Pricing in the Internet

Heikki Almay

Nokia Telecommunications
P.O. Box 300, FIN-00045 NOKIA GROUP
heikki.almay@nokia.com

Abstract

The Internet is growing very fast. As the current Internet pricing schemes do not support the planned developments in the Internet service offering, .i.e. QoS schemes and broadband access, also Internet pricing models are under study.

In this paper the environment of the ISP that implements the pricing schemes is explained and the current as well as the potential new pricing mechanisms are discussed.

Basically Internet pricing can be flat rate or it can be based on basic usage parameters such as time and traffic volume. Additionally it is possible to price a connection according to the resources it consumes in the network. These advanced schemes can take into account the state of the network (congestion) and the nature of the traffic flow (burstiness).

1. Introduction

Pricing is a key issue when new types of services are introduced into the Internet. The simple reason is that today the Internet is mainly run by commercial organisations, which obey to the rules of market economy.

Until now pricing in the Internet has not been subject to any specification or regulation work.

This paper focuses on the pricing of connections in the Internet and Internet access. Although the cost of providing value added services (e.g. Internet telephony) is discussed from a traffic point of view, the pricing of value added services is beyond the scope of this paper.

On the pricing of QoS (Quality of Service) in the Internet some research has been done [1], [2], [3]. Most of the approaches are refinements of earlier work done in the ATM/B-ISDN context.

People involved in implementations of applications that require QoS schemes list billing and accounting mechanisms as current deficits and subjects of further work. [4], [5].

In this paper *pricing in the Internet* is viewed from a practical perspective. A pricing scheme alone is not enough. In many cases charging information has to be collected by the network. The used services have to be rated and billed. Additionally it should be possible to explain to the user what he is paying for. Keeping these issues in mind ISP's can start developing their pricing.

The paper first gives an overview of the current ISP pricing and some key trends. With simple examples the environment of an ISP is outlined.

Then the cost of offering differentiated services is discussed. The example service is Internet Telephony. The discussion is about the costs of supporting a new service, not a specific QoS scheme [which can be part of the implementation].

As a third topic the pricing models and their implementation related issues are explainde. Here the focus is on practical pricing models and research done on Internet pricing. Pricing theory is not discussed in detail.

Finally some conclusions are drawn and possible future developments are outlined.

2. ISP and pricing

2.1. ISP framework

In this paper the term ISP is used to describe a company that offers Internet access to residential and/or business customers. Some large companies act both as ISP and backbone provider. In this discussion it is assumed that the ISP does not have a global backbone of its own.

The ISP is connected to one or more international backbone providers and other national networks such as other ISP's and national research networks.

For the international capacity the ISP has to pay to the backbone provider. For the connection between national networks each party pays its share of the costs.

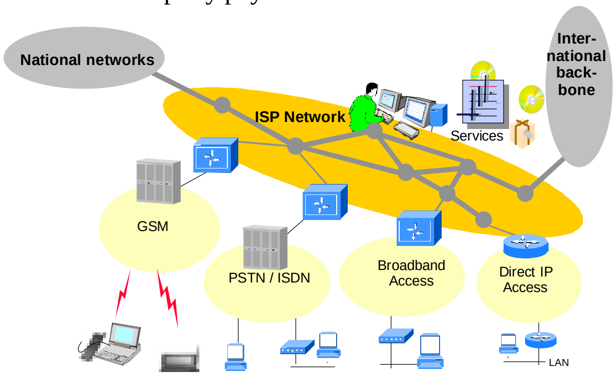

Figure 1: ISP framework

The users are connected to the ISP network using different access technologies. For each of these access types the pricing and the cost structure for the ISP are different.

The ISP also has different customer segments, residential, and different sizes of businesses. The service packages offered to the different customer segments vary. This is also reflected in pricing.

2.2. ISP pricing models

The pricing models used by the ISP's are rather simple. The most common models for residential dial-in customers are:

- Fixed monthly fee
- Fixed monthly fee with time limit + minute based charge for the time exceeding the limit
- Time based charge (premium rate call, no subscription)

The subscription with a fixed monthly fee typically includes e-mail, news and web hosting services.

It is worth to note that during recent years the trend has been towards flat rate pricing.

For business users with direct access to the ISP network the pricing models are less transparent. For dial-in users basically the same tariffs apply as for residential subscribers. Tariffs for business users can include type and bit rate of the access, actual usage, equipment rental etc. It is usually not possible to define them on a per user basis as the ISP's cannot control how many different users actually use one business account.

Often the prices of the ISP services are only a part of the costs that the end user has to pay for Internet access. This is outlined in the following example.

2.3. Example pricing for Internet access

In fig. 2 price examples are given for both the access and the ISP service for residential customers. The price tags indicate monthly charges.

Either the access operator gets all revenue for the access or the income is shared between the access operator and the ISP.

In most European countries an ISP having its access servers connected to the access operators network via ISDN primary rate accesses pays a monthly fee for its lines to the access operator. If the ISP gets a telecom operator license and builds an SS7-interconnect with the access operator, it gets its share of the call charges for all calls made to its Internet access service.

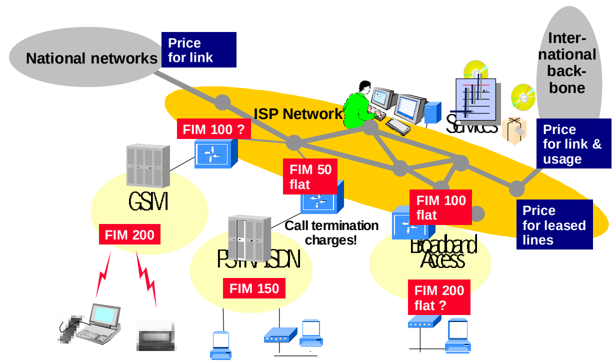

Figure 2: Example monthly prices for residential services & ISP capacity costs

The costs for accessing the ISP are not clearly defined as the telephone bill also covers the costs of voice calls. Internet makes young people keep their fixed line or to buy ISDN. In the example it is assumed that a part of fixed charge is allocated to ISP access.

In the example different price indications are given for mobile, PSTN and broadband IP services. This reflects the current market situation. The basic service is dial-in Internet access. Users with broadband access should be able to use higher bandwidth services than modem users. Without broadband services hardly anyone is willing to pay for the extra access capacity.

There are two ways for the ISP to solve the broadband services issue. It can buy more capacity to the backbone and it can make services available locally (e.g. by extensive use of proxies). Both approaches lead to higher costs.

The price tag for mobile users is high because these users are primarily addressed with value added services (travellers). Example value added services are E-mail notification with SMS and roaming, i.e. Internet access at local tariffs using a local PoP when abroad. Pricing can be time based (e.g. EUnet Traveller).

2.4. Will the ISP business model change

The current business model where residential subscribers pay a flat fee for Internet access bundled with e-mail and web-hosting services to their local ISP is under attack from several directions.

It is increasingly difficult to earn money with conventional "value added" IP services. E-mail services and limited web-hosting services are already available free of charge from various organisations that want to attract people to their sites e.g.

www.altavista.iname.com
www.crosswinds.net
www.geocities.com
www.gnwmail.com
www.hotmail.com
www.fortunecity.com

The service is typically paid by advertisers.

Even Internet access may become free of charge. Retailers (e.g. Tesco in the U.K.) start offering free of charge Internet Access for their customers. Their intention is to get loyal customers to the company portal site and so to push their e-commerce offering.

Because of competition the price level of the basic Internet access is very low. Basic ISP business is in many cases not profitable [6]. As indicated in the example of Fig. 2 a large portion of the costs seen by the user are charged by the telecom operators. This is why ISP and telecom services are increasingly offered by the same organisations. Telcos build and buy ISP's.

The biggest opportunities for ISP's are in access services differentiated by QoS and in charging for content [7]. A survey done by Ovum suggests that an average premium of USD 10 per months is sustainable for QoS and for high quality content.

3. Cost of offering differentiated services

In this chapter the overall costs of offering differentiated services is discussed. The focus here is not on the QoS schemes, but the total cost of the ISP.

IP telephony is not fully representative to all potential applications that could make use of differentiated services. While the learnings from the IP telephony example could rather well apply to other streaming applications such as video, the behaviour of e.g. data VPN should be studied separately.

3.1. Internet Telephony example

One of the services most often mentioned in combination with differentiated service is IP telephony. The costs, pricing and policy of offering Internet Telephony have been analysed [6].

ISP's with Internet Telephony in their service offering have a different cost structure when compared to ISP's without IPT offering (Baseline Scenario). Findings are outlined in table 1.

Table 1: ISP cost Structures [6]

	Baseline Scen. (Without IPT)	*Internet telephony Scen.*
Capital equipment	11%	11%
Transport	24%	28%
Customer service.	26%	26%
Operations	11%	9%
Marketing, sales etc	28%	26%

The importance of transport costs increases as typical dial-in subscribers consume 5 kbit/s. For Internet Telephony usage is 15 kbit/s.

Overall the costs for the Internet Telephony scenario are significantly higher than for the Baseline Scenario. The increase of the individual cost items is shown in table 2.

Table 2: Subscriber cost increases [6]

	Residential Subscribers.	*Business subscribers*
Capital equipment	45%	45%
Transport	75%	75%
Customer service.	44%	44%
Operations	7%	7%
Marketing, sales etc	7%	7%

3.2. Cost structure

In the above example the cost of offering a *differentiated service* for Internet Telephony was not considered. If higher QoS is required, the transport costs rise more than the indicated 75%, unless resources available to non-voice traffic are reduced. In the example Internet Telephony scenario this is not practical as voice traffic is 60% of the total.

It is worth noticing that in the cost structure of an ISP fixed costs dominate. The costs of operating an empty network are essentially the same as the costs of a congested network.

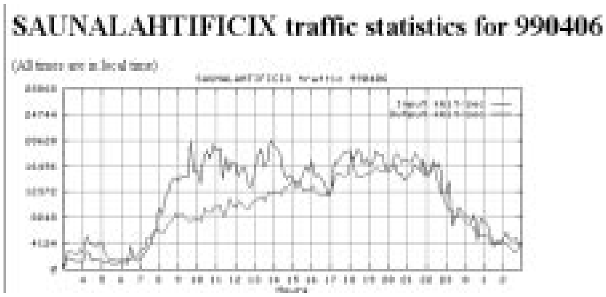

Figure 3: Traffic between Saunalahden Serveri and other Finnish networks

A look at the example traffic charts (Fig. 3 and Fig. 4) reveals that use of the network varies over time. Especially in Fig.4 most of the time a large portion of the network resources remains unused. The ISP capital costs are not affected by traffic volumes but the network capacity built to meet peak traffic.

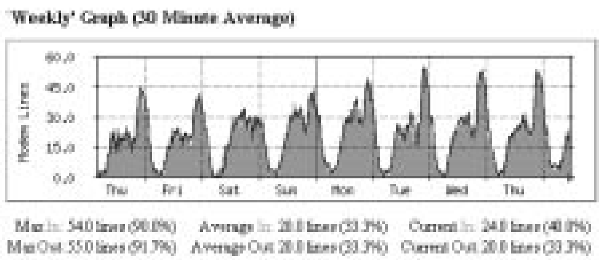

Figure 4: Modem traffic statistics, Cistron, Dutch ISP, 4/1999

In table 3 the world-wide ISP equipment market is shown. Points of Presence (POP) and the backbone network account for 85% of the market. The rest is Network operation centre (NOC), support systems and customer premises equipment (CPE). In these market figures the backbone providers are included. So the figures give a rough indication on the capital costs in providing Internet service (both access and backbone).

Table 3: Segmentation of ISP network market by site of deployment, 1997-1998 (USD) [8]

	1997	1998
Backbone	1,090	1,600
% of total	46%	48%
POP	920	1,230
% of total	39%	37%
NOC and backoffice	210	270
% of total	9%	8%
CPE	140	230
% of total	6%	7%
Total ISP network equipment market	2,360	3,330

4. Pricing models

Price setting contains the following six steps: [9]

- Selecting the pricing objective
- Determining demand
- Estimating costs
- Analyzing competitor´s prices and offers
- Selecting a pricing method
- Selecting the final price

A careful reader might notice from the list that there is no direct link between costs and prices. In stead of carefully looking at the ISP objectives, the demand elasticity and elaborating on mark-up pricing perceived-value pricing only some selected pricing models are discussed. First the rather simple flat rate pricing is analysed. Then schemes for usage based pricing are outlined. Finally some elements of the research work that has been done on more sophisticated pricing schemes are presented.

4.1. Flat Rate Pricing

Flat rate pricing is currently very widely used. It has several advantages:

- It is easy to budget, as the cost of the service is known all the time. This appeals both to residential users and company IT personnel.
- For billing no information from the network is needed. This simplifies ISP operations as no rating of usage information is needed
- Customers do not dispute the bills
- Because of the simplicity flat rate services are very easy to market

Flat rate pricing also has some drawbacks.

- Users are encouraged to use more network resources than they actually need (on-line while idle)
- As usage (minutes/bytes) does not cost, it generates the illusion that the whole Internet is for free – including applications and software
- Flat Rate pricing offers only very basic possibilities for QoS charging [per subscriber].

Flat rate pricing is sometimes said to be “unfair”. An Internet addict preferring round-the-clock real-time video with bit rates that disturb other traffic is paying the same as an occasional surfer.

Now think of two persons renting an apartment. One of them is all the time at home and prefers round-the-clock partying with a lot of friends with a noise level, which annoys the neighbours. The other person is a silent, most of the time away from home. Hardly anyone gets the idea that the noisy person should pay a higher rent than his silent neighbour even though this might be justified from the house owners point of view (e.g. cost of handling complaints, renovation costs, people moving away from the house etc.).

One could say that this is not fair – but we have accepted that in housing rental flat rate pricing dominates.

The ISP business does not work exactly like an apartment building, but it should be noted that the cost difference between a heavy Internet user and the occasional surfer is not necessarily as big for the ISP as it seems. In the IP Telephony example in chapter 2 the costs that can be directly affected by increasing traffic volume and on-line time are capital equipment and transport costs. In the Baseline Scenario they account for only 35% of the ISP total cost. Of the two cost types only transport costs may directly be affected by traffic volumes. For both capital equipment and transport the capacity of the network [peak traffic] is more relevant from a cost point of view.

If we use the monthly price of FIM 50 (example in section 2) as an estimate for the ISP costs, we can state that the fairness issue boils down to one question: How much effort is it worth to allocate an average monthly sum of FIM 17.5/user in a more just way?

The above example is valid for the pure ISP. For operators having both ISP- and backbone operations the transport costs also boil down to capital costs [of the backbone network]. This means in practice that for these players an even lower portion of the total cost is depending on the user traffic.

4.2. Usage based pricing

Some years ago usage based pricing was very common among ISP's for dial-in users. In Finland most ISP's changed to flat rate pricing 1997-1998. End 1998 even T-Online, the most prominent European ISP with usage based pricing included two "free hours" into their monthly fee.

The simplest variant of usage based billing is using time as a measure. One variant is to offer a premium rate telephone number for Internet access without separate subscription. Charging and billing are the responsibility of the PSTN operator. The second option is that the ISP bills for the used minutes. For this the user has to be a registered customer of the ISP.

For dial-in customers using conventional Internet applications time is a good measure of usage as traffic is sporadic and bit rates are low. Time based pricing gives dial in users an incentive to log out in stead of timing out. For the ISP this saves costly modem ports.

For broadband subscribers (e.g. ADSL or Cable Modem) time as a measure of usage is applicable only in special cases. The basic assumption is that the users are always online.

In the ADSL case pricing can be based on sent and received traffic volume (packets, bytes). This information is available in the Remote Access Node where the user PPP-session is terminated and it is carried e.g. over RADIUS to the AAA-server. From there the data can be made available to the billing system and its rating engine which determines the price of the sessions.

In systems with shared media (e.g. most Cable Modem systems) determining the origin of a packet is an issue.

Some ISP's are looking for finer pricing mechanisms such as destination pricing. The idea is to put a price tag on those packets that are leaving or entering the ISP network. So the user has to pay for the traffic that causes the ISP actual costs. With special software (e.g. Cisco Netflow Collector) running on selected routers the additional information on individual packets can be obtained.

Usage based pricing steers user behaviour away from wasting network resources. It is suggested [10] that without usage based pricing there might actually be a higher congestion combined with a lower usage [of a network resource]. The usage is lower if congestion sensitive users do not use the resource.

With more sophisticated pricing schemes some ISP's may try to shift demand e.g. to virtual local communities.

The usage based pricing mechanisms described work on a per packet or per byte basis. They do not take into account the nature of the traffic flows in question. To deliver a bursty hard real-time connection on UDP through a congested network costs per byte as much as any other flow. It is however clear that for the UDP real-time flow to run gracefully much more network resources have to be available than for a traditional TCP connection. So per byte pricing does not necessarily reflect the costs in the network.

Actually one could argue that usage based pricing should only be used during the peak hour, as delivering any traffic during off-peak is inexpensive. As discussed earlier, fixed costs dominate.

Usage based pricing schemes can with some effort be extended to cover QoS schemes by introducing different traffic counters for different CoS.

4.3. Congestion-sensitive Pricing and Smart Market

Research has been done on Internet pricing schemes that are based on the statistical characteristics of the traffic.

The basic assumption for this work is that user demands are increasing exponentially and the network traffic is expected to outstrip available capacity. Under these conditions, efficient bandwidth allocation through statistical multiplexing by itself may not be sufficient to meet user demands. [3]

For characterising bursty connections [originally in ATM connection admission control and dimensioning] the *effective bandwidth* has been introduced. The effective bandwidth of a connection reflects the statistical characteristics of the flow. These properties of a flow can be used to determine how much of the link resources it consumes on a link and if it can be accepted on the link without compromising the level of service of the already present flows. The effective bandwidth has been proposed as a basis for charging [1].

When offering guaranteed services an ISP has to be able to quantify the amount of resources are needed for fulfilling a service level agreement. With the data the network manager can decide how many service level agreements can be supported simultaneously on a link. The information can also be used for pricing. [2]

Congestion-sensitive pricing should supplement flat-rate pricing [3]. Typically limited-term contracts (service level agreements) that reflect the current state of the network are proposed. Probably the theoretically most puristic line of thought is the smart market solution. In this scheme, the users bid the maximum price they are willing to pay to send their message. The highest bid messages are sent first. In any given time interval the lowest bid message that gets sent sets the price for all messages sent. [11]

While the economic theories behind the proposed congestion-sensitive pricing models and the smart market approach are sound, the implementation of the schemes would at least require

- new functionality in the ISP network
- new functionality and a new business model for ISP peering and backbone connections
- new systems for charging and billing
- new price aware user software

New functionality is needed in the ISP network because the price of a connection cannot be determined locally. End-to-end awareness is needed to check for potential congestion. This means that a local intra-ISP implementation of congestion-sensitive pricing is rather meaningless. This again brings along questions of signalling and revenue sharing between ISP´s along the data path.

Charging and billing for congestion-sensitive pricing is not trivial, as prices change uncontrolled over time. The ISP should be able to tell his customers how the costs occurred.

The user should have the possibility to control the money he is spending. In practice this means setting an upper limit to price, but price can be price per byte, price per minute, per hour or per month or per session.

In the following example one method of pricing is presented. It should give the reader a feeling of why user control of the connection price is highly desirable.

4.4. Example: Charging and Accounting for Bursty Connections

In [1] the calculation of a tariff for busty traffic is proposed

$$f(m;M) = a(m) + b(m)M$$

Here M is the [unknown] mean rate of the call, m the users approximation of it. B(h,M) is the effective bandwidth of the connection. f(m;M) is a tangent to B(h,M) at the point M = m. a(m) is a per second price, b(m)M a per bit price. The coefficients are calculated using

$$b(h,m) = (e^{sh} - 1)/s[h+m(e^{sh} - 1)]$$
$$a(h,m) = B(h,m) – mb(h,m)$$

h is the [known] peak rate of the connection, s is a constant.

A numerical example given in [1] with a 100 Mbit/s link and s=0.333 gives following values.

Table 4: A numerical example [1]

	Rate Mbit/s.		*Charge*	
Service type	*Peak*	*Mean.*	*Fixed /s*	*Mbit/s*
1	0.1	0.04	2.7×10^{-4}	1.0
2	2.0	0.02	1.3×10^{-4}	1.4
3	10.0	0.01	1.1×10^{-3}	7.9
	H	M	A(h,m)	B(h,m)

Looking at the mean and peak rates it is hard to determine the charge intuitively. For service type 3 the transfer of a byte is 7.9 times more expensive than for service type 1.

Even though the above charges are fictive, it is quite obvious that a user should be made aware of what types of connections the application are setting up and what costs will occur.

4.5. Who should pay

Once the ISP has set up a pricing scheme where bytes are to be charged a new question pops up. Who should pay for the traffic? Is it always the receiver, always the sender, should both pay, or should there be a way to choose?

At the moment it seams that at least for ADSL access the user will pay for both sent and received traffic when usage based pricing is applied. So actually he pays for using the access line. If this is acceptable in an environment with highly variable prices for different types of flows remains to be seen. It is unlikely that the sender of a high quality video commercial would like the receiver to pay for his ad. But it is as unlikely that non-professional content providers would like to pay for people visiting their home pages.

Clearly the *who should pay* question remains for further study.

5. Conclusions & outlook

Pricing mechanisms in the Internet give ISP's a lot of new opportunities. Pricing of higher QoS could offer ISP´s a way out of the low flat rate price levels of today. There is evidence that some users are willing to pay extra for better service.

The actual winning scheme for QoS pricing is far from clear. It could be a higher flat rate price, a premium on bytes/packets sent using a better CoS. Additionally elements from congestion-sensitive pricing or simpler time-of-day pricing schemes could be added to the scheme for demand shifting and improved resource utilisation. All these measures can be justified.

Three key questions remain unanswered. The first one is, *how to make the new pricing appealing to customers*. What makes this task extra challenging is the common perception that the Internet is for free. Today you just pay the monthly fee for entrance. On user acceptance of new pricing schemes little has been published [12].

The second major question is *how to make it worth the effort*. The discussion in chapter 3 indicates that only a rather small portion of the ISP costs is depending on actual user traffic. Fixed costs dominate.

The third question *who should pay* must be answered when tariff schemes of residential users are based on traffic volumes or flows. A random choice may severely restrict the available content in the Internet or willingness of people to surf the Internet.

Despite of the open questions usage based pricing is likely to come back. At least in the broadband context it is gaining popularity as a tool for differentiation.

Introducing dynamic congestion-sensitive pricing to the real-life network is hardly going to happen soon. The simple reason is, that the users would not understand what they are paying for, neither would the ISP know why he is charging any given price.

What model will win on the market place remains to be seen. We should keep in mind that the cost structures of all players are not equal. If network traffic really should outstrip available capacity and make congestion-sensitive pricing attractive probably the service providers with own international backbone networks are the last to suffer. So the competing service offers might look like the two outlined in the figure below. One ISP optimising the use of the network, possibly building pricing incentives for shifting traffic to off-peak hours, QoS schemes for high-end users, the other just adding capacity – and probably charging a premium for the high quality network.

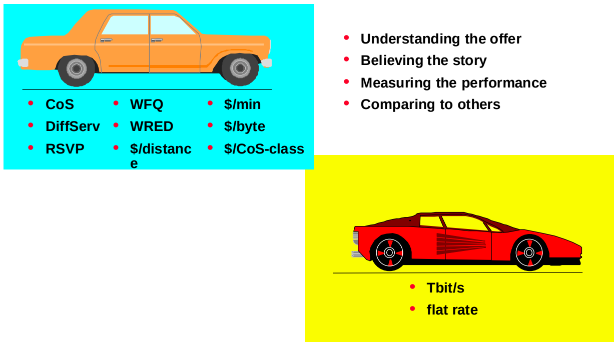

Figure 5: Perceived value &pricing

It is quite easy to make guess that the simple alternative will be more appealing for a large portion of the market.

References

[1] Kelly, F., Charging and Accounting for Bursty Connections, MIT Workshop on Internet Economics, March 1995, http://www.press.umich.edu/jep/works/KellyCharg.html

[2] Courcoubetis C., Kelly F. P., Weber R., Measurement-based usage charges in communication networks. Statistical Laboratory Research Report 1997-19, University of Cambridge. http://www.ics.forth.gr/proj/race/publications/index.html

[3] Kalyanaraman S. et al, Dynamic Capacity Contracting: A Framework for Pricing the Differentiated Internet, submitted to First

International Conference on Information and Computation Economics, May 1998, http://www.ecese.rpi.edu/Homepages/shivkuma/research/papers-rpi.html
[4] Schulzrinne Henning, Rosenberg Jonathan, Internet Telephony: Architecture and protocols – an IETF perspective, Computer Networks 31/3, 11.2.1999, p. 251-253
[5] Toga James, Ott, Jörg, ITU-T standardization activities for multimedia communications on packet-based networks: H.323 and related recommendations, Computer Networks 31/3, 11.2.1999, p.219 – 222
[6] Mc Knight Lee, Leida Brett, Internet Telephony, Costs, pricing and policy, Telecommunications Policy, Vol 22/7, p.555 – 569, 1998
[7] Accessing the Internet: The Challenge for ISPs and Telcos, Ovum Report, July 1998
[8] Datamonitor, Equiping ISPs, Datamonitor Plc, London, UK, December 1997
[9] Kotler, P., Marketing Management, Analysis, Planning, Implementation and Control, 7th edition, Prentice Hall, 1991, 0-13-563479-2, p.474-506
[10] Pricing congestible network resources, MacKie-Mason, J.K.; Varian, H.R., Selected Areas in Communications, IEEE Journal on, Volume: 13 7 , Page(s): 1141 –1149
[11] Parenteau Bernard, Rishe Naphali, Internet Pricing and Priorization, Community Networking Proceedings, 1997 Fourth International Workshop on , Page(s): 93 -101
[12] Edell Richard, Varaiya Pravin, Demand for quality-differentiated network services, Proceedings of the 36th Conference on Decision & Control, San Diego, California, USA December 1997