MAGAZINE

Economic Framework for Pricing and Charging in Digital Libraries

J. Sairamesh
Computer Science Department
Columbia University
New York, NY, 10027
jakka@cs.columbia.edu
ramesh@ics.forth.gr

C. Nikolaou
University of Crete and ICS-FORTH
Heraklion, Crete
GR 711 10
nikolau@ics.forth.gr

D. F. Ferguson
IBM T.J. Watson Research Center
Hawthorne, New York, NY, 10532
dfferg@watson.ibm.com

Y. Yemini
Computer Science Department
Columbia University
New York, NY, 10027
yemini@cs.columbia.edu

D-Lib Magazine, February 1996

ISSN 1082-9873

Abstract

We consider commercial Digital Libraries as information economies consisting of several players: authors and publishers who create and sell their collections, suppliers (e.g. computer systems) who provide information storage, indexing and access services, information-agents who provide searching and presentation services, and users who request for services. In such an economic framework, one can envision suppliers and information-agents competing to provide services for information storage, searching, access and presentation. In providing such services, several issues arise; among them are socio-economic and cultural aspects of pricing information objects and Quality of Service (QoS) to access and view these objects. These issues play an important role in allocating resources such as processing time, network bandwidth and buffers, memory, cache and I/O, which are distributed (and owned) among various players in the economy. Using this framework, we present the interactions among the players, service models, pricing and charging/billing mechanisms, and corresponding implementation issues in large digital libraries.

Introduction

Digital Libraries will have a major influence on the design of future information systems. They are the cradle from which future advanced information technologies will emerge to provide "transparent" services to a variety of users. Digital Libraries have already attracted many real-world agents, which participate for various economic reasons. For example, publishers of books are in the process of presenting their collections in digital formats in order to reduce production costs and increase profit margins.

Digital Libraries will house information objects in various media, such as text, audio, video and image, and provide information access services to a variety of users. One can envision a plethora of applications, such as collaborative computing and virtual class-rooms, using such services. Users of these applications, being heterogenous in their preferences, could request for different qualities of service for information access. The preferences of users could be based on the charge to access information.

For example, users browsing through a Digital Library of video clips or image clips (as shown in Figure 1) would prefer different qualities of service based on the charge and the availability of local resources (local software and hardware). In the Figure 1, we also illustrate that fact that current networks have limited bandwidth, and efficient usage of bandwidth is neccessary to support many concurrent accesses to the Digital Libraries. Therefore, one can envision several levels of services being offered by Digital Libraries to access information objects. In the example, we show that a video object can be presented to the user in different formats (MPEG, MPEG-II, etc.).

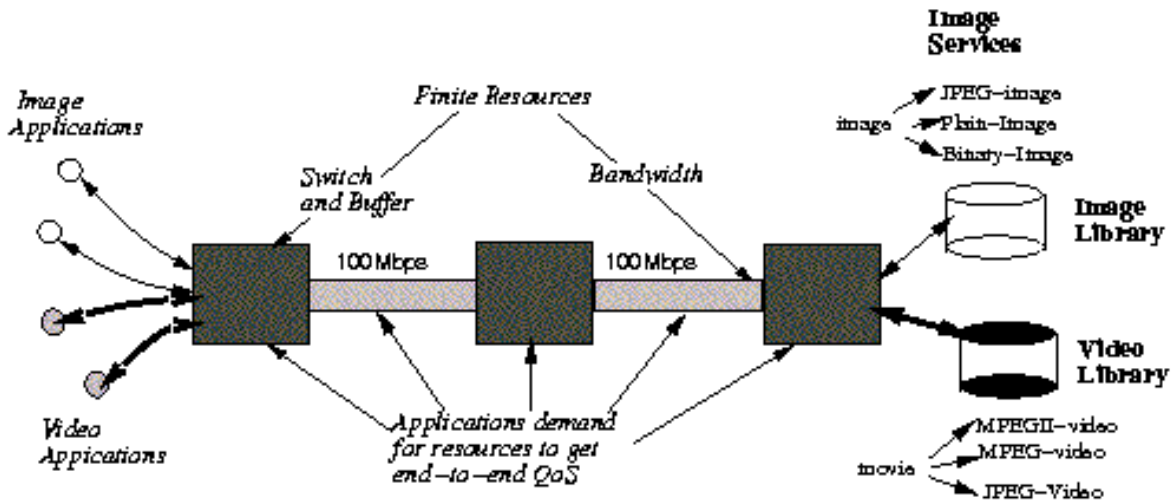

Figure 1.

In designing service mechanisms for accessing information objects in digital libraries, several issues arise. One of them is the issue of pricing services (and multiple service levels) to access information objects, and accounting and billing users for QoS. These are economic issues, however, they play an important role in allocating resources such as processing (CPU) time, network bandwidth, buffer, memory and cache for storing, searching, accessing, transporting, and presenting various information objects in a coherent way to the users. For example, users might want to know the prices for different qualities of service to access an information object, such as a video clip of a lecture based on a specific day (and time).

In order to develop a framework for pricing services and charging/billing users for access, several questions arise

and some of them are as follows: What is the architecture? What is the economic framework? Who are the players in the economy, and what are their interactions? What should be priced? What are the pricing models? What are the practical issues in implementing pricing schemes, and charging frameworks in large distributed Digital Libraries? There are several accompanying system level issues for storage and access of objects. Some of them are as follows: What are performance demands on the local operating systems and distributed system as a whole? What are appropriate scheduling and load-balancing mechanisms? What are the object storage and replication mechanisms?

Economic models for storage and replication of information objects have been investigated by Ferguson et. al and Stonebraker et. al.. Economic models for resource management, in general, have been considered in [1] and [11] . Considerable research in the area of Information economies is being done in [8]. Also, in [8], a collection of references to related works in this area is given. Work in the area of autonomous, profit-seeking agents in market economies is being conducted in [9].

Current networks and systems have limited resources (such as link bandwidth). Therefore, allocation of resources, such as processing time and network bandwidth, can be crucial in providing efficient services. In the future, bandwidth and processing time will be more than sufficient, and information objects will be priced based on the demand for them rather than on the Digital Library system (DLSystem) that houses them. We consider a DLSystem to be a complex computer system that provides storage, indexing, and access services to various information objects. In this paper, we assume that a set of resource allocations make up a service. An information object could also be a resource, but we do not price information objects in our current models, however, we plan to incorporate this in the near future.

Our Approach

Our main approach in pricing and accounting for QoS is an economic model of the players in large Digital Libraries. This involves modeling and understanding the various tasks that are required to support Digital Libraries. There are several reasons why this form of modeling helps in understanding the interfaces between the various elements of Digital Libraries. The main reasons are as follows:

- *Limiting Complexity:* Economic models provide several interesting contributions to resource sharing algorithms. The first is a set of tools for limiting the complexity by decentralizing the control of resources. The second is a set of mathematical models that can yield several new insights into resource sharing problems (such as scheduling in computer systems, object replication in distributed systems).
- *Decentralization:* In an economy, decentralization is provided by the fact that economic-agents attempt to achieve their goals in competitive fashion (selfish fashion). There are two types of such agents, *suppliers* and *consumers*. A consumer attempts to maximize its benefit by obtaining the services or resource allocations that maximize its preferences under a wealth constraint (see [4]). Suppliers compete in order to maximize revenue and minimize costs, which means attract more users and provide the best services.
- *Pricing and Coordination:* Most economic models introduce money and pricing as techniques for coordinating the behavior of agents. Each consumer is endowed with money, which they use to purchase services or resources (see [4]). The price a supplier charges for a resource or a service is determined by its supply and the market demand.
- *Usage Accounting, Billing and Dimensioning*: In using economic models for service provisioning in distributed systems, usage accounting becomes a part of the economy. Suppliers have to keep track of resource usage in order to price resources effectively, and thereby charge/bill the users for access.
- *Administrative Domains:* Large distributed systems and computer networks are spread over several domains. Each domain is typically managed and controlled by an administrator. An economic model of such an environment would imply that administrators advertise their services and prices. Therefore, providing a simple interface to negotiate between domains.
- *Scalability;* A key issue in designing architectures for services in large computer networks and distributed systems is scalability. Models of competition provide--in a *natural fashion*--mechanisms for scaling services

appropriately based on the service-demand and resource-availability.

The *novelty* introduced by the economic approach is not only in modeling large decentralized, autonomous systems, but also in designing efficient systems that account naturally for user QoS requirements, and optimally allocate resources for the various services in Digital Libraries.

Simple Model of Digital Libraries

We consider three kinds of players in the economic model of Digital Libraries: Suppliers (storage and indexing), Agents (who search, transform and present information) and Users. The three players are described below:

- **Suppliers**: Commercial DLSystems buy works from various publishers and authors, and provide services to a variety of users for a fee. The DLSystems not only price information but also price QoS to access the information. Pricing is based on the user demand, and users represent their demands through agents. The prices remain fixed every time-interval, assuming that time is divided into intervals. For example, a time-interval could be a 6 hours in a day. The time-intervals are decided by each DLSystem independently based on the market demand. The suppliers provide storage, access and local indexing services.
- **Information Agents**: These are suppliers that provide *Value-Added Services*, such as transparent searching and presentation of objects to the users. From now on, we refer to Information Agents as agents. The agents charge users for QoS provided. Agents contract services and resources from DLSystems over a period of time, and provide stable set of services to the users. The stability is provided by a flat service fee over period of time (example: half a day) to the users. Agents renegotiate for services from each DLSystem provided it is profitable, else they release resources. The agents compete to provide *One-Stop Shopping* services to the users.
- **Users**: They present complex queries to the information agents, and request for a service in viewing or obtaining the objects. Users could be grouped into query classes based on similar access patterns or just based on service required. Users within a group or class share the costs of the services. Users are free to choose among the various agents based on the QoS and charge.

In Figure 2, we illustrate the economic model with three players. There are many classes (K) of users who obtain services from one or many agents. The agents buy services from one or more of the DLSystems (Library Systems), and provide searching and presentation services to the users for a fee. In the Figure, users are shown to get services only through the agents, however, one can also model users obtaining services directly from the DLSystems, which means that users have to search, on their own, for information about Digital Library services.

Pricing Services: Searching and Presentation

In the economic framework, agents compete for resources and services from DLsystems. They buy information access services ahead of time, and provide a set of integrated services to consumers (users). Similarly, DLsystems which have common information objects naturally compete to attract users and maximize profit. In designing pricing schemes for such services in an economic framework, several considerations need to be taken into account. The main considerations are the following:

- What should be priced? Resources (such as buffer, bandwidth and CPU time) and services, based on performance parameters such as response time should be priced. Services for quality of search can also be priced. Information objects, in general, will be priced, and we plan to study this in the future.
- Prices must be set such that demand equals the supply. This is done to make sure that resources are not over-booked, and service quality is always maintained.
- Prices should reflect user query-request behaviour (load). This is essentialy to control congestion of requests to a single digital library system.
- Prices should not fluctuate rapidly as users can get confused. In practice, prices should be stable for a

reasonably long period of time.
- Prices need to be negotiated ahead of time. This means that arriving users choose the services from agents based on the price information and the service levels.

Agents, post their prices and services offered for all the time-intervals on a certain day. For example, agents might offer low prices for a service between 6.00 A.M. and 4.00 P.M. and offer a higher price for the same service between 4.00 P.M. and 10.00 P.M. Agents, based on factors such as user demand profile and market behaviour, might compete to buy more services from the DLSystems.

Agents have bugdets for buying services from DLsystems. Agents compete among themselves for services from suppliers. Agents charge the users based on QoS requested in searching and presentation. Suppliers price resources/ services based on demand from the users. We assume that time-intervals, where prices are fixed, are set independently by each supplier.

For example, supplier A might have 6.00 A.M till noon as one time period, and from noon till 6.00 P.M. another time period and price structure, and from 6.00 P.M. till 10. P.M. a different set of prices. Supplier B might have fixed prices between 6.00 A.M. and 5.00 P.M and another price structure for 5.00 P.M to 11.00 P.M. The behaviour of the suppliers is an economic one, the change in prices is due to the market demand.

The suppliers operate indepedently in selecting the time periods for stable prices, however, they are driven based on the collective market demand. Similarly, each agent independently selects the time periods for stable prices to the users.

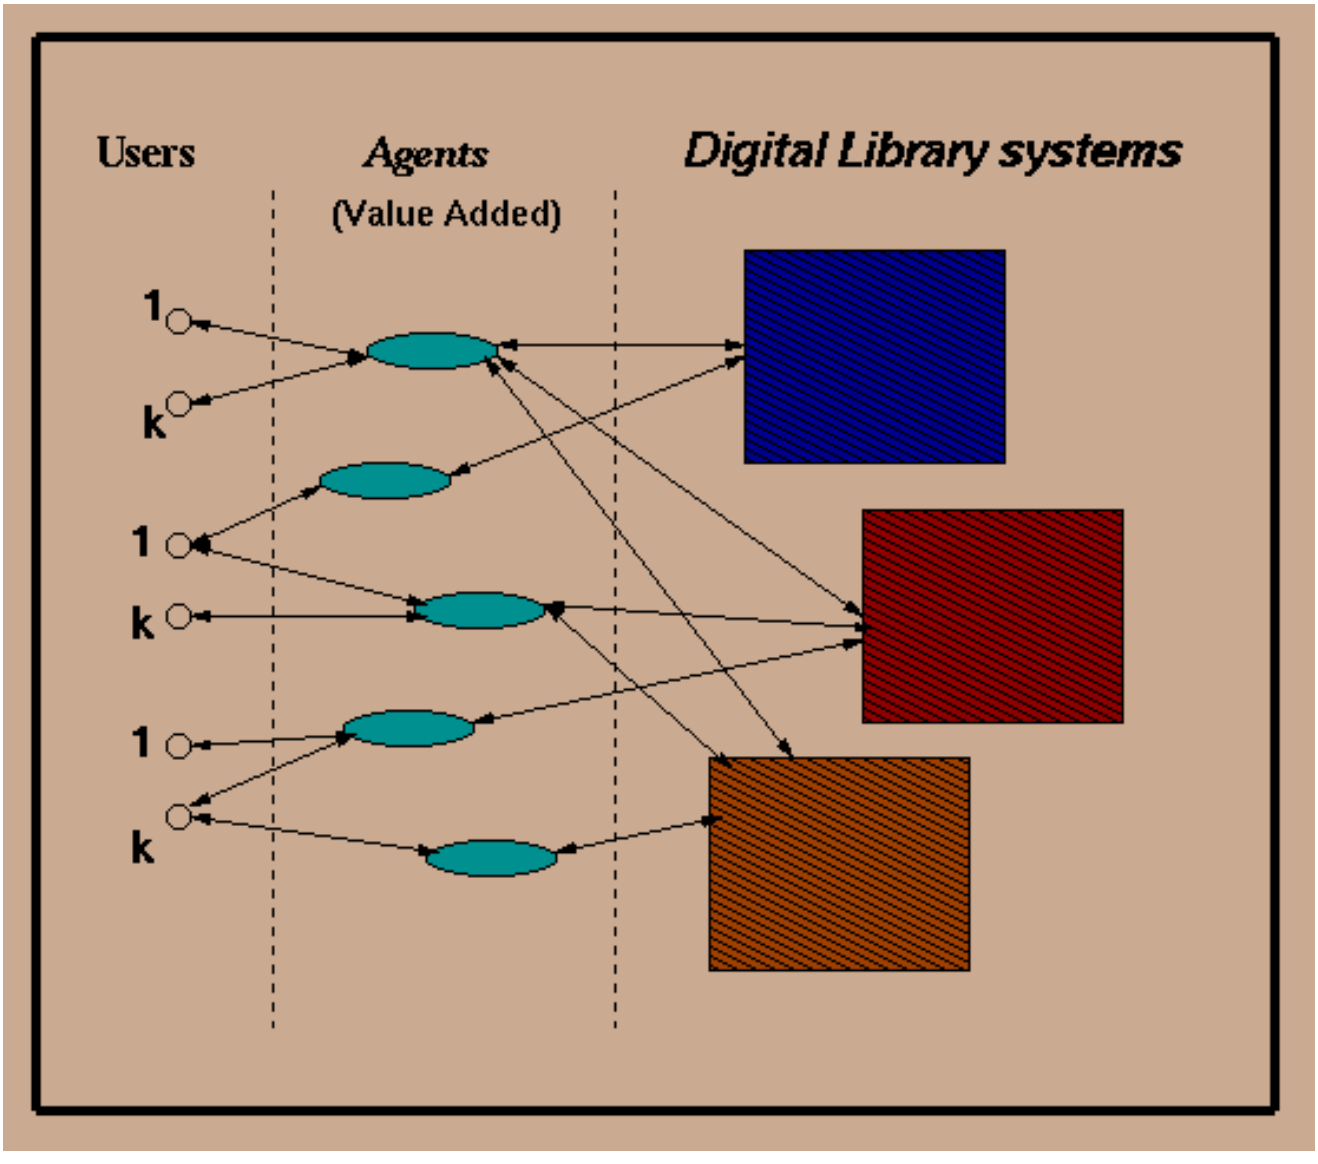

Figure 2.

Price Negotiation:

The suppliers, based on the demand profiles of the agents, decide on times when negotiation of the services and prices is done. For example, the price negotiation times everyday could be 6.00 A.M and 4.00 P.M. for supplier A, and 6.00 A.M and 5.00 P.M. for supplier B. Negotiation of prices and services by the suppliers is done at the beginning of every time period. Prices could be announced just after the negotiation.

In economic terms, the negotiation protocol between the agents and the DLsystems could be done using the *tatonnement* process: agents are endowed with some wealth. Each agent computes the demand from a benefit function and wealth. The benefit function is simply a utility function as described in Ferguson et. al.. The aggregate demand from all the agents is collected by each supplier who then computes the new resource price. If the demand for a resource is greater than its supply, the supplier raises the price of the resource. If there is surplus supply, price is decreased. With the new prices, agents again compute and present their demands to the suppliers. This process continues iteratively till the equilibrium price is achieved where demand equals the supply.

Bidding and auctioning for resources/services are other forms of price fixing. There are several auctioning mechanisms such as the Sealed Bid Auction, Dutch Auction, and English Auction . The basic philosophy behind auctions and bidding is that the highest bidder always gets the resources, and the current price for a resource is determined by the bid prices.

We model the digital library system as a simple computer system with CPU and storage (disks, memory and cache). We assume that storage is very large and therefore the charge is negligible, however, we assume that processing power and network I/O bandwidth is limited, and users are charged for them. Users who wish to get a good response time, pay for network I/O bandwidth and processing time allocations.

Some simple performance based economic models are discussed in [1]. Similarly, models for allocating network bandwidth and buffers for various services are given in [2], [3], [4], [7] and [8]. Our goal is not only to investigate pricing mechanisms, but also design architectures to support various pricing policies and charge users for QoS provided. More details on the modeling and pricing, and performance models (e.g. queueing models) for QoS provisioning (such as response time and video/image quality) is given in [10].

- **Long-Term Pricing**: One can argue that prices should be set on a monthly or weekly basis. However, this can cause potential problems as multiple levels of services cannot be supported flexibly for the various users. For example, users may not want services on all days, instead they would prefer to get services on demand.
- **Trade-offs**: If the agents in the information economy do not exist, then users will have direct access to the DLSystems. In such situations, users will send requests for information at arbitrary points in time. Therefore, guarantees for services and prices can be complex, as several thousands of requests could potentially arrive. If prices are not stable for long periods of time, then users can get confused about choosing the right agent.

Charging, Accounting and Payment

Users will see a charge for access ahead of time, and will have a choice in choosing the right agent for searching and access services to the information objects. Assuming that prices fluctuate slowly, several practical issues of billing and payment can be resolved very easily. There are already a variety of payment models (see [6]), and a

whole new chapter has begun in this area with the multi-industry project initiated by CommerceNet and the W3 Consortium.

Implementation Issues and Platforms.

We are currently investigating ways to implement a charging/accounting framework over NCSTRL , which is an evolving, distributed Digital Library System for Computer Science Technical Reports. NCSTRL currently uses DIENST, which uses WWW framework (HTTP and HTML) for communication and searching, and CGI for interaction with Web browsers. We plan to extend DIENST to support a charging framework as shown in Figure 3. The extensions are going to be done on a testbed of DIENST servers that can support objects such as video clips of lectures. For example, universities and research institutions might sell services to view video lectures for a small fee.

In Figure 3, a distributed name service architecture along with the DIENST architecture is shown. The name service is designed to provide information about services offered by the Digital Library servers (DIENST servers). This information includes service prices and copyright information to access the objects. In the Figure, we also show a Digital Library site (example: campus or an organization), where several WWW based clients query for information from the Digital Library Server.

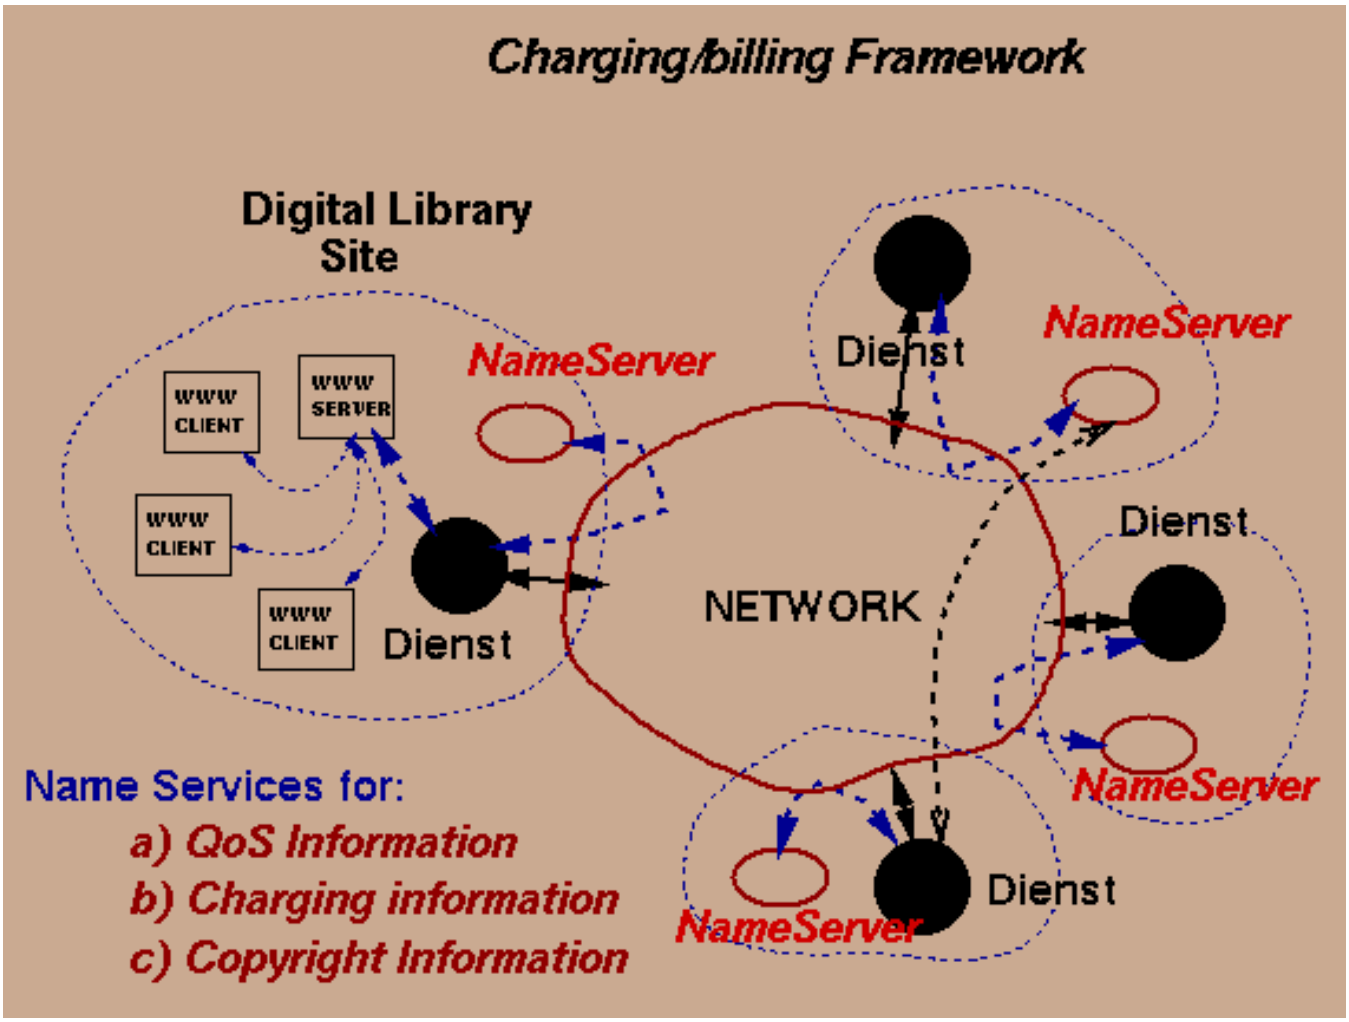

Figure 3.

Our goal is to give to the user, along with the search/query results, the qualities of service offered and the corresponding price for access. For example, if the objects are multimedia objects, then along with the search results come the service levels (MPEG or JPEG quality video) offered by the corresponding Digital Library servers

and the price to access those objects. In Figure 4, we show the model that we are currently building and testing. DIENST currently returns a list of URLs when a user submits a query (keyword). We plan to extend this by adding QoS and price information along with each URL in the list.

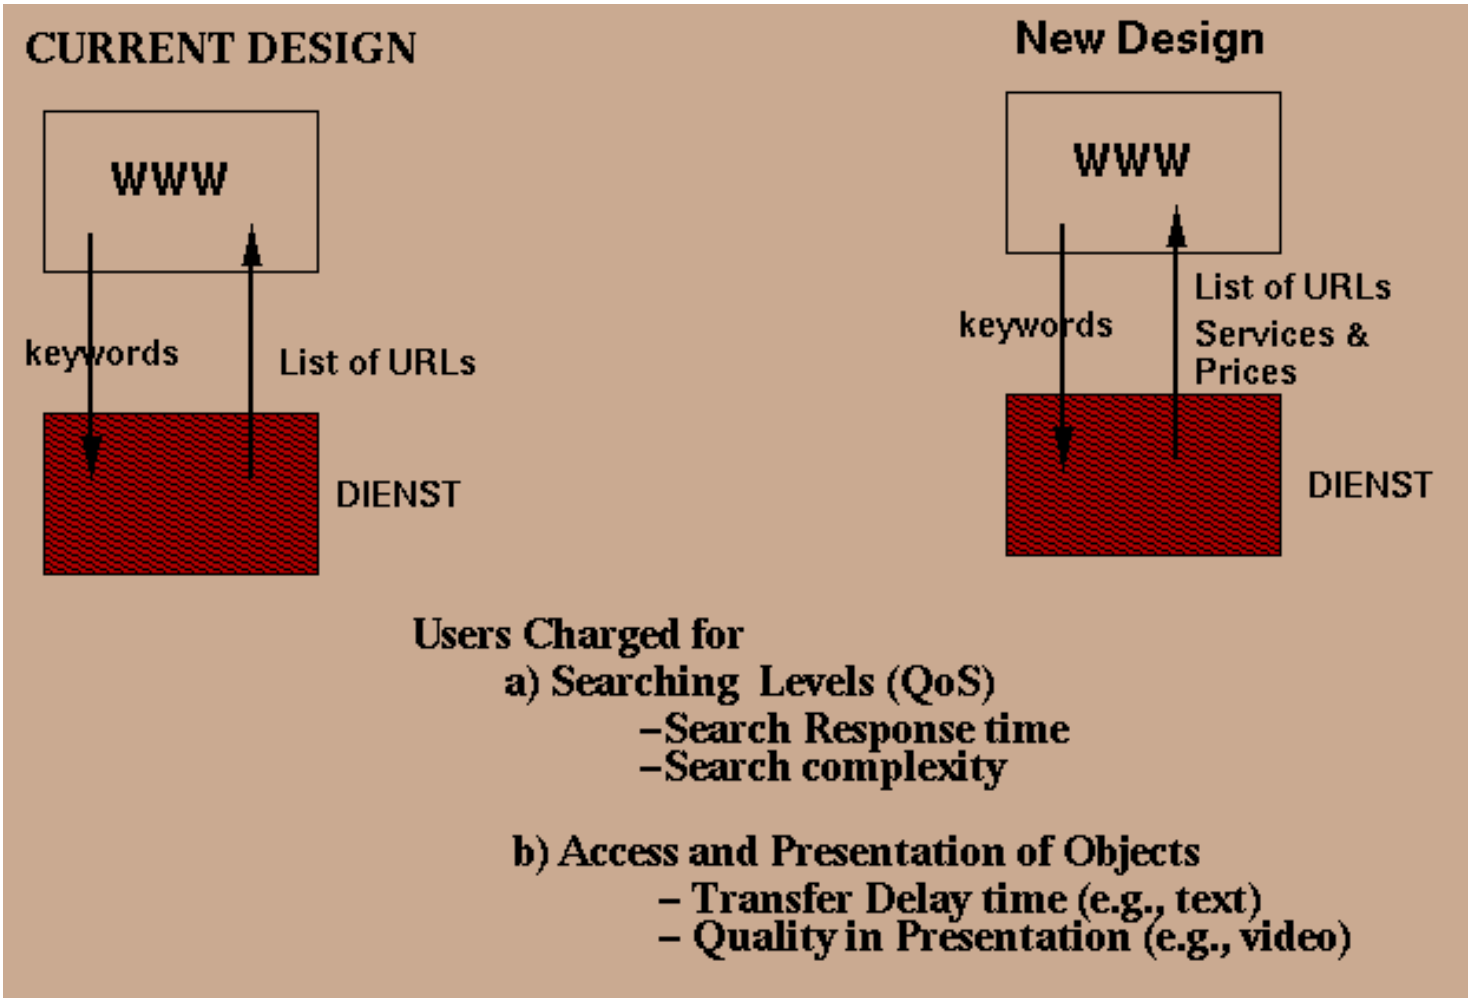

Figure 4.

In Figure 5, we show an example of a payment scheme that we plan to study. The payment is between a stock broker and a consumer, and between the Digital library system and the consumer. The Digital Library contains information about companies (their earnings, stock performance, and other related information) and news papers. The user first queries the Digital Library for stock information such as stock price, earnings and performance plots of a specific company or a collection of companies. The user then decides on stocks to invest and contacts a stock broker to submit the request. The stock broker buys stock for the consumer using the consumer-account in the bank. The interesting issues are in storing and presenting information about companies in Digital Libraries and payment mechanisms.

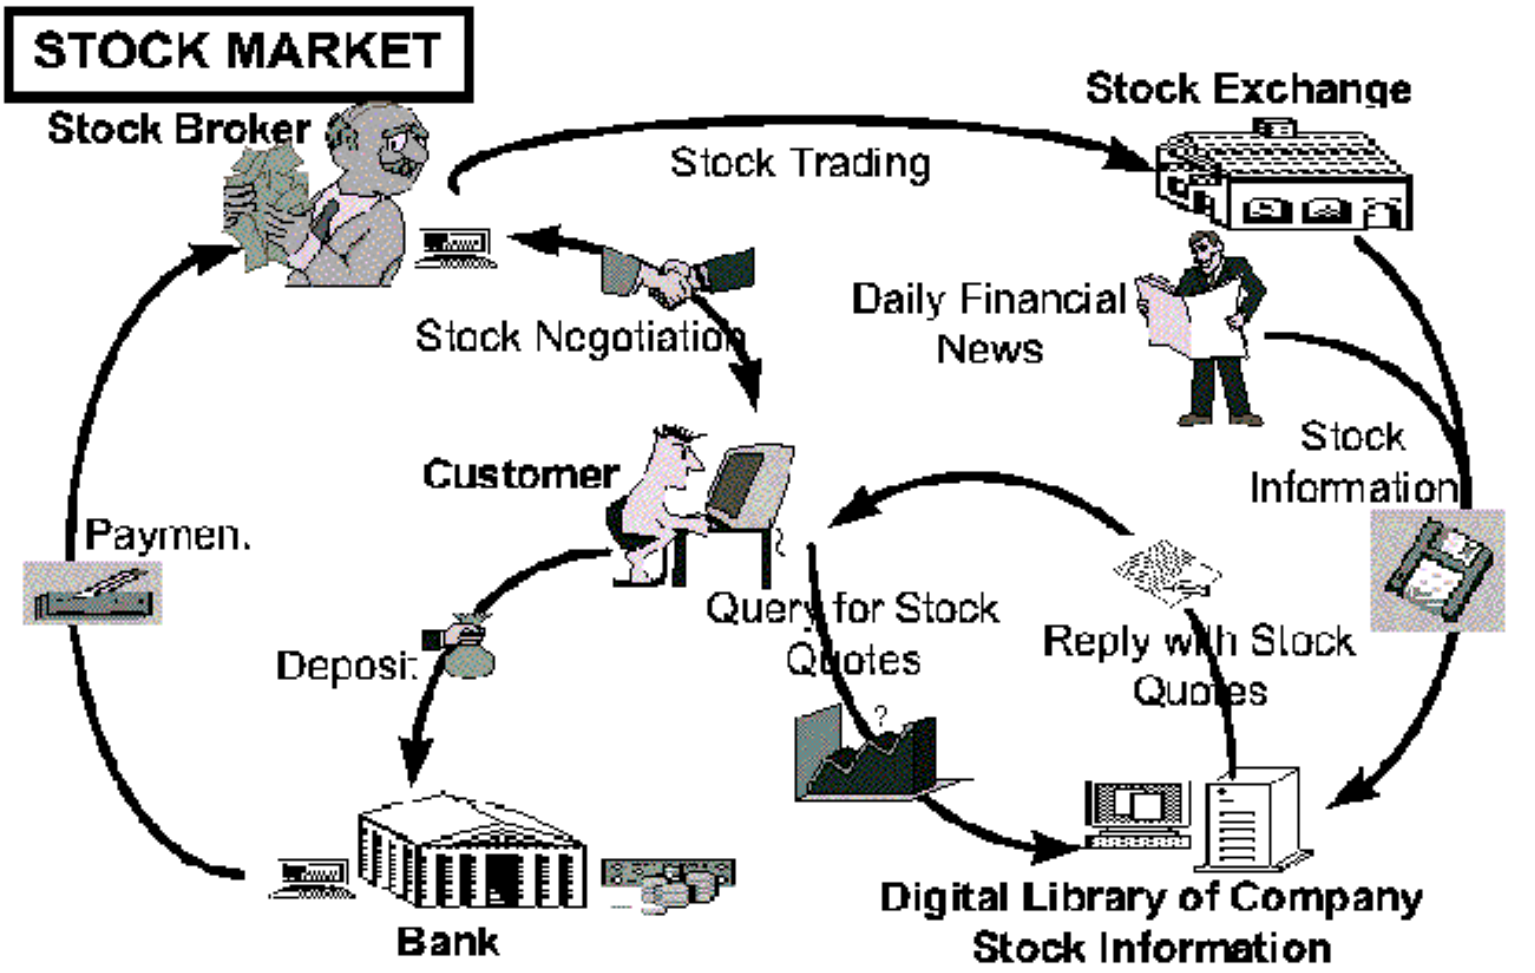

Figure 5.

Future Work

In this story, we presented an economic framework for pricing services and charging/billing users in Digital Libraries. In this framework, we presented a competitive model of interaction among the players of the economy. Our goal is to use such a framework to support various other pricing schemes in order to charge/bill users for services. We also presented an architecture for charging/billing in a Digital Library system using the DIENST technology. Our future work is to study and implement various payment schemes in large Digital Libaries, and to investigate issues of copyright and protection. We also plan to investigate issues of caching and prefetching in large Digital Libraries, where agents sell caching services to consumers.

Acknowledgments

The research work in this paper were supported in part by the Center For Telecommunications Research, Columbia University, New York, in part by the LYDIA project (ESPRIT III P8144), and in part by the CASHMAN project (ACTS AC 039).

References

- [1] D. F. Ferguson, C. Nikolaou, J. Sairamesh, and Y. Yemini, *Economic Models for Allocating Resources in Computer Systems*, Market based Control of Distributed Systems, Ed. Scott Clearwater, World Scientific Press, 1995.
- [2] J. Sairamesh, D. Ferguson, and Y. Yemini, *An Approach to Pricing, Optimal Allocations, and Quality of Service Provisioning in High-Speed Networks*, *Proceedings of the INFOCOM'95*.

- [3] D.F. Ferguson, C. Nikolaou and Y. Yemini, *``An Economy for Flow Control in Computer Networks,''* Proc. of the INFOCOM'90.
- [4] D.F. Ferguson, Y. Yemini and C. NiKolaou, *``Microeconomic Algorithms for Load Balancing in Distributed Computer Systems,''* Proc. of International Conference on Distributed Systems (ICDCS'88). Technical Report
- [5] Service Models and Pricing
- [6] InterPay: Managing Multiple Payment Mechanisms in Digital Libraries by S. Cousins, S. Ketchpel, A. Paepcke, H. Garcia-Molina, S. Hassan, and M. Röscheisen
  to appear in Digital Libraries 95
- [7] Connection Establishment and Pricing
- [8] The Information Economy
- [9] Market-Oriented Programming
- [10] J. Sairamesh, C. Nikolaou, D. F. Ferguson, Y. Yemini, *Pricing Services in Digital Libraries* , Working paper, to be published as a joint Columbia University CS-TR, New York and ICS-FORTH CS-TR, Crete, Greece, 1996.
- [11] C. A. Waldspurger, T. Hogg, B. Huberman, J. Kephart, and S. Stor netta, *Spawn: A distributed computational ecology*, IEEE Trans. on Software Engg., 18(2), February, 1992.

---

hdl://cnri.dlib/february96-sairamesh