Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE DIRECT
Computer Communications 27 (2004) 522–528
computer communications
www.elsevier.com/locate/comcom

Content pricing in the Internet

Burkhard Stiller[a,b,*], Kevin Almeroth[c], Jörn Altmann[d], Lee McKnight[e,f], Maximilian Ott[g,h]

[a]Information Systems Laboratory IIS, University of Federal Armed Forces, Munich, Germany
[b]Computer Engineering and Networks Laboratory TIK, ETH Zürich, Switzerland
[c]Department of Computer Science, University of California, Santa Barbara, CA, USA
[d]University of Berkeley, Berkeley, CA, USA
[e]School of Information Studies, Syracuse University, New York, NY, USA
[f]Massachusetts Institute of Technology MIT, Engineering Systems Division, Cambridge, CA, USA
[g]Semandex Networks, Inc., Princeton, NJ, USA
[h]WINLAB, Rutgers University, Piscataway, NJ, USA

Received 8 August 2003; accepted 8 August 2003

Abstract

Pricing content defines a major challenge for tomorrow's Internet, since existing models appear to be unworkable. However, the provider of content, commercial ones as well as private persons, show in a commercialized world of interactive trade the clear need for content charging systems, which in particular are necessary for electronic content. Those systems generally are beyond the pure technical scope of a system and they go beyond the economic point of view as well. Only an integrated approach, fully understood in their details and technological–economic interactions, will provide the platform for future success. Although traditional manners of content charging will work in the Internet, such as credit card payments or paper bills, the set of constraints and technical possibilities has changed, e.g. a customized newspaper can be provided only electronically. This new approach for pricing content in the Internet needs to cover mechanisms for resource allocation, service fusion, payment schemes, valuation of content, pricing, charging for services and content, and a management system.

Keywords: Internet pricing; E-content; Information network; Service differentiation; Content market

1. Introduction

With the transmission capabilities of today's wired and wireless networking technology, the delivery of services and in particular content to a user becomes an important business model. Although a number of basic principles for supporting dedicated content is well known, a content-based view on pricing such services in the Internet shows a large diversity of opinions and forecasts. It remains a largely unsolved issue, how an overall service architecture for content, its Quality-of-Service support, and economic aspects are to be detailed and how they may be integrated efficiently. In particular, the pricing of content, mainly e-content, challenges IP services and Internet Service Providers as well as the management tools of peer-to-peer networks, services, and content. In addition, those services, which have to support the offer of content in a broad-scale, and probably with a strong one-to-one marketing focus, have to be defined and they must be supported by the network or middleware between the network and the application or customer. While pricing of network services has been worked on, cf. Ref. [2] for an overview and Ref. [6] for market-managed work, Internet tariffing shows the so-called feasibility problem [9].

However, the technical set of content models, transaction models, resource allocation schemes, and networking models determine only one side of the problem. The other side of the problem originates from the fact that new markets arise and they need theoretical as well as practical considerations. It considers customer behavior, utilities and utility functions, valuation of content and services, and

* Corresponding author. Address: Information Systems Laboratory IIS, University of Federal Armed Forces, Munich, Germany.

E-mail addresses: stiller@tik.ee.ethz.ch (B. Stiller), almeroth@cs.ucsb.edu (K. Almeroth), jorn_altmann@acm.org (J. Altmann), lmcknigh@syr.edu (L. McKnight), max.ott@semandex.net (M. Ott).

0140-3664/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.comcom.2003.08.017

economic and technical feedback loops between customers and providers. While the selling of customized content can be managed and supported by charging systems, the challenge lies in the provisioning of differentiated pricing, e.g. based on the customer's preferences, the customer's subscription, or its location.

1.1. Why to price content?

People and machines communicate over the Internet in order to share information of widely varying types. Some of the content is bought and sold, while a great deal of content is at present processed and consumed without charge. The challenge for providers of Internet content who hope to be paid for their content is great. Technical, economic, business, and policy factors, and most importantly user behaviors and preferences, affect demand for the myriad possible types of content. Firms continue to refine schemes and offer new content types for sale via the Internet, while researchers continue to offer suggestions to answer questions.

There is no one-size-fits-all answer to this conundrum, for it depends very much on what precisely is offered for sale, and what alternatives the consumer has. Pricing mechanisms for Internet content must take into account demand factors, or there is not much point—and certainly no profit—in creating supply. The optimal strategy for developers of Internet content pricing mechanisms is, therefore, to seek to become a platform technology, useful for pricing a wide variety of content, according to a wide variety of differing pricing schemes. Easier said than done—of course!

As indicated in Fig. 1, content charging is effected by three different roles. While the customer will utilize content offered from a content provider, the network provider is responsible for the delivery of such content. This tri-lateral situation is different compared to commonly used contracting scheme, where bi-lateral contacts determine the standard approach. In addition, the set of tasks to be performed by each of those three roles mentioned seems to require a minimum level of cooperation, if an economically and technically efficient system is intended. Certainly, this type of cooperation must show a market-driven approach, since otherwise no real incentives will be available to function properly and scale in larger dimensions of customers, geographically distributed network providers, and number-wise as well as content-wise content providers. This paper discusses and provides ideas, migration paths for possible architectures, roles and tasks, and future developments and necessities on pricing content in the Internet.

1.2. Paper overview

The remainder of this paper is organized in three parts. Section 2 introduces an approach to a market for e-content, in which pricing, customer behavior, and resource constraints are discussed. In addition, Section 3 discusses an architectural approach on content charging, where network service differentiation is related to content. Furthermore, Section 4 presents the idea of an information network, where an application-related schema offers an end-to-end characteristic for content. Finally, Section 5 summarizes and presents those topics of Internet content pricing.

2. A market for e-content

Available bandwidth and Internet usage has increased tremendously over the past few years. According to recent surveys, broadband connectivity is accessible to more than 75% of homes in the US alone. Similarly, the number of people accessing the Internet and using it for purchases has been increasing steadily. On the other hand, backbone networks are **over-provisioned to the extent that many backbone links are utilized at no more than 5% of their capacity. Together, these trends suggest that there is an enormous potential for selling multimedia software in the Internet. For example, content providers can sell downloadable CDs over the Internet. E-books have been suggested as a viable alternative to paper-based books. These are examples where the customer downloads the content. One can also envision services based on streaming content. The current Internet has numerous streamed services available today. Most of these services are free, the revenues earned primarily through advertising. However, it has been argued [7] that advertising alone cannot sustain the market. The challenge, therefore, is to envision scenarios where customers pay for content. Our work is focused on the opportunities and challenges of selling e-content and services.

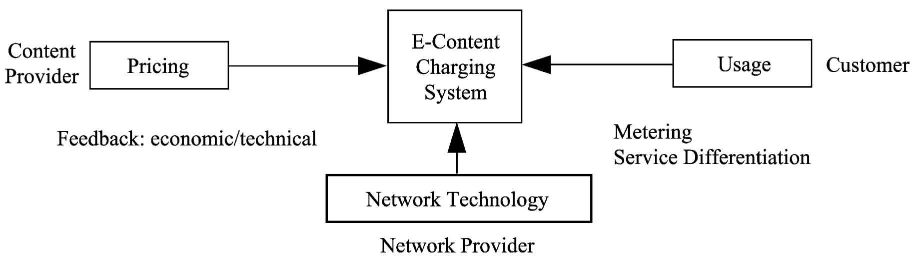

Fig. 1. Content charging—roles, tasks, and major mechanisms.

2.1. Transaction models

In an e-content market, the content provider has two primary transaction models for selling content: (1) quoting a price to the customer and (2) letting the customer quote a price to the content provider. We call the first a Quoted-Price Model. In this model, if the customer agrees to the price, the content providers initiate the download (or streaming) process within a maximum time delay. The second model is similar to a Sealed-Bid Auction. In this model, different customers quote what they are willing to pay for the content of their choice. Based on the available resources and revenue considerations, the content provider clears the market at regular intervals. The download process is initiated for the winning bids as soon as the market is cleared. One can envision specialized strategies like bundling, combinatorial auctions, etc. in each of these transaction models for maximizing the revenue. One can also combine these two models in the form of automated bargaining where the agents for the content provider and the customer bargain over the price.

In this section, we primarily discuss the quoted price model. A quoted price model is adopted in conventional markets for content. For instance, CDs, books, etc. are sold using this model. Similarly, video and DVD rental stores quote a price to rent movies. Providers in such markets have enormous expertise in pricing these forms of content. Strategies evolved in conventional markets cannot be applied in an Internet setting. There are three reasons for this. First, the customer population is diverse. With Internet speeds increasing, it is possible that customers from geographically dispersed regions make requests for content at the same web-site. In contrast, customers in conventional markets are typically local. Second, customer behavior in the Internet can be dynamic. For instance, a sudden news event can affect customer behavior in near real-time. Conventional wisdom on pricing cannot keep up with such dynamic behavior. The third reason is that in an Internet setting, diverse content share the same distribution medium. To illustrate this point, consider the case where a content provider has only enough resources to accept one request. Suppose that there are two requests—one for content A, where the customer is willing to pay $5 and the other for content B, where the customer is willing to pay $10. By rejecting the request for A (by quoting a price greater than $5), and accepting the request for B (by quoting exactly $10), the content provider generates more revenue. Thus, the content provider must intentionally over-price content A in order to increase the returns per unit resource consumed. On the other hand, in a conventional market, it would be counter-productive to intentionally over-price any of the content.

2.2. Static, discriminatory, and dynamic pricing

In an Internet setting, the content provider has enormous flexibility in quoting a price to individual customers. This is because the price quoted to one customer is not visible to other customers. The content provider thus has three choices:

- *Static pricing*. Quote a price which changes infrequently. A static pricing scheme appears to be the most acceptable to customers. Customers know what to expect, and the price is fair to all customers. In fact, if nothing is known about individual customer preferences, and resources are assumed to be infinite, it can be shown that there is a fixed price that maximizes revenue [4]. However, there are some serious problems with a static pricing scheme. First, finding the optimal fixed price is a non-trivial task. Second, the proof that there is a fixed price that maximizes revenue assumes that customer behavior is well defined and temporally invariant. This may not be true in real life. For example, customers are likely to spend more during weekends or other holidays. Customer behavior may vary over short time durations also. Movie theatres take advantage of such time of day variations in customer behavior by charging different prices for matinee and evening shows. Third, with resource constraints, it can be easily shown that the optimal fixed price depends on the currently available resources.
- *Discriminatory pricing*. Quote a price that varies with customers. For example, let two customers request the same content from a content provider at the same time. The content provider can quote different prices to the two customers based on their profiles. If the content provider can guess the customers' valuation correctly, then the content provider can increase revenues. Since customers may pay different amounts for the same content, discriminatory pricing can appear unfair to many customers. While in the short term, discriminatory pricing can maximize revenue, in the long run, the number of dissatisfied customers is bound to increase.
- *Dynamic pricing*. Quote a price that can change with time. Simultaneous requests are quoted the same price irrespective of the individual customers' valuations. Prices may typically change based on current system load, request arrival rate, and other external factors. A dynamic pricing strategy appears to be the best strategy when the content provider does not have any specific information about individual customers. A dynamic pricing scheme is fair to customers who request content at the same time. Moreover, a dynamic pricing scheme can also be used to experiment with the customer behavior and learn how customers react to different prices. For example, the content provider can charge a set of 'test prices' over a period of time and learn what price maximizes revenue. In our research, we have developed strategies for dynamic pricing that can achieve revenues comparable to the maximum expectation.

2.3. Customer behavior

To quote a price that maximizes revenue, the content provider should have an understanding of how customers react to the quoted price. Even in the absence of specific information about individual customers, the content provider can make some generalized observations about the entire customer population. We list these observations below:

- Customers have a finite valuation for content. If the quoted price is greater than this valuation, then customers do not proceed with the transaction. As a result, there exists some finite price above which no customer will purchase the content.
- Customers are rational, i.e. they would prefer a lower price to a higher price.
- The number of customers who will accept a quoted price does not increase with price.

The content provider can thus use a simplified model of customer behavior. The content provider is interested in the fraction of requests that will result in successful transactions. This fraction is a function of the quoted price. Let us assume that the content provider is using a dynamic pricing policy. For a price x, let $f(x)$ denote the fraction of customers who will accept the price. Let x(low) be a price below which $f(x)$ is exceptionally high, and let x(high) be a price above which $f(x)$ is exceptionally low. Then $f(x)$ can be approximated in the domain $[x(\text{low}), x(\text{high})]$ using some non-increasing function. We have developed a family of decreasing functions which depend on a parameter d. Fig. 2 shows this family of non-increasing functions. By experimenting with different prices to observe the fraction of customers who accept the price, and using statistical methods like squared error minimization [5], one can estimate the parameter d, and the threshold prices x(low) and x(high). Once all parameters are known, the content provider can predict the customer behavior and, thereby, find an optimal price for the content.

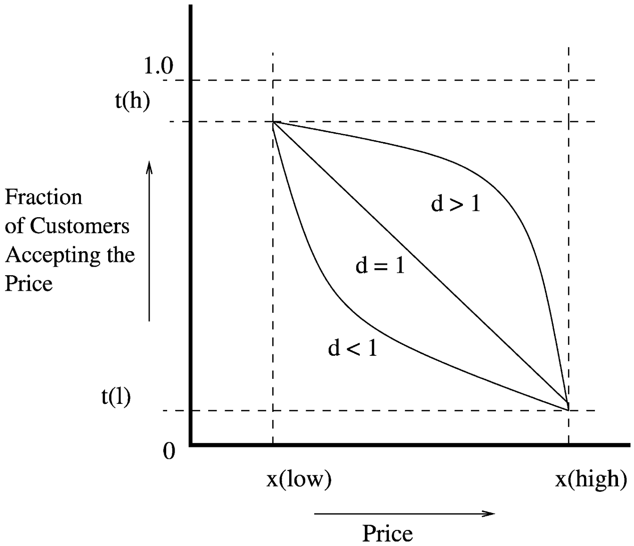

Fig. 2. Customer model.

2.4. Resource constraints

The number of requests that can be served by a content provider is limited by available server and bandwidth resources. Further, the number of new requests that can be served can vary based on the number of requests already being served. To manage resources efficiently, the content provider must quantify the available resources. For simplicity, let us assume that all the customers get the same quality of content, i.e. there is only one level of service. Under this assumption, a simple abstraction would be to model the resources as 'logically equal-capacity' channels. Typical resource constraints in a real-time system include memory, bandwidth and latency. The minimum resources required to serve one request are together labelled as a single channel. When a request is made and the customer accepts the quoted price, the content provider allocates a channel for the entire duration of the download/streaming process. Since the number of free channels can vary at any given time, resource constraints can affect how many requests can be serviced. The resource constraints in the system can be analytically modelled in the revenue maximization problem using the notion of system utilization. System utilization is the relative fraction of time for which system channels are 'occupied'. However, here we use a system utilization to formulate the problem of revenue maximization as constrained optimization problem. In general, when the rate of channel allocation exceeds the maximum channels that can allocated per unit time (i.e. resources are constrained), the price should be increased.

In systems where resources are scarce, as may often happen at times of peak load, a technique called batching can be employed to improve scalability of the system. In batching, requests for the same content are aggregated over a period of time. Then, using a one-to-many delivery mechanism like multicast or satellite broadcast, the content is streamed from the server. Only one channel needs to be allocated to serve multiple batched requests for the same content. The trade-off in this model is the delay incurred by customers before the download/streaming process is initiated. Various batching schemes have been explored for improving the scalability and efficiency of the system. However, none of these approaches consider the problem of how to set the price in a batching scheme. Clearly, the customer's value for the content will decrease with an increasing delay in initiating the download/streaming process. Moreover, long delays will increase customer dissatisfaction. The content provider may solve this problem by bounding the maximum delay as part of the service contract. In our work, we have addressed the problem of pricing when the maximum delay in a batched content delivery system is bounded.

3. A network and content service differentiation architecture

Pricing and charging for information is one of the major issues, which has to be solved, in order to guarantee further growth of the Internet. Most of the current business models of information service providers (i.e. content provider, retailer) on the Internet are only based on revenues from advertisement (i.e. selling advertisement space on their web site) and retailing (e.g. selling books). However, these revenues hardly cover expenses for network connectivity, storage, and server capacity, necessary to run their Internet businesses as narrow-band network services. As recent developments showed, several information service providers had to change their business models or go out of business when their revenues from advertisement went down.

For sure, revenues from advertisement and retailing are not sufficient to cover even higher expenses for providing broadband network services, required for multimedia content. These services consume additional network, server, and storage resources. In order to resolve this dilemma, information service providers have to start charging their customers for the content or for the additional resources consumed.

Currently, only a few information service providers actually charge for their content. The pricing and charging models deployed are very simple. They are subscription-based or transaction-based. Under a subscription-based model, customers are paying a monthly charge for accessing real-time information (e.g. Bloomberg—financial information). A usage-based fee is charged for accessing archived information (e.g. New York Times), if it is a transaction-based model.

3.1. Problem statement

Although there are customers who are willing to pay for information (i.e. multimedia content), it is still not possible to do it. This raises the question what technology is missing in order to charge for content, since most of the necessary infrastructure to deliver rich media is in place already. Customers in metropolitan areas have access to broadband access technology. The Internet backbone is over-provisioned (e.g. the utilization of the Internet backbone is less than 15% and large capacities of black fiber are unused). Software solutions are available to deliver multimedia content to the end-users' homes. Even multimedia content is available, such as live sport events and on-line gaming.

Examining this scenario, it becomes obvious that a simple charging system is missing, which enables service providers to charge for content and the infrastructure resources consumed. This, in turn, would require a future Internet that can provide quality-of-service and allows negotiating charges between sender and receiver of information. With such a charging system in place, new business models are possible. However, the implementation of new business models requires the understanding of pricing of bundled information and infrastructure services, the payment flows between all participating service providers, and the provisioning of quality of network services. It also requires a management system that can handle the exchange of usage charging information, service provisioning information, and traffic data between different providers and end-users.

3.2. Solution proposals

In such an environment, information service providers and their customers must be able to allocate resources on demand from infrastructure providers. This dynamic resource allocation, in turn, enables to sell information services in two different ways. Either the network service charges are bundled into the content service or the end-user pays for the content service and the network service separately. Both approaches help recovering costs for infrastructure expenses.

At the same time, infrastructure service provider (i.e. network service provider, storage service provider, processing service provider) will be able to increase their revenues by providing customized services. Customers, who value the service higher, will receive more resources. Currently, infrastructure service providers can only provide best-effort service, and if they only charge a flat rate, they are not even capable to stop over-usage. The proposed system will prevent this.

The purchase of a video from a video service provider can be illustrated as follows. The end-user purchases the video service on the web site of the service provider by selecting the quality level that he requires. In this example, the end-user chooses 'good' as the quality for the video transmission. Since the video service is bundled with the infrastructure service, the video service provider has to make the necessary arrangement with infrastructure service providers so that the end-user receives the video at the selected quality level. However, the end-user will be able to increase the quality selection either by choosing a higher quality level on the video provider's web site or by selecting a higher priority through his network quality controller [1]. In both cases, the additional charges for the higher quality have to be paid by the end-user. In one case, the extra charges have to be paid to the video service provider and, in the other case, the charges are paid to the access service provider of the end-user.

4. Proposal for an information network

The commoditization of data networks together with the continuous rapid improvement in the underlying technology has turned the telecom business into an economical nightmare [3]. With 'data piping' at the end of the value chain, there are too many parties between a paying customer and the pipe provider. Customers facing services can often be value priced, while anything after that is cost-based, especially if there is a low barrier of entry for providers of a generic service. In fact, technology advances and the accompanying cost reductions reward the newcomer (until the next one comes).

While the discussions between proponents of dumb and smart networks have been going on for a long time, it is often forgotten that end-users really do not care. Their interface to a service is the screen and keyboard. If pictures on the screen are generated inside the terminal, or on a server in a data center, or inside a routing element is of little relevance. In fact, we have seen a trend, at least on the edge, to move from 'blindly shoveling bits' to first 'color' them and then treating them differently. Companies, such as Akamai and Alteon provide services and equipment for routing URLs instead of IP addresses. By shifting the value proposition from bandwidth to 'reducing the user's waiting time' closer to the user, they could charge more.

Taking this further, what are people using the network for? One can argue, that it comes down to exchanging information. What starts as data, becomes useful information to a person if it is relevant. E-mail and Instant Messaging are very useful tools, because they allow us to communicate effectively. Unfortunately, the information used to distribute a message is normally disjoint from the message content itself: "I just found this leaking container labeled X, what should I do". Where should that go. To a mailing list, or newsgroup? Which one and how up-to-date is the subscription list anyway. Even if I know an expert, she may be at lunch. In summary, we need an information network that maintains all useful properties of the underlying data network, but provides new functionalities which are much closer to the (social) networking needs of its users.

4.1. Solution proposal

The packet network functions like the postal service. Every packet is labeled with the address of the recipient. What we need is a network where all 'consumers' declare their interest to the network, and all packets are labeled with a description of their content. Now the network delivers a copy of this packet to everyone whose interest matches the content descriptor. This is not a new idea; many will immediately recognize this as a 'publish and subscribe' (PS) system.

The problem with PS systems is that they have very poor scaling properties. PS systems started as software buses, where every message is broadcast to every consumer where local filters separate the wheat from the chaff. Topic channels based on IP multicast addresses is also a rather limited approach (mapping inherently multi-dimensional data into a single dimension—multicast address range—is tricky and expensive to say the least).

Obviously, we need a solution where knowledge itself is not embedded in the network. Requiring a 'network upgrade' whenever we introduce a new topic is not a workable solution. Fortunately, XML schemas and all the related standards bodies provide us with a framework to design a general framework based on the basic XML schema standard. The application specific schema becomes an end-to-end property. So while the network cannot distinguish a baseball bat from a long-legged bat (M. volans), it does know about how semantic relationships are encoded in taxonomies [10].

The structure of schemas also allows for aggregation (and compaction) of interest profiles, enabling the practical deployment of large-scale information networks. In addition, exposing more information about a packet to the network also allows for additional services, such as giving priority to a 'help, fire' message over the 'hi mum' one. Fully structured messages such as sensor readings also pave the way for providing generic processing resources 'on the way' [8].

5. Summary and outlook

This paper addressed in an overview four following areas: (1) motivation for pricing and charging needs, (2) a market for e-content, (3) a network and service content differentiation architecture, and (4) an information network.

Although many different pricing models exist and Internet charging systems have been developed [11], the content-based charging solution with enough flexibility for all three roles of the customer, the network provider, and the content provider do not exist yet. Therefore, the e-content market proposal has the potential to revolutionize how customers exchange content. A sound business model is a prerequisite for its success. This part identifies those important factors that govern the success of any business model. Thus, the work will provide useful guidelines to content providers. The solution on the differentiation architecture presented still shows several open issues, such as the collection of payments, the split of charges among service provider, and the impact on existing business models. However, regarding the collection of payments, the question is who will collect the charges from the end-user. That can be either the infrastructure service provider, who has already established the customer relationship with the end-user, or the content service provider.

Both approaches have advantages and disadvantages. Considering the business site of this problem, it becomes obvious that the infrastructure service provider as well as the information service provider try to maintain the relationship to the customer, since this is the most profitable part in the value chain. Once the payments are collected, they have to be split among the service providers. The question here is: how can that be done in a fair and efficient way? Finally, if the proposed system will be implemented, it might have a strong impact on existing businesses. For example, imagine such a system, which enables artists to collect royalties for their music directly from the end-user. The music industry, which already struggles to keep control over the distribution of their content, would have to change their business model in order not to become obsolete. Finally, with respect to the information network, the view changed as well, since people pay for information and not data. The value, and in turn the price, of a data packet being distributed by a network is ultimately linked to the value of the contained information to the recipient. What is valuable information to me, may be wasted resources to you, or even to me under different circumstances, such as time and context. However, connecting billions of information producers and consumers at the information level cannot easily be achieved by simply deploying PC or terminal-based applications, or wider pipes. It requires networking at a different plane.

Acknowledgements

Many thanks are addressed to Srinivasan Jagannathan, working on the content of Section 2 with K. Almeroth. In addition, many thanks go to David Hausheer, who collaborated in content management discussions with B. Stiller.

References

[1] J. Altmann, H. Daanen, H. Oliver, A. Sánchez-Beato Suárez, How to market-manage a QoS network, IEEE Conference on Computer Communications, INFOCOM, New York, USA, June 2002.

[2] M. Falkner, M. Devetsikiotis, I. Lambadaris, An overview of pricing concepts for broadband IP networks, IEEE Communications Surveys (2000) 2–13. 2nd Quarter.

[3] R. Googin, How Networking Advances Screwed Up the Economy, http://www.isen.com/archives/0112161.html (16/12/2001).

[4] S. Jagannathan, K. Almeroth, Dynamics of price, revenue, and system utilization, IEEE/IFIP Conference on Management of Multimedia Networks and Services, Chicago, Illinois, USA, October 2001.

[5] S. Jagannathan, J. Nayak, K. Almeroth, M. Hofmann, On pricing algorithms for batched content delivery systems, Electronic Commerce Research and Applications Journal 1 (3/4) (2002) 264–280.

[6] M3I, Market managed multi-service internet, Fifth Framework European Union Project, IST Program, No. 11429, 2002, http://www.m3i.

[7] A. Odlyzko, Paris metro pricing for the Internet, ACM Conference on Electronic Commerce (1999) 140–147.

[8] M. Ott, L. French, R. Mago, D. Makwana, D. Reininger. Semantic Multicast based on XML Routing. http://www.semandex.net/whitepaper/semantic_multicast.pdf (30/01/2002)

[9] P. Reichl, P. Flury, J. Gerke, B. Stiller, How to overcome the feasibility problem for tariffing Internet services. The Cumulus Pricing Scheme, IEEE International Conference on Communications (ICC'01), Helsinki, Finland, June 11–14, 2001.

[10] Semandex, What is Content-based Routing?, 2002, http://www.semandex.net/whitepaper/Semandex_Content_Based_Routing_101.pdf.

[11] B. Stiller, P. Reichl, J. Gerke, P. Flury, A generic and modular internet charging system for the cumulus pricing scheme, Journal of Network and Systems Management 9 (3) (2001).