

# Pricing priority services over DiffServ-enabled transport networks

Christos Bouras, Afrodite Sevasti

*RA Computer Technology Institute - RACTI, Patras, Greece and Department of Computer Engineering and Informatics, University of Patras, Greece*

**Abstract:** The rapid evolution of the DiffServ framework in our days has provided the means for efficient QoS provisioning over contemporary IP networks. The provisioning of services according to the DiffServ framework has in turn raised the requirements for pricing mechanisms that preserve the potential and flexibility of the DiffServ framework. At the same time, such mechanisms should reflect resource usage, allocate resources efficiently, reimburse costs or maximize service provision profits and lead customers to requesting services that will maximize their revenue. This work proposes a policy for pricing based on resource allocation of a particular category of DiffServ-based services for aggregated traffic in the case of transport networks. Our research takes into account the particularities that apply to the case of DiffServ services' provision over transport networks while imposing minimal overload and a-priori estimation of costs.

**Key words:** QoS charging and pricing, DiffServ, traffic profile, Service Level Agreement

## 1. INTRODUCTION

An important issue in designing pricing policies for today's networks is to balance the trade-off between engineering and economic efficiency. In [1], it is emphasized that pricing schemes that determine prices over short intervals in order to maximize economic efficiency may be unrealistic. Instead, schemes where the utility and cost functions are known and valid for duration longer than a connection's duration are recommended. Results from [2] and [3], based on strong evidence of the history of all communication

technologies and users' reactions, support the argument for keeping pricing mechanisms simple, flat if possible, and claim that even the slightest attempt to impose complex, incomprehensible charging will have a substantial negative impact on usage.

Our work focuses mainly on pricing schemes for the DiffServ framework, which seems to gain significant importance in transport networks worldwide ([4], [5]). The evolution in networking that has emerged from the introduction of service differentiation and QoS provision by the IntServ and DiffServ frameworks has affected traditional network pricing and shifted the interest from fixed access and connection fees to usage-based fees. Usage-based fees are considered appropriate to account for congestion costs, differentiated services, QoS provision and other relevant costs for pricing today's connectionless IP networks ([1], [6]).

Usage-based charging was traditionally based on accounting for the traffic flowing within a network, even in packet granularity, and then determining charges by multiplying the pre-determined price per packet with the number of packets transmitted. Later, the 'smart market' approach that was introduced in [7], required customers to declare their willingness to pay by bidding for network resources for each packet sent. This way, each customer was charged for the marginal cost imposed by the transmission of an additional packet during congestion. The clearing price, determined from the bids supplied, was then used together with per-packet accounting to charge the best effort service.

Although usage-based charging is useful to account for service provision within the DiffServ framework, per-packet or per-flow accounting has to be avoided. The DiffServ framework was designed so as to avoid fine granularity, dealing with traffic aggregates and keeping complexity at the edges of network domains. Moreover, the DiffServ approach focuses on pricing service classes characterised by guarantee parameters and therefore a packet belonging to a traffic aggregate receiving QoS assurances must have a different price than that of a packet that belongs to an aggregate served as best-effort. DiffServ pricing is therefore a service class-based pricing. However, the issue still remains: how should traffic belonging to a single service class be charged?

The establishment of long-term contracts between the customer and the service provider, instead of detailed accounting, was proposed in [8]. The contracts contain traffic profiles as approximation of the 'expected capacity' that the customers purchase from the network services' provider and thus are recommended as indication of resource usage by a customer and the basis for charging. The challenge is to direct the customer towards the selection of the traffic profile that best meets his needs and budget, while at the same

time using this traffic profile in order to determine his usage of resources and emerging charges.

In related research, not many proposals exist that describe concrete methods on how should the money be actually collected within a DiffServ enabled network that supports classes with different, either qualitative or quantitative QoS properties. ‘Expected capacity’ pricing, which has already been mentioned, was innovative in the sense that it proposed charging based on long-term contracts or traffic profiles instead of per-packet accounting or transmission duration. The ‘edge pricing’ paradigm, presented in [9], complements ‘expected capacity’ pricing by shifting pricing activities at the ingress points for a domain but still does not provide a detailed solution for pricing of DiffServ-based services.

The theory of effective bandwidths ([10]) has been used for the purposes of exploiting the ‘expected capacity’ framework towards the direction of specific pricing mechanisms and prices’ determination. In [11] and [12] two compatible approaches for charging flows that obey traffic contracts (or Service Level Agreements-SLAs) according to their effective bandwidth are presented. In [13], the effective bandwidth as a means for reservation of resources in a QoS-enabled network is also used, introducing routing parameters in the allocation and pricing equation. However, the assumption that customers’ benefit received from traffic transmitted between nodes  $i$  and  $j$  is independent of traffic between  $i$  and any other node  $h, h \neq j$ , contradicts our claim that there exist negative externalities in a customer’s benefit from traffic transmitted between two network elements in a QoS-enabled environment.

Unlike most of related work, we claim that, for DiffServ-based services, a flat per packet or per transmitted-volume-unit price within a service class is not efficient from an economical and engineering point of view. We propose a pricing scheme that applies to a significant portion of DiffServ-based services, demonstrates engineering and economic efficiency, preserves simplicity in calculation of customers’ charges and effectively reveals the details of service differentiation and QoS provision. Our approach is innovative because it anticipates for externalities hidden in the costs involved and caused by the nature of such DiffServ services and also because it goes all the way up to the determination of actual prices.

After this introductory section, making a brief reference to related research work, the following section outlines the rationale for the implementation of the proposed pricing methodology. Section 3 describes the proposed methodology, having two separate sections for provisioning and pricing bandwidth and buffer space. In section 4 our proposed future work is outlined and the paper is completed with our conclusions.

## **2. RATIONALE**

The case that will be further investigated in this work is that of pricing a high-priority, low latency QoS service for the customers of a transport network. Such services are provided under different names in DiffServ-enabled WANs worldwide and are built according to the Expedited Forwarding Per-Hop-Behavior (EF PHB) ([14]) of the DiffServ framework.

At every level of an inter-domain hierarchy, the EF-based service architecture requires policing of each customer's EF entitled aggregate at domain ingress borders so that all aggregates will conform to certain characteristics, referred to as the EF aggregate's traffic profile from now on. The very convenient feature of EF-based services is that it is possible to estimate the amount of resources (in terms of capacity and buffer space) that the network should reserve in order for the provisioning of the required quality guarantees, by taking into consideration the traffic profiles of the aggregates for each customer.

For the provisioning and pricing of an EF-based service it is also important to point out the service's critical elements. Reliable transmission of data with the least possible end-to-end delay, almost zero packet loss and the minimum possible variation between the end-to-end delay experienced by different packets are the most crucial factors from the customer's point of view. Moreover, in an EF-based service, the provisioning of transmission resources is taken for granted and the focus shifts to the transmission quality obtained. Unlike best-effort services, bandwidth is not a resource under contention: the amount of bandwidth devoted to such a service is always over-provisioned and therefore whether one of the customers restrains himself to less bandwidth than contracted or not does not affect the performance perceived by the customers' community. However, instead of bandwidth, the resource under contention is buffer space. The negative externalities imposed by congestion in best-effort service provision have their analogy to the negative externalities imposed by delay due to buffer occupancy and packets' waiting time in an EF-based service. In economic theory, externalities are referred to as costs (for negative externalities) or benefits (for positive ones) that do not accrue to the consumer of the good ([15]).

A pricing scheme for the EF-based service must charge customers for both the bandwidth they consume and the buffer space they occupy. Since excessive buffer space occupancy has effects to the whole customers' community (negative externalities), the pricing scheme must lead each customer to select the amount of buffer space that he will buy from the provider in such a way that the negative externalities imposed by that amount of space are compensated (included in the price for this buffer

space) and the customer does not have to shape his traffic, in an effort to reduce that amount, more than he can endure.

In order to emphasize these statements, an indicative comparison to the DiffServ SLAs' pricing proposal made in [12] is provided. This proposal makes an approximation of a flow's effective bandwidth with a token bucket traffic profile of  $(r, b)$  as that of an on-off flow with 'on' rate  $h$  and bases charges on the amounts of effective bandwidth consumed. This approach provides incentives to customers to reduce their peak rate ( $h$ ) without necessarily reducing their demand for buffer space ( $b$ ).

Contrary to the latter approach, what we are proposing is a distinction between the costs imposed to customers for the rate of their token bucket traffic profiles and the costs imposed to customers for the depth of their token bucket traffic profiles. The approach of our work provides a means to charge for the provision of an EF-based service treatment to the aggregates of EF-traffic injected to a transport domain (TD) from its adjacent customers by quantifying the externalities imposed by the depth of the token bucket of each customer's traffic profile, in the quality metric of delay.

Our proposal is not the first one that introduces perceived packet delay in the pricing equation, see also [16] and [17]. However, the context in which delay is used by our proposal, as a means to help users reach a reasonable decision on their required resources, is different from these approaches. For example, in [16] delay is not used as one of the parameters that affect the selection of a customer's traffic profile.

For simplicity of the analysis that will follow, we assume the use of a strict priority scheduler both at the ingress and at the core and egress nodes of the TD. Our analysis will focus on a pricing methodology that will not aim at maximizing the provider's revenue, applying thus on policies exercised by non-profitable providers and having as a goal (apart from expenses reimbursement) to maximize social welfare. It will be one of the targets of our future work to extend this methodology towards scenarios of providers' profit maximization.

### 3. PRICING THE SLAS

Over-provisioning and careful dimensioning can be intuitively assumed to guarantee the required transmission rate and low end-to-end delay for the EF traffic aggregates traversing a TD. In such a situation, the utility function of customers is dependent upon the equivalent capacity that each aggregate perceives and the quality metrics guaranteed (end-to-end delay, jitter and packet loss).

We depict by  $p(S_i)$  the costs that a customer has to pay for purchasing from the TD an SLA with the  $S_i = (r_i, b_i)$  token bucket profile. For ensuring reimbursement of costs for provisioning of the EF-based class ( $c_{EF}$ ), we initially propose that the TD provider partitions  $c_{EF}$  into the cost of the transmission medium  $c_R$  and the cost of equipment  $c_B$  invested in the provision of the EF-based service. Thus

$$\sum_{i \in \{\text{set of SLAs offered}\}} p(S_i) \geq c_{EF} \Rightarrow \sum_{i \in \{\text{set of SLAs offered}\}} p(S_i) \geq c_R + c_B \quad (1)$$

It is important to emphasize that the pricing mechanism proposed should aim at restricting the customer's demands in such a way that, at the equilibrium, each customer's revenue is maximized, without equation (1) being violated.

The customer's actual needs for QoS should lead to the negotiation with the TD provider of an appropriate  $S_i$  and guaranteed end-to-end delay ( $D$ ) for all the customer's EF packets. The provider should engineer its infrastructure so that once traffic contracts are signed with all customers, the provisioned transmission rates and the common to all customers end-to-end delay bound guarantee is ensured.

Provisioning and charging for transmission rate

According to the approach of [18], the TD provider can guarantee a worst-case end-to-end delay bound to all its customers, provided that the ratio of the TD links' capacity to be devoted to the EF traffic injected to the TD is bounded as follows:

$$a < \min_l \frac{P_l}{(P_l - C_l)(h - 1) + C_l} \quad (2)$$

where  $C_l$  is the capacity of each link of the TD, assumed constant  $\forall l, l \in TD$  and equal to  $C$  and  $P_l$  is the maximum rate with which the EF traffic aggregate (emerging from the merging from EF aggregates upstream) is injected at each TD node. Under the assumption of strict priority scheduling  $P_l = d_l C$ , where  $d$  is the fan factor for link  $l$  of the TD. Also  $h$  is the maximum number of hops within the TD that a customer's EF traffic can traverse.

It is recommended by [18] that the TD provider chooses  $\alpha$  such as it is much less than the quantity of the right part of (1). Also, the prerequisite of over-provisioning upon which the provision of an EF class is based, requires that, if  $N$  is the set of customer aggregates routed through this node, for every node  $n$  of the TD it holds that:

$$\sum_{i \in N} r_i \leq aC \Rightarrow a \geq \frac{\sum_{i \in N} r_i}{C} \quad (3)$$

In the worst-case scenario, when all the customers' aggregates are routed through one or more nodes of the TD's core, if  $k$  is the total number of customers for TD, then  $i$  is replaced by  $k$  in (3).

Thus  $a$  has to be equal or larger than the sum of all the customer profiles' token bucket rate divided by the capacity of TD's links. Assuming that each customer will ask for the highest  $r_k$  possible, the proposed mechanism has to turn up with a set of acceptable  $r_k$  values and corresponding prices for the customers so that one or more values for  $a$  can exist, according to

$$\frac{\sum_k r_k}{C} a \ll \min_l \frac{P_l}{(P_l - C)(h - 1) + C} \quad (4)$$

From (4), it is clear that the range  $\{a_{\min} \dots a_{\max}\}$  within which  $a$  can vary is quite limited. In fact  $a_{\max}$  is constantly bounded by the right part of the inequality of (4) which is constant for a certain topology. The TD provider can only vary the selection of a value for  $a$  below  $a_{\max}$  according to the total EF capacity he wishes to sell to his customers. For the rest of this section, we will assume that the TD provider selects a value for  $a$  so that (4) holds and that this value of  $a$  is not negotiated with TD's customers whatsoever. Maintaining a constant value for  $a$  for the rest of our analysis helps in isolating the charging for EF traffic methodology from its side effects on the rest of the traffic that crosses TD. Non-EF traffic will thus be served by  $(1 - a)C$  capacity on each link of TD and will not be affected by any kind of distribution or re-balancing of the resources devoted to EF traffic due to the charging scheme proposed.

After the selection of  $a$ , the TD provider has to distribute a total of

$$r_{tot} = \sum_k r_k = aC \quad (5)$$

EF capacity among his customers. Under the model that this work addresses, all of TD's customers are transport domains themselves, the EF traffic aggregates of which have emerged as the result of aggregation of hundreds or thousands of EF micro-flows. Based on this observation, the TD provider during the pricing mechanism's initialisation phase is suggested to distribute  $r_{tot}$  to his customers in a fair way according to

$$r_i = \frac{C_{access}^i}{\sum_i C_{access}^i} r_{tot} \quad (6)$$

In this way, each customer  $K_i$  receives a share of the EF capacity available according to the capacity ( $C_{access}^i$ ) of his access link to TD.

After the mechanism's initialisation phase, we propose re-negotiation phases of all the contracted traffic profiles simultaneously over long-term intervals. During re-negotiations, each customer will be able to base his new traffic profile's  $r_i$  value selection for the next period on statistical data for the utilization of the rate value allocated to him in the elapsed period. This data can directly be retrieved by the statistics of a token bucket policer of the customer's aggregate in the point of entrance to the TD. Long-term re-negotiation phases will allow customers to evaluate their needs for resource provisioning based on solid, single-dimensional measurements and request the corresponding resources from the TD provider. We claim that this model will demonstrate fluctuations in the beginning, leading to more stable distribution of resources after a number of re-negotiations. Fluctuations are also possible when a new customer will require EF services from the TD provider.

In terms of charging the provided EF rates for each phase, the TD provider is proposed to adhere to the conclusions of relative work in order to calculate the utility perceived by customers for different values of provided transmission rates. According to this ([12], [19]), the price should be a concave, increasing function of the rate provided. Therefore, TD is proposed to charge the EF capacity provided to each customer according to

$$P_{r_i} = \log_x r_i \quad (7)$$

The base  $x$  will be determined separately for each service provision phase by solving

$$\sum_i \log_x r_i \geq c_R \quad (8)$$

where  $c_R$  is the first factor of the right part of (1). This way, since rate provision does not include externalities, each rate ( $r_i$ ) provided is charged with a price strictly proportional to its share of the total available rate space  $r_{tot}$ . After obtaining the value of  $x$ , the TD provider can go back to (7), and announce to each customer the price he will have to pay for the rate  $r_i$  provided to him in the upcoming operation phase.

Provisioning and charging for burstiness



After the selection of  $a$ , the provisioning of resources for servicing EF traffic throughout TD is possible, by configuring all nodes' strict priority schedulers to provide a service rate of at least  $R = aC$  to the EF traffic on all TD links. According to [18], the end-to-end delay is then bounded by

$$D = \frac{h}{1 - ua(h-1)} \left( u \frac{b_{tot}}{C} + \Delta \right), \text{ where } u = \max_l \frac{P_l - C_l}{P_l - aC_l} \quad (9)$$

with  $\Delta = ct$  under the assumption of strict priority scheduling.

At this point the TD provider can negotiate with customers the other part of the customers' SLAs, that of the token bucket depth in correlation with the delay guarantees provided by the TD. The TD provider uses (9) in order to estimate the total buffer space that can be distributed among his customers. In the case of a TD with  $a = 0.05$ ,  $h = 8$ ,  $d_l = 4$ ,  $MTU = 4700bytes$  and  $C = 622Mbps$  the bound on end-to-end delay provided to all customers for different  $b_{tot}$  values is provided in Table 1.

Table 1. Bounded end-to-end delay in a TD with  $b_{tot}$  sized buffers for EF traffic

$b_{tot}$ (packets)	10	20	30	50	100	150	200
$D$ (ms)	7.9	14.88	21.86	35.82	70.72	105.62	140.52

It is apparent that according to the current TD's size and specifications there is a limited amount of total buffer space that can be distributed to its customers. The customers must thus be prompted by the bucket depth charging policy of the TD provider to restrain themselves from selecting large values for  $b_i$  by the fact that this will penalize themselves and others in terms of the delay perceived by their packets. Also the TD provider has to distribute  $b_{tot}$  among his customers so that

$$\sum_{i \in N} b_i \leq b_{tot} \quad (10)$$

where  $N$  is the set of all customers. The latter holds because in the worst-case scenario where all customers' aggregates at some point are routed through one node of the TD's core and all aggregates' bursts coincide inside the buffer space of this node, the node must have enough buffer space to place packets, so that no packets are dropped.

It is at this point that the 'smart market' approach already presented in the introductory section applies. As already mentioned, in the case of EF-based services, resources (i.e. buffer space) must be distributed to those who value them most and distribution has a direct impact on all customers (the end-to-end delay guaranteed by TD). The TD provider announces the end-to-

end delay that can be guaranteed to customers, emerging from (9), and the customers place bids on the available buffer space ( $b_{tot}$ ) in order to obtain a share. The clearing price for a buffer position ( $P_b$ ) is set at the point where the sum of demands for buffer space, starting to add from the higher-bids' demands, reaches the amount of available buffer space  $b_{tot}$  (see Figure 1).




Figure 1. The 'smart market' for a buffer space market

In this way, customers who have valued more a buffer position taking into consideration the end-to-end delay guarantee  $D$ , receive a larger portion of the available buffer space, or in other words obtain SLAs with larger  $b_i$  values than customers who placed lower bids. Of course each customer will be notified of the cost he will have to pay for buffer space when signing a token bucket  $(r_i, b_i)$  SLA as equal to

$$P_{b_i} = b_i * P_b \quad (11)$$

However, if at this point,

$$P_b * b_{tot} \leq c_B \quad (12)$$

where  $c_B$  is the second factor of the right part of (1), the TD provider is notified by the proposed mechanism that his investment in buffer space for the provision of the EF service exceeds the customers needs and it is indicated that the provider does some re-dimensioning.

In a real-life scenario, it is envisaged that the TD provider will distribute the available buffer space  $b_{tot}$  during the initialisation phase according to intuitive bids placed by customers or in a fair manner, since no real-use data will be available. At the moment of re-negotiations, instead of speculating for the future, the customers will be able to place bids on the available buffer

space based on the statistics of a token bucket policer  $(r_i, b_i)$  applied to their aggregates as they enter the TD for the elapsed period. Again, fluctuations will be observed in the first phases or when a new customer will require EF services from the TD provider. However, since the ‘smart market’ and bidding are proven to successfully integrate externalities in goods provision costs, it is envisaged that in equilibrium, the buffer space will be distributed to those who value it most and are willing to compensate for the delay their bursts might cause to others.

#### 4. FUTURE WORK-CONCLUSIONS

As already mentioned, our future work will focus on the case of the analysis of the systems behaviour through re-negotiation phases. We will also focus provider’s profit optimisation and the careful investigation and on investigating the effects of our approach to the customers’ utility functions.

We also aim at dealing with the case of pricing services based on the Assured Forwarding PHB (AF PHB), as defined within the DiffServ framework. Auctioning mechanisms could be of particular use for the pricing of services built upon the DiffServ AF PHB where bidding can be done either among flows of the same AF class but different precedence class, or among members of different AF classes. These issues will be further investigated in our future work, in order to determine whether the research work on auctioning mechanisms for allocating prioritised resources could be exploited for pricing AF PHB-based service classes.

The pricing mechanism proposed in this work is based on traffic profiles that the customers negotiate with a TD provider and concludes on prices announced to customers prior to the service provision interval. In the case of EF-based services, which is under consideration here, the traffic profiles of the customers comprise a sound representation of the utility that each customer finds on the service. At the same time, traffic profiles are used by the TD provider in order to dimension the EF-based service and allocate the resources used by it. The proposed pricing mechanism uses the traffic profiles of customers as the intermediate between each customer and the provider. In this way it reflects both the customers’ revenue from the EF-based service provided and the costs for the service provisioning that the TD provider undertakes. Moreover, the proposed pricing mechanism takes into consideration the in-elasticity in demand for transmission rate that applies in the case of the customers of a backbone transport domain and efficiently allocates the available buffer space to those customers for which accommodation of their bursts is more valuable. Finally, the proposed mechanism provides indications of the quality that will be provided to

customers (in terms of end-to end delay), in order to assist them in the qualitative valuation of the service they will receive and express accurately their needs for resources.

## 5. REFERENCES

- [1] M. Falkner, M. Devetsikiotis and I. Lambadaris, 'An Overview of Pricing Concepts for Broadband IP Networks', IEEE Communications Surveys & Tutorials, Vol. 3, No. 2, 2000
- [2] A. Odlyzko, 'Internet Pricing and the History of Communications', Computer Networks 36, Elsevier Science B.V., pp. 493-517, 2001
- [3] R. Edell and P. Varaiya, 'Providing Internet Access: What We Learn from INDEX', IEEE Network, 12(5), 1999
- [4] Abilene QoS Planning Group, 'Abilene Premium Service Test Program', 2000 (found at <http://www.internet2.edu/abilene/qos/>)
- [5] M. Campanella, 'Implementation Architecture specification for the Premium IP service', Deliverable D2.1-Addendum 1, SEQUIN Project (IST-1999-20841), 2001
- [6] L. A. DaSilva, 'Pricing for QoS-Enabled Networks: A Survey', IEEE Communications Surveys & Tutorials, Vol. 3, No. 2, 2000
- [7] J. MacKie-Mason and H. Varian, 'Pricing the Internet', in 'Public access to the Internet', Brian Kahin and James Keller, editors, Prentice Hall, New Jersey, 1995
- [8] D. D. Clark, 'A model for cost allocation and pricing in the Internet', In L. W. McKnight and J. P. Bailey, editors, 'Internet Economics', MIT Press, 1996
- [9] S. Shenker, D. Clark, D. Estrin and S. Herzog, 'Pricing in computer networks: Reshaping the research agenda', ACM Computer Communication Review, Vol.26(2), pp.19-43, 1996
- [10] F. Kelly, 'Notes on effective bandwidths', In F. P. Kelly, S. Zachary, and I. Zeidins, editors, 'Stochastic Networks: Theory and Applications', Oxford University Press, pp. 141-168, 1996
- [11] C. Courcoubetis, F. P. Kelly, and R. Weber, 'Measurement-based charging in communications networks', Technical Report 1997-19, University of Cambridge, 1997
- [12] C. Courcoubetis and V. Siris, 'Managing and Pricing Service Level Agreements for Differentiated Services', in Proc. of IEEE/IFIP IWQoS'99, UCL, London, UK, 1999
- [13] J. K. MacKie-Mason, 'A Smart Market for Resource Reservation in a Multiple Quality of Service Information Network', Technical Report, University of Michigan, September 1997
- [14] V. Jacobson, K. Nichols and K. Poduri, 'An Expedited Forwarding PHB', RFC 2598, 1999
- [15] T. Henderson, J. Crowcroft and S. Bhatti, 'Congestion Pricing: paying your way in communication networks', IEEE Internet Computing, September-October 2001, pp. 85- 89
- [16] C. Parris and D. Ferrari, 'A resource based pricing policy for real-time channels in a packet-switching network', Technical report, ICSI, Berkeley, CA, 1992
- [17] H. Jiang, I. Sidhu and S. Jordan, 'A Pricing Model for Networks with Priorities', 34th Annual Allerton Conference On Communication, Control and Computing, Illinois, 1996
- [18] A. Charny and J.-Y. Le Boudec 'Delay bounds in a network with aggregate scheduling', in Proc. of First International Workshop on Quality of future Internet Services (QofIS), Germany, 2000
- [19] P. Thomas, D. Teneketzis and J. MacKie-Mason, 'A Market-Based Approach to Optimal Resource Allocation in Integrated-Services Connection-Oriented Networks', in Proc. Of Fifth INFORMS Telecommunications Conference, Boca Raton, Florida, 2000