

2003

Economics of a Collaborative Peer-to- Peer Infrastructure for Content Distribution

Mohamed Hefeeda

Philipp Afeche

Bharat Bhargava

Purdue University, bb@cs.purdue.edu

Report Number:

03-015

Hefeeda, Mohamed; Afeche, Philipp; and Bhargava, Bharat, "Economics of a Collaborative Peer-to- Peer Infrastructure for Content Distribution" (2003). *Computer Science Technical Reports*. Paper 1564.
<http://docs.lib.purdue.edu/cstech/1564>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact epubs@purdue.edu for additional information.

**ECONOMICS OF A COLLABORATIVE PEER-TO-PEER
INFRASTRUCTURE FOR CONTENT DISTRIBUTION**

**Mohamed Hefeeda
Philipp Afeche
Bharat Bhargava**

**CSD TR #03-015
May 2003**

Economics of a Collaborative Peer-to-Peer Infrastructure for Content Distribution

Mohamed Hefeeda
Department of Computer Sciences
Purdue University

Philipp Afeche
Kellogg School of Management
Northwestern University

Bharat Bhargava
Department of Computer Sciences
Purdue University

Abstract

Distributing digital media contents to a large number of users in a cost-effective manner is a challenging task for the content provider. Traditionally, the content provider either deploys a set of high capacity servers or contracts a content delivery network (CDN) to transport contents to users. The first approach requires a significant investment to set up and administer the servers. While in the second approach the CDN charges the provider for every megabyte served. For large files such as movies, both approaches burden the budget of the content provider. For instance, the limited deployability of video on demand services may be attributed, in part, to the cost factor.

Motivated by the success of the peer-to-peer (P2P) paradigm in the last few years and by the immense number of the often underutilized end systems connected to the Internet, we propose a collaborative P2P infrastructure for cost-effective content distribution. The objectives of this position paper are: (1) To highlight the economic potential of a content distribution system built on top of a P2P infrastructure, and (2) To identify the key research problems that need to be addressed in order to realize this economic potential.

1 Introduction

Consider a content provider who is interested in distributing digital media contents (e.g., music files, movies, documents) to a set of potential clients. Currently, two approaches are being used to distribute contents to clients: direct and third-party. In the direct approach, the content provider deploys and manages a

set of servers with capacity commensurate to the expected demand from clients. To enhance performance (in terms of, e.g., short delay and small loss ratio) a third-party is involved to “deliver” contents to clients. This third-party is known as a Content Delivery Network (CDN). Content delivery networks, such as Akamai and Digital Island, deploy thousands of servers at the *edge* of the Internet. (Akamai deploys 10,000+ servers.) These servers (also called caches) are installed at many POPs (point of presence) of major ISPs such as AT&T and Sprint. The idea is to keep the contents close to the clients, and hence traffic traverses fewer network hops. This reduces the load on the backbone networks and yields a better service. The CDN typically caches the contents at many servers and redirects a client to the most suitable server. Proprietary protocols are used to distribute contents over caches, monitor the current traffic situation over the Internet, and directs clients to caches. Cost-effectiveness is a major concern in both of these approaches, especially for distributing large files such as movies. For instance, CDN charges the content provider for every megabyte served, which might be acceptable for relatively small files such as web pages with some images.

We envision a collaborative content distribution infrastructure centered around the peer-to-peer (P2P) paradigm. Instead of deploying powerful caches at many locations, the P2P model relies on resource contribution from peers (client machines). Every peer may contribute a little, but there is an enormous number of them. The P2P approach strives to push the contents even closer to the clients: contents are obtained from fellow peers within the *same* network domain. The

collaborative P2P model can be used in two settings. First, it can serve as a *substrate* through which content providers disseminate contents to clients by employing and aggregating resources from participating peers. In this case, content providers should motivate peers to contribute resources to the system. Second, it can be used as an infrastructure for a cooperative sharing of resources and contents among peers. A cooperative file-sharing system is an example for the second case. In this case, incentive mechanisms should be developed to ensure fair contribution and consumption of resources.

The collaborative P2P model has the potential to create substantial value in a cost effective manner, compared to a system where no such sharing occurs. However, realizing this potential is a challenging task. First, how to optimally create and disseminate information resources in a network among a large number of distributed participants with stochastic and dynamically changing demand and supply is a hard enough problem, even for a central "system manager" with complete information about the preferences and full control of actions of all participants. Second, this problem is further complicated by the fact that participants are autonomous and self-interested economic agents (individuals or firms) whose own incentives and objectives are typically not aligned with those of the overall system. Furthermore, participants have private information about their own preferences and other important variables that impact the system behavior. Hence, the success of collaborative information sharing requires mechanisms that coordinate the actions of its participants such that they increase the value of the system.

Before we present the research problems, we describe the similarities and differences between the economic issues in collaborative information systems and transmission services in data networks.

1.1 Economics of Data Transmission versus Collaborative Information Systems

The design and pricing of data transmission services in networks emerged in the mid-1990s with the commercialization of the Internet (see for example [2, 7, 10, 12, 16, 17, 21]). It typically focuses on how to design price-service mechanisms that optimally allocate a capacity-constrained and therefore congestion-prone network to given customer demands for transmission services. Data transmission services share some

of the economic characteristics of collaborative information sharing: both involve the transmission of various media types with heterogeneous quality requirements over congestion-prone computer networks, connect many independent, spatially dispersed and self-interested users who are sensitive to quality of service and whose consumptions create externalities. However, it is important to emphasize that collaborative information sharing poses some fundamentally new and different challenges. (1) In data transmission pricing, users are typically only consumers of the resources. By contrast, in collaborative information sharing, the *peers are both consumers and suppliers of content and resources*. (2) Even though the Internet is managed by independent service providers, its supply infrastructure is nevertheless much more aggregated than is the case in collaborative information sharing where *each* participant who supplies system resources is an independent agent. In this sense, P2P systems are "completely decentralized"; their available infrastructure and the location and quantity of the information they carry are subject to the decisions of a much larger number of smaller entities than in data transmission systems. (3) The work on data transmission pricing focuses on the provisioning and allocation of a *given* set of resources to a *given* set of demand functions for data transmissions between a *given* source and destination. By contrast, in collaborative information sharing, the set of available resources and their capacities are *functions* of peers' decisions, the demands for transmissions are *functions* of how much content is made available for sharing, and the source-destination(s) pairs are *functions* of peers' sharing choices, network locations and the mechanism for matching supplying peer(s) to client peers.

2 Research Problems

The goal of our research is to *design and test economic mechanisms that yield highly performing collaborative information sharing systems*: they should be at least economically viable, and ideally economically attractive, for all participants. To that end, we need to develop a systematic understanding of how system behavior depends on (1) the technological properties of the system, (2) the incentives of the economic agents that control their resources, and (3) the mechanisms that are in place. Specifically, the following research problems need to be resolved in order to realize the collaborative

model for content distribution.

2.1 Defining Economic Performance Objectives

First of all, the economic performance objectives of the collaborative system have to be clearly defined. These may include: (1) maximize the *system benefit*, or aggregate utility minus cost. This aggregate metric typically results from more specific objectives, including, (2) large variety (selection) and high recency of available information content, (3) high quality of service delivery, (4) high consumption (sharing) levels of available information, and (5) low incremental cost of deploying the sharing infrastructure. In addition, it may be important to evaluate the *distribution of benefits* among client peers, supplying peers, and other relevant participant groups.

2.2 Mechanism Design for Information Network Provisioning

We define network provisioning as the process of creating and distributing information resources among supplier peers in preparation for client requests. This definition bears some similarities with the replication problem addressed by Choen and Shenker [6]. The authors propose optimal replication strategies that minimize the expected search size in an unstructured P2P environment. They prove that replicating objects in proportion to the square-root of their query rates yields the minimum expected search size for locatable items. However, the replication strategies assume “full co-operation” from peers, in the sense that a peer voluntarily commits some of its capacity to the system and follows the prescribed protocol for replicating objects. Peers of this nature fall in the *obedient* nodes category in the classification given by Shneidman and Parkes [20] and Feigenbaum and Shenker [9]. The obedient nodes along with faulty nodes—those that may stop working (fail stop), drop messages, or act arbitrarily (Byzantine)—are the typical nodes used in the distributed systems literature. In contrast, nodes in P2P systems are found to be economically *rational* or utility-maximizing agents [20, 9]. Unless properly incentivised, nodes may deviate from the protocol or not participate at all. In the replication problem, if peers are not paid for sharing data, they may not have an incentive to cache data for subsequent sharing. If they do

get paid, their decision to cache may depend on more subtle factors such as their “forecast of demand” for the cached materials. In a multi-product environment, a peer may prefer to only cache what she perceives to be the most popular and therefore the most lucrative content, even though this may not be in line with the overall system objective.

This makes the provisioning problem more challenging because we need to: (1) understand peers behavior and more specifically their valuation of their own capacity, (2) study how the system performance depends not only on the aggregate capacity but also on the heterogeneous contributions from individual peers, (3) analyze the interaction between the provisioning and matching algorithms, and (4) design an incentive mechanism to integrate the above issues into a provisioning algorithm that optimizes a system-wide objective function.

The fact that agents are distributed and have a considerable amount of private information also raises the question of how centralized the mechanism should be. Adding the critical computational tractability issue, the problem becomes a distributed algorithmic mechanism design (DAMD), whereby the agents, the relevant information, and the computation of the mechanism are distributed. Examples of DAMD in the data transmission setting include [8] and [9].

2.3 Mechanism Design for Matching Client and Supplier Peers

The problem of matching client requests with supplier peers can be viewed as a complex routing problem. Unlike in standard routing, here the destination is to be endogenously determined as a function of the incentives of clients (e.g., price, content and service quality), suppliers (e.g., current vs. future revenue opportunity, resource consumption of fulfilling request) and those of the overall system (e.g., value of request vs. negative externalities on service quality of competing requests).

2.4 Mechanism Design for Joint Provisioning And Matching

The next task is to study how the provisioning and matching mechanisms interact and to search for mechanism pairs that perform well together. This search will give special consideration to mechanisms that link

a peer's ease of access to others' resources to her willingness to share her own.

2.5 Study of P2P System Macro Structure and Behavior

The provisioning and matching mechanisms are likely to impact how a network forms and behaves. This raises interesting fundamental questions about the relationship between the macro behavior of the system and the behavior of its individual components. From this perspective, collaborative information sharing systems can be viewed as complex systems [18], whose study is creating significant interest in various branches of the physical and social sciences. A number of interesting open questions can be asked about the macro structure and function of P2P systems that emerge in such decentralized fashion: How do they compare to those that are obtained by centralized design? What can be said about their dynamics? Do they converge to certain structural and functional patterns? If so, how stable are these patterns and how "organized" do they appear to be? These questions may well uncover interesting links between the structure of collaborative information sharing systems and their economics: for example, how concentrated are the peers and is the content in a system that has reached steady-state? Does the resulting structure of supply and demand reflect a high degree of market power for a small number of large peers, or do all peers control a comparable amount of resources? How do the answers depend on the mechanism design and on key features of the participants?

2.6 Comparison of P2P Systems with Key Alternatives

The objective of this task is to systematically compare the economic and operational performance of P2P systems with those of key alternatives, including content distribution networks. A comparison framework that identifies key economic and operational performance metrics as well as metrics that categorize the market and network environment need to be developed. The economic and operational performance metrics may include: expected provider profits, cost effectiveness, customer benefits, system stability, and service quality. While the geographical dispersion of potential users, concentration of capacity bottlenecks in the delivery infrastructure, and nature of data to be distributed can

be used to categorize the market and network environment.

3 Related Work

The economic aspects of peer-to-peer systems have received little attention so far. Previous research appears to mainly focus on the free riding problem, whereby only a small fraction of peers contribute resources into the system. Free-riding has been shown by [1] (through a measurement study) and [11] (through game-theoretic analysis). Using a model of user behavior and empirical data collected from OpenNap networks, [3] shows that free-riding increases with the size of the P2P network, a known phenomenon in the public goods setting [13, 19]. Free riding threatens the future of P2P systems by stifling the growth of the system capacity and the variety and volume of sharing. Researchers of [11] and [1] advocate the use of payment mechanisms that motivate the peers with incentives to contribute to the system. In [11], the authors construct a game theoretic model of P2P systems and study user equilibria under different payment mechanisms that the system offers the peers, such as micro-payments, points-based, and rewards for sharing. By contrast, [15] focuses not on peers' incentives but on those of the central authority in file-sharing services such as Napster that are centrally managed. They propose how peers should pay the central entity to motivate it to make clients aware of their content. To prevent the problem of content piracy, [14] proposes a system architecture that uses economic incentives instead of tamper resistance protocols and motivates users to keep the content within the subscription community.

In summary, existing economic studies of P2P systems give *partial*, mostly qualitative, *insights* into some of the incentives that drive peers' behavior in collaborative information sharing, and they explore some *partial solutions*. However, since they abstract away important technological features of the network environment within which peers operate, these studies are not equipped to provide a systematic understanding of how a P2P system behaves depending on the technical and economic mechanisms that govern its operation.

4 Conclusion

In this position paper, we introduced the idea of a collaborative infrastructure for content distribution, which relies on aggregating resource contributions from the participants in the system. We argued that the collaborative model has the potential to create substantial value in a cost-effective manner. We also presented the research challenges facing the model. These challenges mainly include designing incentive-compatible algorithms for: replicating information resources among suppliers, matching clients with the appropriate suppliers, and the interaction between the replication and the matching. They also include characterizing the macro structure behavior of the collaborative model and comparing it versus its alternatives.

References

- [1] E. Adar and B. Huberman. Free riding on Gnutella. *First Monday*, 5(10), October 2000.
- [2] P. Afeche. Pricing data network transport services: Review and directions, 2003. Under review for publication in *Handbook on Economics of Electronic Commerce*.
- [3] A. Asvanund, K. Clay, R. Krishnan, and M. Smith. An empirical analysis of networks with externalities in p2p music-sharing networks. Technical report, H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, July 2002. Working Paper.
- [4] S. Barnett and G. Anido. A cost comparison of distributed and centralized approaches to video-on-demand. *IEEE Journal on Selected Areas in Communications*, 14(6):1173–1183, August 1996.
- [5] S. Chan and F. Tobagi. Distributed servers architecture for networked video services. *IEEE Transactions on Networking*, 9(2):125–136, April 2001.
- [6] E. Choen and S. Shenker. Replication strategies in unstructured peer-to-peer networks. In *Proc. of ACM SIGCOMM'02*, pages 177–190, Pittsburgh, PA, USA, August 2002.
- [7] R. Edell and P. Varaiya. Providing internet access: What we learn from INDEX. *IEEE Network Magazine*, 13(5):18–25, September/October 1999.
- [8] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmission. *Journal of Computer and System Sciences*, 63:21–41, 2001.
- [9] J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *Sixth International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIALM'02)*, Atlanta, Georgia, September 2002. ACM Press.
- [10] R. Gibbens and F. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35:1969–1985, 1999.
- [11] P. Golle, K. Leyton-Brown, and I. Mironov. Incentives for sharing in peer-to-peer networks. In *Proc. of Second workshop on Electronic Commerce (WELCOM'01)*, Heidelberg, Germany, November 2001.
- [12] A. Gupta, D. Stahl, and A. Whinston. A stochastic equilibrium model of internet pricing. *Economic Dynamics and Control*, 21:697–722, 1997.
- [13] J. Hindriks and R. Panes. Free riding on altruism and group size. Technical report, University of London, 2001. Working Paper.
- [14] B. Horne, B. Pinkas, and T. Sander. Escrow services and incentives in peer-to-peer networks. In *Proc. of ACM Conference on Electronic Commerce (EC'01)*, Tampa, Florida, October 2001.
- [15] A. Kim and L. Hoffman. Pricing Napster and other Internet peer-to-peer applications. Technical report, George Washington University, 2002.
- [16] J. MacKie-Mason, P. Thomas, and D. Teneketsis. A market-based approach to optimal resource allocation in integrated-services connection. *Operations Research*, 50(4), July-August 2002.
- [17] J. MacKie-Mason and H. Varian. Pricing congestible network resources. *IEEE Journal of Selected Areas in Communications*, 13(7):1141–49, September 1995.
- [18] J. Ottino. Complex systems. *AIChE Journal*, 49(2):292–299, February 2003.
- [19] T. Palfrey and H. Rosenthal. Participation and the provisioning of discrete public goods: a strategic analysis. *Journal of Public Economics*, 24:171–193, 1984.
- [20] J. Shneidman and D. Parkes. Rationality and self-interest in peer-to-peer networks. In *Proc. of 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03)*, Berkeley, CA, USA, February 2003.
- [21] D. Songhurst. *Pricing Congestible Network Resources*. Elsevier, Amsterdam, 1999.