
Bootstrapping the Long Tail in Peer to Peer Systems

Bernardo A. Huberman¹ and Fang Wu²

¹ HP Labs, Palo Alto, CA 94304 bernardo.huberman@hp.com

² HP Labs, Palo Alto, CA 94304 fang.wu@hp.com

1 Introduction

The provision of digitized content on-demand to millions of users presents a formidable challenge. With an ever increasing number of fixed and mobile devices with video capabilities, and a growing consumer base with different preferences, there is a need for a scalable and adaptive way of delivering a diverse set of files in real time to a worldwide consumer base.

Providing such varied content presents two problems. First, files should be accessible in such a way that the constraints posed by bandwidth and the diversity of demand is met without having to resort to client server architectures and specialized network protocols. Second, as new content is created, the system ought to be able to swiftly respond to new demand on specific content, regardless of its popularity. This is a hard constraint on any distributed system, since providers with a finite amount of memory and bandwidth will tend to offer the most popular content, as is the case today with many peer-to-peer systems.

The first problem is naturally solved by peer to peer networks, where each peer can be both a consumer and provider of the service. Peer to peer networks, unlike client server architectures, automatically scale in size as demand fluctuates, as well as being able to adapt to system failures. Examples of such systems are Bittorrent [4] and Kazaa, who account for a sizable percentage of all the use of the Internet. Furthermore, new services like the BBC IMP, (<http://www.bbc.co.uk/imp/>) show that it is possible to make media content available through a peer-to-peer system while respecting digital rights.

It is the second problem, that of an adaptable and efficient system capable of delivering any file, regardless of its popularity, that we now solve. We do so by creating an implementable incentive mechanism that ensures the existence of a diverse set of offerings which is in equilibrium with the available supply and demand, regardless of content and size. Moreover, the mechanism is such that it automatically generates the long tail of offerings which has been shown

to be responsible for the success of a number of online businesses such as Amazon or eBay [2]. In other words, while the system delivers favorite mainstream content, it can also provide files that constitute small niche markets which only in the aggregate can generate large revenues.

In what follows we describe an efficient incentive mechanism for P2P systems that generates a wide diversity of content offerings while responding adaptively to customer demand. Files are served and paid for through a parimutuel market similar to that commonly used for betting in horse races. An analysis of the performance of such a system shows that there exists an equilibrium with a long tail in the distribution of content offerings, which guarantees the real time provision of any content regardless of its popularity. In our case, the bandwidth fraction of a given file offered by a server plays the role of the odds, the bandwidth consumed corresponds to bettors, the files to horses, and the requests are analogous to races.

An interesting consequence of this mechanism is that it solves in complete fashion the free riding problem that originally plagued P2P systems like Gnutella [1] and that in milder forms still appears in other such systems. The reason being that it transforms the provision of content from a public good into a private one.

We then analyze the performance of such a system by making a set of assumptions that are first restrictive and are then relaxed so as to make them correspond to a realistic crowd of users. We show that in all these cases there exists an equilibrium in which the demand for any file can be fulfilled by the system. Moreover this equilibrium exhibits a robust empirical anomaly which is responsible for generating a very long tail in the distribution of content offerings. We finally discuss the scenario where most of the servers are bounded rational and show that it is still possible to achieve an optimum equilibrium. We conclude by summarizing our results and discussing the feasibility of its implementation.

2 The system and its incentive mechanism

Consider a network-based file exchange system consisting of three types of traders: content provider, server, and downloader or user. A content provider supplies—at a fixed price per file—a repertoire of files to a number of people acting as peers or servers. Servers then selectively serve a subset of those files to downloaders for a given price. In a peer-to-peer system a downloader can also, and often does, act as a server.

If the files are typically large in size, a server can only afford to store and serve a relatively small subset of files. It then faces the natural problem of choosing an optimal (from the point of view of maximizing his utility) subset of files to store so as to sell them to downloaders.

Suppose that the system charges each downloader a flat fee for downloading any one file (as in Apple's iTunes music store), which we normalize to

one. Since many servers can help distribute a single file, this unit of income has to be allocated to the servers in ways that will incentivize them to always respond to a changing demand.

In order to do so, consider the case where there are m servers and n files. Let b_{ij} be the effective bandwidth of server i serving file j , normalized to $\sum_{i,j} b_{ij} = 1$. Also, denote the bandwidth fraction of file j by $\pi_j = \sum_k b_{kj}$.

Suppose that when a downloader connects to the system, it starts downloading different parts of the file simultaneously from all available servers that have it. When it finishes downloading, it will have received a fraction of the file j

$$q_{ij} = \frac{b_{ij}}{\sum_k b_{kj}} = \frac{b_{ij}}{\pi_j} \quad (1)$$

from server i . Our mechanism prescribes that *the system should pay an amount q_{ij} to server i as its reward for serving file j .*

Now consider the case when server i 's reserves an amount of bandwidth b_{ij} as his "bid" on file j . Because we have normalized the total bandwidth and the total reward for serving one request both to one, the proportional share allocation scheme described by Eq. (1) can be interpreted as redistributing the total bid to the "winners", in proportion to their bids. Thus our payoff structure is similar to that of a pari-mutuel horse race betting market, where the π_j can be regarded as the odds, the bandwidth corresponds to bettors, the files to horses, and the requests are analogous to races (Fig. 1). It is worth pointing out however, that in a real horse race all players who have placed a bet on the winning horse receive a share of the total prize, whereas in our system only those players that kept the "winning" file and also had a chance to serve it get paid. In spite of this difference it is easy to show that when rewritten in terms of expected payoffs, the two mechanisms behave in similar fashion.

3 The solution

3.1 Rational servers with static strategies and known download rates

In this section we make three simplifying assumptions. While not realistic they serve to set the framework that we will utilize later on to deal with more realistic scenarios. First, every server is rational in the sense that he chooses the optimal bandwidth allocation that maximizes his utility, whose explicit form will be given below. Second, every server's allocation strategy is static, i.e. the b_{ij} 's are independent of time. Third, we assume that each file j is requested randomly at a rate $\lambda_j > 0$ that does not change with time, and these rates are known to every server.

Consider a server i with the following standard additive form of utility:



Fig. 1. A parimutuel horse betting board.

$$U = \mathbb{E} \left[\int_0^{\infty} e^{-\delta t} u(t) dt \right], \quad (2)$$

where $u(t)$ is his income density at time t , and $\delta > 0$ is his future discount factor. Let X_{j1} be the (random) time that file j is requested for the first time, let X_{j2} be the time elapsed between the first request and the second request, and so on. According to our parimutuel reward scheme, server i receives a lump-sum reward b_{ij}/π_j from every such request, at times $X_{j1}, X_{j1} + X_{j2}$, etc. Thus, the server i 's total utility is given by

$$U = \sum_j \frac{b_{ij}}{\pi_j} \sum_{l=1}^{\infty} \mathbb{E}[e^{-\delta \sum_{k=1}^l X_{jk}}] \equiv \sum_j \frac{b_{ij}}{\pi_j} u_j. \quad (3)$$

which amounts to each server receiving a utility proportional to the fraction of the file that he serves. Notice that the sum of expectations in Eq. (3) (denoted by u_j) can be calculated explicitly. Because the X_{jk} 's are i.i.d. random variables with density $\lambda_j^{-1} \exp(-\lambda_j x)$, it can be calculated that $u_j = \lambda_j/\delta$. If we let $\lambda = \sum_j \lambda_j$ be the total request rate and $p_j = \lambda_j/\lambda$ be the probability that the next request asks for file j , then we can also write $u_j = \lambda p_j/\delta$. Plugging this back into Eq. (3), we obtain

$$U = \frac{\lambda}{\delta} \sum_j \frac{p_j b_{ij}}{\pi_j}. \quad (4)$$

Since we assume that server i is rational, he will allocate b_{ij} in a way that it solves the following optimization problem:

$$\max_{(b_{ij})_{j=1}^n \in \mathbb{R}_+^n} \sum_j \frac{p_j b_{ij}}{\sum_k b_{kj}} \quad \text{subject to} \quad \sum_j b_{ij} \leq b_i. \quad (5)$$

where b_i is the total upload bandwidth of user i . Thus we see that the servers are playing a *finite budget resource allocation game*. This type of game has been studied intensively, and a Nash equilibrium has been shown to exist under mild assumptions [8, 11]. In such an equilibrium, the players' utility functions are strongly competitive and in spite of a possibly large utility gap, the players behave in almost envy-free fashion, i.e. each player believes that no other player has received more than they have.

3.2 Rational servers with static strategies and unknown request rates

We now relax some of the assumptions made above so as to deal with a more realistic case.

It is usually hard to find out the accurate request rate for a given file, especially at the early stages when there is no historical data available. Thus it makes more sense to assume that every server i holds a *subjective belief* about those request rates. Let p_{ij} be server i 's subjective probability that the next request is for file j . Then server i believes that file j will be requested at a rate $\lambda_{ij} = \lambda p_{ij}$. Eq. (5) then becomes

$$\max_{(b_{ij})_{j=1}^n \in \mathbb{R}_+^n} \sum_j \frac{p_{ij} b_{ij}}{\sum_k b_{kj}} \quad \text{subject to} \quad \sum_j b_{ij} \leq b_i. \quad (6)$$

which is still a finite budget resource allocation game as considered in the previous section.

It is interesting to note that when m is large, b_{ij} is small compared to $\pi_j = \sum_k b_{kj}$, so that π_j can be treated as a constant. In this case, the optimization problem can be well approximated by

$$\max_{(b_{ij})_{j=1}^n \in \mathbb{R}_+^n} \sum_j \frac{p_{ij} b_{ij}}{\pi_j} \quad \text{subject to} \quad \sum_j b_{ij} \leq b_i. \quad (7)$$

Thus, user i should use all his bandwidth to serve those files j with the largest ratio p_{ij}/π_j .

This scenario (7) corresponds to the so-called *parimutuel consensus* problem, which has been studied in detail. In this problem a certain probability space is observed by a number of individuals, each of which endows it with their own subjective probability distributions. The issue then is how to aggregate those subjective probabilities in such a way that they represent a good

consensus of the individual ones. The parimutuel consensus scheme is similar to that of betting on horses at a race, the final odds on a given horse being proportional to the amount bet on the horse. As shown by Eisenberg and Gale [6], an equilibrium then exists such that the bettors as a group maximize the weighted sum of logarithms of subjective expectations, with the weights being the total bet on each horse.

Moreover a number of empirical studies of parimutuel markets [5, 7, 9] have shown that they do indeed exhibit a high correlation between the subjective probabilities of the bettors and the objective probabilities generated by the racetracks. Equally interesting for our purposes is the existence of a robust empirical anomaly called *the favorite-longshot bias* [5, 7, 9]. The anomaly shows that favorites win more frequently than the subjective probabilities imply, and longshots less often. This anomaly enhances the long tail, which is populated by those files which while not singly popular, in aggregate are responsible for a large amount of the traffic in the system.

3.3 Rational servers with a dynamic strategy

We now consider the case where the rate at which files are requested can change with time. Because of this, each server has to actively adjust its bandwidth allocation to adapt to such changes. As we have seen in the last section, user i has an incentive to serve those files with large values of p_{ij}/π_j . Recall that $\pi_j(t)$ is just the fraction of total bandwidth spent to serve file j at time t , which in principle can be estimated from the system's statistics. Thus it would be useful to have the system frequently broadcast the real-time π_j to all servers so as to help them decide on how to adjust their own allocations of bandwidth.

From Eq. (1) we see that, by serving file j , user i 's expected per bandwidth earning from the next request is $p_j q_{ij}/b_{ij} = p_j/\pi_j$. Hence a user will benefit most by serving those files with the largest " p/π ratio". However, as soon as a given user starts serving file j , the corresponding p/π ratio decreases. As a consequence, the system self-adapts to the limit of uniform p/π ratios. If the system is perfectly efficient, we would expect that $p_j/\pi_j = \text{constant}$. Because p_j and π_j both sum up to one, this implies that $\pi_j = p_j$, or $\sum_k b_{kj} = \lambda_j/\lambda \propto \lambda_j$. In other words, the total bandwidth used to serve a file is proportional to the file's request rate.

This result has interesting implications when considering the social utility of the downloaders. Recently, Tewari and Kleinrock [10] have shown that in a homogeneous network the average download time is minimized when $\sum_k b_{kj} \propto \lambda_j$. This implies that in the perfectly efficient limit, our mechanism maximizes the downloaders' social utility, which is measured by their average download times.

Since in reality a market is never perfectly efficient, the above analysis only makes sense if the characteristic time it takes for the system to relax back to uniformity from any disturbance is short. As a concrete example,

consider a new file j released at time 0, being shared by only one server. Suppose that every downloader starts sharing her piece of the file immediately after downloading it. Because there are few servers serving the file but many downloaders requesting the file, for very short times afterwards the upload bandwidth will be fully utilized. That is, during time dt , an amount $\pi_j(t)dt$ of data is downloaded and added to the total upload bandwidth immediately. Hence we have

$$d\pi_j(t) = \pi_j(t)dt. \quad (8)$$

which implies that $\pi_j(t)$ grows exponentially until $\pi_j(T) \sim p_j$. Solving for T , we find

$$T \sim \log \left(\frac{p_j}{\pi_j(0)} \right). \quad (9)$$

Thus the system reaches uniformity in logarithmic time, a signature of its high efficiency.

3.4 Servers with bounded rationality

So far we have assumed that all servers are rational, so that they will actively seek those files that are most under-supplied so as to serve them to downloaders. In reality however, while some servers do behave rationally, a lot of others do not. This is because even a perfectly rational server sometimes can make wrong decisions as to which files to store because his subjective probability estimate of what is in demand can be inaccurate. Also, such a bounded-rational server can at times be too lazy to adjust his bandwidth allocation, so that he will keep serving whatever he has, and at other times he might simply imitate other servers' behavior by choosing to serve the popular files. In all these cases we need to consider whether or not the lack of full rationality will lead to equilibrium on the part of the system.

As a simple example, assume there are only two files, A and B. Let $p = \lambda_A/\lambda$ be file A's real request probability, and let $1 - p$ be file B's real request probability. Suppose the servers are divided into two classes, with α fraction rational and $1 - \alpha$ fraction irrational, arriving one by one in a random order. Each rational server's subjective probability in general can be described by an identically distributed random variable $P_t \in [0, 1]$ with mean p . Then with probability $\mathbb{P}[P_t > \pi(t)]$ he will serve file A, and with probability $\mathbb{P}[P_t < \pi(t)]$ he will serve file B. In order to carry out some explicit calculation below, we consider the simplest choice of P_t , namely a Bernoulli variable

$$\mathbb{P}[P_t = 1] = p, \quad \mathbb{P}[P_t = 0] = 1 - p. \quad (10)$$

(Clearly $\mathbb{E}[P_t] = p$, so the subjective probabilities are accurate on average.) It is easy to check that under this choice a rational server chooses A with probability p and B with probability $1 - p$.

On the other hand, consider the situation where an irrational server chooses an existing server at random and copies that server's bandwidth allocation. That is, with probability $\pi(t)$ an irrational server will choose file A.³

From these two assumptions we see that

$$\mathbb{P}[\text{server } t \text{ serves A}] = \alpha p + (1 - \alpha)\pi(t), \quad (11)$$

and

$$\mathbb{P}[\text{server } t \text{ serves B}] = \alpha(1 - p) + (1 - \alpha)(1 - \pi(t)). \quad (12)$$

The stochastic process described by the above two equations has been recently studied in the context of choices among technologies for which evidence of their value is equivocal, inconclusive, or even nonexistent [3]. As was shown there, the dynamics generated by such equations leads to outcomes that appear to be deterministic in spite of being governed by a stochastic process. In the context of our problem this means that when the objective evidence for the choice of a particular file is very weak, any sample path of this process quickly settles down to a fraction of files downloaded that is not predetermined by the initial conditions: *ex ante*, every outcome is just as (un)likely as every other. Thus one cannot ensure an equilibrium that is both optimum and repeatable.

In the opposite case, when the objective evidence is strong, the process settles down to a value that is determined by the quality of the evidence. In both cases the proportion of files downloaded never settles into either zero or one.

In the general case that we have been considering, there are always a number of servers that will behave in bounded rational fashion and a few that are perfectly rational. Specifically, when $\alpha > 0$, which corresponds to the case where a small number of servers are rational, the $\pi(t)$ will converge to p in the long time limit. That is, a small fraction of rational servers is enough for the system to reach an optimum equilibrium. However, it is worth pointing out that since the characteristic convergence time diverges exponentially in $1/\alpha$, the smaller the value of alpha α , the longer it will take for the system to reach such an optimum state.

4 Conclusion

In this paper we described a peer-to-peer system with an incentive mechanism that generates diversity of offerings, efficiency and adaptability to customer

³ This assumption can also be interpreted as follows. Suppose a downloader starts serving his files immediately after downloading, but never initiates to serve a file. (This is the way a non-seed peer behaves within Bittorrent.) Then the probability that he will serve file j is exactly the probability that he just downloaded file j , which is $\pi_j(t)$.

demand. This was accomplished by having a pricing structure for serving files that has the structure of a parimutuel market, similar to those commonly used in horse races, where the the bandwidth fraction of a given file offered by a server plays the role of the odds, the bandwidth corresponds to bettors, the files to horses, and the requests are analogous to races. Notice that this mechanism completely solves the free riding problem that originally plagued P2P systems like Gnutella and that in milder forms still appears in other such systems.

We then analyzed the performance of such a system by making a set of assumptions that are first restrictive but are then relaxed so as to make the system respond to a realistic crowd. We showed that in all these cases there exists an equilibrium in which the demand for any file can be fulfilled by the system. Moreover this equilibrium is known to exhibit a robust empirical anomaly, that of the *favorite-longshot bias*, which in our case generates a very long tail in the distribution of offerings. We finally discussed the scenario where most of the servers are bounded rational and showed that it is still possible to achieve an optimum equilibrium if a few servers can act rationally.

The implementation of mechanism is feasible with present technologies. The implementation of a prototype will also help study the behavior of both providers and users within the context of this parimutuel market. Given its feasibility, and with the addition of DRM and a payment system, it offers an interesting opportunity for the provision of legal content with a simple pricing structure that ensures that unusual content will always be available along with the more traditional fare.

Acknowledgements.

We benefited from discussions with Eytan Adar, Tad Hogg and Li Zhang.

References

1. Eytan Adar and Bernardo A. Huberman. Free Riding on Gnutella. First Monday October (2000).
2. Chris Anderson. The long tail. http://longtail.typepad.com/the_long_tail/ (2005).
3. Jonathon Bendor, Bernardo A. Huberman and Fang Wu. Management fads, pedagogies and soft technologies. <http://www.hpl.hp.com/research/id1/papers/fads/fads.pdf> (2005).
4. Bram Cohen. Incentives build robustness in Bittorrent. Working Paper, Workshop on the Economics of P2P Systems (2003).
5. L. Coleman. New light on the longshot bias. *Applied Economics*, 36(4), 315–326 (2004).
6. Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The parimutuel method. *Annals of Mathematics Statistics*, 30(1), pp. 165–168 (1958).

7. M. Gramm and D. H. Owens. Efficiency in Pari-Mutuel betting markets across wagering pools in the simulcast era. *Southern Economic Journal*, 72(4), 926–937 (2006).
8. L. Shapley and M. Shubik. Trade using one commodity as a means of payment. *Journal of Political Economy*, Vol. 85:5, 937–968 (1977).
9. Richard H. Thaler and William T. Ziemba, Parimutuel betting markets: race-tracks and lotteries, *Journal of Economic Perspectives*, Vol. 2, No. 2, pp. 161–174 (1988).
10. Saurabh Tewari and Leonard Kleinrock. On fairness, optimal download performance and proportional replication in peer-to-peer networks. *Proceedings of IFIP Networking 2005*, Waterloo, Canada (2005).
11. Li Zhang. The efficiency and fairness of a fixed budget resource allocation game. *ICALP* (2005).