

Price Issues in Delivering E-Content On-Demand

SRINIVASAN JAGANNATHAN

and

KEVIN C. ALMEROOTH

Department of Computer Science

University of California

Santa Barbara, CA 93106-5110

{jsrini,almeroth}@cs.ucsb.edu

The explosive increase in Internet bandwidth and usage opens a vista of opportunities to sell multimedia-rich software and services using the Internet. Once *e-content* is created, the cost of replication is negligible. Customers can download the e-content immediately after online transactions. Alternately, the content provider can stream the content to the customers. A sound business model is necessary for the success of such an enterprise. In this paper, we examine the determinants of revenue for an Internet based on-demand content delivery service. The determinants of revenue are: transaction model, pricing strategy, customer behavior, distribution resources, and competition. We briefly describe each of these factors and discuss how they relate to revenue. Our belief is that by better understanding how these factors affect revenue, content providers can develop services that generate more revenue while also being more compelling to users.

1. INTRODUCTION

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 U.S. 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 not more than 5% of their capacity. Moreover, the multimedia capabilities of computers are becoming better and cheaper. Together, these trends suggest that there is an enormous scope for selling multimedia software

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2002 ACM 0000-0000/2002/0000-0018 \$5.00

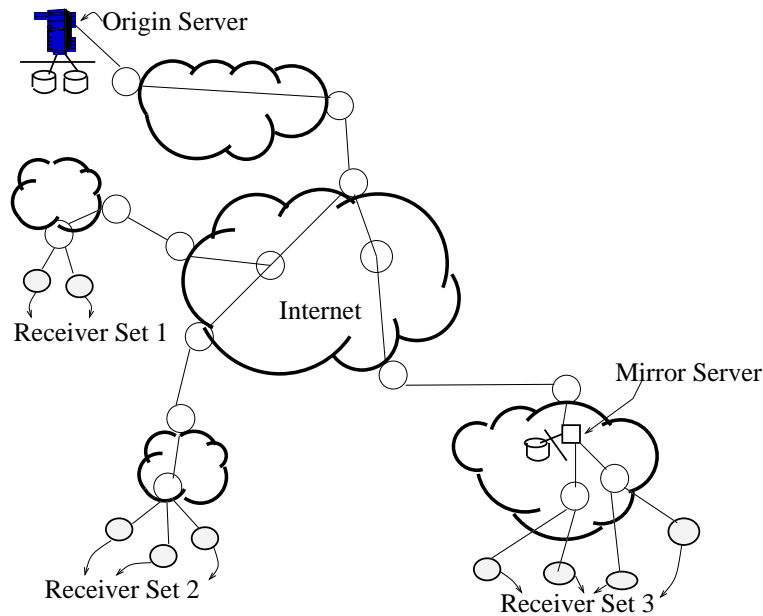


Fig. 1. A Content Delivery Architecture

in the Internet. For example, content providers can sell downloadable CDs over the Internet. E-books have been suggested as a viable alternative (on fast multimedia computers) 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, as Odlyzko [Odlyzko 2000] argues, advertising alone cannot sustain the market. The challenge therefore is to envision scenarios where customers pay for the content. In this paper, we discuss a simple generic architecture for selling on-demand e-content and services and outline issues relevant to e-content commerce.

2. CONTENT DELIVERY ARCHITECTURE

Users' Internet experience can be broadly classified into two categories: *connectivity* and *content*. Connectivity pertains to getting access to the Internet. Content refers to the actual data received/sent by the user. Content presupposes connectivity and adds value to the latter. A clear distinction between connectivity and content is important because the providers of either service can be different. For instance, consider the architecture in Figure 1. It shows an origin server, a mirror server and many consumers. The content provider (servers) as well as the consumers are connected to the Internet. For this connectivity, both parties interact with an Internet Service Provider (ISP). In addition to paying the ISPs for connectivity, consumers pay the content provider for the content. Such payments may be one-time transactions, or a subscription for the content delivery service over a period of time. This is similar to the connectivity model where consumers pay the phone company for telephony connectivity and the ISP for Internet connectivity.

While connectivity pricing has been studied in great detail, content pricing has not received much attention. Our focus therefore is on the issues related to content pricing. We specifically focus on content that is delivered on-demand, i.e., content delivery (or streaming) starts immediately after (or within a short interval of) the transaction.

In the architecture illustrated in Figure 1, receivers in Set 1 and Set 2 receive content from the origin server, while the receivers in Set 3 receive content from the mirror server. There are two possible scenarios for the content delivery transaction. In the first scenario, the customers purchase the content from the content provider and are then redirected to the nearest mirror server. In the second scenario, the customers purchase the content directly from one of the mirror servers. The second scenario is preferable to content providers because of two reasons. First, it allows the content provider to charge a price based on resources available at the local server. This is especially important for services like video-on-demand (VoD), where the number of videos that can be simultaneously streamed is restricted by available server and bandwidth resources. Second, this approach is preferable when the mirror servers are geographically dispersed. For instance, the same content can be sold in Europe as well as in the North America. At a European server, the price will be in euros, and at a North American server, in dollars. The prices can be different at the mirror servers reflecting local customer preferences. This transaction scenario also gives customers added flexibility in choosing the mirror server of their choice. For instance, during evening hours in Europe, when one can expect servers to be over-loaded, customers can choose servers in North America or Asia, where servers may not be as loaded.

In the content delivery architecture outlined above, some resources may be reserved at each server for free content samples and other promotional activities which help the customer search for content. We also assume that the content provider uses sophisticated content protection mechanisms like e-signatures and data encryption when the content is in transit. This helps prevent data theft.

Once the initial infra-structural investments are made, the maintenance costs (for bandwidth, servers etc.) are fixed. This is a reasonable assumption because servers incur fixed costs and bandwidth is typically bought at a flat rate. In addition, once the digital content is created, the cost of replication is negligible. In this case, profit maximization is equivalent to revenue maximization.

3. 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 initiates 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 paper, 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 always 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, different 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 has to 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 because they do not share the distribution medium. In the same example, now suppose the content provider could accept two requests instead of one. In this case, quoting exactly \$5 for *A* and \$10 for *B* would increase the revenue.

4. 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 [Jagannathan et al. 2001]. 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 also vary over short time durations. Movie theaters 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. For example, in the hypothetical example discussed in the previous section, we observed how varying the price of content *A* based on available resources increased revenue.

- 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 information about each customer. If the content provider can guess how much each customer is willing to pay, 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. Another problem with discriminatory pricing is that it may be illegal in many jurisdictions to charge different prices¹. Further, in case of content theft, the content provider has limited legal choices. If the thief is caught, the thief can possibly get away with paying only the lowest price for that content.
- 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 [Jagannathan et al. 2001; Jagannathan and Almeroth 2001a; 2001b].

5. CUSTOMER BEHAVIOR

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

- Customers have a finite valuation for the content. If the quoted price is greater than this valuation then the 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.

¹In conventional markets, mail-in rebates are used as a form of discriminatory pricing.




Fig. 2. Customer Model

—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, say more than t_h and let x_{high} be a price above which $f(x)$ is exceptionally low, say below t_l . Then $f(x)$ can be approximated in the domain $[x_{low}, x_{high}]$ using some non-increasing function. In our work [Jagannathan et al. 2002] we propose a family of decreasing functions which depend on a parameter δ described as follows.

$$f(x) = \begin{cases} t_h & , \quad 0 \leq x < x_{low} \\ (t_h - t_l) \left[1 - \left(\frac{x - x_{low}}{x_{high} - x_{low}} \right)^\delta \right] + t_l & , \quad x_{low} \leq x \leq x_{high} \\ t_l & , \quad x > x_{high} \end{cases} \quad (1)$$

Figure 2 illustrates the 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 [Jagannathan et al. 2002], one can estimate the parameter δ , and the threshold prices x_{low} and x_{high} . Once all the parameters are known, the content provider can predict the customer behavior and thereby find an optimal price for the content. In the pricing scheme we developed, the customer reaction is continuously monitored, and the price is varied at regular intervals².

²Temporal price variations are an inherent feature in many commodity markets.

6. 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”. In our work on pricing [Jagannathan et al. 2001; Jagannathan and Almeroth 2001b; 2001a; Jagannathan et al. 2002], we use system utilization to formulate the problem of revenue maximization as a constrained optimization problem. In general, when the rate of channel allocation exceeds the maximum channels that can be 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 tradeoff in this model is the delay incurred by customers before the download/streaming process is initiated. Various batching schemes have been explored [Dan et al. 1994; Almeroth et al. 1997; Chan and Tobagi 1999] 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 process³. 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 [Jagannathan et al. 2002], we address the problem of pricing when the maximum delay in a batched content delivery system is bounded.

7. LEVELS-OF-SERVICE

A content provider can sell the same content at different qualities or levels of service. This is necessary in the Internet because of heterogeneous connection speeds. For instance, a customer using a 56 kbps modem may request content at a “lower”

³The customers’ reaction may differ based on whether the content is streamed or downloaded in the background.

level-of-service than say a customer using a DSL connection. In such an example, the prices of the content at different levels-of-service are not significantly correlated. However, the price of content at a lower level-of-service should in general be lower than at a higher level-of-service. One may also think of a scenario where a customer has the flexibility to choose from a variety of levels-of-service, based on the minimum quality he desires and the money he is willing to spend. In this scenario, the prices of the content at different levels-of-service are tightly coupled. The customer's choice of the level-of-service depends not only on the minimum quality that desired but also on how higher levels-of-service are priced. For example, suppose that a customer is willing to pay at most \$10 for a streamed video at 64 kbps. If the price of the video at 64 kbps is \$5 and at 128 kbps is only \$9, then the customer has an incentive to choose the higher level-of-service. However, the content provider is possibly losing one dollar in the transaction, since two requests at 64 kbps could have been served using the same amount of resources. The prices of each level-of-service should therefore be chosen with care.

The key feature in a market with multiple levels-of-service is that the maximum number of requests that can be served depends on the relative fraction of resources allocated for each level-of-service. Since this depends on the relative fraction of requests for each level-of-service and also on the price, it is difficult to quantify the resources available to the content provider. One heuristic that can be used is to observe the relative fraction of resource allocation for each level-of-service over a period of time. One can then formulate an approximation to the system utilization and use it in the constrained optimization problem.

8. PRICES IN A COMPETITIVE MARKET

There are many factors that suggest that the e-content market is not perfectly competitive. The nature of the competition also depends on the relationship between content owners and content providers. There are two possible scenarios.

In the first scenario, the content provider is the *only* distributor of the content. This is the case when either the content provider owns the content or has exclusive license to distribute the content. The market is clearly monopolistic. Our work on monopolistic markets is directly applicable in this case. Notice however, that there may be indirect competition from "similar" content from other content providers. For instance, consider two competing video-on-demand services, each owned by competing, and equally famous, Hollywood studios. The studios have monopoly over their content. However, how the customers value their content indirectly depends on the prices set by the competing studio for similar content.

In the second scenario, the content owner has licensed more than one content provider to distribute the content. Assuming that the quality, speed of delivery and other service parameters are identical among the content providers, the market would still be imperfectly competitive because of the market recognition of the individual providers. For instance, Internet web site access patterns have been observed to follow a Zipf-like distribution [Breslau et al. 1999], i.e., most accesses are restricted to a few highly popular web sites⁴. In a market scenario this would imply

⁴In a Zipf-like distribution, the i^{th} most popular web site will have a frequency of accesses proportional to $\frac{1}{i^\theta}$, where θ is some small constant

that brand recognition is very important. Such a trend suggests that competition may not affect prices, at least for a few more years. Customers may typically visit a chosen few web portals like Yahoo! or AOL, for purchasing content. It is also likely that customers may not compare prices at different content providers before making a choice. Our work on monopolistic markets is still applicable in this case.

There has been considerable research on agents that search for the best price for a product [Bakos 1997; Chavez and Maes 1996; Greenwald and Kephart 1999; Tsvetovaty et al. 1997]. Such tools can in the future possibly be modified to search for prices of similar content. Today, not many customers use tools to search for best prices. But one can expect that with time, customers will be more Internet savvy and use agents to search for the best price. The dimensions of competition in such markets need to be investigated in greater detail.

9. SUMMARY AND CONCLUSION

In this paper we introduced the idea of e-content markets. We discussed a simple, yet deployable content delivery architecture and described transaction scenarios for purchasing content. We outlined the quoted-price model for content transactions and introduced static, discriminatory and dynamic pricing schemes. We described how resource constraints can affect revenue and why expertise from conventional markets cannot always be applied in an Internet setting. We proposed how dynamic pricing schemes can be used to learn customer behavior by experimentation. We also suggested how constrained resources can be efficiently utilized using schemes like batching. We then generalized the resource model to enable multiple levels-of-service. Finally, we discussed the nature of competition in e-content markets.

The e-content market has the potential to revolutionize how customers exchange content. A sound business model is a prerequisite for its success. Our work identifies the important factors that govern the success of any business model. Our work will thus provide useful guidelines to content providers.

REFERENCES

- ALMEROOTH, K., DAN, A., SITARAM, D., AND TETZLAFF, W. 1997. Long term channel allocation strategies for video applications. In *IEEE Infocom*.
- BAKOS, J. Y. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science* 43(12).
- BRESLAU, L., CAO, P., FAN, L., PHILLIPS, G., AND SHENKER, S. 1999. Web caching and zipf-like distributions: Evidence and implications. In *Infocom*. 126–134.
- CHAN, S. AND TOBAGI, F. 1999. On achieving profit in providing near video-on-demand services. In *Proceedings of the 1999 IEEE International Conference on Communications (ICC'99)*.
- CHAVEZ, A. AND MAES, P. 1996. Kasbah: an agent marketplace for buying and selling goods. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*.
- DAN, A., SITARAM, D., AND SHAHABUDDIN, P. 1994. Scheduling policies for an on-demand video server with batching. In *ACM Multimedia*. San Francisco, California, USA.
- GREENWALD, A. AND KEPHART, J. 1999. Shopbots and pricebots. In *Sixteenth International Joint Conference on Artificial Intelligence*.
- JAGANNATHAN, S. AND ALMEROOTH, K. C. 2001a. An adaptive pricing scheme for content delivery systems. In *Global Internet Symposium*. San Antonio, Texas, USA.
- JAGANNATHAN, S. AND ALMEROOTH, K. C. 2001b. The dynamics of price, revenue and system utilization. In *Management of Multimedia Networks and Services*. Chicago, Illinois, USA.
- ACM SIGecom Exchanges, Vol. 3, No. 2, May 2002.

- JAGANNATHAN, S., NAYAK, J., ALMEROOTH, K., AND HOFMANN, M. 2001. E-content pricing: Analysis and simulation. Tech. rep., University of California Santa Barbara. available at <http://www.nmsl.cs.ucsb.edu/papers/ECONTENTPRC.ps.gz>.
- JAGANNATHAN, S., NAYAK, J., ALMEROOTH, K., AND HOFMANN, M. 2002. On pricing algorithms for batched content delivery systems. Tech. rep., University of California Santa Barbara. available at <http://www.nmsl.cs.ucsb.edu/papers/BatchingPrc.ps.gz>.
- ODLYZKO, A. 2000. The history of communications and its implications for the Internet. available at <http://www.research.att.com/~amo/doc/networks.html>.
- TSVETOVATY, M., GINI, M., MOBASHER, B., AND WIECKOWSKI, Z. 1997. Magma: an agent-based virtual market for electronic commerce. *Applied Artificial Intelligence*.