A Theoretical Framework for Analyzing Satellite-based Web Multicasting Services

Marios Dikaiakos
Department of Computer Science
University of Cyprus
P.O. Box 20537, 1678 Nicosia, Cyprus
mdd@ucy.ac.cy

Abstract

In this paper we study Web multicasting, a service offered by satellite operators to Internet Service Providers around the world. This service employs satellite connections for disseminating periodically Web content to regional and "institutional" WWW caches. We propose a theoretical framework for Web multicasting and formalize the notions of Utility and QoS perceived by customers of Web multicasting services. We explore two alternative charging schemes, Usage- and Subscription-based pricing, and propose a framework for negotiating the provision of the Web multicasting service between a satellite operator and its potential customers. We use this negotiation framework to compare theoretically the two pricing schemes at hand. We show that a given level of QoS can be guaranteed under Subscription-based pricing at a cost at least as low as under Usage-based pricing. Our theoretical framework can be used further by satellite operators and prospective customers to define the content provided through Web multicasting, estimate its quality, negotiate its price and assess its overall effectiveness.

1 Introduction

World Wide Web usage represents the largest single source of traffic on Internet and is expected to grow further with the rise of Internet usage [2, 17] and the advent of new Web-based applications. The increased WWW usage has resulted in heavy workloads on popular Web servers and local networks. Currently, these loads are difficult to meet; in the future, if Web use continues to grow as fast, systems and networks face the danger of being eventually overwhelmed. As a way of coping with increased Web-loads, network operators have adopted and deployed Web caching hierarchies on proxy servers residing at local or regional networks [5, 28].

The emergence of proxy servers as Web caches has raised numerous research issues related to proxy server performance, effective caching and efficient caching architectures (e.g., see [5, 26, 20, 18, 10, 14, 11, 28]). Furthermore, the wide-scale deployment of hierarchical and co-operative Web caches opens new grounds for the development and implementation of cache-based techniques to sustain adequate levels of Web-performance, like prefetching [14] and content dissemination [6, 15]. Our conjecture is that, besides the expected performance advantages, such techniques represent a promising field for the exploration of emerging schemes for pricing and charging Internet-content [21, 22, 16, 12].

Recently, satellite networks have been adopted as an alternative for TCP/IP provision, by ISPs that seek to establish access connectivity to global Internet, backbone operators that wish to extend their terrestrial networks anywhere in the world, etc. [1, 3, 4]. Satellite networks are also employed to expand the performance gains achieved by Web-caching hierarchies deployed on Internet. In particular, satellite operators provide *Web multicasting services*, which consist of periodic multicasts of Web content to subscribed clients. Typical customers are Internet Provides seeking to enrich their cache hierarchies with Web content, without overloading their terrestrial links.

Web multicasting works as follows: satellite operators use their terrestrial connections to Internet backbones in order to "pull" a collection of WWW-objects into their multicast servers. Terrestrial multicast servers send the collected of content to the satellite through an up-link channel; this content is subsequently broadcast ("pushed") to authorized subscriber-organizations through the satellite's down-link channel. A subscriber stores received content into an *institutional* Web-cache, which is typically the "parent" in a Web-caching hierarchy established on top of proxy servers such as Squid [28].

In summary, a satellite operator (*"content-distributor"*) broadcasts Web content to various client-organizations around the world (typically ISPs). This process is repeated periodically throughout a day, materializing a *periodic push* scheme for information dissemination (see the taxonomy in [7]). The employment of satellite broadcasting for the dissemination of Web content to institutional caches is called *Web multicasting* via satellite. This service is currently deployed by major satellite operators around the world (e.g., see [1, 3, 4, 19]), and is adopted by client-organizations that purchase this service as a means for prefetching Web content through existing, under-utilized satellite links. Clients redistribute the content to their user-base through established Web-caching schemes.

In this paper, we introduce a theoretical framework to study issues pertinent to Web multicasting and formalize the notions of *Utility* and *Quality of Service* perceived by clients of Web multicasting services. We explore two charging schemes, *Usage-* and *Subscription-based pricing*, and propose a framework for negotiating the provision of the Web multicasting service between the satellite operator and its potential clients. We use this negotiation framework to compare theoretically the two pricing schemes at hand.

2 Multicasting Web Content via Satellite

The basic premise behind Web multicasting services is that the prefetched content covers adequately the interests of subscribers, improves the hit-ratio of installed Web-caches and, therefore, relieves overloaded terrestrial TCP/IP connections [25]. The soundness of this premise and the overall feasibility of the proposed approach depend on a number of issues, such as:

- The *"profiles"* of potential subscribers, which represent their information interests, the size of their customer-base, the level and cost of their terrestrial Internet-connectivity, cultural and language issues, etc.
- The way clients perceive and formalize the utility they expect to receive with the adoption of Web multicasting. Note that utility is a measure of the "pleasure" a client derives from the consumption of a particular service or good.
- The charging schemes proposed by content-distributors and the negotiation framework that can be established between distributors and clients to reach flexible and mutually profitable pricing mechanisms.
- The choice of content multicast to subscribers. Content selection should take into account client utility and distributor costs.
- The scheduling of data broadcasts.

The main thrust of our work is to establish a theoretical framework that takes into account the issues mentioned above, and can be used by satellite operators and potential customers to define the content provided by the Web multicasting service, estimate its quality, negotiate its price and assess its overall effectiveness.

Related Work

The development of wireless and satellite networks, and the expanding availability of assymmetric high-bandwidth links have created a lot of interest on issues related to data broadcasting. A large number of projects have examined various aspects of information dissemination over broadcast channels. Most of the schemes explored in the litrature deal with the tuning of information-dissemination systems to better serve immediate user-requests either through improved multicast-scheduling algorithms, "lighter" multicast architectures, or optimized caching schedules [27, 7, 8, 23]

In contrast, our approach looks into the case where the broadcast channel is used simultaneously with terrestrial links and in conjunction with Web-caching hierarchies deployed. This is also the focus of the work presented in [24, 25]. Our approach differs from this, however, in a number of ways: Firstly, we focus on schemes for periodic rather than continuous prefetching. Secondly, we study the application of Web multicasting for prefetching Web-content to groups of users that belong to different backgrounds. So, instead of considering a single, unified user population, we consider multiple ones with possibly different characteristics (language, culture, size) and requirements.

Last, but not least, we focus on the modeling of issues that determine Web multicasting, beyond its infrastructure requirements: Utility, Quality-of-Service, Pricing, and service-negotiation between provider and customers.

3 A Theoretical Formulation of Web Multicasting Services

3.1 Service Definition

As stated earlier, a content-distributor multicasts content according to a simple periodic schedule. On every multicast, all client-organizations receive identical information. These assumptions correspond to the actual configuration of emerging satellite services that multicast WWW data on an international scale [19].

For the content-distributor to choose Web content appropriately, we assume further that it collects *profile information* from each client regularly; a profile represents the

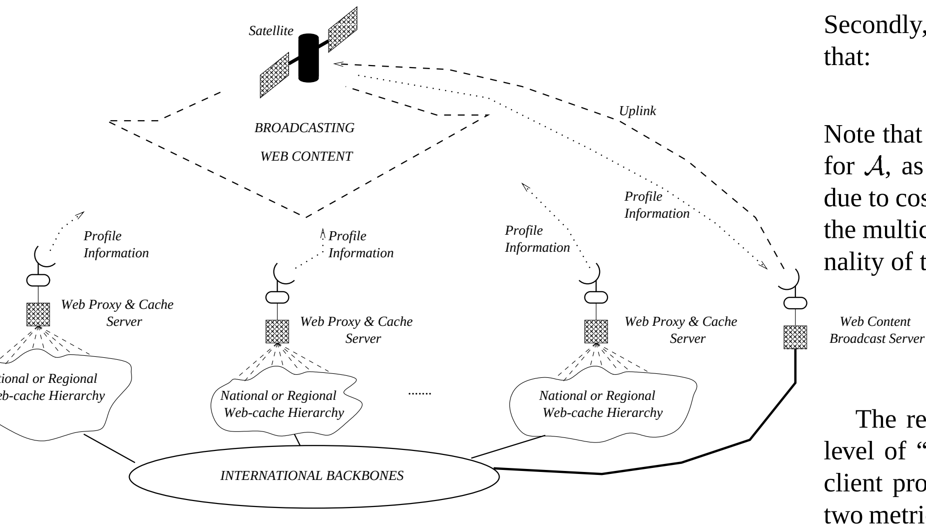

Figure 1. Multicasting content to Web-caches.

most recent information-needs of a client. Based on client-profiles, the distributor can select the content to be pulled from the Web and stored on the broadcasting server for the subsequent transmission (see Figure 1).

For a theoretical formulation of the Web multicasting service, we assume that the multicast operator has $\mathcal{M}$ clients. We represent the *URL-profile* of each client-organization i with a set A_i:

$$A_i = \{a_{i,j} \mid j = 1, \ldots, n_i\}, \; i = 1, \ldots, \mathcal{M} \qquad (1)$$

where $a_{i,j}, j = 1, \ldots, n_i$ correspond to "popular" URLs in the user community of organization i. Notably, different clients may have profiles differing as widely as the interests of an ISP clientele in Cyprus and a regional network user-base in Brasil; that is, they may differ both in terms of their size (n_i) and content ($a_{i,j}$'s) [9]. In practice, the URLs of an A_i-set can be extracted from the URL-traffic captured by the institutional cache of i.

Web-multicasting Service Definition: Based on the contents of the A_i's, the Web multicasting service corresponds to a set $\mathcal{A}$ of URL addresses that the multicast operator will disseminate to its subscribers. $\mathcal{A}$ is called the *multicast profile* and is defined as follows:

$$\mathcal{A} = \{\alpha_k \mid k = 1, \ldots, N\} \qquad (2)$$

where α_k's are the URLs disseminated to the Web multicasting subscribers.

The multicast profile should comply with two fundamental conditions. First, the elements of $\mathcal{A}$ should be chosen amongst the elements of the A_i's, so that:

$$\mathcal{A} \subseteq \bigcup_{i=1}^{\mathcal{M}} A_i. \qquad (3)$$

Secondly, $\mathcal{A}$ should provide some "cover" to all A_i's, so that:

$$A_i \cap \mathcal{A} \neq \emptyset, \; \forall i \in \{1, \ldots, \mathcal{M}\}. \qquad (4)$$

Note that the union of all A_i's would be an obvious choice for $\mathcal{A}$, as it satisfies conditions (3) and (4). Nevertheless, due to cost considerations we assume that the cardinality of the multicast profile should be much smaller than the cardinality of the union of A_i's, i.e.:

$$\|\mathcal{A}\| \ll \|\bigcup_{i=1}^{\mathcal{M}} A_i\|. \qquad (5)$$

The required multicast profile should possess a certain level of "similarity" between the multicast profile and all client profiles (A_i's). To gauge this similarity, we define two metrics that can be used to assess the relevance between two URL-profiles.

Definition 1 (Resemblance) *Let A and B be two URL-profiles with: $A = \{a_i \mid i = 1, \ldots, n_A\}$, and $B = \{b_i \mid i = 1, \ldots, n_B\}$, where the a_i's and b_i's correspond to URL addresses. Then, the* resemblance *between A and B is defined as follows:*

$$res(A, B) = \frac{\|A \cap B\|}{\|A \cup B\|}.$$

In practice, the resemblance of two profiles A and B represents the portion of the overall pool of elements of A and B belonging to both A and B. We can easily prove the following properties for resemblance: $0 \leq res(A, B) \leq 1$, $res(A, A) = 1$, $res(A, B) = res(B, A)$, and, if $A \cap B = \emptyset$ then $res(A, B) = 0$.

Definition 2 (Coverage) *Let A and B be two URL-profiles with: $A = \{a_i \mid i = 1, \ldots, n_A\}$, and $B = \{b_i \mid i = 1, \ldots, n_B\}$, where the a_i's and b_i's correspond to URL addresses. Then, the* coverage *of set A by set B is defined as follows:*

$$cov(A, B) = \frac{\|A \cap B\|}{\|A\|}.$$

In practice, the coverage of profile A by a set B represents the percentage of A's elements that belong to B. We can easily prove a number of basic properties for coverage: $0 \leq cov(A, B) \leq 1$, $cov(A, A) = 1$, $cov(A, B) \neq cov(B, A)$, and, if $A \cap B = \emptyset$, then $cov(A, B) = 0$.

3.2 Pricing and Quality-of-Service

To select the URLs of the multicast profile $\mathcal{A}$, the multicast operator should aim at satisfying the utility requirements of all clients that adopt the satellite-based Web-content-dissemination service. The formalization of utility, however, depends, among other things, upon the pricing model agreed between the multicast operator and its clients. Here, we suggest two simple pricing models and explore how their adoption affects client utility and the calculation of $\mathcal{A}$ from A_i's.

Subscription-based pricing: *To receive the Web multicasting service, clients of the content-distributor pay a fixed, monthly subscription fee covering leased satellite equipment and the periodic data feed.*

Usage-based pricing: *To receive the Web multicasting service, clients of the content-distributor pay a standard fee, covering leased satellite equipment, and a monthly fee proportional to the amount of bytes they receive from the satellite.*

In both models, it is assumed that each client-organization i has adequate storage capacity for storing the broadcast content in its institutional cache. Furthermore, the institutional cache of i can discard content not deemed of interest to its user-base, i.e., not belonging to A_i.

Under Subscription-based pricing, each client achieves optimal utility when receiving a selection of URLs that provide a maximal coverage of its profile (A_i); in other words, the client seeks the maximization of $cov(A_i, \mathcal{A})$. Under Usage-based pricing, each client seeks to minimize the amount of useless information received and charged, i.e., $\mathcal{A} - A_i$, in addition to maximizing the coverage of its profile. This is equivalent to maximizing $res(A_i, \mathcal{A})$.

These considerations dictate the client's perception about the quality of the proposed service. Therefore, we model the Quality-of-Service (QoS) offered by the multicast operator as follows:

Quality-of-Service: *The Quality of Service offered by a multicast operator to its client i is represented as a function $Q(\tau, A_i, \mathcal{A})$, where A_i is the URL profile of i, $\mathcal{A}$ is the multicast profile disseminated by the operator, and τ is the pricing model agreed between the operator and its clients. For Usage and Subscription-based pricing, Q is defined as follows:*

$$Q(\tau, A_i, \mathcal{A}) = \begin{cases} res(A_i, \mathcal{A}), & \textit{where } \tau = \textit{Usage-based pricing} \\ cov(A_i, \mathcal{A}), & \textit{where } \tau = \textit{Subscription-based pricing} \end{cases} \quad (6)$$

From Equation (6), we can easily see that $0 \leq Q(\tau, A_i, \mathcal{A}) \leq 1$.

4 A Framework for Negotiating Web-multicasting Services

The models presented in the previous section enable the satellite-operator to establish a framework of negotiation with its clients about the provision of the Web multicasting service. This framework entails three dimensions:

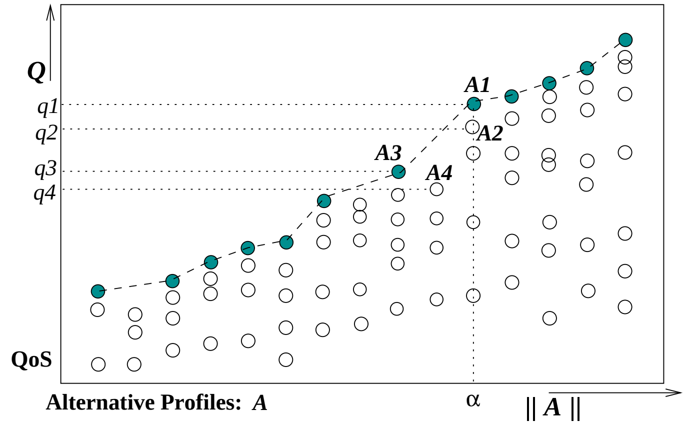

Figure 2. A Negotiation Framework for Web multicasting.

- The definition of the service provided, which we model by the multicast profile $\mathcal{A}$.
- The Quality-of-Service, which we model according to definition (6).
- The price tag paid by a particular client for a given service and service-quality.

In an ideal situation, the operator and each client "negotiate" in order to reach a service agreement: following the collection of URL-profiles A_i from the clients, the satellite operator calculates a number of alternative service-provisions in terms of alternative multicast-profiles $\mathcal{A}$. Each alternative multicast profile corresponds to a different Quality-of-Service and is offered at a cost determined according to the pre-agreed pricing scheme.

Figure 2 represents the space of alternative multicast profiles proposed by the satellite operator to some client i. Proposed profiles are represented as circular points in a two-dimensional space: the horizontal dimension corresponds to the cardinality of multicast profiles whereas the vertical dimension corresponds to their respective QoS values. Notably, the operator could propose to its client a number of different multicast profiles with identical profile size but with different QoS values. For instance, in Figure 2, profiles $\mathcal{A}_1$ and $\mathcal{A}_2$ have the same cardinality α; $\mathcal{A}_1$, however, offers an improved QoS over $\mathcal{A}_2$ since $q_1 > q_2$.

Each multicast profile is offered by the satellite operator at a particular price. We assume that, under the pricing schemes introduced earlier, multicast profiles of the same cardinality $||\mathcal{A}||$ have the same cost; furthermore, that the more URLs are broadcast via the satellite, the higher the cost of the satellite service is. In other words, we make the following conjecture:

Conjecture 1 *For any two multicast profiles* $\mathcal{A}$ *and* $\mathcal{B}$ *proposed by the satellite operator to its clients, if* $\|\mathcal{A}\| \leq \|\mathcal{B}\|$ *then* $price(\mathcal{A}) \leq price(\mathcal{B})$.

In summary, a client can choose among a set of triplets that define the Web multicasting service in terms of a proposed multicast profile, its quality, and its price. It is up to the client to agree upon the particular service deemed satisfactory. Taking into account the remarks above, it is not difficult to see that from the range of proposed multicast profiles of Figure 2, a client is expected to negotiate for the "purchase" of only a small subset of profiles that we call **candidate profiles** and are marked as dark circles. The client has no reason to consider other profiles: for instance, profile $\mathcal{A}_2$ would be rejected since $\mathcal{A}_1$ offers a better QoS ($q_1 > q_2$) at the same price. Moreover, $\mathcal{A}_4$ would be rejected because profile $\mathcal{A}_3$ offers a better QoS ($q_3 > q_4$) at a price that is no worse than $\mathcal{A}_4$'s (since $\|\mathcal{A}_3\| < \|\mathcal{A}_4\|$). Candidate profiles are defined formally as follows:

Definition 3 (Candidate Profile) *A multicast profile* $\mathcal{A}$*, proposed by a satellite operator to some client* i*, is called* **candidate profile** *if and only if, for any other proposed profile* $\mathcal{B}$ *such that* $\|\mathcal{B}\| < \|\mathcal{A}\|$*, it is:* $Q(\tau, A_i, \mathcal{B}) < Q(\tau, A_i, \mathcal{A})$.

With these remarks in mind, it is not difficult to establish the following conjecture and prove Lemma 1.

Conjecture 2 *Among the range of multicast profiles that are proposed by a satellite operator to some client, the client will be willing to consider for purchase only candidate profiles.*

Lemma 1 *For a client* i*, the Quality-of-Service of candidate profiles is monotonically increasing with respect to the candidate-profiles' cardinality. In other words, for any two candidate profiles* $\mathcal{A}$ *and* $\mathcal{B}$ *such that* $\|\mathcal{A}\| < \|\mathcal{B}\|$*, it is:* $Q(\tau, A_i, \mathcal{A}) < Q(\tau, A_i, \mathcal{B})$.

Proof: By contradiction, directly from Definition 3 and Conjecture 2.

Service Configuration through QoS-guarantees

It is impractical to run separate, automated negotiations between the operator and its clients, each time the operator has to construct a multicast profile. Such an approach would require significant computation and communication resources and might not result to a single multicast profile satisfying all clients. Therefore, to make things simpler, the multicast operator can incorporate client considerations in a *service contract* proposed to potential clients. According to this contract, the satellite operator undertakes the responsibility of continuously broadcasting a *candidate multicast profile* that provides all clients with a minimum, guaranteed Quality-of-Service level. This *QoS-guarantee* is offered to each client i through a *quality factor q*, which is accepted by both sides in the service contract. The quality factor defines the minimum guaranteed Quality-of-Service level offered by the operator to all clients, through the following inequality:

$$Q(\tau, A_i, \mathcal{A}) \geq q \tag{7}$$

The utility requirements of the clients are accommodated in this contract through the quality factor q. Under such a scheme we can prove the following theorem:

Theorem 1 *Let* $\mathcal{A}_s$ *and* $\mathcal{A}_u$ *be two candidate multicast profiles of minimum cardinality that provide all clients with the Quality-of-Service guarantee* q *under Subscription and Usage-based pricing, respectively. Then:* $\|\mathcal{A}_s\| \leq \|\mathcal{A}_u\|$.

Proof (by contradiction): We assume that:

$$\|\mathcal{A}_s\| > \|\mathcal{A}_u\| \tag{8}$$

Given that $\mathcal{A}_s$ is a *minimum-cardinality* candidate profile under Subscription-based pricing, for any other candidate profile $\mathcal{B}$ with cardinality less than $\|\mathcal{A}_s\|$, there would be at least one client for which the QoS provided by $\mathcal{B}$ would be less than q, under Subscription-based pricing. This remark holds for $\mathcal{A}_u$ as well, according to our assumption (8). Therefore:

$$\exists j : cov(A_j, \mathcal{A}_u) < q \tag{9}$$

From the definition of $\mathcal{A}_u$ and (6), however, it is:

$$\forall i, res(A_i, \mathcal{A}_u) \geq q \tag{10}$$

Furthermore, from Definitions 1 and 2 of *Resemblance* and *Coverage*, we can easily see that:

$$\forall i, cov(A_i, \mathcal{A}_u) \geq res(A_i, \mathcal{A}_u) \tag{11}$$

Hence:

$$(10), (11) \Rightarrow \forall i, cov(A_i, \mathcal{A}_u) \geq q,$$

which is a direct contradiction to inequality (9). Consequently, assumption (8) is wrong and therefore we conclude that $\|\mathcal{A}_s\| \leq \|\mathcal{A}_u\|$.

What this theorem shows, in combination with Conjecture 1, is that if the satellite operator and its clients accept the negotiation scheme presented earlier, a given level of the QoS-guarantee can be established under Subscription-based pricing at a price *at least as low* as under Usage-based pricing.

Besides the satisfaction of client-utility, however, the multicast operator is expected to pursue the maximization of profit it receives from the deployment of the Web-multicasting service. Under Subscription-based pricing, the operator's "income" is constant for a given number of client organizations. Therefore, we can assume that the operator seeks to minimize its collection and distribution costs in its selection of multicast content, while at the same time maintaining the QoS-guarantee agreed with its customers. We model the operator's costs with $\gamma \times ||\mathcal{A}||$, a value proportional to the total number of Web-objects disseminated, i.e., to the cardinality of $\mathcal{A}$. It should be noted that modeling distributor's costs proportionately to $||\mathcal{A}||$ is only an approximation as this does not take into account the *byte size* of objects.

The operator's income and benefits are proportional to $||\mathcal{A}||$, under Usage-based pricing. Consequently, we assume that the multicast operator seeks to send more content when selecting its multicast profile $\mathcal{A}$, that is to increase $||\mathcal{A}||$. Nevertheless, $||\mathcal{A}||$ cannot be increased up to $|| \bigcup_{i=1}^{\mathcal{M}} A_i||$; in most cases, such an increase could violate the QoS-guarantee described by definition (6) and inequality (7), and/or exhaust storage and networking resources of the operator.

It should be noted that issues such as the selection of content, the profitability of the service for a varied number of clients, etc. are beyond the scope of this paper.

5 Conclusions and Future Work

In this paper we studied Web multicasting, a periodic "push" scheme that uses satellite links to disseminate information to WWW-caches worldwide. This information dissemination takes place under a service offered by satellite-network operators to subscriber ISPs around the world. Satellite-based dissemination is combined with hierarchical caching schemes deployed by the ISPs, providing prefetched Web-content to WWW-caching hierarchies.

In this context, we introduced a novel theoretical framework that takes into account Utility, Pricing and Quality of Service for Web multicasting. Within this framework we propose the notion of QoS-guarantees that can be adopted by satellite operators and prospective customers to determine the content of the Web multicasting service. Based on our modeling, we proved that the multicast operator can guarantee under Subscription-based pricing, a Quality-of-Service at least as good as under Usage-based pricing, at the same or lower cost. This conclusion provides a basis for preferring the Subscription-based pricing scheme for Web multicasting services established upon the negotiation framework introduced here. Upon selection of a particular pricing scheme, our framework can be used for deriving content-selection algorithms (see for instance [13]), estimating the advantages gained by the satellite mutlicasting approach over traditional solutions, analyze the effects of customer profiles on service cost, etc.

References

[1] Broadcast Satellite Services. http://www.isp-sat.com.

[2] GVU's WWW User Surveys. http://www.gvu.gatech.edu/gvu/user_surveys/.

[3] INTELSAT. http://www.intelsat.com/products/internet/atint.htm.

[4] Skycache. http://www.skycache.com.

[5] M. Abrams, C. Stanbridge, G. Abdulla, S. Williams, and E. Fox. Caching Proxies: Limitations and Potentials. In *Fourth International World Wide Web Conference*, December 1995. http://www.w3.org/Conferences/WWW4/.

[6] D. Aksoy, M. Altinel, R. Bose, U. Cetintemel, M.J. Franklin, J. Wang, and S.B. Zdonik. A Framework for Scalable Dissemination-Based Systems. In *Proceedings of the 1997 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages & Applications (OOPSLA '97)*, pages 94–105. ACM, 1997.

[7] D. Aksoy, M. Altinel, R. Bose, U. Cetintemel, M.J. Franklin, J. Wang, and S.B. Zdonik. Research in Data Broadcast and Dissemination. In *Proceedings of the First International Conference on Advanced Multimedia Content Processing, AMCP '98, Lecture Notes in Computer Science*, pages 194–207. Springer Verlag, 1999.

[8] D. Aksoy and M. Franklin. Scheduling for Large-Scale On-demand Data Broadcasting. In *Proceedings of the 1998 IEEE Infocom Conference*. IEEE, 1998.

[9] V. Almeida, M. Cesario, R. Fonseca, Jr.W. Meira, , and C. Murta. Analyzing the behavior of a proxy server in light of regional and cultural issues. In *Proceedings of the Third International WWW Caching Workshop*, June 1998.

[10] A. Bestavros, R. Carter, M. Crovella, C. Cunha, A. Heddaya, and S. Mirdal. Application Level Document Caching in the Internet. In *Proceedings of IEEE SDNE'96: The Second International Workshop on Services in Distributed and Networked Environments*. IEEE, June 1995. http://www.cs.bu.edu/~best/res/papers/sdne95.ps.

[11] P. Cao and S. Irani. Cost-Aware WWW Proxy Caching Algorithms. In *Proceedings of the 1997 USENIX Symposium on Internet Technology and Systems*, pages 193–206, December 1997.

[12] C. Courcoubetis and V.A. Siris. Managing and pricing service level agreements for differentiated services. In *Proceedings of the 7th IEEE/IFIP International Workshop on Quality of Service (IWQoS'99)*, 1999.

[13] M. Dikaiakos and A. Stassopoulou. Content-Selection Strategies for the Periodic Prefetching of WWW Resources via Satellite. *Computer Communications*, In Press, 2000.

[14] L. Fan, P. Cao, W. Lin, and Q. Jacobson. Web Prefetching between Low-Bandwidth Clients and Proxies: Potential and Performance. In *Proceedings of the 1999 Sigmetrics Conference on Measurement and Modeling of Computer Systems*, pages 178–187, 1999.

[15] M.J. Franklin and S.B. Zdonik. "Data in Your Face": Push Technology in Perspective. In *SIGMOD 1998, Proceedings ACM SIGMOD International Conference on Management of Data, June 1998, Seattle, Washington, USA*, pages 183–194. ACM Press, 1998.

[16] A. Gupta, D. Stahl, and A. Whinston. The Economics of Network Management. *Communications of the ACM*, 42(9):57–63, September 1999.

[17] D. L. Hoffman, W.D. Kalsbeek, and T. P. Novak. Internet and Web Use in the U.S. *Communications of the ACM*, 39(12):36–46, December 1996.

[18] B. Liu, G. Abdulla, T. Johnson, and E. Fox. Web Response Time and Proxy Caching. In *WebNet '98*. AACE, November 1998.

[19] C. Makris. Personal Communication. Cyprus Telecommunications Authority, August 1999.

[20] C. Maltzahn and K. J. Richardson. Performance Issues of Enterprise Level Web Proxies. In *Proceedings of the 1997 Sigmetrics Conference on Measurement and Modeling of Computer Systems*, pages 13–23. ACM, 1997.

[21] A. Odlyzko. The economics of the Internet: Utility, utilization, pricing, and Quality of Service. Technical report, AT&T Labs - Research, 1998.

[22] A. Odlyzko. Paris Metro Pricing for the Internet. In *Proceedings of the ACM Conference on Electronic Commerce (EC-99)*, 1999.

[23] C.H. Papadimitriou, S. Ramanathan, and P.V. Rangan. Multimedia Information Caching for Personalized Video-on-Demand. *Computer Communications*, 18(3):204–216, March 1995.

[24] P. Rodriguez and E.W. Biersack. Bringing the Web to the Network Edge: Large Caches and Satellite Distribution. In *Proceedings of WOSBIS '98, ACM/IEEE MobiCom Workshop on Satellite-based Information Services*, October 1998.

[25] P. Rodriguez and E.W. Biersack. Prefetching Web Documents into Large Caches using a Satellite Distribution. *ACM Special Issue of the Journal on Special Topics in Mobile Networking and Applications (MONET)*, To appear, 2000. http://www.eurecom.fr/˜rodrigue/.

[26] A. Rousskov and V. Soloviev. On Performance of Caching Proxies. In *Proceedings of the 1998 Sigmetrics Conference on Measurement and Modeling of Computer Systems*, pages 272–273. ACM, 1998.

[27] C.J. Su, L. Tassiulas, and V.J. Tsotras. Broadcast scheduling for information distribution. *ACM/Baltzer Journal of Wireless Networks Technology and Systems*, 5(2):137–147, 1999.

[28] D. Wessels and K. Claffy. Evolution of the NLANR Cache Hierarchy: Global Configuration Challenges. Technical report, NLANR, October 1996. http://www.nlanr.net/Papers/Cache96/.