Pirates and Samaritans: a Decade of Measurements on Peer Production and their Implications for Net Neutrality and Copyright

J.A. Pouwelse, P. Garbacki, D.H.J. Epema, H.J. Sips

Parallel and Distributed Systems, Delft University of Technology, The Netherlands
P.O. box 5031, 2600GA Delft

Abstract

This study traces the evolution of commons-based peer production by a measurement-based analysis of case studies and disusses the impact of peer production on net neutrality and copyright law. The measurements include websites such as `Suprnova.org`, `Youtube.com`, and `Facebook.com`, and the Peer-to-Peer (P2P) systems Kazaa, Bittorrent, and Tribler. The measurements show the two sides of peer production, the pirate side with free availability of Hollywood movies on these P2P systems and the samaritan side exhibited by the quick joining of 400,000+ people in a community to organize protests against events in Burma. The telecommunications and content industry are disrupted by this way of peer production. As a consequence, revenues of both industries are likely to suffer in the coming years. On the other hand, innovative P2P systems could win the battle on merit over classical distribution technologies. As a result, a continuation is expected of both legal actions against P2P and possible blocking actions of P2P traffic, violating net neutrality. It is argued that this hinders innovation and causes a large discrepancy between legal and user perspectives. A reform of copyright laws are clearly needed, otherwise they will be unenforceable around 2010.

Key words: P2P, collaboration, commons-based peer production, copyright

1 Introduction

Now that one billion people are interconnected through the Internet, their combined creative force outweighs that of any company. In the Internet, the

Email address: `j.a.pouwelse@tudelft.nl` (J.A. Pouwelse, P. Garbacki, D.H.J. Epema, H.J. Sips).

roles of producers and consumers are beginning to blur and merge, a development that is best described by the term *prosumer* (producer-consumer) (Toffler 1980). The rise of the prosumer is changing the business landscape. Yochai Benkler has described this phenomenon as "the third mode of production", thereby separating it from the property-based and contract-based models of firms and markets (Benkler 2002). Benkler uses the term *commons-based peer production* to denote the creation of output by prosumers. Generally, in literature all concepts and examples of peer production focus on what is defined here as the *samaritan* side of peer production. However, there is also a dark and disruptive side of peer production, i.e. the *pirate* side (Section 2).

In this article, measurements are presented of enabling *Peer-to-Peer (P2P) platforms*, that allow large groups of prosumers to engage in peer production. A key concept of this technology is its decentralized nature, resulting in a lack of any single person, any authority, or any central computer server that is in charge of the system, hence the term P2P. The measurements characterize the development of P2P platforms over the last decade (Section 3) and cover a diverse range of topics such as TV channel switching behavior, P2P file sharing, friendship within Internet communities, and reputation systems.

The measurements show that peer production is growing in four directions: i) more people are using products created using peer production, ii) the creative output of peer production is increasing, iii) the sophistication of P2P platforms is improving swiftly, and iv) the output types are expanding (e.g., P2P money lending (Manjoo 2006)).

This study addresses two aspects of peer production. The first aspect is to understand and quantify the mechanisms underlying peer production, identify strengths, weaknesses, and conditions that enable it to flourish. The second aspect deals with the influence of peer production on convergence, net neutrality, copyright law, and society at large.

2 The disruptive power of commons-based peer production

Peer production can have a big impact on existing business models. The unleashing of the collective power of large groups of users engaging in peer production is a disruptive change for the content, the telecommunications, and other industries involved. The offerings of a company can in some instances even be provided at *no cost* by a group of organized volunteers. In this section, evidence is given of the disruptive power of peer production.

The first of example is the website `Wikipedia.org`, where users collaborate to create an encyclopedia that can be used and downloaded for free. The

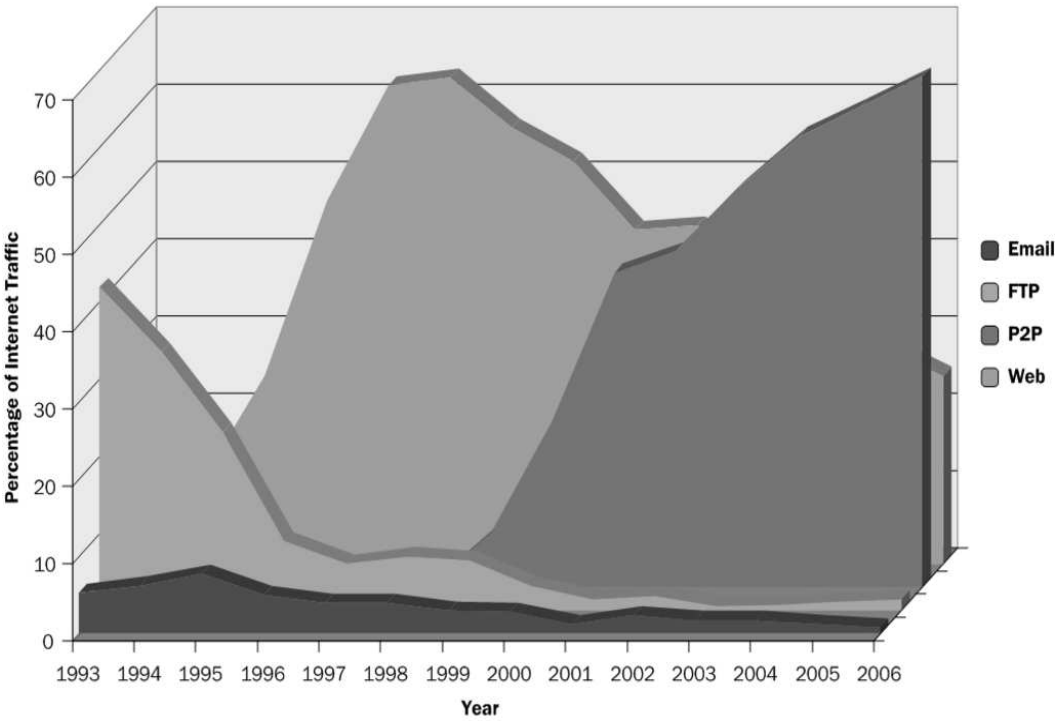

Fig. 1. Over a decade of Internet backbone traffic, source: **Cachelogic.com.**

quality of this encyclopedia has been assessed as being roughly equal to that of the *Encyclopedia Britannica Online* (Giles 2005). Significant media coverage followed the publication of this assessment, and it met with serious criticisms; for example, it was called "wrong and misleading".[1] Wikipedia is one of the top-10 most visited websites in the world, while `Britannica.com` is located in the lower parts of the top-5000, according to monitoring company `Alexa.com`. This illustrates the disruptive effects of peer production.

The second example shows the disruption caused by P2P technology, which can be seen as a natural companion of peer production. P2P replaces the traditional model of an Internet user as a consumer of central server resources with the prosumer model of a user that contributes both content and hardware resources, thereby eliminating the need for central servers. In this respect, P2P file sharing can be viewed as peer production where prosumers pool content (Pouwelse et al. 2005).

Internet measurements confirm that P2P file sharing is the killer application when bandwidth consumption is considered. CacheLogic (`cachelogic.com`) investigated the types of traffic on the Internet backbone during the period 1993–2006, of which the results are shown in Figure 1. In the early 1990s, the dominant Internet protocol was FTP. In mid 1990s, FTP traffic was overtaken

[1] `http://www.nature.com/nature/britannica/eb_advert_response_final.pdf`

by web browsing. The emerging P2P technology almost doubled its share in the Internet backbone each year, becoming traffic wise the most popular protocol after only four years of its existence. In 2006, P2P traffic was responsible for *over two thirds* of all Internet traffic, surpassing web browsing by a factor of almost 3. Over 71% of all this P2P traffic consists of video, as show by a measurement in Germany by `ipoque.com`. According to a Nielsen UK study (Nielsen 2006), "over two-thirds of Britons online have watched some form of video content via the Internet". Both streaming video and downloading where found to be popular. Yet another measurement by `Ellacoya.com` found that web traffic consisted of 36% HTTP streaming video. From these various measurements, it can be concluded that numerous Internet users watch video online, it dominates Internet traffic, and is transported by either HTTP and P2P protocols.

Why do people use online video instead of existing TV broadcast infrastructure? Internet video is often free of charge, is not tied to broadcast schedules (on-demand), is frequently removed of interruptive advertisements, and in the case of P2P may not yet (or not anymore) be available in theaters.

To quantize the disruptive power of peer production, the availability of top Hollywood movies is presented on the leading P2P portal of 2004, the website `suprnova.org`. On the website users could find a wealth of content and trigger a P2P download with a single click. The system reached a popularity of around 800,000 unique visitors per day. By using website crawling techniques, detailed records have been obtained of all files published through Suprnova (Pouwelse et al. 2005). The measurements capture the entire lifetime of Suprnova, from its creation in early 2003 until its (forced) shutdown near the end of 2004. The collected dataset contains the filenames that have been with the names of the top-1000 Hollywood movies of those years. This matching is non-trivial as filenames are only in rare cases literal copies of the official movie titles. For this matching task special software has been developed that employs various naming unification rules (e.g., 'part II' equals '2'), a scoring function to find the most suitable match from various alternatives, and string matching based on the Damerau-Levenshtein distance. [2]

Figure 2 shows the availability of top Hollywood movies on the Suprnova portal. The results span the time period from 2003 until end of 2004. The horizontal axis presents the official movie release date in US cinemas. The vertical axis shows the date at which the file matching the movie title was injected into Suprnova. The figure provides insight into both the number of Hollywood titles that were available for download and the time of availability. The correlation of the movie release time and P2P injection time is visible as a fuzzy line pattern from the lower left corner towards the upper right corner.

[2] `http://en.wikipedia.org/wiki/Damerau-Levenshtein_distance`

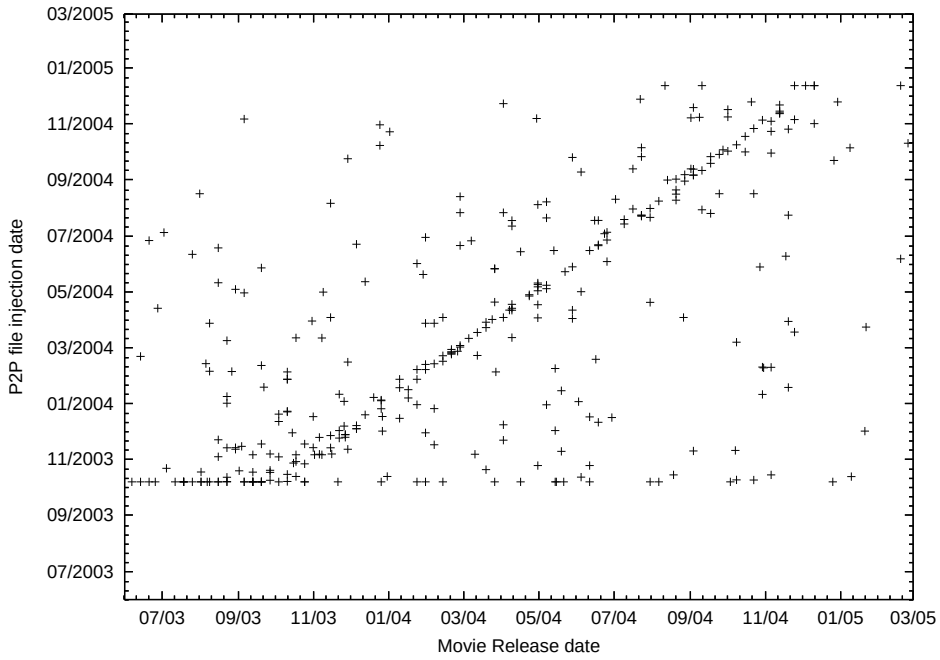

Fig. 2. Comparing the Cinema release and P2P availability of top Hollywood movies.

Prior work examined a sample of 183 movies and found that 77% of them were leaked by industry insiders (Byers et al. 2004). The measurements make clear that for a significant number of titles, users no longer needed to wait for either the screening on TV or even the release in theaters, as P2P availability precedes them. It is estimated that the cost of operating Suprnova was less than $500 per month. From these figures it can be concluded that prosumers collaborating to distribute video content is proving to be effective, popular, cost-efficient, and disruptive to existing business models.

3 The evolution of P2P platforms

In this section, a measurement-based analysis of the evolution of P2P platforms is presented. A P2P platform is defined as "an enabling mechanism for human interaction and cooperation at an unbounded scale that lacks central points of authority and is helped by mutual donations of computer resources". The analysis takes the form of a sequence of seven case studies of P2P platforms that each represent a step in the evolution of these platforms. A careful study of this evolution leads us to believe that effective P2P platforms have the following four cardinal features:

(1) The ability to distinguish good from bad contributions;
(2) A regulation mechanism for computer resources;
(3) Good mechanisms for group communication;

(4) A sense of community.

As to the first of these features, a single *contribution* is considered the elemental building block of collaboration made by a single prosumer. A contribution can have the form of, for instance, a comment of a few lines, a quality rating, or a complete video clip. However, some contributions contain errors or are simply spam. The second feature is to ensure that sufficient computer resources are available to run the P2P software and store the contributed data. The third feature relates to creating the communication channels for prosumers to craft their contributions. The fourth feature is the requirement of a shared vision and culture to ensure long-term viability. A sense of community is vital to overcome differences of opinion. Communities also provide a sense of belonging and incentives to motivate volunteers. This fourth cardinal feature has been studied in the related context of collective action problems and group formation (Cornes and Sandler 1996).

3.1 Early user-generated content and moderation

P2P platforms have a history of about ten years, starting with the creation of the website `Slashdot.org` in 1997. This website is the first large-scale case of user-generated content and user-controlled moderation. Slashdot has been for many years now a very popular source of technology-oriented news. It not only just presents news, but above all it has pioneered in building a self-regulating community of prosumers around the news by means of a discussion area. The novelty here was that this discussion area is moderated by the users themselves.

Slashdot users have the ability to freely attach *news comments* to news items. Such comments contain remarks, enhancements, and insights on the corresponding news items. Slashdot uses a rating system based on voting and reputation to determine a *score* for each news comment. This score ranges from -1 (inferior) to 5 (insightful). Users are randomly given the ability to moderate comments by decreasing (-1) or increasing ($+1$) scores. The initial comment score is 1 for registered users and 0 for anonymous comments.

In May 2000, we have measured Slashdot moderator responsiveness for 30 news items which received 4,250 news comments subjected to 1,400 moderations. Figure 3 shows the average length of the moderation process, from insertion until receiving the final score. The initial score of almost 80% of the comments never changed. For the others, it took on average a mere 37 minutes before insightful comments got their first score increase. The results show that this collaborative moderation process is swift and very responsive. Within hours the good and the bad content is separated, as indicated by the

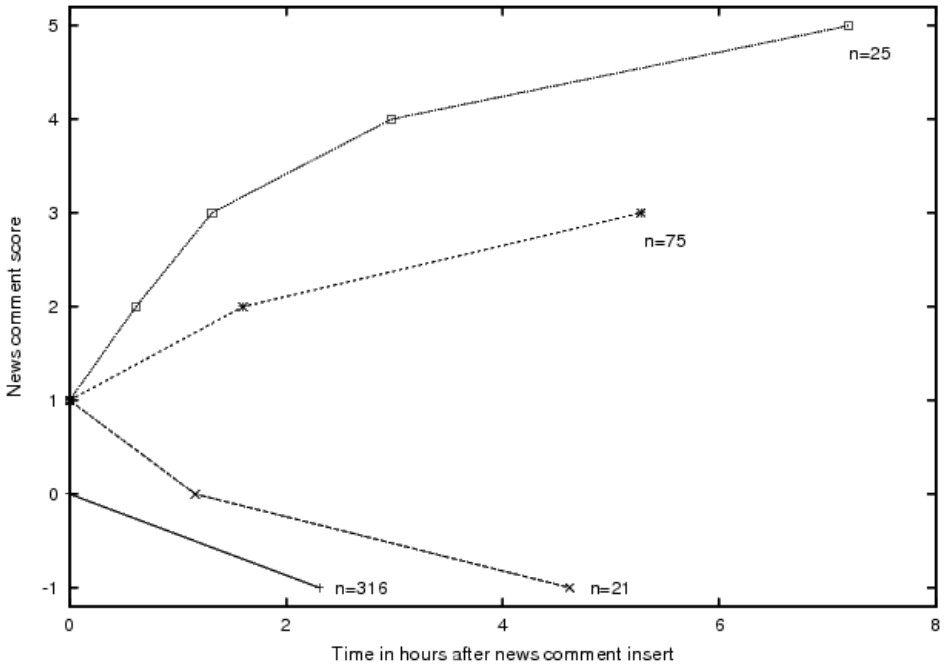

Fig. 3. The responsiveness of the Slashdot rating system.

quick identifications of anonymous (start value=0) spam comments (−1). Note that Slashdot as the first stage in the evolution of P2P platforms lacks two of the four cardinal features: Slashdot users cannot form groups or communicate, collaborative tasks are still basic, and the resources needed to operate the system are provided by revenues from advertisements.

3.2 Public ownership of user-generated content

Slashdot comments are generated by the users, but owned by the site owners. The second point in the evolution of P2P platforms represents a move towards community ownership of content. Content is placed in the public domain with a copyright license similar to the Open Source software license. The most sophisticated example by 1999 is the `Musicbrainz.org` website.

The Musicbrainz website is the home of a community that builds a music encyclopedia. Musicbrainz collects information about artists, their works, country of origin, relationships between the artists, song and album release dates, etc. Musicbrainz as of May 2008 contains the names of 380,000 artists and 6,710,000 music track releases. Most importantly, over 389,000 volunteers have inserted content into this database and corrected (spelling) errors using an advanced voting and moderation system.

When compared to Slashdot it is clear that the way the content moderation process works in Musicbrainz has more complexity and sophistication. The

Table 1
Different amount of donated computer resources in Kazaa (number of users=556).

Percentage of users	Kazaa labels	Sharing ratio R
38.7%	None, Microscopic, Very Tiny	$R < 0.04$
24.6%	Minimal, Low	$0.04 \leq R < 0.5$
14.2%	Medium, High	$0.5 \leq R < 2$
22.5%	Guru, Deity, God	$2 \leq R$

Musicbrainz system exhibits basic forms of all four cardinal features of a P2P platform: a voting process, a fund for collection of voluntary donations for website hosting, a mailing list dedicated to building consensus, and a shared vision as outlined in their "social contract" (Musicbrainz 1999).

3.3 Exploiting donated computer resources

The Kazaa P2P file sharing system launched in March 2001 can be considered the third point in the evolution of P2P platforms (Leibowitz, Ripeanu, and Wierzbicki 2003). Kazaa enables the exchange of multimedia files between prosumers without any server.

Early 2003 we have conducted a series of measurements investigating the regulation of bandwidth in the Kazaa network. As a general networking rule, every downloaded file must be uploaded by another user. A sharing ratio value equal to 1.0 is achieved by users who donate the same amount of bandwidth as they consume. Most users consume significantly more bandwidth than they contribute. This problem is known as "freeriding" (Adar and Huberman 2000). Table 1 shows the sharing ratios for a sample of 556 Kazaa users (ratio of uploaded versus downloaded bytes). A small minority of prosumers (22.5%) in Figure 1 labeled in the Kazaa user interface as 'Deity' and 'God' are responsible for most of the bandwidth donations. Due to freeriding the system performance determined by file download times is significantly degraded. The Kazaa case study shows the need for a robust mechanism which ensures a sufficient amount of resource donations. In the ideal case, a leak-free resource economy should be created where each prosumer must donate an amount of resources equal to its consumption of the collective resource. This problem can again be seen as a separation of the good from the bad (e.g. freeriders). Kazaa represents an initial step towards distributed regulation of donated computer resources, adhering to the second cardinal feature of effective P2P platforms.

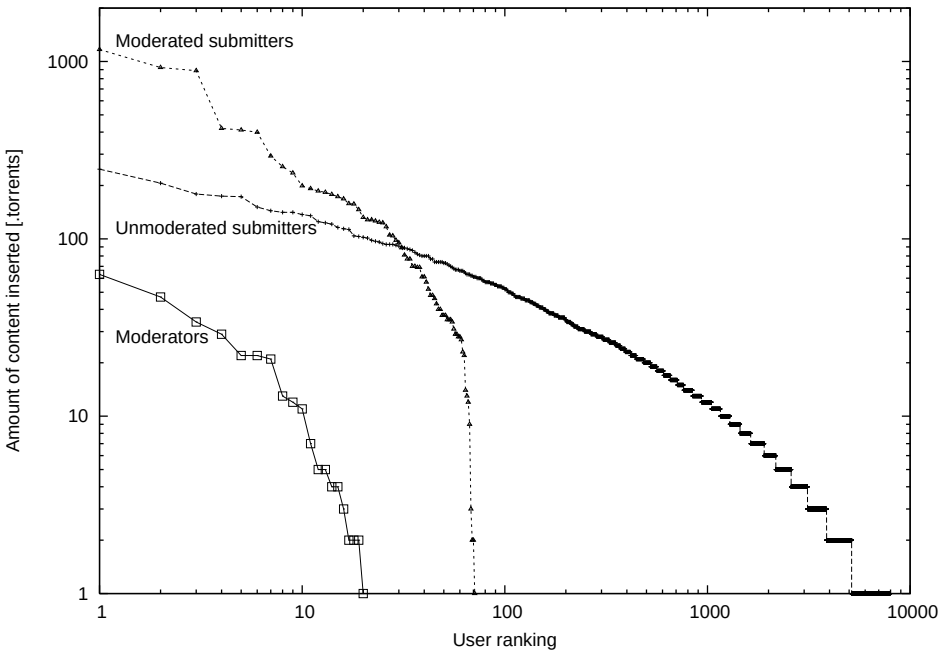

Fig. 4. The contributed content by three types of Suprnova volunteers.

3.4 Regulating bandwidth donations

The freeriding problem encountered in Kazaa is addressed in the design of the Bittorrent (Cohen 2003) protocol. Bittorrent is a P2P file download protocol that provides an incentive to donate bandwidth. Bittorrent's incentive mechanism is built on a data bartering economy. More precisely, Bittorrent employs a variant of tit-for-tat mechanism to restrict the content download rate based on the value of the current contribution. With tit-for-tat users are exchanging (bartering) local content for the content they wish to obtain from other users.

Bittorrent is only a content download protocol. The content search functionality required for a complete file sharing solution, has to be provided by means of external services. A well-known example of such service is the Suprnova site already described in detail in Section 2. Bittorrent/Suprnova represents a step forward from a monolithic approach towards a two-layered solution. It allowed independent evolution of the content transfer protocol and the content browse & search functionality.

The decentralized Kazaa system quickly suffered from spam and pollution. Suprnova took a step back in the evolution and introduced a group of mirrored web servers with a prosumer-based moderation system. To keep the Suprnova content clean, new submitted content was first manually inspected by *moderators*, who weeded out fake content, content with low perceptual quality, and content with incorrect naming. An ordinary user who submits content is called a *moderated submitter*. To lower the burden on the modera-

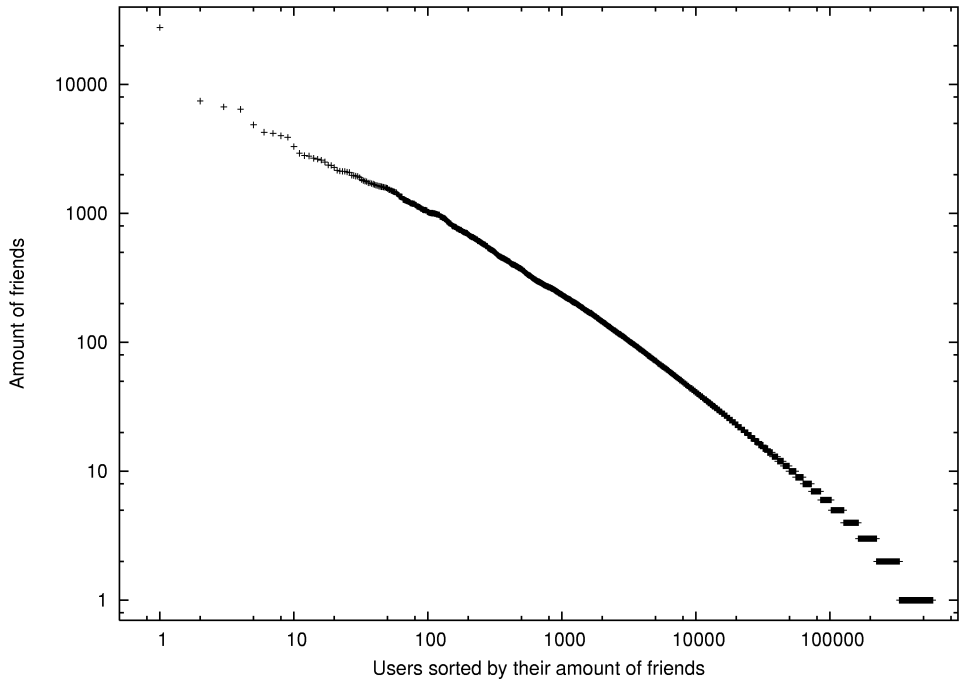

Fig. 5. Online friendships in the Youtube Internet video community.

tors, a user who frequently injects correct content is promoted to the rank of *unmoderated submitter*, and is allowed to directly add content. Unmoderated submitters can request from the existing moderators a promotion to the moderator status. Figure 4 shows the amount of content published by the three types of users extracted from the data collected from the Bittorrent/Suprnova measurements. The horizontal axis presents the users, ranked by the amount of content they inserted. The vertical axis shows the amount of content they inserted into Supernova.

The Bittorrent/Suprnova case study shows the difficulties of building a scalable and spam-free content sharing system. Bittorrent addresses the second cardinal feature with an upload incentive, but this is not fraud-proof (Piatek et al. 2007).

3.5 Community interaction and culture

Currently the largest website for publishing and serving video clips, `Youtube.com`, represents the fifth point in the evolution of P2P platforms. By all accounts, the Youtube architecture is a step back in the evolution of P2P platforms. All functionality in the Youtube system is provided through a single website that runs proprietary server software. Furthermore, similarly to Slashdot, user contributions have restrictive copyright clauses. The architectural and copyright limitations of Youtube are compensated by the richness of its community building functionality.

With the help of a supercomputer we have crawled over 5+ million webpages of Youtube, starting in the summer of 2006. This crawl contains detailed information on numerous users such as their age, country of origin, last online time, list of their favorite video clips, and their Youtube friends. Figure 5 shows for 592,900 users of Youtube how many friends they have. The horizontal axis shows the various users sorted by their number of friends (vertical axis). The leftmost user shown in this picture has 27,716 "friends".

The scale of the social network captured in the measurements indicates that the third and fourth cardinal feature of P2P platforms has clearly advanced in Youtube to a superior level of sophistication. The simplicity and ease-of-use of Youtube is one of the driving forces behind its success.

The cost of operating the vast amount of Youtube servers must somehow be recuperated with targeted advertisements. Google acquired Youtube for $1650 million and the current and future revenues are the focus of much speculation.[3] It remains to be seen if this approach can compete with P2P file sharing and can sustain new costly features such as HDTV support.

3.6 Peer production of functionality

The website `Facebook.com` takes P2P platforms to a new level, and constitutes the sixth point in the evolution of P2P platforms. Facebook is a website that enables social networking, special interest communities, and rich multimedia content. The key element is that Facebook allows a prosumer to improve and expand the functionality of the platform by creating a "prosumer application".

All available prosumer applications can be seen through a "browse applications" menu. A user can install prosumer applications with only a few mouse-clicks. During the installation procedure a user grants access rights to the prosumer application. The new application can then access the list of friends, send messages, and can add itself to the Facebook homepage of the user. On 27 September 2007 there where 4626 of such applications available, with 6049 applications four weeks later. This growth of 355 applications/week indicates that Facebook has a flourishing community that actively expands it. A key finding is the existence of "application spam" on Facebook. For instance, the application entitled *Will you be a miltimillionaire?* forces the user to send a personal message to his/her friends, asking them to also install this application. Application spam tries to spread virally across a social network, gain access to personal information, derive income from (targeted) advertisements, while offering no real functionality.

[3] `http://techdirt.com/articles/20070827/121830.shtml`

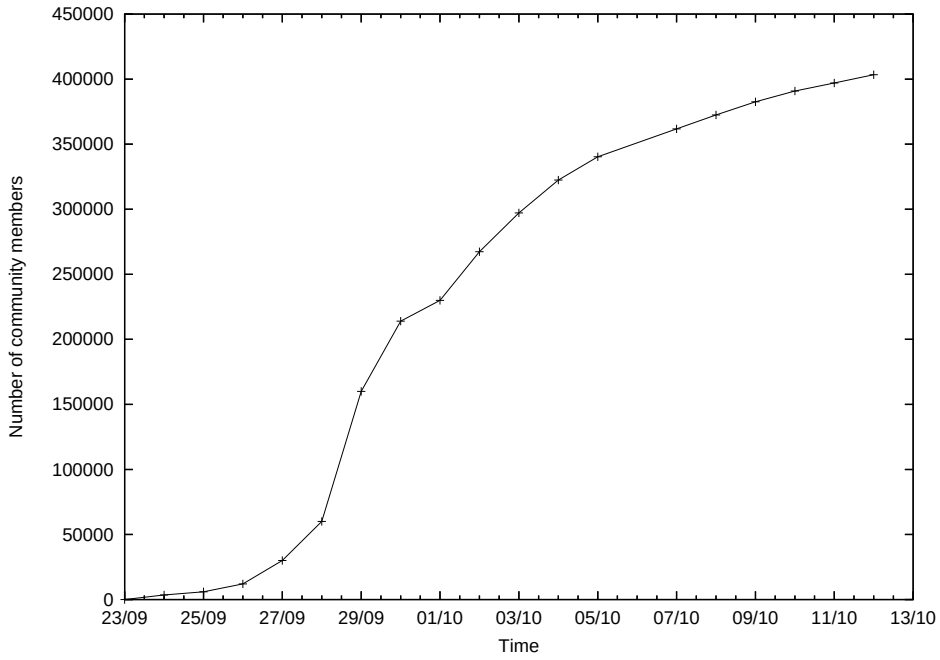

Fig. 6. Growth of the Burma community on the website Facebook.

The observations indicate that peer production has expanded from origins in technology-oriented communities (e.g., Slashdot) to other branches of society. The usage dynamics of P2P platforms in other branches of society is illustrated by the Burma uprising example. The Facebook community called *Support the Monks' protest in Burma* grew to over 400,000 members in less than 20 days. Figure 6 shows the speed of this community growth. The horizontal axis shows the number of days after the community was founded and the vertical axis shows the community size. The maximum growth occurred on 29 September 2007, when 100,000 new members joined in a single day. It can be concluded that Facebook represents a step in the evolution of prosumer communities with impact on society.

3.7 Reputations and decentralized communities

The P2P file sharing system called Tribler (Pouwelse et al. 2008) can be positioned as the seventh point in the evolution of P2P platforms. Tribler includes an explicit reputation system to calculate the "goodness" of others. Reputation systems such as used on the auction website `Ebay.com` have proven to be difficult to realize in a P2P setting (Sarkio 2006, Kamvar, Schlosser and Garcia-Molina 2003). Tribler is the first system to deploy a decentralized reputation system. In its design aspects, Tribler is the complete opposite of Youtube and Facebook as it lacks any central points, is non-profit, uses Open Source licenses, and exploits an open model for software development. Central to Tribler are its social features, its backward compatibility with Bittorrent,

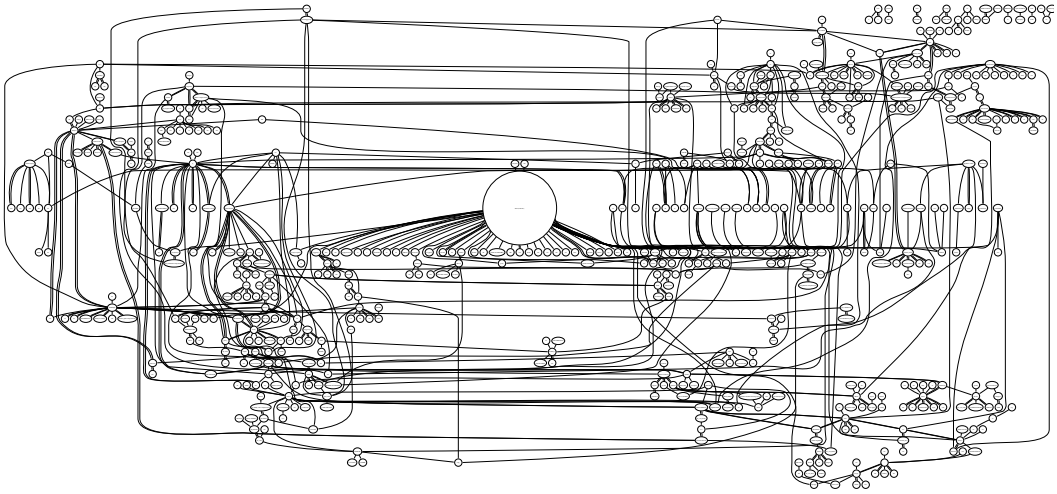

Fig. 7. Bandwidth contribution graph of encountered peers within Tribler.

and fast keyword search for content.

The Tribler reputation system is founded on the "goodness" of other peers in terms of reliability, content preference, and bandwidth sharing. The goodness information of a peer is tied to a non-spoofable public-key identifier and stored in an embedded database, called the *MegaCache*. A Tribler peer constantly explores the P2P network for fresh content and people by making a new connection every few seconds using an epidemic protocol (Pouwelse et al. 2008). For every encountered peer the similarity of their download taste is determined using collaborative filtering techniques. Peers with similar taste are "better" and re-contacted regularly to accelerate content discovery (called semantic clustering). Every Tribler peer tracks bandwidth contributions of others using a protocol called *BarterCast*.[4] to allow the identification of freeriders (Adar and Huberman 2000). Using BarterCast each peer shares with others which peers have given him upload bandwidth. Peers with no prior knowledge of each other thus can use the feedback of the peers they know to assess the trustworthiness of each other.

Figure 7 shows a graphs of such discovered bandwidth contributions as gathered by an Tribler peer, running the unmodified Tribler software. This data was accumulated over a six week period after BarterCast was taken into production within Tribler. The 690 nodes in this figure represents peers, the edges content transfers. The big central node represents the local gathering peer itself. The 52 edges originating from the local peer are due to direct bartering to 52 Tribler peers within Bittorrent swarms. These 52 direct bartering relations are locally observed, all other edges are based on information that is potentially fraudulent. Together these edges form a web-of-trust that can be used to estimate the "upload reputation" of any peer using algorithms such

[4] `http://www.tribler.org/BarterCast`

as PageRank, EigenTrust, and MaxFlow (Sarkio 2006, Kamvar, Schlosser and Garcia-Molina 2003). Every Tribler peer uses the data from such figures to calculate using the MaxFlow algorithm if a peer is possibly freeriding. The density and the coverage of Figure 7 is an indication of the strength of this approach.

Although far from perfect, Tribler is to the authors knowledge the most sophisticated approach with no central servers in actual use. It is a step forward compared to the Kazaa approach and the Bittorrent tit-for-tat algorithm, as both have been shown to be inefficient and insecure (Piatek et al. 2007). It can be noted that all four cardinal features of P2P platforms are present in Tribler. Due to the academic purity of the architecture there are no central points that could limit Tribler scalability.

3.8 Research challenges, trends, and future developments

The evolution of P2P platforms shows that decentralization and trust are the two prime research challenges. Decentralized P2P platforms consistently lag behind in sophistication when compared to central web-server-based approaches. The basic problems of the decentralized approach are converting unreliable and untrusted donated computer resources into a reliable and trustworthy P2P platform and gathering scattered information using software running locally. The decentralized approach is thus intrinsically more complex. However, it is still the preferred approach due to the inherent scalability of decentralized systems. For instance, centralized Wikipedia requires continuous donation rallies to remain functional.[5] It is most likely that a continuous struggle will remain to retain decentralization while improving sophistication. The second challenge recurring in all presented case studies is separating the good from the bad. A distributed reputation system is needed to identify good prosumers in terms of both content contributions and computer resource donations (Sarkio 2006, Kamvar, Schlosser and Garcia-Molina 2003). This would provide significant efficiency gains and allow peer production to flourish further. Figure 7 showed the first operational distributed reputation system. Expanding and utilising this work will be challenging. A hardened reputation system would even permit an "Internet currency" and marketplace for computer resources and services[6] and prevent cheating (Piatek et al. 2007).

The two major developments for P2P platforms that will be ready for mass usage in 2009-2010 will be outlined in the following. The first innovation is HDTV Video on Demand service. This requires improvements to the second cardinal

[5] `http://wikimediafoundation.org/wiki/Personal_Appeal`
[6] `http://news.bbc.co.uk/2/hi/technology/6971904.stm`

feature of P2P platforms (regulation mechanism for computer resources) and further deployment of high speed Internet access such as ADSL2.

The second innovation is anonymous downloading, uploading, and injection of content using a *darknet*. A darknet inhibits both Internet censorship and enforcement of copyright law. The `freenetproject.org` has in 2000 already produced a darknet, but it was slow, difficult to use, and offered little content. Darknets struggle with the second cardinal feature of P2P platforms. Full anonymity costs both extra bandwidth and is difficult to combine with enforcement of resource contributions. By 2010 darknets should be able to offer the same performance as traditional P2P software by exploiting social networking. No effective legal or technological method currently exits to stop darknets, with the exception of banning general-purpose computing (Zittrain 2007). Technologies such as *secure computing* and DRM are convincingly argued to be unable to stop darknets (Schecter, Greenstadt and Smith 2003).

Opinions on the future of peer production range from "a short lived joke" to "the amateur collective will thus supplant the professional institution as the engine of common culture".[7] The predicted cause of decline is that real money will start to reward contributors, but this study has found no evidence to support this view. One likely outcome is that both views are correct. Contributions and resource donations will remain voluntary, but generate a form of "social credit" that can be passed around and utilized elsewhere. It is very likely that "social credits" or an "Internet currency" founded on a hardend reputation system would create an ecosystem in which both prosumers and businesses can flourish.

Current levels of peer production seem durable and sustainable. For years there has been continued growth in usage, creative output, sophistication, and scope. It can be foreseen that peer production will be the dominant content creation model on the Internet by roughly 2010-2012 if trends continue.

4 Discussion on peer production, net neutrality, and copyright law

The illustrative case studies in the previous section provide insight into the growth and further disruptive potential of peer production. In the next sections, a broad view of peer production is taken and linked to historical precedents of price discrimination and 1,000 years of history in property rights.

[7] `http://www.roughtype.com/archives/2006/07/jason_calacanis.php`

Table 2
Revenue per MByte of traffic for various services.

Revenue per MByte	Service
$1000	Wireless texting
$10	Wireless voice
$0.1	Wireline voice
$0.01	Residential Internet
$0.0001	Backbone Internet

4.1 Convergence, net neutrality, and innovation

A single converged infrastructure can now be used to offer services such as TV, voice, and data. The Net neutrality debate revolves around not favoring one service or service provider over another and not charging service providers extra. The rich history of price discrimination which goes back centuries is often neglected in this debate. Numerous historical precedents of price discrimination are linked by Andrew Odlyzko in (Odlyzko 2008) to the current net neutrality debate, combined with general observations of the telecommunications industry. This section extends this work and enhances it with peer production insight.

Key to the existing practice of price discrimination is Table 2 which shows service revenues. The revenue levels of wireless texting are many orders of magnitude higher than transporting bulk Internet backbone bits. Keeping this structure in place is deemed essential by Internet providers. "The key lesson here is that *legacy* service providers resist the pressure to become mere bit pipes" (Crowcroft 2007).

The telecommunications industry main driving forces seems to be duplicating the unique revenue levels of wireless texting, causing three structural industry inefficiencies. The industry i) fails to focus on its core connectivity business and is not trying to satisfy proven consumer-demand in a cost-efficient manner; instead billions are invested and often wasted ii) in developing unproven new services and iii) technologies to adjust the network to applications.

Evidence for the first cause of inefficiency is provided by the pioneering French broadband provider Free.fr.[8] The industry is claiming that net neutrality would prevent them from building "horrendously" expensive broadband networks. However, broadband networks have proven to be inexpensive to build, which can be called *the French lesson*. Free.fr is offering unlimited high-speed Internet access at a fraction of the cost of competitors. Key is that they still

[8] `http://en.wikipedia.org/wiki/Free_(French_ISP)`

obtained the highest earning before interest and taxes in the EU telecommunications industry. Free.fr combines unbundled lines, dark fiber, vertical integration, their own DSLAMs, a simple network structure with a disruptive flat-fee of 30 Euro per month for 28 Mbps broadband, IPTV, and unlimited international VOIP calls. To reduce cost they build and fine-tune their own set-top boxes and developed a simple backend billing system. [9]

The second cause of telecommunications industry inefficiency is the continuous attempt to compete with end-user service providers using expensive inferior technology and focus on client/server content delivery while ignoring P2P. The lack of successful innovations in the area of end-user service development by the telecommunications industry has been analyzed in detail (Odlyzko 2004). They failed to predict and exploit innovations such as Youtube and Facebook. Companies such as Comcast, Cox, and others even block P2P as measured by project Glasnost. [10] The "excessive usage of broadband" by P2P users is said to interfere with the service provided to others. This is again an example of industry failure to see what lies ahead in terms of consumer-demand and required network capacity. All the top three revenue sources of Table 2 are under threat of being disrupted due to upcoming innovations in P2P platforms. P2P platforms offering (wireless) voice and texting are maturing and are feared by the industry. [11] Future dual-mode GSM-wifi handsets with low-cost municipal wireless Internet would likely win in a fair competition on merits.

The third cause of inefficiency is the industries deep rooted love for historical circuit switching with quality guarantees and resentment of best-effort packet delivery from the Internet-era. Experiences over the past decades have indicated that it is more economical to adjust applications to best-effort networks instead of building smart networks with quality guarantees. Peer production, P2P technology, and dumb networks form a natural combination offering light-speed interconnectivity at the center, intelligence at the edges, and scalability to include all humans. Detailed studies such as (Yuksel et al. 2007) show that overprovisioning a best-effort network requires extra network bandwidth when compared to using "flows" (e.g. circuits) with QoS. However, this study and others often fail to accurately model the economics of saving some bandwidth. Overprovisioning may still be economically sound with a consistent yearly growth rate of Internet traffic of 50-60% [12] and inherently more expensive QoS equipment. In short, keeping it dumb is probably smarter.

The success of Free.fr and the public outcry over Comcast P2P blocking shows that transparency and competition work. It is unfortunate that companies are

[9] http://gigaom.com/2007/12/21/xavier-niel-free-fr/
[10] http://broadband.mpi-sws.mpg.de/transparency/results/
[11] http://www.itu.int/osg/spu/presentations/2007/kelly-melody-challenges-opportunities-of-VoIP-1-march-07.pdf
[12] http://www.dtc.umn.edu/mints

allowed to falsely advertise with "unlimited" broadband and can point to vague *fair use* policies to restrict GByte usage and block P2P. To conclude, creating transparency and real competition in the broadband and telecommunications market as a whole are more likely to improve economic welfare then an attempt to enforce net neutrality.

4.2 Content industry, copyrights, and compulsory licensing

A transformation in the content industry due to peer production is a very likely scenario. The content industry is seeing an assault from four directions; i) eroding of copyright fundamentals, ii) the pirate side of peer production, iii) changing attitudes towards copyright, and iv) improvements to darknets.

Peer production and detailed historical analysis is undermining the (economic) justification of current copyright legislation. The dominating view is that without statutory protection of cultural products there would be market failure and under-production. However, peer production is increasingly replacing markets with commons and offers a wealth of competitive products. A historical analysis of over 1,000 years land enclosure and commons property reveals that the right to be included in the benefit stream has recently moved to the background and hurts innovation (Ford Runge and Defrancesco 2006). An economic analysis reveals that EU legislation contains three assumptions which are not thta valid anymore (Towse 2005). These assumptions are that there is no clash between content publisher interest and author interest, that only publishers are capable of entrepreneurship, and that firms innovate instead of creators. The value added by publishers is diminishing due to P2P platforms, as every talented creator can now independently reach an audience of millions and can pioneer a new business model.

Figure 2 and related work (Byers et al. 2004) show that copyrighted works are widely available on P2P. This is the key driver for expansion of the US copyright statute. The US copyright statute is "biased towards the copyright industry groups who have largely written them to serve their interests" (Samuelson 2007). A proposed "copyright czar" at the US White House level would expand these laws further (Ahrens 2008). These laws are increasingly disconnected from the real world where, for instance, over a third of PC's worldwide have a P2P client installed. [13] New legal actions against P2P users in the US and the recent award of $111 million in damages in the Torrentspy P2P case indicate the extreme different viwepoints from a legal perspective and a common user perspective. Support for this legislation is eroding. A study of over 6000 people in Finland showed that the legislative maximum

[13] `http://arstechnica.com/news.ars/post/20071227-one-third-of-pcs-prefer-limewire.html`

penalty for file sharing is not matching the participants sense of justice.[14] By 2010 copyright laws will become fully unenforceable due to the continued mass usage of file sharing, lack of countermeasures, and availability of *darknets* (Section 3.8).

The choice for the future is between protection of failing business models with stronger laws & enforcement efforts versus legalizing existing practices and de-criminalization of millions. The experience with P2P shows that a switch to the latter is needed and new and sustainable Internet-compatible business models are explored. By creating a new *digital superdistribution right* P2P can be legalized for non-profit prosumer usage in combination with a royalty payment system and compulsory license. Possible compulsory license approaches are examined in (Gervais 2003). Many people agree that copyright reforms are needed (Samuelson 2007) and that the rights of the commons "to be included" should be restored (Ford Runge and Defrancesco 2006).

It is hoped for that peer production will further flourish, laws & business models adapt, and this wild-west period ends. Unfortunately, one scenario is that the telecommunication and content industries will join to halt "innovation at the edge" by blocking competing P2P services and litigation, thereby essentially making them gatekeepers of information by utilizing the same deep-packet inspection technologies that some goverments use to keep the Internet "clean".

5 Conclusion

In this article, the evolution of fully distributed *peer production* on the Internet has been outlined and supported by measurement data and empirical analysis. The website `Slashdot.org` from 1997 represents the starting point of this evolution. On Slashdot it is observed that it takes less then an hour before insightfull comments are recognized by the community. With `Youtube.com` the concept of user-generated videos became mainstream and the concept of a friend is taken to an extreme. The top-10 Youtube users with the highest amount of friends together boast 72,866 "friends". Careful study of such numbers and technology developments leads us to believe that the cardinal features for effective platforms for peer production (P2P platforms) are: separation of good and bad contributions, regulation of computer resources, group communication, and community building. The measurement studies included in this paper indicate that peer production is growing in both usage, creative output, sophistication, and scope. The largest academic measurement to date on the pirate side of peer production shows that copyrighted works are widely

[14] `http://inhiit.blogspot.com/2007/09/p2p-survey-results.html`

available for free download and enforcement of copyright is likely to become impossible by 2010.

The evolution of peer production shows that value creation and innovation is increasingly moving away from the telecommunication and content industries towards "the edges". The current main sources of revenue of both industries will be difficult to protect and uphold. The past decade of evolution shows the power of volunteers with self-governance and loose formation of hierarchical structure based on merit. The next decade is likely to show further growth as we begin to understand how to organise millions into a fully self-managing collective.

References

Adar, E. and Huberman, B. (2000) Free riding on Gnutella, Technical report, Xerox PARC.

Ahrens, F. (2008) House Bill To Create Anti-Piracy Czar Advances, Washington Post, pp. D01, May 1.

Benkler, Y. (2002) Coase's Penguin, or, Linux and The Nature of the Firm, Yale Law Journal, No. 112, pp. 369-446.

Borch, N. (2005) Social Peer-to-Peer for Social People, In proceedings of 1st Conference on Internet Technologies and Applications.

Byers, S., Cranor, L., Cronin, E., Kormann, D. and McDaniel, P. (2004) Analysis of Security Vulnerabilities in the Movie Production and Distribution Process, Telecommunications Policy Journal, No. 28, pp. 619-644.

Cohen, B. (2003) Incentives Build Robustness in BitTorrent, `bitconjurer.org/BitTorrent`.

Cornes, R. and Sandler, T. (1996) The Theory of Externalities, Public Goods, and Club Goods, Cambridge University Press.

Crowcroft, J. (2007) Net neutrality: the technical side of the debate: a white paper, Computer Communication Review, No. 37, pp. 49-56.

Ford Runge, C. and Defrancesco, E. (2006) Exclusion, inclusion, and enclosure: Historical commons and modern intellectual property, World Development, Vol. 34, No. 10.

Gervais, D.J. (2003) The Price of Social Norms: Towards a Liability Regime for File-Sharing Intellectual Property Law Journal, Vol. 12, No. 1, pp. 39-74.

Giles, J. (2005) Internet encyclopaedias go head to head, Nature, Vol. 438, No. 7070, pp. 900-901.

Kamvar, S., Schlosser, M. and Garcia-Molina, H. (2003) The EigenTrust algorithm for reputation management in P2P networks, In proc. of 12th WWW conference.

Leibowitz, N., Ripeanu, M. and Wierzbicki, A. (2003) Deconstructing the Kazaa Network, 3rd IEEE Workshop on Internet Applications, San Jose, CA.

Manjoo, F. (2006) The virtual moneylender, Salon.com, `salon.com/tech/feature/2006/05/22/prosper/index.html`

Musicbrainz (1999) Social contract, `musicbrainz.org/doc/SocialContract`

Nielsen (2006) Teens the biggest spenders on online content, N`www.nielsen-netratings.com/pr/PR_110706_UK.pdf`

Odlyzko, A. (2008) Net neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets, unpublished manuscript, `http://www.dtc.umn.edu/~odlyzko/doc/net.neutrality.pdf`

Odlyzko, A. (2004) Telecom dogmas and spectrum allocations, `http://www.dtc.umn.edu/~odlyzko/doc/telecom.dogmas.spectrum.pdf`.

Piatek, M., Isdal, T., Anderson, T., Krishnamurthy, A. and Venkataramani, A. (2007) Do incentives build robustness in BitTorrent?, In proceedings of NSDI.

Pouwelse, J. Garbacki, P., Epema, D. and Sips, H, (2005) The Bittorrent P2P File-sharing System: Measurements and Analysis, 4th International Workshop on Peer-To-Peer Systems (IPTPS'05), Ithaca, NY.

Pouwelse, J. Garbacki, P., Wang, J., Bakker, A., Yang, J., Iosup, A., Epema, D., Reinders, M., van Steen, M. and Sips, H. (2008) Tribler: A social-based peer-to-peer system, Concurrency and computation: Practice and experience, Vol. 20, No. 2.

Samuelson, P. (2007) Does copyright law need to be reformed?, Communications of the ACM, vol. 50.

Sarkio, K. (2006) Distributed Reputation Management in P2P Based Virtual Communities: Tailored Trustworthiness Estimations, Licentiate Thesis, Helsinki University of Technology.

Schecter, S., Greenstadt, R. and Smith, M. (2003) Trusted Computing, Peer-To-Peer Distribution, and the Economics of Pirated Entertainment, In 2nd Workshop on Economics and Information Security.

Toffler, A. (1980) The Third Wave, Collins Books, London, UK.

Towse, R. (2005) Economics and copyright reform: aspects of the EC directive, Telematics and Informatics, No. 22, pp. 11-24.

Yuksel, M. et al. (2007) value of supporting class-of-service in IP backbones, International Workshop on QoS.

Zittrain, J. (2007) Saving the Internet, Harvard Business Review.