consequences for their impacts on networks and economics, especially in the case of video.

Peer-to-Peer Techniques

There are two basic processes for acquiring rich media content[23] over the Internet: file transfer and streaming. File transfer involves delivering the entire file over the network to the user's computer. Once the transfer is completed, the file can be played, copied, or transferred elsewhere.[24] In streaming, the user receives a small initial segment of the file, which is stored locally in a buffer file. The file begins playing from the buffer, while the next segment is transferred over the network. In this manner, the user has the experience of hearing or viewing the entire file almost immediately. However, the file is only transferred as fast as it can be played back, and only the segment of the file being played is stored on the user's computer, making it impossible to copy or replay the file locally.[25]

From a pure network standpoint, file transfer is more efficient. Because the file is not actually being played across the unpredictable network link, there is more tolerance for delay and lost packets. Pieces of files can be transferred out of order, up to the speed of the network connection, rather than being limited to the playback speed of the file. These capabilities are particularly well-suited to a Peer-to-Peer architecture, in which information flows between heterogeneous and intermittently connected nodes at the edge of the network rather than centrally managed servers. The major Peer-to-Peer applications, including Napster, Gnutella, FastTrack, and BitTorrent, are all file-transfer systems.[26]

Beyond that, the Peer-to-Peer services differ in their technical architecture. None of the currently-popular systems employ Napster's central directory. Some use dynamically-created "supernodes," which turn users with high-quality connections into temporary directory nodes for other users. Other systems, such as Gnutella, relay requests from one node to another, until the request finds a directory including the desired file. BitTorrent, described in greater detail below, further distributes the directory function through the use of multiple "trackers," which keep track of pieces of files.

Every major Peer-to-Peer system has its strong and weak points. Some scale well to large numbers of simultaneous users (popularity); some compensate well for the inherent unreliability of Peer-to-Peer network nodes (availability); some offer higher file transfer speeds (download performance); some allow files to remain available for long periods of time (content lifetime); some offer content (such as movies) soon release (content injection time); and some are resistant to accidental or deliberate uploading of files with incorrect names or corrupted contents (pollution level). A recent study compared five leading Peer-to-Peer systems along these axes, showing that no single factor accounts for popularity.

The FastTrack platform, which is the basis of KaZaA and several other Peer-to-Peer software applications, has until recently been the most popular Peer-to-Peer service. eDonkey has enjoyed greater success in Europe and Asia than in the USA.[27] This is partially because eDonkey's queue-based scheduling and distributed downloading capabilities are useful for transferring video and other larger files. Such files are a larger proportion of the Peer-to-Peer mix outside the USA, where broadband speeds and penetration are greater.

As already noted, the distinguishing characteristic of video files is their immense size. Downloading performance is the biggest hurdle to obtaining such files across the network. Though video files have been available on Peer-to-Peer systems since the beginning, they typically represented a very small fraction of available files and network traffic. Performance simply wasn't reliable enough to make transferring large video files such as movies worthwhile (Table 4.2).

Table 4.2 Characteristics of the five most popular Peer-to-Peer systems

Peer-to-Peer system	Strong points	Weak points
FastTrack (KaZaA)	Popularity, availability, content lifetime	Pollution level
Overnet (eDonkey)	Popularity, content lifetime	Download performance
BitTorrent	Popularity, download performance, content injection time, pollution level	Availability, content lifetime
DirectConnect	Download performance, content lifetime	Availability
Gnutella	Download performance	popularity, pollution level

From J.A. Pouwelse, P. Garbacki, D.H.J. Epema, and H.J. Sips, "A Measurement Study of the BitTorrent Peer-to-Peer File-Sharing System," at 7, preprint available at http://www.pds.ewi.tudelft.nl/reports/2004/PDS-2004-003/pdsreport.html

BitTorrent and Swarming

Even with a fast broadband link and a sophisticated Peer-to-Peer platform, downloading an entire movie or other long video clip in one fell swoop is difficult. Real-world transfer speeds on Peer-to-Peer services are significantly below the peak download speed of the broadband connection, generally in the range of 20–50 kilobytes per second. At those rates, a large video file may take many hours, even days, to transfer completely. The chances that the originating node will be online and reachable that entire time are small. And if a transfer is interrupted in the middle, it may be impossible to pick it up again.

The solution is what is sometimes called swarming technology.[28] Swarming breaks up large files into many small pieces.[29] When a user wishes to download the file, rather than pulling it all from a single source, the system locates and downloads the pieces from many different locations in parallel. When more than one user attempts to download a file at the same time, the downloaders simultaneously upload pieces of the file to each other.[30] Thus, instead of choking a node hosting a popular video file, a swarming system automatically distributes the file transfer load.

Swarming is a key element of BitTorrent, which was first released in mid-2002. This explains why it is used so heavily for video and large software files.[31] The BitTorrent architecture involves three components: .torrent files, trackers, and user nodes. The .torrent files, which are accessed through ordinary Web servers, provide basic information on the file to be obtained, but not the actual content itself. They also include pointers to trackers, which are a form of directory server. A tracker maintains information about who has what pieces of the relevant file. Using the tracker, the BitTorrent software begins downloading and uploading pieces of the file among other nodes. At least one node must function as a seed, which means it has a complete copy of the file, to verify integrity. The other nodes are in the process of downloading, and may only have a small portion at any given time.

Unlike other popular Peer-to-Peer services, BitTorrent does not directly provide search functionality. In order to obtain a BitTorrent file, a user must locate it through other means than the BitTorrent software itself. The most common way of doing so is through Suprnova.org, an independent website which maintains a moderated list of new BitTorrent files.

BitTorrent addresses a significant limitation of other Peer-to-Peer systems known as free riding. Free riding is a significant problem with some Peer-to-Peer systems, notably Gnutella; one study in 2000 found that 70% of users shared no files at all, and merely downloaded.[32] BitTorrent users do not have the option not to share files. The system also incorporates a bartering mechanism.[33] In other words, users who upload are rewarded with the ability to download more rapidly, while those who do not upload are punished with limited download capacity. Bram Cohen, the developer of BitTorrent, calls this “leech resistance.”[34]

The BitTorrent software is open source. This allows developers to build their own client software based on the protocol, and several have. Developers are incorporating BitTorrent into different kinds of applications, including video syndication feeds. Commercial software developers are also evaluating BitTorrent as a technology platform for video distribution.

If widespread video Peer-to-Peer distribution does continue to take off, especially for applications other than unauthorized sharing of copyrighted material, BitTorrent or BitTorrent-like swarming technology will likely be part of any popular application. Faster broadband connections, especially in the upstream direction, will make it easier to transfer entire video files, but without a bartering system to limit traffic imbalances, any system will have difficulty scaling. The market appears to be moving in this direction. eDonkey, already one of the most popular Peer-to-Peer file-sharing platforms, especially outside the USA, has added BitTorrent-like swarming capability called Horde in recent versions.

It is important to recognize that BitTorrent and the Horde-enabled eDonkey are recent arrivals. Though released in 2002, BitTorrent was originally used primarily for software distribution. In essence, it has been less than a year since well-designed tools became available for effective video Peer-to-Peer transfers. And those tools are far from perfected. BitTorrent lacks a straightforward user interface and integrated search functionality, while eDonkey can't yet match the performance of BitTorrent's swarming implementation. As software to distribute Peer-to-Peer video reliably and simply becomes more widely available, video Peer-to-Peer traffic will grow. Any predictions at this stage of the market must necessarily be speculative. Given the growth of video Peer-to-Peer traffic in the past year, though the question is when, not if, today's usage level will expand.

Video Peer-to-Peer Usage Scenarios

The other half of the video Peer-to-Peer equation is the demand side. Good tools can only go so far without killer apps to drive usage. In the case of video Peer-to-Peer, there has been relatively little consideration of novel usage scenarios. Though the economic and legal discussion around video Peer-to-Peer centers on trading of commercial video content such as movies and television programs, this is not likely to be the only substantial form of video Peer-to-Peer activity.

As with the technology and network impacts, the precise timing and contours of popular usage scenarios is hard to predict. Video Peer-to-Peer is still in the early stages of development. Many entrepreneurs and even established companies have failed miserably in predicting usage patterns for new forms of media, which video Peer-to-Peer represents. So any predictions should be taken with a grain of sale. Nonetheless, it is possible to sketch out some likely developments, given the capabilities that will soon be available and known user demand.

Video as a medium is different from audio. The richness and emotional impact of video lends itself to different experiences. Just compare the number of people who shoot home movies or otherwise use video recorders to the number who make personal audio tapes. A song can be played in the background, while engaged in other activities, while most video content requires full attention. And, although movies as a form of pre-recorded commercial content are somewhat analogous to record albums, a great deal of commercial video content, such as television news and sporting events, is live. All these examples suggest that, as video becomes a more prominent part of the Peer-to-Peer world, different applications may predominate than the song-sharing that dominates audio Peer-to-Peer usage.

Four primary classes of application are likely to drive utilization of video Peer-to-Peer: sharing of pre-recorded video files; distribution of personal video among families and friends; dissemination of "do it yourself" entertainment and news content; and monitoring and sensor applications. The first, which dominates video Peer-to-Peer today, is a fixture of the broadband wired Internet. The second and fourth will, to a great extent, grow out of the wireless Internet, including both local-area WiFi wireless "hotspots" and wide-area cellular data networks.

Video File Sharing

Peer-to-Peer platforms are used to distribute copyrighted commercial video files such as movies, music videos, and television episodes. The pattern here is similar to the sharing of music files, though so far the level of activity has been smaller. Because video files are much larger than songs, they take much longer to download. Thus, in the same amount of time, a user can obtain fewer video than audio files. Moreover, large video files such as movies are still difficult to download reliably, though with faster broadband pipes and systems such as BitTorrent, downloading movies is becoming more feasible.[35]

Still, there is evidence to suggest that a reasonable amount video file-sharing is taking place. A Jupiter Research survey found that 15% of European Peer-to-Peer users download at least one movie per month.[36] A study by the Motion Picture Association of America found that one in four American Internet users interviewed said they had downloaded movies via the Internet, and 60% of Korean users said they had.[37] These numbers may sound frightening for the movie industry, but keep in mind that one movie download per month is likely to be small relative to both the equivalent music-sharing activity and the number of films the user views legitimately during that period. Moreover, evidence from music sharing has been inconsistent on the critical question of whether Peer-to-Peer file sharing leads to reduced record sales.[38]

If the music experience is any guide, video Peer-to-Peer file sharing is likely to become increasingly popular as time goes on. Efforts by the record industry to limit Peer-to-Peer file sharing through legal actions, including suits against end-user uploaders and software developers, have not significantly curtailed Peer-to-Peer activity.[39] The limiting factors on video Peer-to-Peer trading – principally, the size of the files – will diminish over time. Broadband access providers will offer faster transmission speeds, and software based on video-friendly technology such as BitTorrent will becoming increasingly widely distributed and easy to use. On the other hand, the appeal of downloading commercial video content will be limited until that content can more easily be viewed on a television screen, as opposed to a computer monitor.

Because it so closely tracks the audio Peer-to-Peer experience, file sharing is the least interesting application of video Peer-to-Peer. Whether content owners are successfully able to thwart unauthorized video file trading through lawsuits or licensed downloading services is an important economic and social question, but one being actively considered in relation to audio file sharing. The novel question is whether video Peer-to-Peer will produce something new and different. Even within the realm of file sharing, there are reasons to believe it will.

First off, not all video file sharing involves copyright violations. BitTorrent is widely used for sharing of high-quality audio concert recordings.[40] Many bands have given permission for fans to record and swap concert bootlegs. A network known as etree.org is dedicated to using BitTorrent to share such legal concert recordings, using lossless compression mechanisms that provide higher-quality audio than MP3 or similar formats.

In addition, some content owners are exploring the possibility of using Peer-to-Peer systems to distribute their content, subject to digital rights management restrictions to prevent unauthorized distribution. The BBC has openly discussed the possibility of using Peer-to-Peer platforms for distributing rich-media program guides.[41] Independent film developers may want to release their movies onto Peer-to-Peer networks to build demand, much as some independent bands have with their songs.[42] Peer-to-Peer distribution makes even more sense for small content producers of video than for music, because the content creators can distribute trailers or segments of their works rather than the whole film.

The same platforms will be used for large non-video files, principally software. BitTorrent is used for unauthorized distribution of games and other software. However, as with video, not all the activity involves copyright violations. Some distributors of the Linux operating system now use BitTorrent as a regular method for making new software versions available. Linux distributions can be several hundred megabytes in size, and when new versions are released, high download requests can clog originating servers. By using Peer-to-Peer services, the distributors save on hosting and bandwidth costs to get the software out to users.

Life Sharing

Video is a window on personal experiences. Sharing with someone what you *see* is so much richer than sharing what you hear that it forms a qualitatively different experience. A video clip of you or what you see around you, whether live or transferred asynchronously is a piece of your life. As such, it is likely to be of some interest to your friends and family, but less so to the general public. Personal video communications are not the same as phone calls, as four decades of failed videophone efforts show. Nonetheless, there may well be markets for video Peer-to-Peer life sharing. Content that is of interest to a small audience becomes significant in the aggregate when there are enough originators.

For the first time, many million people now have digital video camcorders, with unit sales exceeding 10 million in 2004.[43] This eliminates one of the hurdles to Peer-to-Peer life-sharing: the capability of encoding video. Because video life sharing is largely a personal experience, it does not require professional quality production values, which are expected for most of the music files shared on Peer-to-Peer networks. All that is necessary is a reasonable-quality recording, which is now widely affordable.

The second element is to get those recordings onto the Internet. Digital video editing software is now inexpensive and easy to use, but that still requires an additional step to transfer the content from a camera to a storage device such as a hard drive or rewriteable DVD drive. Two developments are likely to change the equation: networked video recorders and video cameraphones. A variety of short-range unlicensed wireless technologies, including WiFi and WiMedia, could be employed to transfer data directly from a video camera to a home network server, from which it could be uploaded over the Internet. Or, with a webcam or other built-in video capture device, an existing PC, media, or gaming device becomes a networked video origination point.

The bigger impact may come from the continued proliferation of camera-enabled mobile phones. The mobile phone is the world's most popular personal computing platform, with over 1.5 billion users worldwide. Annual handset sales worldwide exceed 500 million units annually. Handset vendors are now adding applications, data networking, and cameras to phones, turning single function mobile phones into digital smartphones. An estimated 200 million cameraphones will be sold in 2004, compared to roughly 50 million digital cameras and 60 million film cameras (excluding single-use models).[44] That number is expected to grow to over 600 million a year by 2008.[45] Thus, in less than a decade, there will be over one billion users carrying networked digital cameras around with them at all times.

Since video is, in essence, simply a series of still photographs, the basic hardware for a still cameraphone can also handle video. Many existing cameraphones offer video modes, and as resolutions, storage, and wireless network capacity increase, a greater share of cameraphones are likely to offer this capability. The resolution and other capabilities of cameraphones are a function of hardware performance and miniaturization. The information technology industry has developed fabulous expertise in applying standardized processes to improve performance and reduce costs steadily over time. Higher unit sales mean lower per-unit costs. Whatever the price-performance of a video cameraphone today, it is therefore a safe bet that the standard device a year from now will be better, cheaper, or both. High-resolution video cameraphones selling for less than $200, or much less when tied to service plans, are inevitable by the end of the decade.

For the first time, therefore, a substantial and growing audience exists with all the fundamental capabilities necessary for wireless life sharing: video encoding capability, direct network connectivity from the device, and a broadband data pipe. Exactly when and how quickly personal video sharing takes off will depend on other factors, including the pricing and ease-of-use that service providers offer.

What is clear is that people love to share their personal experiences, either with social networks of family and friends, or with anyone who cares to view them. The proliferation of personal World Wide Web pages is testimony to this fact. So are the sales of still and video cameras. As Andrew Odlyzko has shown, "communications" applications involving user-generated content have consistently outpaced services delivering professionally-created content to a passive audience.[46] From giving grandparents who live far away a glimpse of their grandchildren to keeping a running diary of your European trip to showing your doctor what that bump on your arm looks like to immortalizing the time you ran into a celebrity, there are countless situations in which people will want to share their experiences.

Much of this content will only be of interest to a small circle of recipients. That makes it ideal for a distributed Peer-to-Peer environment, with the video files flowing directly between individuals. Making this experience seamless will require new software tools that build on social networking services such as Friendster, Orkut, and LinkedIn. Already, a startup called Ludicorp has developed a social networking service, Flickr, designed around digital photo sharing, and sold it to yahoo! Similar services for video, tied together with other networking features, will be the glue for video Peer-to-Peer life sharing.

Some life-sharing is not, strictly speaking, Peer-to-Peer. Video streamed live from a camera or cameraphone to a server, and then viewed, involves a client-server broadcast model. Apple's iChat DV integrates a webcam and software for real-time video conferencing. Services such as SightSpeed, Viditel, and Convoq support reasonable-quality videoconferencing for both consumer and business applications over ordinary broadband connections. Instant messaging services, such as AOL's AIM and ICQ, Yahoo! Messenger, and Microsoft's Windows Messenger, are adding video chat functionality.[47] As gaming consoles such as the Microsoft Xbox and Sony PlayStation morph into multi-function digital hubs with integrated broadband and media capabilities, they are also likely to serve as real-time video communications endpoints.

Distributed Media

Where life-sharing is episodic and usually directed to a narrow social circle, distributed media involves aggregation of content for and by a wider audience. Until now, media has been centralized. The high fixed costs of producing and distributing content gave an advantage to large entities, such as television broadcasters and cable system operators, who packaged the programming for viewers. Though the media has become more fragmented and somewhat more interactive in recent years, it still follows the same basic template.

Now, thanks to the proliferation of devices and networks described in the preceding section, the ability to create high-quality content is within the reach of a far greater number of people. Drazen Pantic, who created the Internet department of pioneering Radio B92 in Serbia, explained the potential for distributed media in a recent manifesto:

> "Today, everyone has access to the latest high quality consumer electronic devices. Every cell phone has the ability to capture images, even movies. Once people begin to use these devices to record the significant events of their lives, there is no way to prevent them from slipping cameras into any location. When sensitive material is captured in digital form, it takes on a life of its own. Circulating across the Internet, it becomes a fact in itself."[48]

Imagine, then, the opening ceremonies of the 2008 Olympic Games in Beijing.[49] Tens of thousands of spectators will crowd into the Olympic stadium, and it is a safe bet that many of them will have video-enabled mobile phones.[50] What they see and hear will be available instantly to many millions of potential viewers, who are no longer limited to the official broadcasts of the games. And this will be the case for unexpected events as well. For any major breaking news story, there will be dozens if not hundreds of potential journalists on the ground, if they care to take on that role.

Another nascent form of video-based distributed media is video-blogging. Blogs, or weblogs, are online personal commentary or diary sites, organized in a series of date-stamped "posts." The vast majority of blogs today are text-based Web pages. However, there is a small but growing community of bloggers whose posts take the form of video. Videoblogs are not Peer-to-Peer today, in that the content is uploaded to a Web server and downloaded from that central location. However, as with all video content, it is possible that in the future the files will be transferred directly between end-user nodes.

The key element in the chain is syndication. Blogs have been the primary driver for the growth of syndication protocols such as RSS and Atom.[51] Using the simple syndication protocols, blog software automatically tags each post with standardized metadata: information such the subject matter, author, and time of posting. Software packages called aggregators can read that information and automatically pull in the latest posts from dozens or hundreds of blogs their users subscribe to. Experiments are underway to integrate syndication and video. Some involve linking RSS with BitTorrent, so that an aggregator would automatically download video content using BitTorrent's efficient Peer-to-Peer technology.

Syndication and its corollary aggregation are the glue that can turn a cornucopia of content around the network into something approaching a unified media experience. Users can subscribe to feeds, either representing content creators they find interesting or new information automatically retrieved that meets certain criteria. Once the aggregator pulls in that latest material, it can be viewed or organized in many different ways, because it is already marked with standard metadata. The full value chain for this new media form hasn't been formed, but companies are working on many of the piece-parts.

The end-product will be a sort of personal TV on steroids. Users will have the ability to select from a vast array of programming that meets their needs, which they can view whenever they choose. Companies such as Tivo[52] and Akimbo are taking baby steps in this direction with their digital video recorders. Akimbo's service, which just launched, allows users to select from a huge library of programming on the Internet and download it directly to a hard drive on the Akimbo box, from which it can be played on a television.

Wireless video Peer-to-Peer distribution may also converge to some degree with interactive and on-demand television. Many vendors and service providers are working on TV-over-broadband offerings that would include far more program choices and flexibility than existing cable and satellite systems. However, these are still centrally managed networks. The content is delivered to the user from a remote server, rather than Peer-to-Peer.

In the interim, video is making its way from traditional broadcast and cable networks to other kinds of devices. Services such as MovieLink offer downloadable movies directly to personal computers. Sprint's MobiTV streams broadcast video directly to mobile phones. And a small startup called Slingbox is building a device that bridges the gap between TV and Peer-to-Peer.[53] The $199 device, scheduled to go on sale before the end of 2004, plugs into a set-top box or digital video recorder. It converts the program signal to digital form, resizes it for delivery to a wireless device, compresses it, and sends it out over the Internet. The user can then watch the program on a mobile phone or handheld computer with a suitably fast wireless connection.

Sling Media is a tiny startup, and it may well face legal obstacles from television and movie industry as it launches its product. Conceptually, though, the Slingbox shows the potential for a new mode of media distribution, one in which content is reflected Peer-to-Peer even if originally delivered through centralized systems. Just as Tivo and other digital video recorders are challenging the advertising-driven economics of television despite still relatively limited sales, the Slingbox model could have significant disruptive impact whether or not the company succeeds in building a business.

Monitoring and Sensors

The final category of video Peer-to-Peer usage may ultimately be the largest in terms of bits, even though much of the content will never actually be viewed. With networked video cameras becoming increasingly cheap and widely available, many opportunities for monitoring will become apparent. Sprint PCS now offers a service called EarthCam mobile, which allows users of certain handset models to view streaming video from any Internet-connected webcam.[54] Security and traffic monitoring are two obvious applications, but there will be many more. Thanks to the WiFi unlicensed wireless protocol, it is now easy to deploy networked video cameras even where wired Internet connections are unavailable.

Some video monitoring scenarios overlap with the life sharing. Networked video-capture devices, such as webcams, can be used for real-time monitoring of family-related activity. The classic example is the so-called "nanny cam," which lets a parent look in on a babysitter or day care center watching his or her children. Unlike the life-sharing applications described above, most of the video delivered in these scenarios will never be viewed. It need not be. With inexpensive networked video cameras widely deployed, the mere possibility than what they record will be of interest will be enough.

In time, more intelligent software will be developed for automatically categorizing, filtering, aggregating, and searching this mountain of video. Excerpts from real-time streams could then be transferred automatically over Peer-to-Peer connections, avoiding the wasted bandwidth of sending the entire stream across the network. Imagine an oil company that wishes to monitor the condition of its equipment deployed in remote locations around the globe. Or a medical service that monitors older people who wish to continue living at home.

As with the life-sharing applications, most content generated through monitoring devices will be of immediate interest to a small audience.[55] However, because the content is generated with little or no human effort, there will be a tremendous amount of it flowing through the network. A Peer-to-Peer architecture, which avoids central repositories on services within the network, would seem to be the only logical approach to take for such information.

Peer-to-Peer Impacts on Networks

What will all this video Peer-to-Peer activity mean for networks? To answer this question, it is useful to divide the network into three segments: private corporate or campus networks (intranets), access networks, and transport networks.

Intranets are feeling strain from Peer-to-Peer traffic because of their relatively limited capacity, especially through the edge gateways connecting it to the public Internet. University networks have been especially hard-hit, because students are heavy users of Peer-to-Peer file-sharing. Many university network administrators have implemented bandwidth caps, terms-of-service restrictions, or other limits to corral Peer-to-Peer file-sharing, because of the network capacity costs rather than objections to potential copyright violations. Video Peer-to-Peer activity will likely increase the strains on these networks. Campuses and businesses that build intranets do so to support the demands of their users; they are not primarily in the network access business.

Access networks also face significant issues from Peer-to-Peer traffic. Capacity on these networks is at a premium. Most broadband operators engineer their networks with significant contention rates, frequently ten to one or more. In other words, the capacity coming into their network is many times smaller than the total theoretical download speed they offer their users. This approach saves on unnecessary capacity investment. It is feasible because not all users access the network at the same time, much of the time they are online they are not actually requesting data, and the access providers consumer terms of service do not guarantee throughput rates.

Wireless Internet connections are particularly ill-suited to handle large volumes of video Peer-to-Peer traffic. Cellular data networks offer substantially lower speeds and reliability than wired broadband access networks, due to the difficulty of sharing capacity over the air. Even newer wide-area wireless data services only offer top speeds of about 200 kbps, and operators frequently enforce monthly caps on data transfers. Thus, although a video-enabled cameraphone is an ideal device for capturing and uploading personal video for Peer-to-Peer distribution, such an application is unlikely to spread widely in the current environment.

This may change as performance of wireless data networks improves, though a more likely path is through alternate distribution mechanisms that do not tax the wide-area wireless networks. Automatic synchronization of media files between the phone and server is one option.[56] Another is to offload the video content from the wide-area network onto a local-area wireless link such as a WiFi hotspot. Dual-mode WiFi/cellular handsets are now coming on the market, and major wireless operators such as Verizon Wireless and T-Mobile operate significant hotspot networks. With wide-area network capacity at a premium, the operators would have incentives, either through service restrictions or pricing, to encourage Peer-to-Peer over WiFi delivery for the video content created through their phones.

Another alternative is to transfer information directly from one phone to another, Peer-to-Peer. With mobile phones increasing in technical sophistication, they could incorporate mesh networking technology to route traffic from device to device, avoiding the central network. Such a configuration would be particularly useful for applications involving video transfers within a limited geographic area.

The problem of video Peer-to-Peer usage is less acute on backbone transport networks. Peer-to-Peer traffic represents about 20% of traffic on Internet backbones, still large but substantially less than its share of access traffic.[57] Transport networks generally are over-provisioned, in contrast to access networks, because the cost of adding additional capacity to fiber-optic backbones is relatively low.[58] The greatest capacity constraint in the backbone is over international links, especially those involving undersea cables or satellite connections. Despite vast capacity increases during the telecom bubble, the data capacity available across the oceans is still far less than in-country for the developed world.

Fortunately, Peer-to-Peer traffic is relatively local. This is so for technical reasons, encouraged by the Peer-to-Peer software itself, but also because of the nature of the content. Popularity of media content often differs from country to country, for language and other cultural reasons. As noted, popularity of Peer-to-Peer file-sharing software also tends to be regional, with eDonkey more popular in Asia and Europe, and FastTrack in the USA. Such regional fragmentation means the Peer-to-Peer traffic itself is more highly concentrated on national or regional backbones, rather than spanning limited international connections.

Symmetry

Beyond the sheer volume of bits, Peer-to-Peer traffic stresses networks by confounding established traffic patterns. Network architects engineer systems based on assumptions about how and when traffic flows. A system used for one-way video broadcasting of a limited set of content, for example, may employ caching servers or a technique called multicasting to eliminate redundant flows of the same traffic. An electronic data interchange (EDI) network for business trading partners, by contrast, will be optimized for relatively symmetric, relatively unique traffic.

Video Peer-to-Peer traffic is inherently more symmetric than the Web and rich media broadcast content that most broadband networks have been optimized for, and its usage patterns differ from the baseline in other significant respects.

Some network traffic flows roughly equally in both directions. The telephone network as a whole, for example, has this characteristic. On average, people make and receive about the same number of calls. In local cases, though, this symmetry does not hold. Ticket agencies, for example, receive many more calls than they make, while outbound telemarketing call centers have the reverse pattern.

When networks designed for one type of traffic encounter a new distribution, they can experience economic and technical problems. Thus, when local phone companies first experienced high levels of dial-up Internet access in the mid-1990s, they complained that the increased number of long, outbound calls to ISPs forced them to make unplanned investments to add ports to their local switches.[59]

Though dial-up Internet connections are exclusively upstream from the perspective of the phone network – people call their ISP, not the reverse – Internet access traffic itself has historically been primarily downstream. Users of the Web request content from Websites. They rarely operate servers which send content out to others.[60] The asymmetry of Internet traffic traditionally tended to *increase* as file sizes grew. Text-based emails are largely symmetric, static graphics are largely sent downstream from Websites, and rich media content is almost always received in one-way broadcast mode.

Broadband access providers have architected their networks to take advantage of this asymmetry. Asymmetric access networks allow providers to save on capacity, improving downstream performance at a lower cost. One of the fundamental properties of information theory is Shannon's Law, which postulates a maximum information carrying capacity for a given communications link. Introducing asymmetry allows communication in one direction to exceed the apparent Shannon's Law limit. Moreover, asymmetric networks simply do not require the same investment in upstream capacity. This distinction is particularly important for cable modem systems. Cable networks were built for television, which is almost exclusively a downstream application.[61] Cable operators have had to spend money upgrading their networks for upstream capacity, which still comes at a premium.

Asymmetric broadband networks offer other benefits to access providers. They may be able to charge premium rates for specialized video, audio, and gaming content that flows down from their servers to their users. Even if they cannot, they have much more control for traffic engineering purposes over traffic that originates in their network or flows in through a limited number of peering points, compared to that originating at a large number of individual users' edge machines.

With Peer-to-Peer traffic, every user is potentially an originator as well as a recipient of content. This is true even when the content is something, such as an episode of *The Sopranos* that end users only intend to view. Most Peer-to-Peer file-trading applications either encourage or require users to upload as well as download. BitTorrent, for example, has a built-in mechanism to incentivize symmetric utilization. This addresses performance issues that hobbled other Peer-to-Peer systems, including early versions of Gnutella.

Moreover, some Peer-to-Peer traffic *is* user-created. This is an important distinction between video and other forms of Peer-to-Peer traffic. Very few ordinary people record and distribute their own music. Quite a few, however, take photographs and shoot home movies. With the plummeting cost of digital video cameras, video-enabled mobile phones, personal computers capable or running powerful video editing software, and storage devices such as rewriteable DVDs, more and more end-users have the ability to be content creators as well as consumers. Thus, video Peer-to-Peer traffic is likely to involve even more upstream activity than music. This poses a dilemma for broadband access providers in the USA, who have uniformly deployed asymmetric access networks.[62]

The symmetry of video Peer-to-Peer traffic will primarily be an economic issue for access networks. Backbone transport networks have always been largely symmetrical, because they aggregate traffic among different kinds of access networks. For every downstream-heavy broadband access service there is an upstream-heavy server farm. And, as noted above, Peer-to-Peer traffic represents a significantly smaller share of traffic on the backbone than on access networks.

Usage Patterns

As discussed above, video Peer-to-Peer file transfers can take many hours to complete. This extended period of activity contrasts with most Web and email sessions, which are relatively short-lived and require active participation by the user. Peer-to-Peer software, especially when used for video, is often left online for extended periods when the user himself or herself is not present. Consequently, the time of day in which the Peer-to-Peer activity takes place does not necessarily track the daylight or work hours the way other Internet traffic does. This confounds traffic engineering metrics that network operators use to allocate capacity.

In addition, users of Peer-to-Peer file-sharing systems typically only download a particular file one time. Once they have the file on their computer, assuming it is complete and functional, there is no reason to download another copy. A movie today will be the same movie next week. The Web is different. The home page of CNN.com six hours from now may be very different than the same page right now. Users download the "same" web page many times, either because it has in fact changed, or to determine whether it has done so. Therefore, traffic patterns on the Web are driven largely by the speed of changes to content. On Peer-to-Peer file-sharing networks, they are driven by addition of new items.[63]

Because of these various differences, Peer-to-Peer usage does not follow the "power law" distribution that marks Web traffic and many other network phenomena.[64] The most popular objects on KaZaA are significantly *less* popular than a power law would predict.[65]

Again, these variations from standard Internet traffic impose costs on network operators by throwing off their models for allocating capacity. Service providers also engineer peering points with other networks in the most efficient configuration, based on traffic flows. As more and more traffic reflects video Peer-to-Peer content rather than traditional Web content, these peering point allocations will also become less accurate.

Likely Consequences and Responses

Video Peer-to-Peer will impact network usage in two ways: by changing what users do, and by changing traffic loads. Both will create incentives for service providers to respond.

Historical network-stressing applications

The rise of video Peer-to-Peer represents the fourth instance that a shift in online usage patterns has created significant new stresses for network operators. In every previous case, the Internet has been up to the challenge. The response to the prior strains did, however, lead to changes in Internet architecture and to Internet economics.

The first situation was the rise of the dial-up Internet in the mid-1990s. The Net was not the first online service, but its scope exceeded anything before. The Internet prior to the emergence of commercial services providers and the World Wide Web was primary designed for non-commercial academic and research users. The growth of dial-up ISPs and the Web created two primary stresses on the extant Internet architecture: telephone switch congestion and peering congestion.

Dial-up users call ISPs through the public switched telephone network, which was engineered for analog voice traffic. The average voice call is 3–5 min, but the average Internet connection is far longer. Phone companies complained to the FCC that dial-up traffic was imposing significant costs on them and threatening to degrade service for other telephone users.[66] In the core of the network, the bottleneck was the limited number of locations, known as Network Access Points (NAPs), where major ISPs exchanged traffic. These open, multi-lateral peering points suffered increaseing performance problems, leading the largest backbone ISPs to move to bilateral private peering and utilize traffic engineering techniques to route traffic more efficiently. There was also a major effort to construct exchange points outside the USA, limiting the traffic that had to traverse slow and expensive international links.

The second time when Internet infrastructure was inadequate to handle growing demand came in the late 1990s. Internet usage, especially World Wide Web and e-commerce activity, was ramping up explosively and globally. Early broadband deployments added to the load. Even with the sophisticated traffic engineering strategies of the backbone carriers, performance began to slow.

The problem was that any traffic forced to traverse the network to a remote Web serve faces delays, either en route or because the originated Web server was overloaded. The more popular the site, the more serious the problem. The response was the development of content delivery networks (CDNs), the most prominent of which was Akamai. CDNs function as network-wide distributed caches. Popular files are served from caching servers close to the user, rather than from the origin server. CDNs shifted revenue flows among Internet service providers and equipment vendors. They also created a new, albeit distributed, Internet point of failure. If Akamai's network goes down, as parts of it did in early 2004, it is as if the Internet failed.[67]

The growth of Peer-to-Peer networks for music file sharing, beginning with Napster, was the third major stress on the network. Though Peer-to-Peer networks generated a vast amount of traffic, there have been few reports of significant network performance impacts across the public Internet. This is likely due to the fact that Peer-to-Peer arrived on the scene in the midst of a frenzied overexpansion of long-haul network capacity, fueled by the Internet and telecommunications bubble. As fast as Peer-to-Peer file sharing ate up bandwidth, new bandwidth was going into the ground. The one area where Peer-to-Peer is having an impact is on last-mile broadband networks. Phone and cable companies have used terms of service and technically-enforced speed limits to prevent users from extensively sharing files. Nonetheless, many broadband ISPs complain that a small number of users are responsible for a significant percentage of their bandwidth utilization.

And now video Peer-to-Peer seems poised to eclipse all of them, at least in terms of absolute traffic loads. The good news from this historical survey is that the Internet has withstood all the prior deluges. As eminent a figure as Bob Metcalfe, the inventor of the ubiquitous Ethernet networking protocol and founder of network equipment vendor 3Com, predicted that the Internet would collapse under the load of the first growth spurt. He was forced to literally eat his words. In each case, the solution has been a combination of new technology and "throwing bandwidth at the problem."

Responses to Video Peer-to-Peer Traffic

Network operators facing the growing flood of video Peer-to-Peer traffic can and do take several steps to respond. Some of these involve network engineering. For example, the symmetric nature of Peer-to-Peer traffic is likely to cause broadband access providers to peer directly, rather than feeding traffic to peered transport providers.[68] So far, though, the most common responses involve restrictions on user behavior. For example, some broadband providers today limit users' ability to operate servers. These provisions have been used against heavy users of Peer-to-Peer file trading software. Service providers also can enforce caps on upstream traffic to kick heavy video Peer-to-Peer users off their networks.

Deep Packet Inspection and Blocking/Filtering

A new class of "deep packet inspection" hardware promises to identify Peer-to-Peer traffic directly, allowing service providers to exclude it entirely or throttle down capacity available to these applications relative to others. The difficulty up to now has been that Peer-to-Peer traffic does not use standard port numbers, which would allow it to be distinguished from Web or email traffic. The only way to identify Peer-to-Peer traffic is to analyze packets at the application layer, rather than the lower network layers where switches and routers typically operate. Doing so requires hardware able to read packets at extremely high rates of speed, which has only recently become feasible.

Service providers have other reasons to deploy deep packet inspection. The FCC has tentatively concluded that managed voice over IP (VoIP) and broadband access services are subject to the wiretapping obligations of the Communications Assistance to Law Enforcement Act (CALEA). If formally adopted, as seems likely, this requirement will obligate service providers to make their networks amenable to wiretapping of VoIP calls. To do so, however, requires knowledge about which traffic is VoIP – information available at the application layer.

Moreover, broadband access providers may voluntarily deploy deep packet inspection gear for other reasons. Classifying services at the application level potentially allows broadband providers to offer differentiated value-added services and enhance security. It also could be used to identify and either block or degrade third-party VoIP traffic. Though major broadband providers have so far disclaimed any intention of doing so, they may have economic incentives to tilt the scales in favor of their own voice offerings, absent regulation to the contrary.[69] An article in mid-2004 quoted a deep packet inspection vendor who stated that his company was in trials with major cable broadband operators, and that third party VoIP services "raped" the access providers networks.[70] Cisco's acquisition of P-Cube suggests the leading data networking hardware vendors are not ignorant of the potential demand for packet inspection technology.[71]

Service Provider Opportunities

Instead of seeing video Peer-to-Peer as purely a negative development, service providers could exploit it to develop new revenue streams. The "if you can't beat 'em, join 'em" strategy has already been employed on university campuses with regard to music sharing. In a 1-year period beginning in mid-2003, more than 20 universities struck licensing arrangements to grant their students access to Peer-to-Peer music downloads, subject to a monthly fee.[72] In these cases, however, there is proven demand for the application, and the universities are simply looking to manage their networks in a way that avoids both legal complications and excessive costs.

There are, however, reasons to believe that broadband service providers will look for ways to offer video Peer-to-Peer services. As more user activity by broadband subscribers reflects video Peer-to-Peer and related applications, service providers will have incentives to capture more of the revenues from those activities. Broadband services generally involve flat monthly rates with upstream and downstream bandwidth caps. Service providers do not benefit directly when users send more traffic; in fact, they may see such usage as a net loss, because it requires them to provision more capacity. If service providers could realize incremental revenue from video Peer-to-Peer transfers, they would have incentives to invest in the necessary network capacity to support them.

As noted above, there are reasons to believe that a smaller share of video Peer-to-Peer activity will involve unauthorized distribution of copyrighted material than is the case for music. The video Peer-to-Peer content that represents personal "life sharing" activities, distributed media, and monitoring, doesn't raise the intellectual property concerns that dominate audio file sharing. Yet these applications face the same network constraints.

Service providers can offer their customers enhanced service quality, ease-of-use, and additional features such as archiving that would enhance any of these usage scenarios. To the extent users have the opportunity to choose between access providers whose terms of service and affirmative offerings constrain their ability to engage in video Peer-to-Peer activity and those provide premium services tailored to video Peer-to-Peer, market incentives may help create a situation favorable to video Peer-to-Peer expansion. For distributed media, there is an opportunity for aggregation, filtering, and billing service providers who package content and make it available to users for a fee.

Service providers could deploy caches or "superpeers" within their networks to make video Peer-to-Peer file transfers more efficient, while offering their users software packages to take advantage of them. By keeping video

Peer-to-Peer transfers more local, such a strategy would also reduce capacity demands in the service providers' networks, thus reducing costs.

Conclusions

The enduring popularity of Peer-to-Peer file-trading for music, despite intensive legal efforts and licensed music distribution alternatives such as Apple's iTunes, shows that once Peer-to-Peer platforms achieve critical mass, they are virtually impossible to stamp out. In the case of video, the rising sales of video cameras and broadband connectivity seem destined to create the conditions for substantial new applications. Video Peer-to-Peer will place dramatic new demands on data networks regardless of what type of content it carries. And the opportunity to transfer non-commercial content will create new business opportunities and usage shifts.

Given the early stage of video Peer-to-Peer activity, especially in the USA, precise economic predictions are difficult to make. Many factors will influence future developments, including the pace of broadband rollouts for truly high-capacity connections (at least 10 Mbps, and ideally at least 100 Mbps, in both directions); the influence of disruptive actions by non-traditional participants in the networking world, including Apple, Sony, Nokia, and Microsoft; and progress on standardization of short-range high-speed wireless links between media-capable devices. Still, the question is *when*, not *if*. Video Peer-to-Peer is here to stay.[73]

Notes

1. The Buggles, Video Killed the Radio Star (Polygram 1980). This song was the first music video played on the MTV cable network when it launched in 1981.
2. Peer-to-Peer is a technical architecture in which individual nodes such as end-user personal computers connect with one another directly, rather than to a central switch or server. Peer-to-Peer systems can be used for many functions other than transferring files. For example, Peer-to-Peer architectures are being employed for distributed computing, storage, electronic commerce, search, and business collaboration.
3. Questions about the legal propriety and responses to Video Peer-to-Peer distribution, including matters of intellectual property enforcement, are outside the scope of this paper. However, it is worth noting that the legal standard for contributory infringement by a piece of hardware or software that can be used for copyright violations is whether it has "substantial non-infringing uses." *Sony Corp v. Universal City Studios*, 464 U.S. 417 (1984). If a higher percentage of video Peer-to-Peer traffic is in fact a non-infringing activity, then that may affect the outcome of future legal challenges.
4. Peer-to-Peer systems are also difficult to track because they do not use standard port numbers in communicating over the Internet. Port numbers are supposed to represent different applications; for example, the World Wide Web employs port 80. Not only do different Peer-to-Peer systems not use the same port number, but each client may employ a range of different numbers. This variation is partly for technical reasons of penetrating firewalls and ensuring reliable connectivity, and sometimes for purposes of obscuring activity from network operators or content owners.
5. *See* Andrew Packer, The True Picture of Peer-to-Peer Filesharing, *available at* http://www.cachelogic.com/press/CacheLogic_Press_and_Analyst_Presentation_July2004.pdf, at 12 (CacheLogic Presentation).
6. *See infra* text at note 67.
7. CacheLogic Announces New Internet Data Analysis Platform, Provides Exclusive Data on Worldwide Peer-to-Peer Usage, press release, July 15, 2004, *available at* http://www.cachelogic.com/news/pr040715.php.
8. *See* T. Karagiannis, A. Broido, N. Brownlee, K. Claffy, and M. Faloutsos, "File Sharing in the Internet: A Characterization of Peer-to-Peer Traffic in the Backbone," UC Riverside Technical Report, 2003, at 11–12.
9. *See* CacheLogic Presentation, *supra* note 5, at 9.
10. Approximately 70% of bandwidth at one cable broadband access provider measured by equipment vendor P-Cube was attributable to Peer-to-Peer. *See* Approaches to Controlling Peer-to-Peer Traffic: A Technical Analysis, P-Cube Technical White Paper, *available at* http://www.p-cube.com/doc_root/products/Engage/WP_Approaches_Controlling_Peer-to-Peer_Traffic_31403.pdf, (P-Cube White Paper) at 4.

11. J.A. Pouwelse, P. Garbacki, D.H.J. Epema, and H.J. Sips, A Measurement Study of the BitTorrent Peer-to-Peer File-Sharing System, *preprint available at* http://www.isa.its.tudelft.nl/~pouwelse/bittorrent_measurements.pdf, at 13–14.
12. CacheLogic Presentation, *supra* note 5, at 12.
13. Adam Pasick, File-Sharing Network Thrives Beneath Radar, Reuters, November 3, 2004, *available at* http://uk.news.yahoo.com/041103/80/f5x2i.html
14. Exact sizes vary, even for the same original file, based on the compression scheme used and, for video, the resolution of the original. For example, DVDs start at higher resolution than analog formats such as VHS (videotapes) or NTSC (over-the-air television).
15. *See infra* Part V. On the other hand, since a user may view several dozen Web pages in a day, but only download one movie. As bandwidth increases, though, the amount of video content it is reasonable to obtain will increase.
16. Krishna Gummadi, Richard Dunn, Stefan Saroiu, Steven Gribble, Henry Levy, and John Zahorjan, "Measurement, Modeling, and Analysis of a P2P File-Sharing Workload," http://www.cs.washington.edu/homes/gribble/papers/p118-gummadi.pdf, at 2.3.1. KaZaA is a Peer-to-Peer application based on the FastTrack platform. In the USA and much of the world, FastTrack was the most popular Peer-to-Peer file-sharing platform in the years between the shutdown of Napster and 2004, when video-oriented services surpassed it.
17. Organization for Economic Cooperation and Development, "P2P Networks in OECD Countries," pre-release of a section from the OECD Information Technology Outlook 2004, http://www.oecd.org/dataoecd/55/57/32927686.pdf.
18. Peter Lyman and Hal Varian, "How Much Information: 2003," *available at* http://www.sims.berkeley.edu/research/projects/how-much-info-2003/internet.htm.
19. As of January 2004, the USA ranked 11th in per-capita broadband penetration. *See* Jim Hopkins, "Other Nations Zip by USA in High-Speed Net Race," USA Today, January 18, 2004.
20. Regional Characteristics of Peer-to-Peer, Sandvine White Paper, *available at* http://www.sandvine.com/solutions/pdfs/Euro_Filesharing_DiffUnique.pdf, at 5. ("In short, there is more demand amongst European Peer-to-Peer users for video content.")
21. Ramayya Krishnan, Michael Smith, and Rahul Telang, The Economics of Peer-To-Peer Networks, September 2003 draft, at 2.
22. The court's rationale was that Napster provided the "site and facilities" for infringement. *A&M Records v. Napster*, 239 F.3d 1004, 1019 (9th Cir. 2001). A subsequent case involving Peer-to-Peer services without that central directory initially came out the other way. *MGM Studios v. Grokster*, 259 F.Supp.2d 1029 (C.D. Cal. 2003), *affd.*, 380 F.3d 1154 (9th Cir. 2004), but was later reversed by the Supreme Court. *MGM Studios v. Grokster*, 125 S. Ct. 2764; 162 L. Ed. 2d 781 (2005). Moreover, Napster's reliance on a central database meant that once Napster the company was shut down, Napster the service disappeared as well. This is not the case with the newer, more distributed Peer-to-Peer networks.

23. Rich media refers to content that is not static, in other words, sound and video recordings. These objects are inherently sequential and subdividable, which is what allows them to be delivered through streaming. Peer-to-Peer networks are also used to distribute non-rich media, such as photos and software applications, which are outside of the scope of this paper.
24. Digital rights management may be used to limit access to the file or prevent it from being transferred to other users.
25. There are hybrid approaches which provide some of the benefits of file transfer with the immediacy of streaming. For example, a file can be streamed but stored for later playing locally once the entire file has been streamed.
26. The major category of Peer-to-Peer streaming applications are instant messaging applications such as Yahoo! Messenger and Apple's iChat AV, which offer video chat capabilities.
27. Karagiannis et al., *supra* note 8, at 6–7.
28. Variations of swarming technology are used in several research projects. Some commercial ventures, such as OpenCOLA and Onion Networks incorporated swarming. However, BitTorrent is the first implementation of swarming to become popular as a means of distributing files across the Internet.
29. *See* Pouwelse et al., *supra* note 11 (describing BitTorrent architecture and performance).
30. Bram Cohen, Incentives to Build Robustness in BitTorrent, May 22, 2003, *available at* http://bittorrent.com/bittorrentecon.pdf. The standard size is one quarter of a megabyte. *Id.* at 2.
31. Recent versions of eDonkey include a swarm download system called Horde. At this stage, BitTorrent's implementation appears to be more effective, based on anecdotal user reviews and takeup rate. BitTorrent is growing more quickly than eDonkey, and it is used almost exclusively for video and other large files. However, it is possible that eDonkey (or some other platform) will surpass BitTorrent in the future.
32. Eytan Adar & Bernardo Huberman, Free Riding on Gnutella, First Monday, October 2000, *available at* http://www.firstmonday.dk/issues/issue5_10/adar/.
33. The algorithm is based on the well-known "tit for tat" strategy from game theory.
34. Interview Responses From BitTorrent's Bram Cohen, Slashdot.org, June 2, 2003, http://interviews.slashdot.org/interviews/03/06/02/1216202.shtml?tid=126&tid=185&tid=95.
35. The social experience of the two media types is also different. People often listen to music in the background, and they listen to the same song or album many times. Movies, by contrast, require full attention, and are often only viewed once. Irrespective of the technical differences, people typically own more songs than they do movies.
36. BBC News, Films "Fuel Online File-Sharing," July 15, 2004, *available at* http://news.bbc.co.uk/1/hi/technology/3890527.stm. The highest level was in Spain, where 38% of users admitted to downloading movies at least once a month. The rate in the USA was 12%; half the rate claimed in the MPAA study, but still significant.

37. Movie and Software File Sharing Overtakes Music, New Scientist, July 12, 2004.
38. *See* Felix Oberholzer and Koleman Strumph, The Effect of File Sharing on Record Sales: An Empirical Analysis, *available at* http://www.unc.edu/~cigar/papers/FileSharing_March2004.pdf.
39. *See* Thomas Karagiannis et al., Is Peer-to-Peer Dying or Just Hiding, *available at* http://www.caida.org/outreach/papers/2004/Peer-to-Peer-dying/Peer-to-Peer-dying.pdf, at 1. ("In general we observe that Peer-to-Peer activity has not diminished. On the contrary, Peer-to-Peer traffic represents a significant amount of Internet traffic and is likely to continue to grow in the future, RIAA behavior notwithstanding.")
40. For example, ETree (http://bt.etree.org/index.php) is a site specifically for high-quality recordings of concerts by artists who expressly permit their fans to record and share their performances.
41. Lucy Sheriff, BBC Ponders Peer-to-Peer Distribution, The Register, February 17, 2004, http://www.theregister.co.uk/2004/02/17/bbc_ponders_Peer-to-Peer_distribution/.
42. The appeals court in the Grokster case pointed to one prominent example, the band Wilco, in justifying its conclusion that the Peer-to-Peer software had substantial non-infringing uses. *Metro-Goldwyn-Mayer Studios, Inc. v. Grokster Ltd.*, 380 F.3d 1154, 1161 (9th Cir. 2004).
43. Reuters, Canon to Pull Out of Analog Camcorders, May 12, 2004, http://reuters.com/newsArticle.jhtml?type=technologyNews&storyID=5115158 (stating that Canon expects sales of 2.6 million digital camcorders in 2004, representing 20% of the global market).
44. Nokia's Ollila Says Global Mobile-Phone Users Reach 1.7 Billion, Bloomberg News, November 3, 2004 (estimating sales of 200 million cameraphones in 2004). Camera Phone Sales to Outstrip Film, Digital, PhoneContent.com, August 14, 2003, http://www.phonecontent.com/bm/news/gnews/camera.shtml (comparing cameraphone sales to film and digital cameras).
45. Dinesh C. Sharma, Study: Pretty Picture for Camera Phone Sales, CNet News.com, March 11, 2004 *available at* http://news.com.com/2100-1041_3-5172377.html.
46. Andrew Odlyko, Content is Not King, *First Monday*, February 2001, *available at* http://www.firstmonday.dk/issues/issue6_2/odlyzko/.
47. A study by the Pew Foundation, based on data from February 2004, found that 42% of American Internet users, representing 53 million adults, used instant messaging. *See* Pew Internet & American Life Project, How Americans Use Instant Messaging, September 1, 2004. Five percent of those users, or nearly three million people, claim to use instant messaging to send music or video files. Even higher percentages, including 21% of respondents of age 18–27, and 12% of those of age 28–39 said they used streaming audio or video to see or hear people they instant messaged.
48. Drazen Pantic, Anybody Can Be TV: How P2P Home Video Will Challenge the Network News, http://journal.planetwork.net/article.php?lab=pantic0704.

49. I first heard this scenario from Jonathan Schwartz, President of Sun Microsystems, at the Supernova conference which I organized in June 2004.
50. China is the largest mobile phone market in the world, and with its vast population, it lead seems certain to grow in the years ahead.
51. RSS stands for either "really simple syndication" or "rich site summary." There are several variations of RSS, and the competing Atom protocol is similar in its basic approach. Both are built on the extensible markup language (XML).
52. Tivo purchased a startup called Strangeberry that was developing technology to integrate Internet content with television. Kim Gerard, Saving TiVo, Business 2.0, September 2004.
53. Elizabeth Corcoran, Shifting Places, *Forbes*, August 2004.
54. Sprint Launches New Live Webcam Application for Mobile Handsets, http://144.226.116.29/PR/CDA/PR_CDA_Press_Releases_Detail/1,3681,1112170,00.html
55. The exception is information such as live traffic pictures. Such content is a perfect fit for the syndication/aggregation model of distributed media.
56. Mike Masnick, Taking The Upload Out Of The Camera Phone Process, The Feature, http://www.thefeature.com/article?articleid=100917&ref=2725973, July 24 2004 (describing a startup called Cognima that offers mobile synchronization software for cameraphones).
57. Karagiannis et al., *supra* note 8, at 12.
58. Moreover, the locality of much Peer-to-Peer traffic means that a relatively small share traverses long-haul network backbones. Gummadi et al., at 5.
59. Kevin Werbach, Digital Tornado: The Internet and Telecommunications Policy, FCC Office of Plans and Policy Working Paper No. 29, March 1997. It is not clear that this argument was justified. The FCC declined invitation to impose per-minute access charges on the ISPs, which the phone companies argued would compensate them for their increased costs. In time, complaints about switch congestion died down.
60. The commercial Website operators typically use specialized Web hosting providers who offer high-capacity outbound links under a significantly different pricing framework than that available to individual end-users.
61. Cable TV networks have a limited upstream channel for telemetry and basic interactive functions, but this alone is insufficient for any Internet usage.
62. Karagiannis et al. at, *supra* note 8, 11 ("Technologies such as DSL and cable modem were quite sufficient when downstream throughput was the main concern. Their attractiveness will fade and their market share dwindle if alternative broadband technologies are deployed that offer comparable upstream and downstream performance.").
63. Gummadi et al.
64. Albert-Lazlo Barabasi and Reka Albert, Emergence of Scaling in Random Networks, Science, October 15, 1999, at 509.
65. Gummadi et al., at 2.3.3.
66. *See* Digital Tornado, *supra*.

67. Paul Roberts, Akamai outage hobbles Google, Microsoft, others, Infoworld, June 15 2004, http://www.infoworld.com/article/04/06/15/HNakamaioutage_1.html.
68. Karagiannis et al., *supra* note 8, at 11.
69. *See* Tim Wu and Lawrence Lessig, Ex Parte Submission in CS Docket No. 02-52, *available at http://faculty.virginia.edu/timwu/wu_lessig_fcc.pdf*
70. Eric J. Savitz, “Talk Gets Cheap: Internet telephony is bad news for the Bells, but maybe great news for the cable guys,” Barron’s, May 24, 2004.
71. *See* Reuters, Cisco to Buy P-Cube for About $200 Million, August 23, 2004.
72. John Borland, College Peer-to-Peer use on the decline? ZDNet, August 24, 2004, *available at* http://zdnet.com.com/2100-1104_2-5322329.html
73. Kevin Werbach is an Assistant Professor of Legal Studies and Business Ethics at The Wharton School, University of Pennsylvania. Email: kevin@werbach.com.