Dynamic Capacity Contracting:

A Framework for Pricing the Differentiated Services Internet

Shivkumar Kalyanaraman†

Assistant Professor

Electrical and Computer Systems Engineering

Rensselaer Polytechnic Institute

Troy, NY 12180

Tel: (518) 276 8979

E-Mail: shivkum@ecse.rpi.edu

T. Ravichandran†

Assistant Professor

Lally School of Management and Technology

Rensselaer Polytechnic Institute

Troy, NY 12180

Tel: (518) 276 2035

E-Mail: ravit@rpi.edu

Randy Norsworthy

Professor

Lally School of Management and Technology

Rensselaer Polytechnic Institute

Troy, NY 12180

Tel: (518) 276 6649

E-mail: randy@norsworthy.net

Submitted to

First International Conference on

Information and Computation Economics

May 5, 1998

†The names of the authors are in no particular order. This project is a team effort where both authors contributed equally.

Abstract

The exploding growth of the Internet has lead to a general deterioration of service levels experienced by its users. The IP best effort model that is currently in use is inadequate to provide guaranteed service levels. Internet service providers are seeking scalable methods for providing differentiated services to customers who are willing to pay more. This paper proposes a scheme called dynamic capacity contracting for providing differentiated services to users. Technical and economic issues pertaining this scheme are explored.

The central idea of dynamic capacity contracting is a network that could, based upon congestion

monitoring mechanisms, raise prices and vary contract terms dynamically. This is achieved through a software entity called bandwidth broker which liaisons between the users and Internet service providers. The bandwidth broker, based on user prescribed constraints (such as budgets), negotiates "soft contracts" with the service providers on an on-going basis. Soft contracts would expire within a time span and should be renewed to maintain a level of service. This scheme builds on the concepts of expected capacity (Clark, 1995) and smart markets (MacKie-Mason & Varian, 1993), presents a technical architecture for its implementation.

Key Words: Electronic Commerce, Internet Pricing, Internet Economics, Networking.

Dynamic Capacity Contracting:

A Framework for Pricing the Differentiated Services Internet

Introduction

Over the past ten years, the Internet has grown rapidly and diffused across a diverse user population (Rai, Ravichandran & Sammadar, 1998). The development of the word wide web and the burgeoning commercial use of the web have lead to a proliferation of data intensive applications including real time audio and video. The resultant explosion in network traffic is expected to create congestion problems, resulting in a general deterioration of service levels experienced by users.

So far, capacity provisioning has sustained the growth in network traffic. For example, the backbone network capacity has been upgraded from 56 kilobits per second to almost 2.5 gigabits per second. However, user demands are increasing exponentially and the network traffic is expected to outstrip available capacity (Commerce Report, 199x). Under these conditions, efficient bandwidth allocation through statistical multiplexing by itself may not be sufficient to meet user demands. It is necessary to share bandwidth in a more controlled manner incorporating factors such as application requirements, network and economic efficiencies. Many have suggested that responsive pricing schemes can help achieve economic efficiencies.

Several schemes have been proposed to price the Internet and its various domains (MacKie-Mason & Varian, 1995; Clark, 1995; Gupta, Stahl & Whinston, 1997). However, there has been little experience in implementing and studying these schemes in the production Internet. A major impediment in doing so is the minimalist "best effort" service model of the IP protocol which does not provide a standard mechanism to specify packet forwarding behaviors other than the "best-effort" service utilizing the statistical multiplexing efficiencies of packet switching. For example, there is no way to give preferential service to packets of higher-paying customers, other than buying a new leased (non-statistically multiplexed) line.

However, this scenario is rapidly changing as the Internet Engineering Task Force (IETF) is standardizing two approaches to support service differentiation. The first approach is called Integrated Services ("int-serv") (Braden, Clark & Shenkar, 1994). In this service framework, a signaling protocol called RSVP reserves network resources based on user defined service parameters. The second approach called Differentiated Services ("diff-serv") (Nichols & Carpenter, 1998) is expected to provide scalable service discrimination without the need for per-flow state signaling at every hop. While the two approaches can coexist and interoperate, it is expected that the latter approach (diff-serv) will be the choice of ISPs and backbone internetwork providers. We believe that the differentiated services architecture has significant practical implications for Internet pricing as it provides a mechanism to implement scalable service discrimination. In this paper, we focus on developing a pricing framework that utilizes the advanced traffic management features offered by this architecture.

The rest of the paper is structured as follows: first, we summarize the diff-serv architecture and highlight its traffic management features that are relevant to Internet pricing; next, we outline a flexible framework for implementing a range of pricing schemes within the differentiated services architecture; we go on describe a pricing scheme that offers a logical transition from the flat-pricing that is prevalent today to a dynamically contracted, purely congestion-sensitive pricing; finally we briefly describe our on-going research in this area and offer some concluding remarks.

Developments in Internet Traffic Management

Differentiated Services Architecture

Differentiated services ("diff-serv") is an new approach being standardized at the IETF (Nichols & Carpenter, 1998) to provide scalable service discrimination in the Internet without the need for per-flow state and signaling at every hop. Differentiated services allow the service model of Internet to change very quickly from the current minimalist "best-effort" model to a multi-service model where available resources are shared in a more controlled manner. The power of this model lies in its promise to allow both economic and scalable technical means to achieve service differentiation.

Figure 1 depicts the differentiated services architecture. In this architecture, there are two types of routers: interior routers and edge routers. The interior routers are equivalent to the core routers in the current Internet architecture. They perform highly optimized packet forwarding functions and classify packets based on the IP header field. Both these functions are executed on the critical packet forwarding code path. The packet forwarding function involves looking up routing tables to determine which "next-hop" (i.e. router port) to send the packet to. The classification function sets per-hop behaviors (PHBs) such as queuing, scheduling and packet dropping options that define how packets are handled at the router output port. The per-hop behaviors are chosen from a small set of predefined ones (64 is the proposed maximum). Interior routers do not participate in signaling functions. They however do build routing tables based upon a routing protocol running in the background.

On the other hand, edge routers participate in signaling, admission control, traffic conditioning and accounting. Edge routers are usually away from the core of the diff-serv domain. As a result, they deal with a smaller set of flows than the interior routers. This allows them to allocate computational resources for functions other than basic forwarding and classification. Edge routers negotiate service level agreements (SLAs) with users (or providers) on a per-flow or a per-aggregated-flow basis. The SLA negotiation is the equivalent of the signaling and admission control functions done in ATM (or in the Integrated Services architecture using RSVP). Once the SLA is negotiated, the edge routers use it as that basis for traffic conditioning (policing, marking, shaping) and accounting the actual traffic. Specifically, the ingress edge router marks the DS-byte in the IP header (formerly the IP TOS octet) of incoming packets indicating what per-hop behavior they should get at the interior routers. When these packets go through the interior routers, they get the per-hop behavior (PHB) as specified in the DS-byte. This architecture allows a rich set of services to be offered using a small set of per-hop behaviors. Furthermore, packets inside the diff-serv domain carry enough information to allow classification at interior routers.

Service discrimination in the Internet can be achieved through other technologies such as Asynchronous Transfer Mode (ATM) networks and the integrated services. However, these technologies differ from diff-serv in two important ways. First, diff-serv allows "better-than-best-effort" service differentiation as opposed to using resource reservations to provide service guarantees. In the traditional best effort model unlimited number of users share a single service class. On the other, the quality of service model (QoS) requires per flow, end-to-end resource reservations. The better-than-best effort model lies between these two service alternatives. In this model, a small set of classes (PHBs) is provided in the interior routers. These classes are shared between users based on their service contracts.

Second, diff-serv is easily scalable because each packet carries the information necessary for classification at the interior routers. This avoids storage and maintenance of per flow states at every interior router. For these reasons, it is expected that the diff-serv architecture will be used by ISPs and other top-level service providers in the Internet.

Differentiated Services Implementation

The technology changes required to implement the differentiated services architecture are simple. These include 1) upgrading interior routers to recognize and support a set of per-hop behaviors indicated in packet headers and 2) upgrading edge routers with traffic conditioning mechanisms which would mark per-hop behaviors on packet headers based upon service level agreements (SLA). The SLAs could be set up through network management mechanisms though more sophisticated signaling and control architectures are currently being studied.

For example, Nichols, Jacobson and Zhang (1997) proposed a control architecture that included a bandwidth broker. The bandwidth broker would be an entity that is aware of the policies concerning the user and liaisons between the user and the provider. Presumably, its functions would include (re)negotiating service level agreements (SLAs) and policy-based shaping/marking/aggregating of outbound user traffic.

The differentiated services model combined with the proposed signaling architectures allows economic concepts such as "price-based service differentiation" and "contracting" to be incorporated within IPv4. As a result, we could expect several Internet pricing experiments based on these models. The influence of economic ideas on diff-serv is not surprising because the baseline framework for capacity contracting ("expected capacity") was proposed by Clark (1995), an active participant in both the engineering and economics communities.

Other Developments in Traffic Management

There are other developments in traffic management that could potentially impact how the Internet is priced. The Multi-protocol Label Switching (MPLS) group at the IETF (Callon et al, 1997) has been designing mechanisms to specify explicit routes in the core network (or MPLS domain). This functionality could be used to semi-automatically engineer paths to perform load balancing, or preferentially route packets of high-value customers. This technology is targeted at core routers to radically improve forwarding performance and to provide traffic engineering options.

Recently, Kalyanaraman (1998) proposed a feedback mechanism to monitor congestion at the edges of the differentiated services network. In this mechanism an interior router instead of dropping packets during congestion, sets a bit in the packet header. This mechanism extends on-going efforts at IETF to develop an explicit congestion notification (ECN) technology and allows edge routers to use such notifications to monitor congestion in various paths in the domain. Specifically, the congestion signals would be noticed by the egress edge router and conveyed to the corresponding ingress edge router(s). The ingress edge router could use this information to apply controls on the flows experiencing congestion. Alternatively, the monitored congestion information from various paths (such as those set up using MPLS) could be functionally composed to create a spot price for new flows bidding for access to this domain.

Thus far we focused on the technological developments in traffic management. The remaining part of the paper focuses on introducing pricing and economic controls in differentiated services internetwork that take advantage of the advanced traffic management and engineering features outlined above.

Issues in Pricing the Differentiated Services Internet

Differentiated services affects the cost structures of the network service providers. The cost structure of providers at a particular level depends upon the price they pay to higher level providers (Srinagesh, 1997). Since all levels of providers can now offer a range of services, it is possible for a provider to buy services to better fit its needs. Furthermore, these adjustments to services can be done dynamically. This is likely to make the system economically more efficient.

Differentiated services also provides an elegant framework to implement internet pricing schemes. Several proposals to price the Internet have been put forth so far. It is recognized that the sustained growth of the Internet is in part attributable to the low flat-rate pricing for basic access and content. However, as the average traffic load on the infrastructure increases, it is clear that with flat rate pricing, low volume users subsidize high volume users. Also, under congestion, there is no way for the customer to demand better service from the provider. ISPs have recognized this and have imposed larger fees on users who stay connected for long periods (IBD, 1998a; 1998b). One caveat is that this scheme is currently intended to allow better connection multiplexing of the limited modem pools at the access point, rather than conserving raw bandwidth in the core provider networks.

Usage based pricing has been proposed as a solution to the problems outlined above (MacKie-Mason & Varian, 1995). We believe that usage-sensitive pricing is most effective only during congestion and only for a class of customers who seek "better-than-best-effort" performance. Hence, it is important that one focuses on "congestion-sensitive" pricing instead of "usage-sensitive" pricing.

It is also known that most users prefer flat pricing and only a subset of users are interested in the option of paying more for better performance (Anania & Solomon, 1997). This behavior has been observed in other contexts such as using tolled highways. For example, the Orange County highway 91 charges a toll based on the time-of-day and the relative congestion in the highway (IBD, 1998c). In other words, a larger toll will be charged if many cars are currently using the highway. It is found that a majority of the highway users prefer the flat-priced alternative highways and risk running into congestion. But, a few users willing to pay a congestion sensitive price for the privilege of using an uncongested highway. The system delivers its promise of avoiding traffic jams seen on the alternative highways that do not impose such tolls.

In summary, we believe that congestion-sensitive pricing if provided should supplement flat-rate pricing. In other words internet pricing schemes should allow users to transition between flat-rate and congestion-sensitive pricing depending on their needs. Moreover, it is imperative that mechanisms are put in place to ensure that users paying congestion-sensitive prices can expect to experience better performance.

We now examine the key issues involved in implementing pricing schemes in the diff-serv architecture.

Implementation Issues

A key constraint in the diff-serv architecture is the clear distinction between interior routers and edge routers. Interior routers cannot do signaling and accounting; they are optimized for forwarding. Therefore, these routers cannot participate in pricing related functions directly. As a result, one cannot expect to implement a pricing scheme that expects symmetric accounting/processing by all routers in the Internet. The policy and accounting functions have to be concentrated at the edges of the diff-serv domain.

In order to realize efficiencies in implementing pricing schemes we suggest 1) avoiding volume measurements as input to pricing and 2) using limited-term contracts. Measurement of per-user volume is cumbersome when many users are involved. On the other hand, if there are limited users volume measurement does not correlate well with costs because of higher peak rates of usage (Kelly, 1997). A similar point is made by Clark (1995).

Limited term contracts have a specified life during which users get value-added services by paying a fixed price at the time of entering into the contract. Since contracts are not for an infinite time span, provisioning efficiencies can be gained by avoiding resource reservations. Term-contracts would also be attractive from a users' standpoint as they can be invoked based on needs.

Another implementation issue is the providers' capability to assure better service to customers who pay more. A user is typically assigned to the priority class that he contracts for. In a given priority class, there could be congestion. In order to maintain the service levels during congestion, it is necessary to temporarily discourage new users from entering a priority class. This can be achieved by appropriate price signaling to users. The price signaling should be based on congestion, which is transient and path-specific. New traffic engineering and congestion monitoring mechanisms allow congestion-sensitive costing. It is possible for an edge router to send most traffic through explicit paths instead of routing them based on traditional protocols [Callon et al, 1998). This simplifies congestion detection as traffic monitoring needs to be done only on these explicit paths. Recently, Kalyanaraman (1998) proposed a mechanism to collect statistics on explicit paths using feedback within a diff-serv domain. This provides a basis for computing congestion prices. It also allows price computations to be done at the edges of the diff-serv domain.

We now turn to describing a framework for pricing the differentiated service internet.

Dynamic Capacity Contracting

Clark (1995) proposed that users should pay for the privilege of expecting better service. This privilege is captured in the form of an "expected capacity" contract. While Clark (1995) proposed long term capacity contracts, we believe that a dynamic contracting framework would be more effective in implementing congesting sensitive pricing in the diff-serv architecture. In such a framework, users would negotiate short-term contracts with the providers. The price for these contracts could be determined based on recent congestion levels in the network.

Let us consider an example of a simple service contract. Let us assume that the interior routers can support 8 priority classes and two levels of drop within each class. These 16 mechanisms (or "per-hop behaviors") can be coded within 4 bits. Now, contracts can be structured using these 16 mechanisms in combination with time and volume. For example, a contract can be as follows: 2KB of data every second will be marked as "IN" in class "i" and any excess traffic will be marked as "OUT" in class "i". This service allows sending upto 2KB of data every second on a priority basis within class i. Any excess data will be sent in class i but with a lower probability of delivery. In this example, we have used the volume dimension ("2KB") and the time dimension ("per second") to define the service. This is an example of a "better-than-best-effort" service where both the user and provider have more control of how to share the limited bandwidth resources.

Dynamic capacity contracting can be implemented using a bandwidth broker. A bandwidth broker is a component which is aware of the user's policy and serves as the intermediary between the user and provider (see figure 2). In our framework the bandwidth broker could be an enterprise that provides users with connectivity to and choice amongst service providers. This allows connection costs to be optimized. In the next version of IP, IPv6 (Huitema, 1998) it is possible for any network interface to have multiple addresses. The transport protocol could decide to route each packet through a different provider. Therefore the "bandwidth broker" functionality could be implemented in every home PC which supports IPv6.

The bandwidth broker would negotiate term contracts with providers on behalf of the users. It could also maintain user costs to be within a specified budget similar to the Expenditure Controller Interface suggested by Danielsen & Weiss (1997). We note that the providers can vary their advertised term prices based upon monitored congestion, or time-of-day controls. A negotiated contract expires automatically after the term is over - this is consistent with the "soft state" approaches popular in several Internet protocols like DHCP (Comer, 1996) and RSVP [Braden et al, 1997).

Computing Congestion-Based Prices

As discussed earlier, traffic engineering and congestion monitoring mechanisms allow the consolidation of accounting at the edge routers. In our framework, each edge router maintains a table of available contracts, available terms and a set of associated prices. Each feedback of path statistics will update this table. In other words, the price for a given term contract is a function of the congestion statistics on various traffic engineered paths. Specifically, if S_i is the set of statistics on a path i (say, average throughput (T_i) and average loss rate (L_i), we first have a function f that maps these statistics to a price for the path i.e. $P_i = f(S_i)$, where $S_i = (T_{I,} L_i)$. Next, the spot price for access is a composition of these current path prices, i.e. $P = g(P_1, \dots P_n)$. The price P here indicates the price for an entry in the table of available contracts.

One may expect to have the same functions (f and g) or different for updating other table entries. These functions could factor in the length of the contract and the volume or rate of data proposed to be sent at a higher priority. Note that the volume or the rate of data is used to set up the traffic conditioning entities such as the token bucket for policing or enforcing the contract. We do not measure actual volume and price *a posteriori*. Excess volume is sent at lower priority by the traffic conditioner. The price table could also be updated when the number of active contracts crosses critical thresholds.

The bandwidth broker queries simply result in table lookups for the current advertised prices. In addition, the edge router would 1) maintain a list of active contracts, 2) update the list based upon a low granularity timer and 3) purge obsolete contracts. Each active contract would include traffic conditioning components, which would mark the per-hop behaviors (PHBs) in packet headers based upon the current status of the contract. Thus, we believe that a fairly efficient implementation of dynamic capacity contracting can be realized for pricing "better-than-best-effort" services.

A Sample Pricing Scheme

As discussed earlier, it will be required for providers to offer flat rate access. This service does not have any assurances whatsoever - it is a pure best-effort service. However, dynamic capacity contracting allows providers to optionally offer value-added services on top of flat-rate services. Specifically, a user could choose to dynamically contract "better-than-best-effort" service for a limited term.

A provider could have a set of classes (say 1 to N) with priority scheduling (where 1 is the highest priority). Further, it might make sense to split each class into two subclasses. Let us say we have subclasses (1a, 1b), (2a,2b) etc ... Note that in general, the provider need not split every class.

One subclass of every class (say the "a" subclass) could be sold as a flat-priced service to satisfy a large number of users who want flat rate. The only "admission control" is the economic one - the flat price for 1a would be much larger than 2a, or any variable price in 2b and so on. The higher flat price is because 1a is given priority over 2a, 2b. The key point is that we do not assure anything during congestion (seen by this class) for this "a" subclass of service. The other subclass, "b", can be charged based on congestion.

A naive implementation of subclasses "a" and "b" would use two queues. However diff-serv allows a more efficient implementation using the equivalent of an in/out bit for each class. All packets for "a" subclass are marked "out", while packets from the "b" subclass can be marked "in" or "out" depending upon whether they are in-profile or out-of-profile.

Summary

In this article, we first summarized new developments in Internet traffic management. We then proposed a "dynamic capacity contracting" framework to integrate these new technology elements to allow implementation of pricing policies. We also proposed a sample scheme which splits the available bandwidth into classes, offers a flat price for best-effort handling in each class, and allows users to elect for limited-term, dynamically contracted "better-than-best-effort" services. We believe that this is a logical transition from today's flat priced system towards a completely congestion-sensitive pricing system. In our scheme the user can elect to go for either system.

We have implemented a diff-serv model, which simulates packet level dynamics of Internet traffic aggregated and sent through diff-serv classes. We are currently working on defining the pricing functions f and g and incorporating dynamic capacity contracting in our simulation. Our future work will involve studying alternative price policies tightly integrated into a well-modeled internetwork.

References

Anania, L. A & Solomon, R. J (1997) Flat - The minimalist price, Internet Economics, Eds McKnight & Bailey, MIT press, 1997.

Braden, R Zhang, L, Berson, S, Herzog, S & Jamin, S (1997) "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification," Internet Request For Comments (RFC) 2205, September 1997. Available from http://www.cis.ohio-state.edu/htbin/rfc/rfc2205.html

Braden, R, Clark, D & Shenker, S (1994) "Integrated Services in the Internet Architecture: an Overview," Internet Request For Comments (RFC) 1633, June 1994. Available from http://www.cis.ohio-state.edu/htbin/rfc/rfc1633.html

Carpenter, B (1996) "Architectural Principles of the Internet". Internet Request For Comments (RFC) 1958, June 1996. Available from http://www.cis.ohio-state.edu/htbin/rfc/rfc1958.html

Callon et al (1997) "A Framework for Multiprotocol Label Switching", <draft-ietf-mpls-framework-02.txt>, November 1997 Available from http://www.ietf.org/internet-drafts/draft-ietf-mpls-framework-02.txt

Clark, D (1997) Internet cost allocation and pricing, in Internet Economics, Eds McKnight & Bailey, MIT press, 1997.

Comer, D (1996) "Internetworking with TCP/IP, Vol 1" 3rd Edition, Prentice Hall, 1996.

Danielson, K & Weiss. M (1997) User control and Ip allocation, Internet Economics, Eds McKnight & Bailey, MIT press, 1997.

Floyd, S (1994) "TCP and Explicit Congestion Notification" ACM Computer Communication Review, V. 24 N. 5, October 1994, p. 10-23. Available from http://ee.lbl.gov/nrg-papers.html

Gupta, A, Stahl, D. O & Whinston, A, B (1997) Priority Pricing of Integrated Services Networks, Internet Economics, Eds McKnight & Bailey, MIT press, 1997.

Huitema, c (1998) "IPv6: The new Internet protocol," Prentice-Hall, 1998.

IBD (1998a) "Internet Firms Tinker with Service Pricing," Investor Business Daily, 16 April 1998.

IBD (1998b) "End Internet's Free Lunch," Investor Business Daily, 9 April 1998.

IBD (1998c) "Freeways or Feeways", Investor Business Daily, 17th April 1998.

Kalyanaraman et al (1998) "One-bit Feedback Enhanced Differentiated Services Architecture", draft-shivkuma-ecn-diffserv-00.txt>, March 1998 Available from http://www.ietf.org/internet-drafts/draft-shivkuma-ecn-diffserv-00.txt

Kelley, F. P (1997) Charging and accounting bursty traffic, Internet Economics, Eds McKnight & Bailey, MIT press, 1997.

MacKie-Mason, J. K & Varian, H. R (1995) Pricing the Internet, Working Paper, Department of Economics, University of Michigan, Ann Arbor, 1993.

MacKie-Mason, Murphy, L & Murphy, L (1997) Responsive pricing in the Internet, in Internet Economics, Eds McKnight & Bailey, MIT press, 1997.

Mathis, M, Mahdavi, J Floyd, S & Romanow, A (1996) "TCP Selective Acknowledgement Options" Internet Request For Comments (RFC) 2018, October 1996. Available from http://www.cis.ohio-state.edu/htbin/rfc/rfc2018.html

Nichols, K & Carpenter, B (1998) "Differentiated Services Operational Model and Definitions" Internet Draft, <draft-nichols-dsopdef-00.txt>, February 1998 http://www.ietf.org/internet-drafts/draft-nichols-dsopdef-00.txt

Nichols, K, Jacobson, V and Zhang, L (1997) "A Two-bit Differentiated Services Architecture for the Internet", Internet Draft, <draft-nichols-diff-svc-arch-00.txt>, December 1997. Available from http://www.ietf.org/internet-drafts/draft-nichols-diff-svc-arch-00.txt

Rai, A, Ravichandran, T & Sammadar (1998) Modeling the macro-diffusion of distributed IT infrastructures: The case of the Internet, Forthcoming, Communications of the ACM.

Srinagesh, P (1997) Internet cost structures and interconnection agreements, Internet Economics, Eds McKnight & Bailey, MIT press, 1997.