Market Models for P2P Content Distribution *

P. Antoniadis and C. Courcoubetis

Department of Informatics
Athens University of Economics and Business
76 Patision Str., Athens, GR 10434, Greece
{antoniad,courcou}@aueb.gr

Abstract. The new p2p networking paradigms offer new possibilities for content distribution over the Internet. We propose a model that treats peers as independent economic agents buying and selling content and investigate the basic economic properties of such a market managed p2p content distribution network. Initially, we assume that no peer has the content, and there is a substantial initial cost to bring it within the peer group. The bargaining position and hence the price that can be posted by an agent having the content depends on the cost to transport the content to the requesting peer, its value, and the number of other agents providing the same content. We discuss the influence of parameters such as the maximum number of the competitive offers allowed by the system, content popularity, its value to the agents, and the transport costs, taking into account the risk of the first agent incurring the initial content cost.

1 Introduction

In contrast with existing content distribution paradigms such as multicasting, caching and content distribution networks (CDNs), in the p2p model the content is delivered in a fully distributed fashion. Each peer after receiving the content can act as a content provider making the content distribution much more efficient due to the lack of central management. As initially conceived, p2p networks assumed insignificant costs in terms of obtaining and transporting content and deployed minimum control mechanisms.

In such a simple model, an agent that needs the content will always obtain it, and will make it freely available to others. But content may be costly to obtain, since the bandwidth of the network may be restricted at places causing delay costs over particular distribution routes. Furthermore, making content accessible to others reduces a peer's access bandwidth and hence degrades his network access performance. If such costs become significant, the altruistic community spirit may not be enough to sustain the successful operation of the system. Deploying a market mechanism within the peer community may solve these problems and reduce unnecessary waste of resources resulting from peers requesting content for which their value is less than the cost imposed to others. This approach is

* This work is partly supported by EU's Fifth Framework project 'Market Management of Peer-to-Peer Services' - MMAPPS (RTD No IST-2001-34201)

consistent with recent studies of the behavior of popular p2p file sharing applications such as Napster and Gnutella [1] [2] which conclude that free riders consist the majority of the participating peers of these systems. This behavior is allowed where no central management or market mechanisms exist to ensure the fair and economically justified allocation of resources. Recent game theoretic analysis [3] indicates that without some market mechanism in place there exist no suitable incentives for peers to contribute to the p2p community.

Market mechanisms lie somewhere between existing centralized solutions and the current unmanaged p2p systems. Peers are modeled as economic agents having incentives for obtaining revenue from distributing content. We may assume that there is a common cost W to bring the content to a p2p group, resembling to the case of a copyright cost or to the case where there is a substantial cost for transporting the content to any peer from its external source. We also assume that there is some intra-group cost for transporting the content, different in general for each pair of communicating peers and for each service direction. This can model delays or link performance degradation when bandwidth is a scarce resource, as in the case of network access.

Each time an agent requests a specific content we assume there exists a lookup service that provides cost and price information and can control the number of competitive offers allowed for each request. We compare two types of content distribution paradigms:

1. A *restricted competition* model, where at most t agents are randomly selected from the set S of all candidate content providers. This is an oligopoly model, where prices will depend on the number t of competing peers.
2. A *privileged provider* model, where the first agent to acquire the content[1] and hence to incur the initial cost W is granted the privilege to be the unique seller of this content within the group. The above agent has the power of a monopolist that can do personalized pricing.

In this paper we discuss the influence of restricting content availability information to the spreading of the content in the peer community. Granting different market power to agents selling the content affects their expected profits, and hence influences their decisions about buying the content and further reselling it. Allowing a highly competitive market has the positive effect of bringing prices low and resulting in higher social welfare since more agents will finally buy the content. On the other hand, it is not obvious that excessive competition will always benefit the system as a whole. Obtaining information about the quality of the content has always a positive effect by stimulating demand since low prices make it difficult for the agent that quires the content to recover his cost, and hence reduce his motivation to do so. It is hence reasonable to assume that the end result may depend on the size of the system and on the costs of initially acquiring the content and further distributing it within the group compared to the value of the content to the agents.

[1] Since this agent has different characteristics from the rest of the group we will assign him index 0.

We propose a simple economic model that takes into account all the above parameters and results in tractable analytic solutions of the underlying games. It can be used to analyze the sensitivity of various performance metrics such as the social welfare and the expected net benefit of the agents with respect to parameters such as the initial cost, the value and popularity of the content, the available information, the transport costs, etc. These parameters affect the prices and the expected net benefit of the agents which in turn influence their decisions and the resulting content distribution.

2 A simple p2p content distribution market model

Lets assume a group of N peer agents which may request some content, each agent i obtaining value u_i by using the content, where u_i is uniformly distributed in $[0, V]$. We assume that initially the content belongs to some external provider who charges a substantial fee W, high enough relative to V. So, once some agent purchased the content, it makes more sense for this agent to resale the content at lower prices within the group in order to recover his cost and even make some profit, instead of more agents paying this high fee to the content provider. This potential profit allows for agents with $u_i - W < 0$ to take the risk and purchase the content. The larger the size of the group, the greater $\max_i\{u_i - W_i\}$ will be, hence increasing the probability for such a decision. Clearly, the desicion whether to pay this fee W depends on the estimate of the average revenue that can be obtained by reseling it, and on the degree of agents' risk aversion.

On this level there is an interesting game to be played. Each agent waits for another agent to do this first move and pay the fee W. This 'free rider' problem may in some cases result in low probability of purchasing the content. We discuss next the part of the model that describes the internal market of the p2p group, which is used to compute the average revenue obtained by content resale.

An important aspect of our model is the cost associated with moving content around. We denote by c_{ij} the cost for moving content from agent i to agent j. This may be direct communication cost such as delay, information loss or payments to transport service providers, or indirect cost corresponding to performance deterioration of the access service of the agent giving away the content. Such a cost will be charged to the agent purchasing the content in addition to the payment for the content itself. Depending on the situation, c_{ij} may be known a priori only to j, or only to i, or to both. For instance, agent j may get back from the network the identity of agent i and the average throughput of a connection to this agent, while the agent i may only know that there is a potential customer. This uncertainty and asymmetry about the costs is crucial for deriving positive revenues. For instance, if agent j requests the content and there are two competing providers i and k which already have the content, where $c_{ij} = c_{kj} = 0$, then the equilibrium price is zero [4]! On the other hand, if costs are known to all parties involved and $c_{ij} > c_{kj}$, then k will win posting the price $c_{ij} - c_{kj}$. If the distribution of these costs is random with a known distribution, then one may calculate the average revenue from such a transaction. Our model allows

various types of information asymmetry: costs may be random and known only by the customers, or each seller knows only his cost.

There are other important parameters such as content popularity, which can be modeled by the probability of an agent to request the content or by the number of agents that will eventually request the content (the equivalent size of the group with respect to the above content). Reputation may be modelled by associating different values for the same content obtained from different sellers. For instance, the value of agent i may be u_{ij} where j is the seller of the content.

2.1 Calculation of social welfare and expected revenue

We discuss next the market models and the resulting prices and revenues corresponding to different levels of competition. As we show, the degree of competition influences the agents' decisions and the social welfare of the system. Our goal is to convince the reader that our model leads in many interesting cases to analytically tractable solutions, which we hope that can provide better insight to the sensitivity of the performance with respect to the key parameters introduced.

The case of a privileged provider

This is the extreme case of a monopoly, where agent 0 is granted the exclusivity right to resale the content after purchasing it for the initial price W. This agent, given the available information, must choose the optimal price w to charge for the content in order to maximize his expected revenue $\bar{R}_m$ obtained by selling the content to the peer group. Then he must compare this expected revenue to $W - u_0$ and decide if it is worth to him to undertake this effort.

We briefly describe how one may compute the expected revenue from resale. Let c_i be the cost to transport the content to agent i from agent 0. We assume for simplicity that the c_{ij}s are iid with uniform distribution in $[0, C]$. We also assume that the value of c_i is made known only to agent i that requests the content[2] whereas its distribution is known to agent 0. Similarly, the value u_i of the content to agent i is iid and uniform in $[0, V]$. Let w the price that agent 0 uses to respond to the request of agent i. If $u_i < w + c_i$, then agent i will refuse to buy the content resulting in zero revenue for agent 0. The optimization problem faced by agent 0 becomes

$$\max_w R(w) = wP[u > c + w]\,, \tag{1}$$

where for simplicity we have omitted the subscript i. By solving (1) we obtain

$$0 \le C \le 2V/3 \,:\, w^* = V/2 - C/4,\, R^* = \frac{[V - C/2]^2}{4V}\,, \tag{2}$$

$$2V/3 \le C \le V \,:\, w^* = V/3,\, R^* = 2V^2/27C\,. \tag{3}$$

The total average revenue $\bar{R}_m(N)$ is now obtained by multiplying R^* by N.

[2] One may make different assumptions about the disclosure of such information.

To compute the social welfare we must compute the average value of a peer that accepts such an offer, and then multiply it by the average number of such peers. In order to do this, one must compute the conditional distribution of the value of a peer given the fact that he accepted price w^*. The average number of peers that will eventually accept the offer is $NP[u > c + w^*]$. This allows us to compute the social welfare, given that the content is purchased, $SW_m = NP[u > c + w^*]\frac{7C^2+12CV-36V^2}{24(C-2V)}$. Multiplying with the probability that at least one agent will finally decide to initially buy the content, $P_b = [1 - (1 - P[u > W + \bar{R}_m(N)])^N]$, we can compute the expected social welfare $S\bar{W}_m = P_b SW_m$.

The case of restricted competition

In this case we allow t peers to compete for every content request. These are chosen randomly among the peers that have already purchased the content. We consider the illustrative case where $t = 2$. To simplify notation consider the case where peers 1 and 2 compete for a given request by some third peer. Let c_i and w_i, $i = 1, 2$, be the costs and prices corresponding to each of these peers.

Similarly to the priviledged provider case we can calculate the optimal price $w_1^* = \frac{C}{2}$ and the expected revenue $R_1(w_1^*) = \frac{C}{4}P[w_1^* + c < u] = \frac{C(V-C)}{4V}$ of each transaction. [3]

Lets evaluate now the revenue of agent 0. The first time he will resell the content he will use (1) since he is the unique agent that can sell the content. After that he will always use price w_1^*. How much revenue will he make in total? Assuming that agents are randomly chosen among the set S of candidate providers, the probability that agent 0 is chosen in the competing pair is $2/|S|$.

Conditioned on the event that M agents will eventually buy the content, the expected revenue of agent 0 in this duopoly case is

$$\bar{R}_d(M) = \frac{[V - C/2]^2}{4V} + \frac{C}{2}\frac{1}{2}\sum_{k=2}^{M}\frac{2}{k} \approx \frac{[V - C/2]^2}{4V} + \frac{C}{2}[\gamma + \ln M - 1]\,, \qquad (4)$$

where γ the Euler-Mascheroni constant and $M = NP[w_1^* + c < u]$.

Similarly to the case of the privileged provider we can compute the total value generated in this system. For instance, in the case where $C \leq 2V/3$, $SW_d = N\frac{13C^2-12V^2}{24V}$ and $S\bar{W}_d = [1 - (1 - \frac{V-W+\bar{R}_d(M)}{V})^N]SW_d$.

3 Evaluation of results

Below we present some initial results we have obtained comparing the above two cases. The expected revenue of the restricted competition case is significantly reduced in comparison with the priviledged provider case, as it increases logarithmically with N. Thus, the probability that an agent will decide to initially

[3] For simplicity we have assumed that $V >> C$ and hence computed an approximation of (higher than) the optimal price.

take the risk and purchase the content is also reduced. However, if this probability is high enough, when W is close to V, the expected social welfare of the system is significantly higher (see Figure 1) since prices in this case are lower and more agents will finally buy the content.

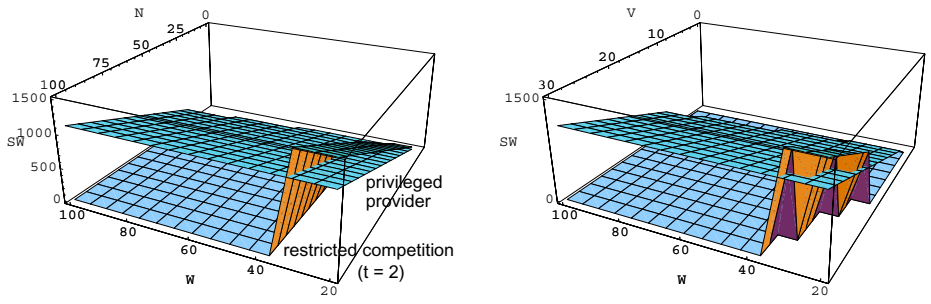

(a) As a function of (W,N) (b) As a function of (W,V)

Fig. 1. Comparison of the expected social welfare between the two cases

4 Conclusions and future work

In this paper we have presented some ongoing work regarding performance evaluation of market managed mechanisms for p2p systems. We have motivated a cost model where simple good will incentives may not be enough to allow content to propagate efficiently within the p2p group. Our model allows for the analysis of the intricacies introduced by a market model, and in many cases produces analytic results. Our next step is to interprete in more depth our results and validate some of the conjectures made in this paper. We believe that complete lack of control may be detrimental in certain cases where substantial costs are involved. The goal of our model is to substantiate this and provide for some quantitative results. We believe that although many of the assumptions regarding cost are of a rather theoretical nature, they capture some essential aspects of the system and will help us understand the role of the key parameters.

References

1. Adar, E., Huberman, B.: Free riding on gnutella. First Monday **5** (2000)
2. Saroiu, S., Gummadi, P., Gribble, S.: A Measurement Study of Peer-to-Peer File Sharing Systems. In: Proceedings of Multimedia Conferencing and Networking, San Jose. (2002)
3. Golle, P., Leyton-Brown, K., Mironov, I., Lillibridge, M.: Incentives for Sharing in Peer-to-Peer Networks. In: Proceedings of WELCOM'01. (2001)
4. Varian, H.: Microeconomic Analysis. Norton (1992)