Peer-to-peer Direct Sales

Adam Wierzbicki
Polish-Japanese Institute of Information Technology
adamw@pjwstk.edu.pl
Krzysztof Goworek
Warsaw University of Technology

Abstract

The article describes and gives an economic analysis of a business model for commercial content delivery networks (CDN) based on the Peer-to-Peer model. The content is stored in the CDN on the hosts of the peers. An user pays for access to the content, and can sell the content to other users as in a direct sales network. The content trade is a free market. Transactions (including billing and accounting) are handled by superpeers, who receive a markup for their services and pay the content provider a gratification for every transaction. The system makes use of reputation mechanisms with a goal contrary to most P2P research: to promote content trading and discourage sharing for free. The article compares the profit obtained by the content provider in a client-server CDN and the P2P CDN, and analyzes the stable-state prices in a P2P CDN.

Keywords

Peer-to-peer, content delivery, direct sales, economic model, free market, reputation

Introduction

Lack of business models for P2P systems has been frequently discussed as a factor that limits commercial applications of the peer-to-peer model. On the other hand, P2P file sharing networks remain popular and continue to develop new technological solutions. The wide use of P2P file sharing can have a real economic impact as well, since it may lead to losses of content provider sales and as a result, to increased prices [1] (although some studies show that there is no such impact [5]).

In this paper, we wish to explore an emerging business model of a commercial P2P content delivery network (P2P-CDN) [22, 23]. We shall attempt to present an economic analysis of such a system, and to compare the profits of content providers in a P2P CDN and a client-server CDN (C-CDN).

We are also interested in seeing whether there exist economic reasons that lead to a distribution of the content for free, or whether such a distribution is not a necessary economic solution and can be prevented by system design.

We will introduce a concrete model of a P2P-CDN. This model follows the design used by most popular P2P file sharing networks (a network with superpeers). We then discuss an economic model of the system and analyze a steady-state solution of this model, discussing the effects of price relations, comparing profits in a P2P-CDN and a C-CDN.

We also analyze systematic mechanisms (such as reputation) that can be used to motivate peers to require payment for the content, instead of distributing it for free. An innovative feature of our P2P-CDN is that it can use the income of a peer as reputation. Such a reputation mechanism has several desirable properties that artificial reputation mechanisms lack. Also, it is used for a different purpose than most reputation mechanisms in the literature of P2P file sharing systems: to encourage trading of content, and discourage sharing of content for free.

The P2P Content Distribution System

In our work, we have assumed that the P2P content distribution system is designed similarly to the most commonly used contemporary file sharing networks. For that reason, we have not considered DHT networks [18,19,9]. Our design resembles closely the BitTorrent network (a superpeer model). It is not the purpose of our work to introduce innovations in P2P system design; rather, we wish to explore the economic foundations for a business model of such a system. In order to discuss economic issues, we required a concrete idea of how such a system would work.

There already exist commercial examples of P2P-CDNs (that use the superpeer model). The Potato System [23] is the closest in spirit to the P2P-CDN considered here; however, it uses a complex licensing model that makes economic analysis difficult.

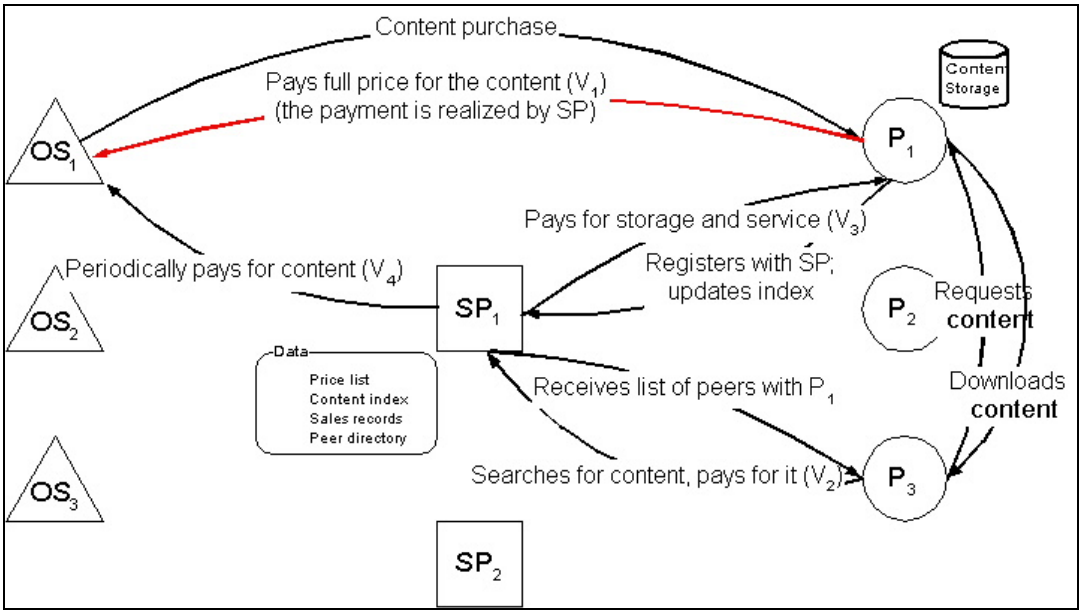

Figure 1 — Model P2P content distribution system; OS_i—Origin Server, SP_i—Super Peer, P_i—Peer

Centerspan [22] is a company that uses Digital Rights Management (DRM) in its P2P-CDN. In our work, we show that increased user participation in legal P2P content distribution is one of the main sources of increased profits for the content provider. We believe that the use of DRM would discourage such participation, besides being ineffective against serious attempts of intellectual property theft. For that reason, we have not considered the use of DRM in this paper, although we do not exclude the possibility.

A schematic view of the P2P content distribution system is presented on Figure 1. Some functions that have been simplified on this diagram (like content purchase and billing) will be described in more detail further on.

The system consists of three main, logically distinct roles: the Origin Server role (also called Content Provider), the Super Peer role and the Peer role. It is worth noting that the OS and SP are functional roles only, therefore they can both be realized by the same physical unit.

Origin Server role. The Origin Servers is a party that provides content that is distributed in the network. It has rights to copy and distribute the content, and is responsible for paying the royalties to the owners of copyrights. Origin Server is interested in distributing the content through a network of peers because it can lead to improvement of QoS without incurring extra costs. OS is paid for the content it sells—it gets a payment for every copy sold directly to a peer as well as a markup for every copy of the content sold by a peer to other peer. There can be many OSs in the market; they compete in the content market by offering wide ranges of content and competitive prices.

Superpeer role. A superpeer (SP) is responsible for bookkeeping in the system and does not trade content. It maintains the following datasets:

- Index of content. A row of the index has the following structure: $\{C_{id}, P_{id}, V\}$ where C_{id} is the id of the content, P_{id} is the id of the peer that offers the content, and V is the price (i.e. the cost of a single download).
- Peer directory. A record of this directory contains $\{P_{id}, P_{addr}, Rep_p\}$, where P_{id} is the id of a peer, P_{addr} is the address of the peer, and Rep_p is current value of reputation of the given peer.
- Billing data. A billing record has the form of $\{P_{s\_id}, P_{b\_id}, C_{id}, V_{PS}, V_{SP}, V_{CP}\}$, where P_{s_id} is the id of seller (a peer or a content provider), P_{b_id} is the id of the buyer, C_{id} is the content id, V_{PS} is the amount to be paid to the selling peer, V_{SP}—superpeer's markup and V_{CP} —the amount to be paid to the content provider.

The last dataset is secret, the others are publicly available.

Peer role. Peers are the clients of the CDN. They can (and are encouraged to) play the role of storage providers. A peer purchases content from an origin server. Content can be purchased in order to get consumed by the peer (which is a client in this case) and/or to become available for other peers. A peer pays OS for downloading the content and it is paid part of the price of the content that is downloaded by other peers from his storage.

Peers join the system by registering with one or more superpeers. Peers can leave the system silently at any

time. To purchase content a peer contacts a superpeer and requests an index of sellers.

Purchase process

The process of purchasing content from a broker peer consists of the following steps:

1. Buyer P_B connects to a superpeer and searches for the content in index held by SP. SP returns a list of peers that store the content. Peers that sell more content will rank higher on the list—peer reputation is a function of income (obtained from transactions brokered by the given SP).
2. The peers can have different content prices and this, along with the value of reputation, influences the choice of content providing peer P_S. Buyer chooses the offer that suits him best and informs the SP of the decision.
3. P_B pays the price for content to SP and sends a request for content to P_S.
4. P_S connects to SP to verify that P_B has paid for the requested content.
5. If SP confirms payment, P_S sends the content to P_B.
6. P_B sends confirmation of the content receive to SP. SP finalizes the transaction by updating its billing record. The payment is divided into three parts: gratification for CP (who is the owner of copyright), SP's markup and P_S's gain.

To avoid the necessity of contact described in steps 4 and 5, one can use the mechanism of nonces. P_S can periodically obtain a list of nonces from SP. P_B receives a nonce from the list after paying for content to SP and sends it to P_S along with the request for content.

The confirmation sent in step 6 can be abused by P_B to prevent payment for downloaded content. To avoid this, SP should follow a two-phase transaction protocol. When the price of content is paid by P_B, SP should keep the money on a separate account until it receives confirmation of correct content delivery. Then, the money is paid to P_S. If P_B complains that delivery is incorrect, SP should request P_S to send the content to SP. SP will then act as a proxy between the peers and forward the content to P_B. SP should also be able to verify content correctness (perhaps using a digital signature). If P_S does not send the content, the money is refunded to P_B. Note that this mechanism is used only for failure situations and should not adversely impact system scalability and efficiency.

Billing, accounting and payment processes

SP is the party responsible for billing the process. It keeps the billing records of transactions between peers. Every transaction between peers must be registered in SP's records. SP calculates the total charges and sends invoices to peers. The payments can be periodic to prevent too frequent exchange of paper or electronic invoices and reduce billing costs. In a system there can be many CPs and many SPs. It is peer who can choose the SP that offers best conditions and SP has an index of CPs it has agreements with. Since peer can be paid for the content that has been downloaded from him, and SP is the broker of these payments, it can pay peer for each download or—as it does to CP—pay on periodic basis (e.g. once a month). No matter which payment system for peers rewarding is used, SP has to keep their records since they are used for calculating their reputation.

Trust and security

Trust management is a central issue in the design of P2P systems [8,12,13]. In the proposed P2P content distribution system, both content providers and peers must trust the superpeer. Let us briefly justify this approach. First, the peer must trust the superpeer to fairly pay for content that the peer sells. The selling peer may trust SP because it can verify billing correctness from its own records and from the prices published by the superpeer.

The content provider must trust SP to correctly account transactions and to fairly pay gratifications for content usage. This can be justified using an audit approach. CP may control a set of peers that are used to audit accounting correctness of an SP. These peers would execute certain transactions, report these transactions to CP and allow CP to verify that SP correctly pays for the audited transactions.

In order to correctly account and bill transactions, SP must have means to authenticate peers [14]. This can be solved using one of several known cryptographic techniques, like PKI certificates or the negotiation of shared secrets between SP and P during the registration of P at SP. The second approach would allow for peer anonymity.

Certain information (like account or credit card numbers) must be transmitted securely in the system. Note, however, that not all transmitted information needs to be encrypted. Existing protocols (such as TLS) allow to implement this requirement.

Economical model of a P2P-CDN

Motivations

The Peer-to-Peer model is proposed for use in a commercial CDN because of some economical benefits that, as the authors believe, all the players receive. OS is encouraged to use a P2P-CDN instead of a C-CDN because of the following benefits:

- Better scalability of the system and higher service availability. Since everyone can easily join the P2P-CDN at no threshold cost (only the content cost has

to be paid), so the network is going to automatically scale to the needs.
- Lower storage and bandwidth costs. The content is stored on extra space on peer hosts.
- Saving revenue by reducing illegal P2P distribution.

A superpeer is a totally commercial market player. Its only motivation is the revenue for providing content and peer indexes, maintaining billing, accounting and payments. A superpeer does not participate in content trade.

As mentioned above, the superpeer role may also be assumed by the content provider, and thus OS and SP may be integrated. A very interesting motivation for the existence of a separate superpeer has been demonstrated by the Potato system [24]. There, the OS is a single artist (independent music author), and SP is a radio station that specializes in independent music. In such a system, SP is also responsible for the promotion of the sold content, and has a separate channel for that promotion (radio). OS is usually too small to launch a commercial content distribution over the internet and to handle all organizational and legal issues.

One of the fundamental assumption of our work is the standard economic assumption of agent rationality: we assume that all agents, including peers, are interested to maximize their profits. A peer is motivated to become a content broker (i.e. a player that resells the content) by the revenue it gets for every download of the content done by other peers. Peers also benefit from increased content availability and quality of service.

Later on, we shall consider the issue of how this assumption can be upheld in contrast with P2P file sharing systems, where peers share content for free. We believe that peers in such systems enter into a complex mutual agreement that they will receive something (other content) for their service. There exist technical, as well as legal reasons (the need for micropayments) why peers choose to share, rather than to resell their content. We believe that in a P2P-CDN these reasons can be removed, resulting in an increased participation of peers in legal content trading.

Prices

Economical characteristics of the described system are similar to characteristics of free market. The parties of every kind compete with each other. The means of the competition are: prices and markups, available content spectrum, bandwidth and quality of service.

The peers can compete with each other by offering competitive prices for their content. A peer that offers lower price for the same content may attract more buyers, but the per-download gain is smaller.

There can be many categories of broker peers. Some peers may decide to offer the most popular content. It is probable that there will be many such peers—so the competition is going to be hard. Others may sell less popular content (niche products). They are not likely to attract as many buyers—so their revenue is going to be smaller, but per-unit gain can be greater since competition will probably not be as hard as in the latter case.

One can distinguish the following areas of competition:

- Content Providers compete with each other by offered content range and prices level.
- Peers who sell content (brokers) compete with each other by content range, download prices, and QoS (bandwidth, availability etc.).
- Content Providers compete with broker peers—every peer who needs a download is free to choose between downloading from a CP or a broker peer.
- Superpeers compete in markup levels, CP market coverage, and broker peer market coverage.

Every party of the market is therefore encouraged to offer best level of service at the lowest possible prices to attract as many clients as possible.

Stable state analysis

Here we assume that the CDN market fulfills all the free market assumptions—the number of players of each kind is large, and they can fully and freely compete with each other. This assumption implicates that in the stable state of the market, no net gain can be achieved by any of the players.

The prices on such stable market are given by the example formulas:

$$V_{CP} = C_{RS} + C_{CP} = C_{RS} + \frac{B_C}{\mathrm{E}\, B_{mth}} C_{CPmonth}$$

where V_{CP} is the single download price offered by CP, C_{RS} is the royalty fee for "first download" (which gives a right to redistribute the content), C_{CP} is the CP operation cost, B_C—volume of content, $\mathrm{E}B_{mth}$—expected monthly download volume, and C_{CPmth} is the monthly operation cost of CP (storage, network access, staff, and marketing).

$$V_{SPtran} = \frac{C_{SPmth}}{\mathrm{E}\, N_{tran}}$$

where V_{SPtran} is the per-transaction price offered by a superpeer, C_{SPmth} is the monthly operation cost of SP, and $\mathrm{E}N_{tran}$ is the expected number of transaction per month.

$$V_{BP} = C_{RE} + V_{SPtran} + C_{BP} = C_{RE} + V_{SPtran} + \frac{B_C}{\mathrm{E}\, B_{mth}} C_{BPmth}$$

where V_{BP} is the download price offered by a broker peer, C_{RE} is the royalty for every P2P download, C_{BP} is the operation cost, B_C—volume of content, $\mathrm{E}B_{mth}$—expected monthly download volume, C_{BPmth}—monthly operation cost of a BP.

We assume here that all the costs are accounted on monthly basis, and that both CPs and broker peers decide to split their operation costs according to volume of the content. The value of C_{BPmth} can be zeroed if peer's hard disk is large enough to keep the offered content (and he does not need the space for any other purposes) and his network access is charged with a flat fee (is independent of traffic volume). If we assume that broker peers do not treat content reselling as a commercial activity, but as a way to lower their own content consumption costs, free market rules will stabilize their markups at a very low level (because there is no market-entry cost other than download cost so anyone can join the market).

Lack of extra costs leads to a very cheap way of content distribution. C_{RS} and C_{RE} are the only fixed costs (but they do not make a cost of distribution). All the other per-unit costs become smaller as the scale of the system increases. If we assume that the peers don't act commercially, which means they don't buy bandwidth and storage in order to provide better services, on a highly competitive market the prices they offer can asymptotically reach $C_{RE}+V_{SPtran}$ level.

Price ratio analysis

As shown above, the price of content in the system generally consists of three parts—CP's costs and copy royalties, SP's costs and peer's cost. The analysis of impact of price levels on probable peer behavior gives interesting results.

The C_{RE}/C_{RS} ratio is especially interesting from our point of view, as it may have serious impact on peer tendencies for free content redistribution. The lower the ratio is, the more a broker peer (BP) is encouraged to resell the content. He/she had to make a serious investment to buy the content from a CP and can sell it for quite a low price (because C_{RS} is rather low). So it is possible that a stable strategy would be: to buy content at price V_1 and then to try to resell it to as many peers as possible, at price V_2, where $V_2 \approx V_1/\mathrm{E}N$, E$N$ being the expected number of buyers. Other peers may decide to buy content from BP, because the price he/she offers is lower than V_1. The drawback is that copyright owner may disagree to use such pricing model (with quite high start payment and rather low per-copy payments). It is probably less attractive to copyright owners—the price of first download cannot be very high (no peer would want to invest in such content), and the per-copy price should yet be considerably lower. It is possible that the total revenue is lower than in a model with C_{RE}/C_{RS}=1. But in the latter model the tendency to illegal content sharing may appear to be very strong, which makes total royalties even lower. So perhaps it is better to accept lower royalties with only a little piracy than potentially high (but less probable) ones. In other words, content providers in our model must make the decision: is it better to have all of nothing or a little of everything?

It can be easily seen than if SP per-transaction price is high (because of relatively high operation costs) the system is unlikely to have wide coverage. Every single download transaction is expensive, which can lead to collapse of the system. We believe that the fixed cost of SP operation can be relatively high, while the incremental costs are low. So there probably exists a "critical mass" of expected transactions number above which they can be cheap—the CDN operation is likely to get cheaper as the network grows.

If peer's costs are relatively large, he/she is discouraged from freely distributing the content, but he/she is also discouraged from buying relatively expensive content from other peers. In such a case, our model can degenerate into a system with a centralized content distribution from OS to all peers.

Peer-to-Peer CDN vs. Centralized CDN

Here we compare the costs and revenues that CPs and peers have in two models: Centralized Content Distribution Network (C-CDN) and P2P Content Distribution Network (P2P-CDN)—the model described above. In order to perform the comparison, we shall assume that the per-download price from the origin server in the C-CDN is the same as in the P2P-CDN (V_{CP}).

The total revenue achieved by a CP in C-CDN is given by the formula:

$$R_C = B_{Cmth} V_{CP}$$

where R_C is the total revenue (in unit of time—a month in our example), B_{month} is the total volume of downloads and V_{CCP} is a per-byte download price. We use the simplest possible tariff model that assumes a constant price per byte of download. In our analysis there are no essential differences between pay-per-download and pay-per-byte models.

In P2P-CDN, CP's revenue is:

$$R_P = hB_{Pmth} C_{RE} + (1-h) B_{Pmth} V_{CP}$$

where R_P is the revenue, B_{Pmth} is the total volume of downloads (per CP), C_{RE}—royalty-fee (for every downloaded byte in the system) and V_{CP}—price for a byte of download from the central server. $h \in \langle 0;1 \rangle$ is the "hit-rate" of broker-peers (i.e. the fraction of downloads that do not involve CP). Example values of h can be obtained from studies of caching of P2P traffic. The closest to our definition of h is the study of [24], where authors used a "passive peer" instead of a cache and reported values of h above 80%. This observation is supported in [25].

The argument that our P2P model encourages users to stay legal, allows using ρ coefficient, where $\rho = \frac{N_P}{N_C} > 1$. Example values of ρ are hard to estimate, but consider that the percentage of US downloaders who paid for music online increased from *8* percent to *22* percent in the first 12 months after the launch of iTunes [21]. This would give a value of ρ=*2.75*.

The costs are given by, respectively:

$$C_C = C_{CP}(B_{Cmth})$$
$$C_P = C_{CP}\left((1-h)B_{Pmth}\right)$$

where C_{CP} is the operation cost (cost of storage, bandwidth and other enterprise costs)—it is a function of downloads volume. If we assume that C_{CP} is proportional to volume of downloads[1].The above give the gain formulas (after simplifying):

$$G_C = B_{Cmth}\left(V_{CP} - C_{CP}\right)$$
$$G_P = h\rho B_{Cmth}\left[C_{RE} + \frac{1-h}{h}\left(V_{CP} - C_{CP}\right)\right]$$

Note that if $\rho = 1,\ h = 0$, then G_C=G_P.

The introduction of the ρ and *h* coefficients simplifies a comparison of content provider profits. We wish to answer the question: when is $G_P \geq G_C$?

The P2P-CDN has profits that are not less than in the C-CDN under the following condition:

$$\boxed{C_{RE} \geq \alpha\,(V_{CP} - C_{CP}) = \alpha C_{RS}}$$

where $\alpha = \frac{1}{\rho h} - \frac{1-h}{h}$. In other words, the royalty for every download must be sufficiently high in relation to the royalty for the first download. However, the obtained relationship allows to draw two conclusions: first, that it is possible to obtain the same profit in the P2P-CDN as in the C-CDN. For our example values of $\rho = 2.75, h = 0.8$ the value of the coefficient $\alpha \approx 0.2$. This means that the royalty need not be prohibitively high (recall our discussion of price ratios that this has an impact on the participation of trading peers in the system). The second conclusion is that as ρ and *h* increase, the royalty can decrease – and this can form a self-enforcing relationship with the number of peers in the system. It is also worth noting that *h* is likely to grow with number of peers—the more broker peers in the system, the more files can be found on their hosts, without need for incurring CP costs.

User costs and profits are given by the formulas:

$$g_C = -V_{CP} - c$$
$$g_P = -V_{CP} - c + \mathrm{E}\,n(V_{BP} - C_{RE})$$

where g_C is the gain of C-CDN user and g_P is the gain of P2P-CDN broker-peer[2], E*n* is the expected volume of downloads of the given content from the peer, C_{RE} is per-byte royalty fee, and *c* is the per-byte storage and bandwidth cost.

If we suppose that $V_{PCP} = V_{BP}$, i.e. a peer sells the content at the same per-download price he/she has bought it, the price formula is:

$$V_{BP} = \frac{n}{n-1}C_{RE} + c + g_P$$

On a stable free market, the net gain is close to zero. Therefore, in a large P2P environment the per-download price of the content gets close to real download costs ($C_{RE} + c$) as $\lim_{n\to\infty}\frac{n}{n-1} = 1$. However, if a peer has got rare content in his/her storage (demand for which is greater than its supply), he/she can easily earn money from reselling it.

Is free content sharing an economic necessity?

One of the questions we have wanted to answer is whether there exist economic reasons for the sharing of content for free in contemporary P2P file sharing systems. The economic analysis of our model shows that in a steady state content would not be distributed for free, but at a price that is determined by operation costs.

A subject that has been ignored by our economic analysis is the cost of financial transactions. If the operation costs for content distribution are low, then the financial transaction cost may be higher than the content price. In our system, we have tried to avoid this case by introducing periodic, aggregated billing. However, in real life this remains a concern. (In the Potato System, the share of financial transaction and code license cost in the total price is three times higher than the share of the other system costs [23]). Additionally, the commercial distribution of content on a P2P basis may require the use of micropayments, which is still a technological obstacle.

Another explanation of free content sharing is that the free market model used in our work is inappropriate. Perhaps a file sharing network is better modeled as an evolving, cooperative strategy in a game of selfish

[1] It may not be the case in reality, but it is convenient for our considerations and does not influence the conclusions

[2] We do not consider the value content represents for the user (utility)—gain is a financial value only.

players (such as the Prisoner's Dilemma). Indeed, many authors in the field have made that assumption [8,10,11]. If that is the case, then perhaps it could be possible to use systematic mechanisms that function as incentives for content trading (and disincentives for free content sharing). Such mechanisms could change the conditions of the game in such a way that an evolution of a cooperative strategy (sharing for free) becomes impossible. In our system, we propose to make use of reputation mechanisms with a goal that is contrary to most work in P2P systems: promote content brokering, discourage sharing for free.

Reputation model

In this section, we describe a reputation mechanism that can be used to promote content brokering by peers (instead of free sharing).

Reputation in the given system should:

- Promote best peers. If reputation were based on content-buyer grades, the system would be sensitive to cheating. A common problem is that peers do not want to issue grades.
- Encourage peers to register in a few local superpeers—registering in more than one SP leads to providing better information for potential buyers. A peer should register in local SPs only, since it makes the network traffic local and reduces unnecessary core network load. Promoting local transfers helps achieve better quality of service as well.
- Discourage peers from distributing the content for free (illegally—we assume the content is commercial and royalties must be paid).

The reputation of a peer in the system is determined by its income. The more money a peer earns for selling content (as shown in the bookkeeping of a superpeer), the higher the reputation of that peer is. The reputation of a peer known to one of his superpeers is based only on income accounted by this SP. This fact can influence the registration decisions of a peer. If a peer decides to stay "loyal" to a single SP, its whole income is accounted only by the one SP, which results in higher reputation. Registering in many SPs could make peer's content visible to more potential downloaders and result in higher total income. However the reputation on each SP is lower which can lead to less downloads.

Traffic can be kept local using reputation if the function Reputation(Income) is convex, so that a peer decides to register with only few superpeers. On the other hand, if a peer registers with only one superpeer, content availability may suffer. So it seems that the relationship between Reputation and Income should have a shape as shown on Figure 2:

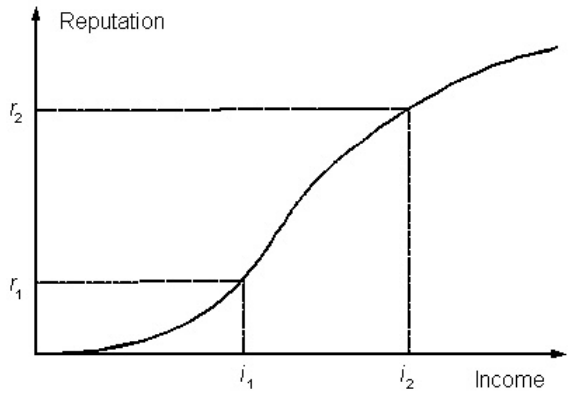

Figure 2 — Desired Reputation(Income) function shape

Reputation based on income reveals information about the real utilities (Quality of Service) obtained by peers during previous interactions. This information is more complete than it could be in a reputation system based on artificial reputation. The main drawback is that a peer gets only the final value of the reputation—it cannot obtain the values of the criteria that influence reputation in which case it could assess the content provider using own weights of quality, price and availability criteria.

As stated before, the reputation could influence the V_{CP} and V_{SPtran} prices. A peer with high reputation can get lower prices for both content and transaction processing. High reputation means that a peer can be trusted by CPs and SPs it had many transactions with. CP and SP can project high level of future revenues from the peer and may want to encourage him/her to stay loyal and buy even more content. Such super-broker peer is likely to care about his reputation level as it gives him serious benefits.

The income accounted by superpeers can be used in one more way. Periodically, all superpeers remove from their directory all peers that have zero income over a configured time period. This mechanism is used in order to prevent peers from distributing the content for free.

Conclusion

We have analyzed the economic foundations of the emergent business model of a commercial P2P content delivery network. An analysis of a free-market model of the system has shown that peers in the system would not share content for free, but at operation costs. A comparison of the content provider profits has shown that the content provider can obtain the same profits in a P2P-CDN, provided that the use of the P2P-CDN model will increase user participation. Our economic analysis shows also that as peer participation increases, the revenues per download paid by every peer can decrease, while the same level of profit is maintained.

We have shown that the P2P content distribution system can be equipped with mechanisms such as reputation based on income, or removal of peers with zero income over a time period. These mechanisms motivate peers to obtain a fair price for their content. In our work, we use reputation with a goal that is contrary to most P2P research: to promote content brokering, and discourage sharing of content for free.

One of the conclusions of our analysis could be that the current, free distribution of content in file sharing networks may be a consequence of inadequate technological and organizational support for payment and accounting of electronic transactions.

Future Work

We consider the following topics the most important for future work:

- An analysis of impact of the income-based-reputation formula and its parameters on the behavior of brokers and buyers.
- An analysis of financial transaction cost on the steady-state solution of the economic model.
- Answering the question of whether micropayments will be needed in our system.
- Investigating whether income-based reputation leads to the creation of small-worlds of bestselling peers.
- Considering the impact of removal of zero-income peers on content availability.

References

1. H. Varian, *The Social Cost of Sharing*, In Workshop on Economics in Peer-to-Peer Systems, 2003
2. P. Antoniadis, C. Courcoubetis, R. Weber, *An Asymptotically Optimal Scheme for P2P File Sharing*, In Second Workshop on Economics in Peer-to-Peer Systems, 2004
3. A. Fuqua, T Ngan, D. Wallach, *Economic Behavior of Peer-to-Peer Storage Networks*, In Workshop on Economics in Peer-to-Peer Systems, 2003
4. M. Adler et al, *Optimal Peer Selection in a Free-Market Peer-Resource Economy*, In Second Workshop on Economics in Peer-to-Peer Systems, 2004
5. F. Oberholzer, K. Strumpf, *P2P's Impact on Recorded Music Sales*, In Second Workshop on Economics in Peer-to-Peer Systems, 2004
6. R. Cole, Y. Dodis. T. Roughgarden, *Pricing Networks with Selfish Routing*, In Workshop on Economics in Peer-to-Peer Systems, 2003
7. Cohen B. *Incentives build robustness in Bittorrent*. In Workshop on Economics in Peer-to-Peer Systems, 2003
8. K. Lai, M. Feldman, I. Stoica, and J. Chuang. *Incentives for cooperation in peer-to-peer networks*. In Workshop on Economics of Peer-to-Peer Systems, 2003
9. Zhao, J. Kubiatowicz, A. Joseph, Tapestry: *An infrastructure for fault-resilient wide-area location and routing,* Technical Report CSD-01-1141, U.C.Berkeley, 2001
10. Tran T. *Reputation-Oriented Reinforcement Learning Strategies for Economically-Motivated Agents in Electronic Market Environments*, PhD Thesis, University of Waterloo. 2003
11. M. G. Zimmermann, Victor M. Egufluz and Maxi San Miguel. *Cooperation, adaptation and the emergence of leadership*. In `Economics with Heterogeneous Interacting Agents', pp. 73-86, A. Kirman and J.B. Zimmermann (eds.), Springer, Berlin. 2001
12. Gmytrasiewicz, P.J. and E. H. Durfee. *Toward a theory of honesty and trust among communicating autonomous agent.* Group Decision and Negotiation 1993.2:237-258
13. L. Mui (2003) *Computational Models of Trust and Reputation: Agents, Evolutionary Games, and Social Networks,* Ph.D. Dissertation, Massachusetts Institute of Technology
14. J. Douceur, *The Sybil Attack,* In Proc. of the IPTPS02 Workshop, Cambridge, MA (USA), March 2002
15. Gnutella/ng, World Wide Web page, http://mangocats.com/annesark/gnutellang/wego_pages.html, 2000
16. Freenet, World Wide Web page, http://freenetproject.org/cgi-bin/twiki/view/Main/WebHome, 2002
17. Mojo Nation, World Wide Web page, http://www.mojonation.net/, 2000
18. Stoica, R. Morris, D. Krager, M. F. Kaashoek, H. Balakrishnan, *Chord: a scalable peer-to-peer lookup service for internet applications,* Proceedings of ACM SIGCOMM'01 Conference, 2001
19. P. Druschel, A. Rowstron, *Pastry*: *Scalable, distributed object location and routing for large-scale peer-to-peer systems*, Proceedings of the 18th IFIP/ACM International Conference on Distributed Systems Platforms (Middleware'01), 2001
20. G. Cohen, *Building a Successful P2P Business Model*, Proc. O'Reilly Peer-to-Peer and Web Services Conference, Washingto, 2001
21. M. Einhorn, B. Rosenblatt, *Peer-to-Peer Networking and Digital Rights Management. How Market Tools Can Solve Copyright Problems*, Policy Analysis, no. 534, February 17, 2005
22. *Mediated Peer-to-Peer (P2P) Enables Viable, Cost Effective Delivery of Digital Entertainment*, Whitepaper, Centerspan, 2002
23. *The Potato System*, Whitepaper, Potato System, http://www.potatosystem.com/info/eng/index.html
24. A. Tagami, T. Hasegawa, T. Hasegawa, Analysis and Application of Passive Peer Influence on Peer-to-Peer Inter-domain Traffic, Proceedings of Fourth International Conference on Peer-to-Peer Computing (IEEE P2P'2004), Zurich, August, 2004
25. A. Wierzbicki, N. Leibovitz, M. Ripeanu, R. Woźniak, Cache Replacement Policies For P2P File Sharing Protocols, European Transactions on Telecommunications, special issue on Peer-to-Peer systems, November 2004, (15:6), 559-569