Pricing and Brokering Services over Interconnected IP Networks

D. Di Sorte, G. Reali

DIEI - University of Perugia
Via G. Duranti 93 – 06125 Perugia - Italy
{disorte,reali}@diei.unipg.it

Abstract

In this paper we consider an IP-based communication platform, characterized by a novel architectural model where barrier-free business and market interactions can be performed. We assume that the network is able to deliver application services which need of network service performance guarantees. In this scenario, we present the concept of the *network commodity* that is traded in the marketplace among network service providers, application service providers and customers. Moreover, we use this concept to define a new usage-based tariff model. Then, we especially focus on the activity of the so-called Network Resource Brokers, which have the goal of finding the end-to-end inter-domain path to deliver an application service that maximizes the users' benefit, in terms of price and perceived service level. In this regard, we present a QoS-and-price based inter-domain routing algorithm, analyze its computational complexity, and show its effectiveness in a selected simulation scenario.

Keywords: *business model, brokering, network commodity, usage-based charging, QoS, inter-domain routing.*

Corresponding Author:

Dario Di Sorte, Ph.D.
Dipartimento di Ingegneria Elettronica e dell'Informazione (DIEI)
University of Perugia
Via G. Duranti, 93 – 06125 Perugia
Italy

Phone: +39 075 585 3630
Fax: +39 075 585 3654
e-mail: disorte@diei.unipg.it

1. Introduction

In this paper, we introduce a new paradigm for network engineering and service distribution, which has been developed in the framework of the IST program within the project WHYLESS.COM [26]. This paradigm is not to be intended as a substitution of the legacy telecommunication paradigms, but rather a viable alternative and an opportunity for networking and service operators to go into business within a competitive environment by means of clear and common rules. The interest in observing such rules is that they may attract customers, which are helped in accessing the telecommunication infrastructure and services through assisted procedures that help them in taking appropriate decisions. A novelty of this paradigm is given by a clear separation of the communication infrastructure from the information content, which is expected to be a consequence of the deregulation tendency of the telecommunication world. Therefore, one of the main challenges of the project is to show that it is worthy to set-up a telecommunication business activity based either on the infrastructure or on the information contents. It is expected that in the future a number of small and medium size proprietary communication networks will be deployed and offered to customers as a viable alternative to the large operated networks. Typical examples are metropolitan networks, campus networks and the plethora of wireless local area networks that will accompany different business activities. This means that potential customers will have a number of different opportunities to obtain services, and the communication arrangement will proceed according to new guidelines. Our goal is to decrease the market entrance barrier to providers and to make the use of their services easier and advantageous for customers.

In particular, we can expect the birth of new entities involved in the overall service deployment that will perform brokering activity. Given the separation of the network infrastructure from the information contents, we assume that the brokering services will be classified similarly. Therefore, our business model includes both Information Brokers (IBs) and Network Resource Brokers (NRBs). To improve the final service for customers and to stimulate competition among providers, brokers must be network entities administratively independent.

Before illustrating their roles in details, we introduce the other main players of the model. They are the Internet Service Providers (ISPs) (or, more generally, Network Service Providers, NSPs), which provide the network infrastructure, and the users, that represent both customers (or end-users) and Application Service Providers (ASPs), which are the providers of application services (browsing, e-mail service, voice and video over IP,…). We assume that each NSP is able to differentiate network services for supporting future multimedia applications adequately, by means of resource reservation capabilities, as suggested in [1] and [2]. The improved network service may be specified in terms of the Quality of Service (QoS) parameters associated with the information transmission. All these concepts will clearly have consequences on the accounting and pricing approaches. If the Internet becomes a multi service network, ISPs will expect additional incomes for providing different QoS levels, and service differentiation will probably imply a tariff differentiation and the introduction of usage-based charges, as suggested in [1]. This new scenario is clearly very attractive for ISPs, since it could provide them more market space. In this sense, usage-based charging is efficient from the economic point of view. From the business point of view, if a user receives an improved network service, it means that he is obtaining something more than a user that receives a bad one, and therefore he has to pay more. In addition, some studies [3] show that customers are willing to pay an additional per-usage charge in order to improve network performance and to avoid performance degradation during network congestion.

In this scenario, the IB is introduced to act as a mediator between the customer and ASPs. On request of a customer, it provides information about which ASPs can provide the desired service. The role of the IB could be very important for both the end-users, when they have to access the network in unfamiliar areas (for instance outside their office/home), and the ASPs, which, through the IB, may present its offer to customers.

For what concerns the NRB, its role is to match, on a per-call basis, ASPs requests and NSPs offers to establish an end-to-end connection able to support a given application service with an adequate QoS level. For long-term Service Level Agreements (SLAs), the set-up of inter-domain routing tables among NSPs can be established off-line. On the other hand, in the case of end-to-end per-call Information Transport Resource (ITR) trading, it is necessary to implement dynamic procedures to establish an end-to-end path each time. For this reason the role of inter-domain routing in the next generation Internet is very important, since it has to control an interconnection of administratively independent networks. Thus, we introduce an independent centralized entity, the NRB, to perform inter-domain routing algorithms. End-to-end services are not built in a nested manner, so that small ISPs have not the necessity of stipulating agreements with larger ISPs and, consequently, they are not dominated by them. In addition, a centralized inter-domain routing approach could permit a more flexible end-to-end service provision by brokering activities without any preliminary agreements among domains.

In the deployment of any QoS inter-domain routing approach, the definition of a standard measure of the QoS, common to each service provider, is necessary. In our approach, we have faced the problem by introducing the concept of network commodity traded in the marketplace and the relevant QoS measure (i.e., the virtual delay). In our approach, the *network commodity* is defined as *the transfer of information units between two end points of the network*, and the measure of the commodity is a function of the QoS parameters. The measure (amount) of the used commodity is related to the QoS parameters associated with the information transmission (end-to-end delay, the delay jitter, and the packet loss probability). We show how the number of commodity units may be determined by using a function $f(d)$ of a specific service descriptor d, called *virtual delay* [4][5]. Since such function gives the number of commodity units associated with the information transmission, it is used by NSPs to define a per-time and/or per-volume tariff model to charge for guaranteed network services [4].

Once such a model is described, the following step is the design of a routing protocol that can be used by the NRB to determine the cheapest path (Minimum Price, MP, inter-domain routing algorithm), among all the available networks (i.e. domains) to connect two end-points while preserving the desired QoS support. Some concepts have been outlined concisely in [5]. In addition to a more detailed description, the value added by this manuscript to the contents of [5] consists of a deeper description of the algorithm and an extension to determine the path that maximizes the users' benefit (Maximum Benefit, MB, inter-domain routing algorithm). The users' benefit is the *customer surplus function* defined as the difference between the amount of money the customer is willing to pay for a given service (*demand curve*) and the cost of the service itself. In addition, we will present a computational analysis of the algorithm and we will show some results obtained from simulations in a network scenario.

The paper is structured as follows. In section 2, we present the network and business reference environment, along with the entities involved and their role in the model. In section 3, we explain the rationales of our business model. In section 4, we present the concept of network commodity and describe the tariff model to charge IP guaranteed services. In section 5, the inter-domain routing algorithms, based on the end-to-end price of the guaranteed network service, are shown, together with the computational analysis and simulation results. In section 6, a typical procedure to provide customers with an application service is described step by step. Finally, section 7 reports some concluding remarks.

2. Network and Business Models

The entities that characterize the business scenario are depicted in Fig. 1 and Fig. 2:

- *User*: this entity models the network users. It indicates:
    - *end-users*, which are private users allowed accessing the network for receiving

specific application services. Typically, an end-user can provide only qualitative information regarding the requested service level without entering in complex technical aspects (e.g., traffic description, QoS parameters) related to the network support. End-users pay a fee to ASPs for the received service;

  - *Application Service Providers (ASPs)*, which provide customers with application services, and charge them accordingly. In order to fulfill the service requests, ASPs must infer the appropriate technical parameters from the customers qualitative description to elaborate the suitable network service request to NSPs. ASPs pay a fee to NSPs, for the network support. They receive a fee from end-users for the application service.

- *Network Resource Managers (NRMs)*: in general, we consider a number of independent NSPs (hierarchically related from the networking point of view), each one managed by an NRM, which provide ASPs with the network infrastructure necessary to set-up a service for the requesting customers. To do this, they receive a fee from them. NRMs are in charge of guaranteeing a set of edge-to-edge services to flows passing through the relevant domain. The techniques used to support QoS (e.g., Differentiated Services [6], Integrated Services [7][2]) inside domains are selected by NSPs arbitrarily. Moreover, each domain is able to provide an external characterization of the traffic handled between any pair of input-output ports of the domain boundary (called Per-Domain Behaviors, PDB, in the Differentiated Services paradigm [22]), in terms of QoS and price parameters;
- *Information Broker (IB)*: this entity has the task to associate the requests of end-users with the offers of ASPs. For this service, it receives a fee from both of them. According to our model, customers may contact an ASP in order to receive a specific service. This is the highest abstraction level to describe business interactions between end-user and ASP; we refer to it as *outer business model*. If necessary, end-users may consult an IB, which acts as mediator between customer requests and offers by ASPs, and provides information about which ASPs can provide the desired service and their “history”, that is the level of satisfaction of previous customers (Fig. 1). The role of IB could be very important when end-users access the network in unfamiliar areas (for instance outside their office/home). Clearly, the presence of lot of brokers can be foreseen in operation. Nevertheless, from the logical point of view our model includes only one IB.
- *Network Resource Broker (NRB):* on request issued by ASPs, it has the task of providing the cheapest path with adequate QoS level to deliver a service for a specific

end-user. It receives a fee for the requested service from both ASPs and NRMs. Also in this case a lot of brokers are necessary in operation, even if in our model they are represented by a single entity. The entities (ASPs, NRMs, and NRBs) and the operations involved in the overall ITR brokering, which is transparent to customers, concern the so called *inner business model*. Such an aspect will be deeply investigated below.

- *Clearing House (CH):* it is a trusting center. It assures that the certified entities involved in transactions are enabled to perform their work, coping with security matters, which are necessary for business over the open platform. Since the operation of the CH is beyond the scopes of this work it will no longer be addressed.

In summary, for the provision of a complete service, two types of contracts must be stipulated: the former between the end-user and the ASP concerning the application service, the latter among the ASP and each NRMs involved in the end-to-end information transfer for the network support. This means that the subjects responsible of the network service payments are just the ASPs, which, in their turn, charge customers for the final service. This model is also a possible solution to the question about the distinction between the charging at network layer (*transport accounting*), which should be transparent to customers, and the charging at higher layers (*content accounting*, which is typically end-to-end). The goal of the former is to charge for the transfer of packets through the network, whereas the goal of the latter is to charge according to the content of packets (e.g., see [8]).

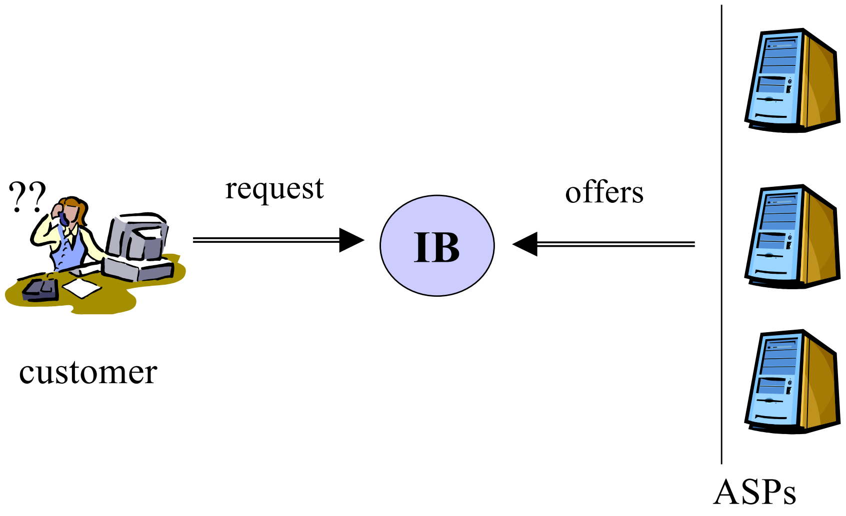

Fig. 1 - Outer business model.

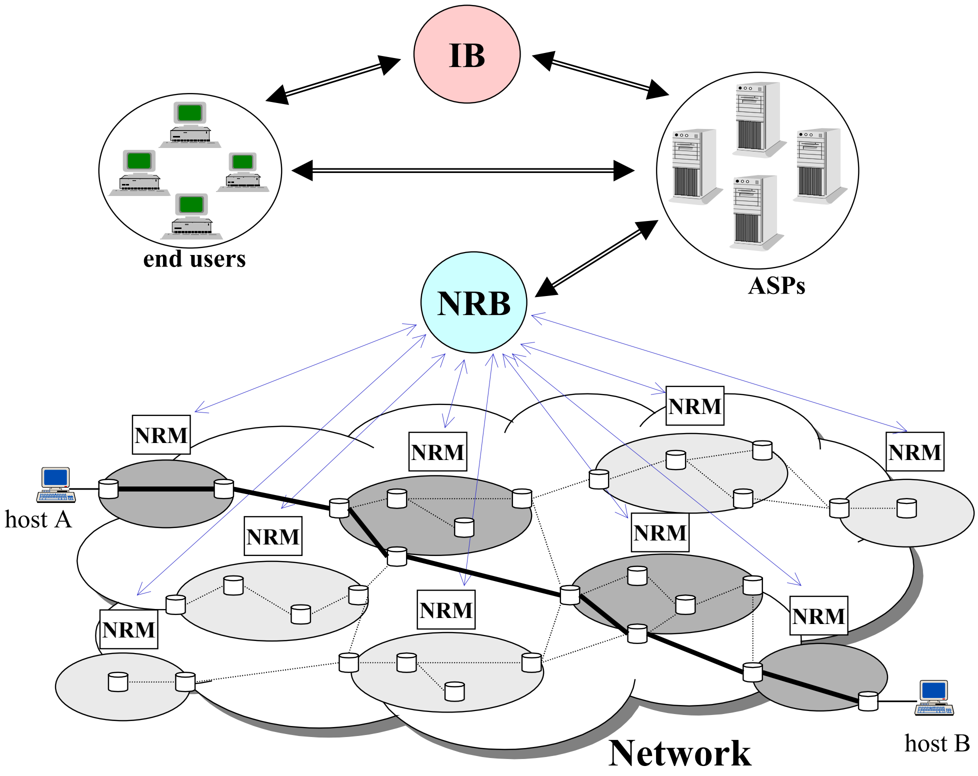

Fig. 2- Reference environment.

In this paper, we are especially interested in network pricing.

As regards the Internet, the most common charging model for the network services is the flat-rate one. According to it, subscribers pay for accessing the network (*access charge*) only, independently of the traffic volume exchanged. Prices usually depend on the speed of the access link. Also the relations among different domains are generally derived from the flat-rate approach.

The advantages of the flat-rate model are that it is simple to understand for customers and to be managed by network operators (it does not need additional accounting architecture!), it incites users to use the Internet freely, thus stimulating its development. At the same time, the flat-rate model has a very important limitation: since it does not provide the possibility to pay to obtain a desired perceived QoS (and also to be paid to provide an added value network service), it is not efficient from the economic point of view.

We essentially agree with the perspective illustrated in [1], where the Author states that "*The introduction of QoS services creates a strong impetus to move to usage-based tariffs, where the tariff is based on the level of use of the network's resources. This, in turn, generates a requirement to meter resource use…*"

For this reason, to overcome the economic inefficiency of the flat-rate pricing model, we propose that each domain adds a usage charge to the flat fee for added value services. The usage charge may be either a function of the traffic volume V exchanged during the communication session, or/and a function of the time T during which the user is active, as explained in section 4. The interested reader can find a deep network, economic and social analysis of pricing models in [9][10][11].

If we consider a flow passing through different domains, it is necessary to determine the total price for the provision of the end-to-end transfer service. According to [12], we assume that the price of the end-to-end service is explicitly given by the sum of the prices charged by domains involved in the end-to-end information transfer.

For the sake of completeness, it is also worth noting that it is also possible assume that pricing procedures are implemented locally, by the access ISPs, rather than in a distributed manner in the network, following the *edge pricing* model [13]. Each ISP charges its own subscribers using the tariffs that fit better the requirements of its own administrative domain, also taking into account the agreements with its neighboring and/or hierarchically higher ISPs.

Our choice has been driven by the following consideration. The network users can negotiate and stipulate agreements not only with the access domain, but also with all the domains involved in the information transfer. This model breaks the basic paradigm of the current Internet regarding bilateral SLAs among neighboring ISPs. The consequence is an increase of the market offer because the user is not constrained by the choices of his access ISP and agreements between ISPs.

3. The rationales of the proposed model

The technical evolution in telecommunications is being accompanied by an important corresponding research activity in the economic field. This activity reflects the same evolutionary speed as it can be observed in the technical fields. Such an activity has generated updated economic and business models, which address the dynamism of the way of accessing the telecommunication services and infrastructure (e.g., [23][24][25]).

In this regard, the definition of the open network architecture is an important concept for the social and economical implications. Essentially, it consists of an integrated model on both the physical network interconnection and the contractual relations among the entities of the business model. The experience of close network architectures, culminated in the definition of huge 3G cellular systems, has shown the weakness of a strong regulation in terms of slow deployment and difficult competition due to high entrance barriers. The reduction of such barriers is one of the main goals of an open network architecture, which has been widely addressed by our network and business model. In our model we deal with economic market entry barriers. Nevertheless it is important to consider that there are also other types of barriers, such as institutional barriers, which could seriously restrict competition. This analysis of this latter type of barriers is beyond the scope of this paper.

As regards the economic market entry barriers, they consist of the costs to be paid before any profit is made. If the essential condition to do business by nationwide communications is to install optical fiber or copper all over the country, this implies a huge block of fixed costs, which is a huge entrance barrier. The reduction of this barrier is possible if we allow “small” companies, which manage local and metropolitan area networks, to gather in an open way, and provide nationwide connectivity. Clearly this is not an easy task. The essential condition is the open way of proceeding. In fact, if we bound this operation by contractual agreements among players, the economy of scale of large operators would likely allow them to control the whole market anyway. On the contrary, if we allow the market player to access an unrestricted market, a real open network may be deployed. Clearly, it is important to define some rules anyway. These rules should allow heterogeneous organizations, which are free to design, implement, and manage these networks, combining forces in a standard and unrestricted way to provide end-to-end connectivity.

The solution that we have proposed and analyzed is based on the logical separation of the supply of telecommunication infrastructure from the delivery of added value services and contents. This separation of contents and infrastructure enables the concept of network resource trading. We propose to implement this trading process via a brokerage of ITR. The reason of this choice is that it allows the construction of end-to-end network connectivity with the following characteristics:

- no commercial agreements among organizations required;
- competition at mere price and QoS level;
- no restrictions on the infrastructure size and quality;
- real independence, without any need of being someway driven by a huge organization;
- no need of bounding customers by specific contracts to compete with other organizations;
- presence of a trusted entity which guarantees the fair competition of players;
- no need of high capital spending for entering the market.

Also from the customer point of view it is possible to identify some advantages. The most important one is that this economic model should prevent the market from being controlled by monopolies, along with their negative consequences. This will improve the social welfare and push competitors to invest to solve problems such as the digital divide.

Another important improvement that can be expected for customers is the simplicity in accessing the ICT (Information and Communication Technology) infrastructure. This simplicity is due to two different reasons. The former is that ASPs mask all technical aspects related to QoS and network support through the interaction with NRB within the inner business model. The latter is that the selection of the ASP may be driven by IBs (outer business model). In fact, currently, customers are surrounded by many devices, which should allow them to access several services. The plethora of services that should be offered within an open network could be confusing, and paradoxically induce customers to access only few of them. Since this could severely limit the effective exploitation of the wide range of different heterogeneous access technologies, ambient intelligence and context-aware solutions are currently being developed within the IST project SIMPLICITY [27]. Our proposal is to decrease this latter user entrance barrier through the adoption of IBs, which are therefore extremely important entities, since they are the glue of the market offers and the user willingness of having improved service.

In summary, the targets mentioned above, in particular the social welfare can be obtained through the proposed solution of an open network, if the implementation, deployment, and operation of it preserve the peculiarities of heterogeneity and independence of the various players. The actions expected by the two brokers, IB and NRB, are aimed at guaranteeing these characteristics in the two business environment in which the market is supposed to be split: network infrastructure and services.

In particular, this paper focuses on the operation of the NRB.

For long-term SLAs, the agreements between NSPs and the set-up of inter-domain routing tables can be established off-line. On the other hand, in the case of end-to-end per-call ITR trading, it is necessary to implement dynamic procedures to establish an end-to-end path each time. For this reason the role of inter-domain routing in the next generation open networks is very important, since it has to control an interconnection of administratively independent networks.

QoS and price inter-domain routing issues have been faced in [17][18], where the Authors present similar solutions, based on SLA Trading Protocol (SLATP) and a modified version of the Border Gateway Protocol (BGP), respectively. For both approaches, the central point is the exchange of information concerning resource allocation, price negotiation and path selection among neighboring domains, without involving other entities other than ISPs, such as brokers. In short, a NSP sells its own edge-to-edge services and buys others. SLA trading is only carried out between neighboring domains. It is possible to build end-to-end transfer services recursively in a nested manner. The price of the service depends on the sum of the costs of each SLA involved. It is important to note that SLA trading, in principle, is not performed on the time scale of a call, but implicitly assumes a given degree of market dynamism.

Moreover, in the deployment of any QoS inter-domain routing approach, the definition of a standard measure of the QoS, common to each service provider, is necessary. In this direction, the lack of standards for SLA is just one of the main problems to overcome. In our approach, we have faced the problem by introducing the concept of network commodity and the relevant QoS measure (i.e., the virtual delay), as explained in the following section. The added value of our pricing approach is just the definition of a tariff model based on this standard and marketable commodity, which accounts of all the technical and market-driven parameters of the transfer service.

Clearly, any solution that is distributed, and therefore that follows the current trend of the IETF, needs the upgrade of the existing protocols and architectures. On the other hand, an independent centralized entity, in charge of performing intra-domain routing, could overcome the main limitation of typical SLA trading in a multi-service scenario. This means that bilateral nested agreements are not necessary for doing business. In this way, as emphasized above, small size ISPs can overcome the necessity of stipulating agreements with a larger one and being dominated by it, as it occurs when end-to-end services are built in a nested manner. Finally, a centralized inter-domain routing approach could permit a more flexible end-to-end service provision by brokering activities without any preliminary agreements among domains. The NRB is an independent entity, and its task is to stimulate competition among service providers, so as to better satisfy customers' requirements.

On the other hand, significant scalability and feasibility constraints must be overcome from a technical point of view. It is not possible to deploy the whole brokerage activity in a single centralized physical entity addressing requests of all network users. It is our opinion that a feasible solution is to create a DNS-like hierarchical structure for a NRB entity (in this sense the system can be considered distributed). Each element of the broker may control a portion of the whole network and is also able to cooperate with the other components. In addition, in an open market, a number of independent NRBs addressing the requests of particular user groups and the offers of a number of NSPs should be expected.

4. Network Commodity and Tariff Model

In this section, we describe the concept of the "commoditization" of the network service, and the pricing law used to charge for IP guaranteed services [4][5], both of which are based on the so-called *virtual delay*, which is a QoS index of the transfer service offered by an administrative domain.

We assume that each flow entering the domain is shaped by a Dual Leaky Bucket (DLB), thus the DLB parameters are the traffic descriptors. As regards the IP networks, the regulated traffic is commonly characterized by a set of parameters, called T_{Spec} (traffic specification), which consists of a token bucket size B_{TS}, a token (sustainable) rate r_S, a peak rate P_S , a minimum policed unit m, and a maximum datagram size M . The presence of the parameters m and M is due to the packet nature of transmission. To simplify the analysis, often in literature a fluid model of the flows is used. The corresponding simplified DLB model, used in this paper, is based on three traffic descriptors only: P_S, r_S, and B_{TS}.

The administrative domain offers a guaranteed service and it is consequently interested in getting revenue from the market. We identify the *network commodity* offered by network operators as the transfer of information units from a point A to a correspondent point B in the network. The QoS of the port-to-port IP network service provided to the specific flow is described by the following *service parameters*: the maximum transfer delay, $D_{\max}$, the maximum delay jitter, D_{jitter}, and the packet loss probability, P_L. We assume that these parameters can be negotiated between network service providers and network users. Other QoS parameters, such as channel reliability, resilience and connection set-up time (*network parameters*), characterize the intrinsic quality of the network, do not depend on the specific flow and cannot be negotiated;

We assume that the service parameters are summarized by the so-called *virtual delay*, *d*. This quantity is a comprehensive and all-inclusive appraisal of the transfer delay, the delay jitter and the loss probability, and characterizes an edge-to-edge service offered by a NSP.

In general, it is possible to assume the virtual delay as:

- a measure of the QoS level of the service computed by the domain from purely technical considerations. Such a measure has validity only within the domain itself. In this regard, a model to compute the virtual delay value can be found in [15];
- a common reference in the market for the QoS level provided by a given network service. In this case, the virtual delay is the basis of inter-domain routing algorithms to minimize the price of a service, while satisfying QoS constraints imposed by a specific application. The measurements of the quality of a service is an interesting research field, the scope of which is to derive the Mean Opinion Score (MOS), which represents the average quality felt by a given number of individuals (see, e.g., [14]). From this point of

view, the virtual delay is a standard measure of the QoS level of a service transmission in the global network market.

In this paper, we assume the latter definition. Our assumption is that the standard law (common to each administrative domain) to associate a virtual delay with a set of service parameters ($D_{transfer}$, D_{jitter}, P_L) is $d = c_1 D_{transfer} + c_2 D_{jitter} + c_3 P_L$, where the constant values, c_i, $i=1,2,3$, depend on the customer sensibility to the specific service considered (e.g., a phone call). Therefore, the virtual delay is assumed to be a standard measure of the QoS level of the service, and it is the basis of inter-domain routing approach described in the next section.

Therefore, a network service is modeled as a hypothetical equivalent service with the virtual delay, d, which gives a measure of the QoS level: the higher the level of the service, the lower the value of d. Moreover, we consider a monotonic, not increasing function of the virtual delay, $f(d)$, which associates the port-to-port transfer of an information unit with a technical measure, expressed in commodity units. Each domain is free to choose the function which best fits its own requirements. The cost (or value) of the transfer of an information unit from a point A to a point B with a virtual delay d is $S(d) = \alpha_{A \to B} f(d)$, where $\alpha_{\mathrm{A}\to\mathrm{B}}$ is the cost of each commodity unit, which depends on

- network parameters. As regards wireless access domains, the specific $\alpha_{\mathrm{A}\to\mathrm{B}}$ value is especially related to the channel reliability, which depends specifically on the error control techniques implemented at link layer;
- the two points A and B (e.g., their distance);
- the policies of the relevant domain.

Clearly, when the commodity is on the market, its price can fluctuate according to factors that are beyond technical considerations. For this reason, we define the price of the transfer of an information unit as the quantity $P(d) = \gamma S(d) = \gamma \alpha f(d) = \beta f(d)$, where γ is a price variation factor that accounts for market fluctuations, and $\beta = \alpha\gamma$ is the market commodity price, that is the price per commodity unit, determined by the NRM. For instance, it could be related to the degree of network congestion (*congestion pricing*). Below, we assume that both the value of β and the QoS level are constant during the connection.

Let T be the duration of a connection and t_0 its starting time. According to [4], the per-call tariff applied to charge for the service offered to a flow entering the domain with an instant bandwidth equal to $B_{ist}(t)$ is

$$Q = \beta f(d) \int_{t_0}^{t_0+T} \max[B_{ist}(t) - B_{res}, 0]dt + \beta f(d) B_{res} T \,.., \qquad (1)$$

where B_{res} is the bandwidth value that a domain charges on a per-time basis. Such a value ranges from zero to the peak rate of the flow. For instance, it could be set equal to the amount of bandwidth reserved for the flow, but, in general, its meaning is only for charging purposes. According to the tariff (1), extra-usage of bandwidth (with respect to the component B_{res}) is charged on a per-volume basis. Thus, the tariff consists of a component depending on the duration time of the connection (*allocation charge*) and a component depending on the amount of traffic volume exchanged (*effective usage charge*). The weights of the two components can be adjusted arbitrarily, by simply varying B_{res}, which, in this view, may be regarded as a tunable knob. Even if the allocation charge ensures a minimum amount of revenue to network operator, on the other hand users would like to be charged according to their actual use of the service (effective usage charge). Consequently, the value of B_{res} can be appropriately chosen by network operators according to their own pricing policy. It is worth noting that if the value of B_{res} increases, the price charged to network users increases as well, unless $B_{ist}(t) \geq B_{res}, \forall\, t$. In this latter case, the tariff remains unchanged.

From (1), it is possible to express the tariff applied by an administrative domain to charge a given network service for a specific input port-output port pair as

$$Q = a_V(d)\widetilde{V} + a_T(d)T\,, \qquad (2)$$

where $\widetilde{V}$ is the amount of traffic volume charged on a per volume-unit basis (see Fig. 3), and $a_V(d)$ and $a_T(d)$ are the per volume-unit charge and the per time-unit charge, respectively.

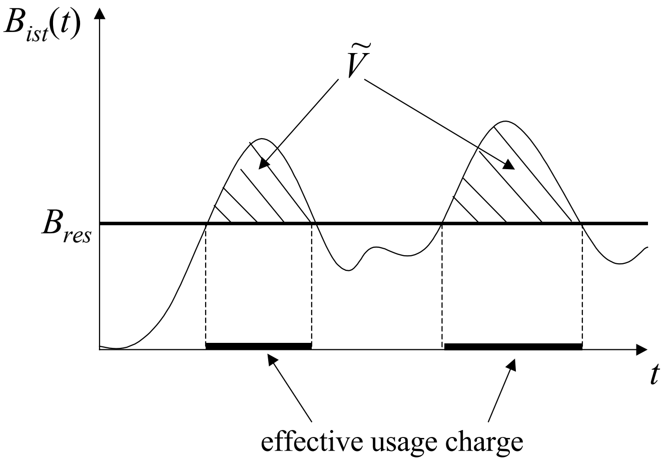

Fig. 3 - Time intervals with effective usage charging, and traffic volume $\widetilde{V}$.

It is worth noting that the highest tariff corresponds to the maximum burstiness of the transmission rate; in fact, it is well known that bursty flows (in particular ON/OFF shaped ones) stress network resources more than flows with a smoothed transmission rate [16].

We would like to express $\widetilde{V}$ as a function of the duration *T* of the connection, so as to identify a cost function in terms of amount of money per time-unit. To this aim, in addition to the DLB parameters (r_S, P_S and B_{TS}) it is necessary to obtain other statistical information on the source flow. In particular, we assume to know two parameters:

- the fraction $0 \le \lambda \le 1$ of the volume transmitted at a rate higher than B_{res};
- the fraction $0 \le \theta \le 1$ of the time when the transmission rate is higher than B_{res}.

In general, the values of λ and θ are functions of the B_{res} value. Clearly, $\lambda(0) = 1$, $\lambda(P_S)=0$, and $\theta(P_S) = 0$. The knowledge of the traffic parameters results from both technical considerations, regarding the flow supporting the application, and statistical knowledge of the behavior of customers.

It is easy to verify that

$$\widetilde{V} = r_S T\lambda - B_{res} T\theta , \tag{3}$$

under the assumption that the average transmission rate of the source is actually equal to the declared r_S value.

This way, from (2) and (3), the tariff can be written as

$$Q = \{a_V(d)[r_S\lambda - B_{res}\theta] + a_T(d)\}T . \tag{4}$$

The cost function in terms of money per time-unit is

$$F(d) = a_V(d)[r_S\lambda - B_{res}\theta] + a_T(d) , \tag{5}$$

that is

$$F(d) = \xi f(d) , \tag{6}$$

where ξ is the amount of money per commodity-unit per time-unit, equal to

$$\xi = \beta(r_S\lambda - B_{res}\theta + B_{res}) . \tag{7}$$

It is important to underline that the functions $\lambda(B_{res})$ and $\theta(B_{res})$ must be estimated off-line by the ASPs. ASPs know the statistical properties of the traffic flows of the application services that they provide. The knowledge typically derives from the DLB parameters used to shape flows, or statistics done by using measurements. By means of this knowledge it is possible to determine the functions $\lambda(B_{res})$ and $\theta(B_{res})$. They can be used for the following reasons:

- to provide NRB with a per time-unit cost function (6) to run the inter-domain routing algorithm, described in the next section;
- to identify the equivalent per time-unit tariff to charge customers. In fact, a simple per-

second tariff, covering network costs, would be easily understandable to end-users. Note that it is also possible to determine an equivalent per volume-unit tariff, which is equal to $G(d) = F(d)/r_S = \xi f(d)/r_S$. The tariff applied must be obviously communicated to the end-user before the beginning of the communication session.

5. NRB activity - Inter-Domain Routing Algorithm

In this section, we describe the brokering activity of the NRB. In our business model, the NRB has the task of controlling: (i) the end-to-end path; (ii) the end-to-end QoS levels of the network service; (iii) the end-to-end price of the network service. In other words, the NRB must perform QoS-and-price based inter-domain routing through a number of different administrative domains, and must choose the best end-to-end path in terms of QoS and price. Therefore, upon a request issued by an ASP, the NRB checks the offer of NSPs (summarized in the so-called e-Tables, where QoS and price parameters characterized the network services are described) and makes use of an inter-domain routing algorithm to maximize the user's benefit (or *customer surplus function*) (Fig. 4).

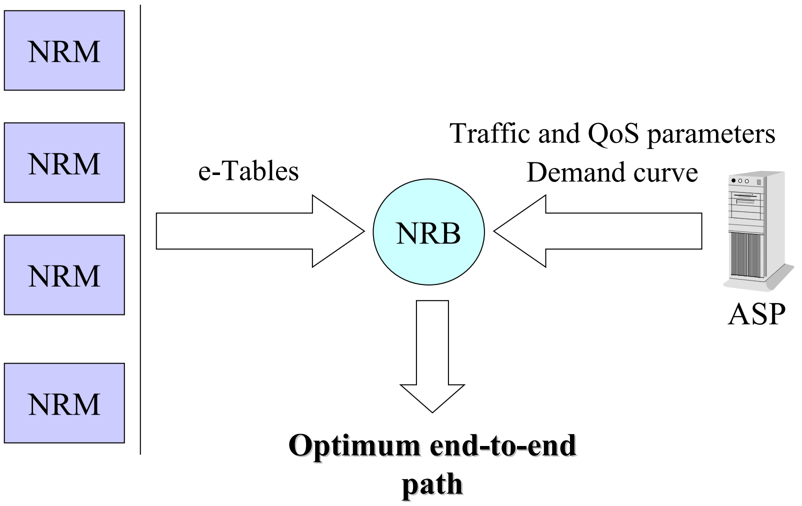

Fig. 4 - Inner business model.

The user's benefit is defined as the difference between the demand curve and the end-to-end tariff charged for the service, given by the sum of tariffs charged by NSPs crossed along the path. As known, the demand curve measures the customer’s sensitivity to the perceived QoS level, and in particular the willingness to pay for a given service performance level. In our approach, the demand curve can be expressed as a function of the value of the end-to-end virtual delay. An example of demand curve is shown in Fig. 5. Without loosing generality, we have assumed that the total end-to-end allowed virtual delay for a given type of application service ranges in the interval [0, 14], and we have partitioned it in three sub-intervals associated with excellent, good and acceptable services, respectively. It is our opinion that a step demand curve is able to satisfactorily represent the users' willingness to pay.

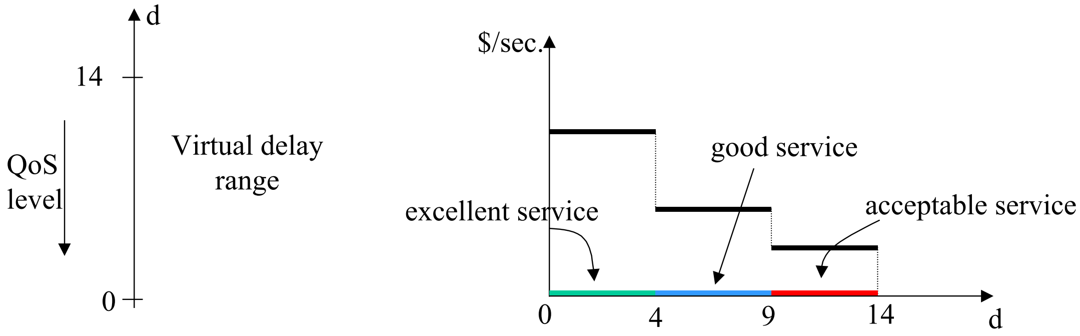

Fig. 5 - An example of demand curve.

5.1 *Algorithm description*

It is necessary to define a metric and a routing algorithm based on that metric to find the path providing end-users with the maximum benefit in terms of QoS perceived and price to pay.

Let F_j be the cost function associated with the generic port-to-port path of a generic domain j. Its value is given by (6). It indicates the equivalent price per time-unit charged by the relevant NRM for the specific flow and the requested QoS, and is computed by the NRB according to parameters provided by the ASP and the NRM. Let $d_{\max}$ be the allowed end-to-end virtual delay, determined by using the specific set of service parameters $(\overline{D}_{\max}, \overline{D}_{jitter}, \overline{P}_L)$, representing the QoS constraints of the specific application to be supported. If a flow is routed through the cascade of N networks, with guaranteed delays $d_1, \ldots, d_N$, the end-to-end delay constraint requires that

$$d_{TOT} = \sum_{j=1}^{N} d_j \le d_{\max} \ . \tag{8}$$

Let us first suppose that the task of the NRB consists in minimizing the cost of the service provided to users, and that each domain is able to provide transfer services with d ranging in the interval [0 $d_{\max}$].

To determine the price, it is necessary to distribute the maximum allowed delay $d_{\max}$ among the crossed domains, since their tariffs clearly depend on the QoS level. It means that it has first to determine the path, then to distribute the total delay among the domains involved, such that the total cost F_{tot} for accessing the service is minimized, that is

$$F_{tot} = \min_{\{d_1,\ldots,d_N\}} \sum_{j=1}^{N} F_j(d_j) \tag{9}$$

with the constraint (8).

One may wonder if relation (9) implies that the minimization can be done on the end-to-end basis only. Generally speaking, if a given path results to be optimum, this does not imply that a portion of it is still optimum for the same delay. This consideration implies that it is not possible to assign a given delay to a specific sub-network without considering all sub-networks of the path together.

This behavior can be easily shown by analyzing the situation illustrated in Fig. 6.

Looking at Fig. 6, it is evident that the preferred path from the user U_1 to the user U_2 is that passing through domains A and C if the target end-to-end delay is small, whereas, if the tolerable delay is large, it is convenient to allocate a small delay to domain C and a larger component to domain B. This simple example shows that even if the shape of the functions is similar but they intersect in some points, it is not possible to define the optimum inter-domain path for all cases. If the two functions do not intersect, then the optimum path is the same for all delay values, and the optimum route from U_1 to U_2 can be determined regardless the target delay value. This also implies that the route is optimum for both hard and soft delay guarantees.

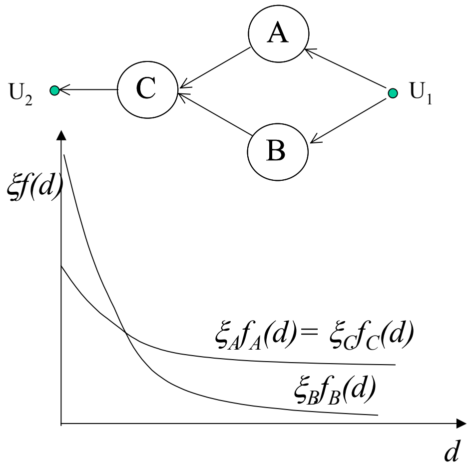

Fig. 6 - Three-domain example with relevant cost functions.

We have developed a procedure that first identifies the optimum inter-domain end-to-end path, and then distributes the overall virtual delay over the involved domains. In the following we show how it is possible to proceed step by step to determine the optimum route.

Let us now consider the situation shown in Fig. 7.

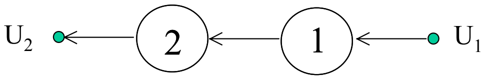

Fig. 7 - Two-domain case.

Relation (9) can be written as follows:

$$\min_{0 \le d \le d_{12}} [\xi_1 f_1(d) + \xi_2 f_2(d_{12} - d)] = g_{1,2}(d_{12}), \tag{10}$$

where d_{12} is the maximum allowed delay from U_1 to U_2. The structure of (10) has appeared very frequently in the literature under the framework of min-plus algebra; it is called min-plus convolution, as seen in [19][20][21]. This operation is associative and commutative. In particular, the cost function $g_{1,2}(d_{12})$ is the min-plus convolution of the functions $\xi_1 f_1$ and $\xi_2 f_2$. Below, we will refer to this operation by using the symbol '*', that is

$$g_{1,2}(d_{12}) = \xi_1 f_1 * \xi_2 f_2 . \tag{11}$$

This way, $g_{1,2}(d_{12})$ gives the optimum delay values $d_{opt,1}$ to be allocated to the domain 1 and $d_{opt,2} = d_{12} - d_{opt,1}$ to be allocated to the downstream domain 2.

In turn, the target delay d_{12} could result from another optimization process. For example, Fig. 8 shows a path through three domains, 1, 2, and 3.

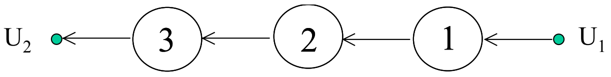

Fig. 8 - Three-domain case.

The optimum delay distribution over the path may be expressed as

$$\min_{0 \le d \le d_{13}} [\xi_3 f_3(d) + g_{1,2}(d_{13} - d)] = g_{1,3}(d_{13}), \tag{12}$$

where d_{13} is the maximum allowed delay to be allocated. It is possible to write (12) as

$$g_{1,3}(d_{13}) = \xi_3 f_3 * g_{1,2} . \tag{13}$$

If we call $d_{opt,3}$ the optimum delay value resulting from (12), the corresponding optimum delay to be distributed over domains 1 and 2 is $d_{12}=d_{13}-d_{opt,3}$. The optimization over 1 and 2 can be done by using (11).

Therefore the total cost function is given by

$$g_{1,3}(d_{13}) = \xi_1 f_1 * \xi_2 f_2 * \xi_3 f_3 . \tag{14}$$

This approach can be extended to N domains by using the function:

$$g_{1,N}(d_{1N}) = \min_{0 \le d \le d_{\max}} \left[\xi_N f_N(d) + g_{1,N-1}(d_{\max} - d)\right] = \xi_1 f_1 * \xi_2 f_2 * \ldots * \xi_N f_N . \tag{15}$$

Let us now consider two different paths, namely path 1 and path 2, which start from a node A and merge to a downstream node B by crossing a number of K and H domains, respectively. Their cost functions are given by K and H min-plus convolutions that produce the cost functions $g_{A,B}^{(1)}$ and $g_{A,B}^{(2)}$. Clearly, if it happens that in a given delay range $g_{A,B}^{(1)} > g_{A,B}^{(2)}$, in that range the path 1 can be discarded, and $g_{A,B} = g^{(2)}_{A,B}$. It may also happen that in a different delay range $g^{(1)}_{A,B} < g^{(2)}_{A,B}$, consequently $g_{A,B} = g^{(1)}_{A,B}$. Therefore, the cost function in the delay range of interest [0, $d_{\max}$], may be associated with different paths. In this case, we can speak of *logical paths*, which, in general, can be composed of different inter-domain physical paths.

Finally, given a target maximum delay $d_{\max}$, the Minimum Price (MP) routing algorithm is defined by the following steps:

1. Starting from the source, which will be referred to as node 1, all the departing inter-domain paths that do not create loops are considered.
2. A metric is associated with each path, obtained by computing the min-plus convolutions of the cost functions of all domains of the path, computed in the range [0, $d_{\max}$]. We are assuming that each domain is able to provide a QoS level equal to d=0.
3. If M_m paths converge towards the same input port of the generic mth domain, they are compared. Since the maximum allowed delay is $d_{\max}$, for each delay value in the range [0, $d_{\max}$], only the path relevant to the minimum cost function survives, and all other paths are discarded. More precisely, the global cost function at the input of the domain is defined in the range [0, $d_{\max}$], and for each value $d \in [0, d_{\max}]$ its value is

$$g_{1,m}(d) = \min_{i=1,\ldots,M_m} \{g^{(i)}_{1,m}(d)\}. \tag{16}$$

4. At the destination domain, referred to as D, the values of the cost functions of the M_D logical paths survived, computed at $d_{\max}$,

$$F_{tot,i} = g^{(i)}_{1,D}(d_{\max}), \quad i=1,\ldots, M_D, \tag{17}$$

   are compared, and the cost function corresponding to the minimum value is selected. Therefore, the problem is reduced to find the optimum logical path $\bar{i}$, the total cost of which is the minimum among the M_D possible solutions, that is

$$F_{tot,\bar{i}} = g^{\bar{i}}_{1,D}(d_{\max}) = \min_{i \in \{1,\ldots,M_D\}} g^{i}_{1,D}(d_{\max}). \tag{18}$$

5. Once the total cost function is built up and the logical path is defined, it is necessary to follow the way back to determine the inter-domain physical path and to distribute the maximum delay $d_{\max}$ over the selected domains. From (15), $d_{opt,N}$, relevant to the last domain, is found. The remaining delay to be distributed is $d_{\max}$-$d_{opt,N}$. Similarly to the determination of $d_{opt,N}$, the value of $d_{opt,N-1}$ can be found, along with the path relevant to the portion of the cost function, previously selected by (16), corresponding to the delay

value $d_{\max}$-$d_{opt,N}$. This way it is possible to return back to node 1 and to terminate the algorithm.

6. Once a virtual delay d_j is allocated to each domain j (j=1,…,N) of the optimum physical path, a portion of the end-to-end service parameters $(\overline{D}_{\max}, \overline{D}_{jitter}, \overline{P}_L)$ is distributed over these domains. The jth domain must guarantee the following set of QoS parameters: $(\delta_j \overline{D}_{transfer}, \delta_j \overline{D}_{jitter}, \delta_j \overline{P}_L)$, where $\delta_j = d_j / d_{\max}$.

It is worth noting that, if the number of domains is high, this approach requires the computation of a potentially high number of min-plus convolutions (see the complexity analysis in the following sub-section), without any guarantees on the possibility of discarding paths. For this reason, it is clear that in this case the definition of a common function $f(d)$ for each domain would be extremely useful. In fact, this would assure that, if two cost functions are generated by the same number of self-convolutions of the same decreasing function, even if multiplied by different coefficients, they never intersect. Moreover, even if the number of convolutions of each path is different, some regularities are preserved, that is the global cost functions do not probably intersect repeatedly. This would simplify the algorithm since it would not be necessary to associate a number of different paths with a cost function at each step, so requesting a lower amount of memory. In any case, it is reasonable to assume that in operation the number of domains could be some unit.

Now, let us extend the MP inter-domain routing algorithm to the MB inter-domain routing algorithm.

Let us suppose that the ASP provides the NRB with the demand curve $U(d)$, describing the cost per time-unit tolerable by a customer, as a function of d, ranging from 0 to $d_{\max}$. The optimization problem can be changed as follows. Now, our goal is to maximize the function *customer surplus* $C(d)$, defined as the difference between the amount of available money $U(d)$ and the cost of the service $g_{1,D}(d)$ (Maximum Benefit, MB, inter-domain routing algorithm).

It represents the benefit perceived by end-users, in terms of price and QoS level.

It is possible to associate each logical path i, i=1,…,M_D, with the relevant value of customer surplus C_i given by

$$C_i(d) = U(d) - g^i_{1,D}(d). \tag{19}$$

The maximum value of $C_i(d)$ is given by

$$C_{i,\max} = \max_{0 \le d \le d_{\max}} C_i(d) = C_i(d_{tot,i}), \tag{20}$$

with $d_{tot,i} \le d_{\max}$.

Therefore, the problem consists in finding the optimum logical path $\bar{i}$, the customer surplus of which is the *maximum* among the M_D possible solutions, that is

$$C_{\bar{i},\max} = \max_{i \in \{1,\dots,M_D\}} C_{i,\max} . \qquad (21)$$

From the optimum value $d_{TOT} = d_{tot,\bar{i}}$, it is possible to find the physical inter-domain path and to distribute d_{TOT} over the N involved domains. In other words, it is necessary to partition the service parameters in the set of values $(a\overline{D}_{transfer}, a\overline{D}_{jitter}, a\overline{P}_L)$, where $a = d_{TOT} / d_{\max}$, according to the factors $\delta_j = d_j / d_{TOT}$ (j=1,…,N), as illustrated above.

We underline that the algorithm shown can be easily extended to the case when domains can not offer QoS guarantees under a given level identified by a minimum value of virtual delay $d_{\min}$. If we consider the cascade of two domains U$_1$ and U$_2$ shown in Fig. 7, where d_{12} is the maximum allowed delay and $d_{\min 1}$ and $d_{\min 2}$ are respectively the minimum amount of virtual delay allocable to U$_1$ and U$_2$ respectively ($d_{\min 1} + d_{\min 2} \le d_{12}$!), relation (10) can be changed as follows: $\min_{d_{\min 1} \le d \le d_{12} - d_{\min 2}} [\xi_1 f_1(d) + \xi_2 f_2(d_{12} - d)] = g_{1,2}(d_{12})$.

In turn, as regards the case with three domains (Fig. 8), relation (12) can be changed as follows: $\min_{d_{\min 3} \le d \le d_{13} - d_{\min 1} - d_{\min 2}} [\xi_3 f_3(d) + g_{1,2}(d_{13} - d)] = g_{1,3}(d_{13})$, where $d_{\min 3}$ is the minimum amount of virtual delay allocable to the third domain.

In conclusion, the routing algorithm is unchanged in its essence. There is to consider only the constraint on the minimum level of QoS that a path is able to provide. In particular, the sum of the minimum virtual delays of relevant domains must be lower than or equal to the value $d_{\max}$.

5.2 Complexity analysis

In order to analyse the complexity of the routing algorithm, we model the topology of the network as an undirected graph. A set of nodes, labelled by n = 1,...,N, may send, receive, and relay data across an intra-domain communication links, represented by arcs. Each one corresponds to an (i, j) pair of different nodes. The presence of an arc (i, j) means that the network can send data from the sender node i to the receiver node j and vice versa.

We label the arcs by the values of cost $\xi_{ij} \ge 0$. Let Q denote the number of paths which connect the source and the destination. Since the number of such paths typically grows up exponentially with the number N of nodes of the network, we will express Q as $O(a^N)$, where a is a constant and N is the number of the nodes of the network.

To determine the complexity of the algorithm, we compute the total number of floating-point operations (flops) to be executed in the worst case. This number is a function of the parameters of the problem. Since the problem complexity has an asymptotical meaning, in order to determine it all non-dominant terms are discarded. A generic algorithm is said to run in $O(g(n))$ time if for some numbers c and n_0, the processing time of the algorithm is at most $cg(n)$, $\forall\, n \geq n_0$.

As illustrated above, the routing algorithm is based on the use of min-plus convolutions, and the general structure of the cost of each path is shown in (15). Essentially the cost value depends on the choice of the commodity function $f(d)$; since the domain of the cost functions are meaningful only in the range between 0 and $d_{\max}$, this is also true for their convolution. This means each convolution domain has to be restricted to the meaningful range $[0,\ d_{\max}]$. It turns out that the cost of each step (i.e., a min-plus convolution computation) is the same; it will be denoted as $C_{M\text{-}P}$. The computational cost of the routing algorithm is given by the theorem below.

Theorem

In the worst case the total cost of the MP/MB routing algorithm is $C_{M\text{-}P} \cdot O(a^N \cdot N)$.

Proof:

In the worst case, the network is very dense, that is every node is adjacent (i.e. connected) to every other node, and every path is composed of up to N domains. Further, at each domain, we have to compare the relevant cost function, defined in the range $[0, d_{\max}]$, with that of the a^N different paths which could converge towards the domain, where a is a constant value.

In summary, in the worst case we must compute $O(a^N \cdot N)$ min-plus convolutions. Then, we must compare the values of the cost functions associated with the Q paths. This step requires Q-1 inequalities, i.e., Q-1 flops (a number negligible which does not influence the asymptotical complexity of the algorithm).

At the end of the algorithm, once we have found the optimum path, we can find the optimal distribution of the virtual delays in few flops (also in this case this number is negligible with respect to the total complexity of the algorithm).

The cost of the operations necessary to extend the MP algorithm towards the MB algorithm is also negligible.

On the basis of the previous observations, we can estimate that in the worst case the total cost of the MP/MB algorithm is given by $C_{M\text{-}P}\cdot$ O($a^N\cdot N$). □

The reader may observe that the theoretical cost in the worst case is extremely high. Even if this analysis is mandatory for this type of algorithms, this result is not representative of the situations that may occur in operation. In fact, being the algorithm proposed for inter-domain routing, in the most common situations the hypotheses of the proof does not apply. In particular it is a matter of fact that a border router is *not* adjacent to all other border routers. Note that in this way the most performance penalising hypothesis is removed. In addition, the number of nodes (i.e. border routers) that is generally considered is not that high, since it is proportional to the number of domains of the inter-network analysed, which is generally in the order of some units or, at most, of some tens. In conclusion we can realistically assume that the time of convergence of the algorithm is fully compliant with the brokering time scale.

5.3 Simulation results

To investigate the basic properties of our proposals, we have selected the 6-domains network topology shown in Fig. 9. The network is made up of two domain types: the "access" one, which includes the connection to end-users, and the "backbone" one, which allows data delivery between access domains. In particular, we consider four access domains (A, B, C, D) and two backbone domains (E, F). Domains are connected by means of bi-directional inter-domain links. We assume that each domain is able to guarantee an edge-to-edge network service between each couple of ports characterized by any value of virtual delay in the interval [1 14].

Our simulator, written in C++, includes both the MP and the MB inter-domain routing algorithms.

Backbone domains have to compete in order to attract traffic and obtain revenues. The results of the competition have been compared according to different strategies, depending on the traffic demand from access domains and the offered network services. The used topology tends to achieve symmetric traffic distribution around backbone domains, in order to provide a fair comparison. Results show the economic success or failure due to different market strategies.

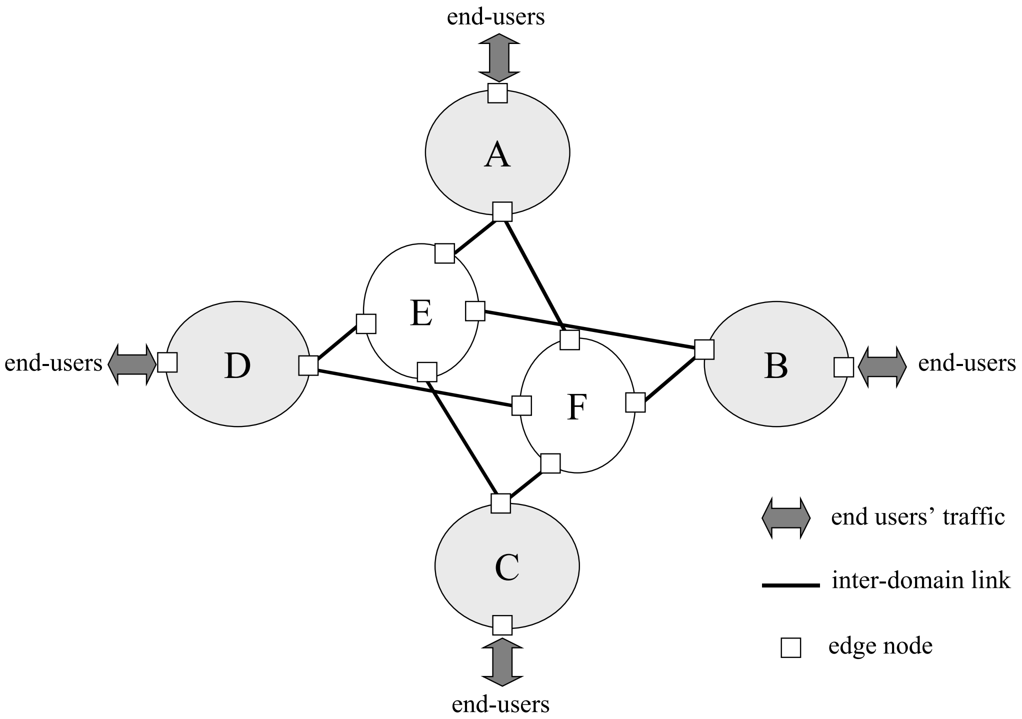

Fig. 9 - 6-domain network topology.

We assume that the network is loaded by homogeneous CBR flows, each one characterized by a bandwidth value equal to 64 Kbps (audio-like source). The traffic entering the network is modeled as a Poisson arrival process with parameter λ, and the call duration is modeled as an exponentially distributed random variable with mean value equal to $1/\tau$. Since the traffic is assumed to be uniformly distributed over the access domains, the call arrival rate at each access domain is equal to $\lambda/4$ on average. In addition, the flows entering a given access domain are uniformly forwarded towards the other three access domains. This implies that the call arrival process from a given access domain to another access domain can be modeled by a Poisson process with parameter $\lambda/12$. The capacity of each internal port-to-port link in the backbone domains is assumed to be equal to 960 Kbps (i.e., 15 calls supported). Since our goal is to study the behavior of backbone domains, we assume that the access domains are not bottlenecks for the end-to-end connections (i.e., there are no blocked calls in access domains). Finally, we assume that all domains charge a per-second tariff in the form $\alpha\gamma B_{res} f(d)$ \$/s. The value of B_{res} is assumed to be equal to 64 Kbps, whereas the choice of the other pricing factors is left to the policy of domains. It is worth noting that in our analysis the tariff charged for a given call remains constant during the call duration.

In the following we analyze three case studies.

- Case study 1

Now we assume that the tariffs charged by domain E are lower than those of domain F. In addition, we have selected the price parameters $\gamma = 1$, $\alpha = \frac{11}{16}10^{-8}$ \$/commodity-unit, and $B_{res} = 64000\,\text{bit/s}$, common to the two backbones. The function which associates a virtual delay value with a measure of the commodity (i.e., commodity-units/bit for a given virtual delay value) is different. In particular, we assume, for virtual delay values $d \in [1, 14]$, for domain E a function $f_E(d) = 2 - (1/14)d$, and for domain F a function $f_F(d) = 2.22 - (1.11/14)d$. Consequently, domain F charges tariffs about 11% higher than the ones of domain E. For instance, for the best service (d=1), the tariff of E is 0.0509 \$/min (i.e., about 0.0848 cent/s), whereas the tariff charged by F is equal to 0.0565 \$/min (i.e., about 0.0941 cent/s). The maximum value of virtual delay for the application is fixed at $d_{\max} = 5$.

We run the MP inter-domain routing algorithm in two load situations, light and heavy. In the former case, the call arrival rate is $\lambda_L = 0.4\,\text{s}^{-1}$ on average and the average call duration is 333 s (i.e., $\tau = 0.003\,\text{s}^{-1}$). In the latter case, the average arrival rate is assumed to be $\lambda_H = 0.7\,\text{s}^{-1}$ and the average duration of a call is equal to 500 s (i.e., $\tau = 0.002\,\text{s}^{-1}$).

Let us first consider the light load case. It is expected that flows are first routed towards domain E. When all the resources of domain E are in use, the inter-domain routing algorithm routes the incoming calls towards domain F.

The global utilization factor and the total revenue of the two backbones are shown in Fig. 10 and Fig. 11, respectively. In the light load case, domain F clearly gets a small amount of traffic and revenue with respect to domain E.

If we charge the network heavily, as described above, the situation changes, as shown in Fig. 12 and Fig. 13. Domain F gets more traffic and more money with respect to the previous case.

At this point, we proceed with the following assumptions: (i) the network is very heavily loaded ($\lambda_{VH} = 0.8$ and $\tau = 0.002\,\text{s}^{-1}$); (ii) domain F increases its own tariff. In particular, the function is chosen equal to $f_F(d) = 4 - (1/7)d$. This means that the tariff of F is twice the tariff of E. Our goal is to show that the total revenue of domain F at the end of the simulation is higher than the one of domain E (Fig. 14), also if domain E gets more traffic (Fig. 15). Obviously, the higher the difference between the tariffs charged by backbone domains, the higher the gap between the revenues of two domains at the end of the simulation. This is due to the fact that even if we have run the MP inter-domain routing algorithm, the willingness of customers to pay, i.e., the maximum acceptable price, has not been considered.

In the following case study, we will introduce the demand curve.

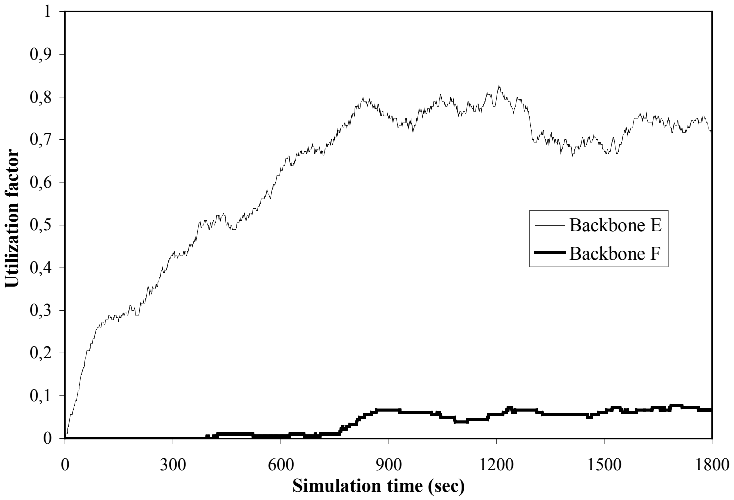

Fig. 10 - Utilization factor of backbone domains: lightly loaded case.

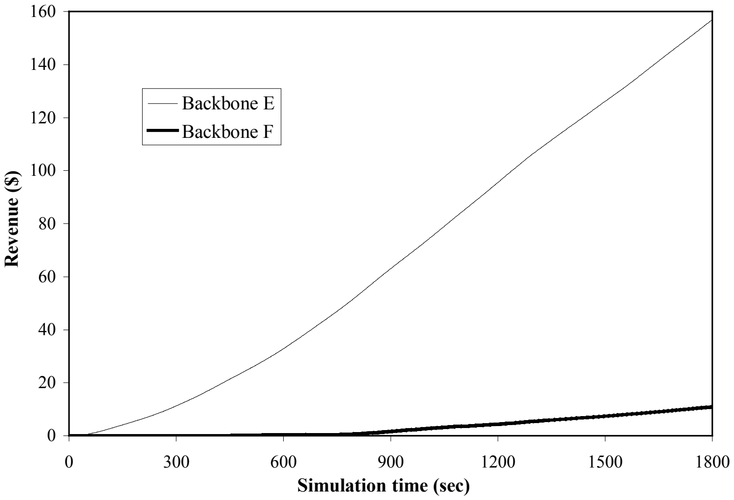

Fig. 11 - Total revenue of backbone domains: lightly loaded case.

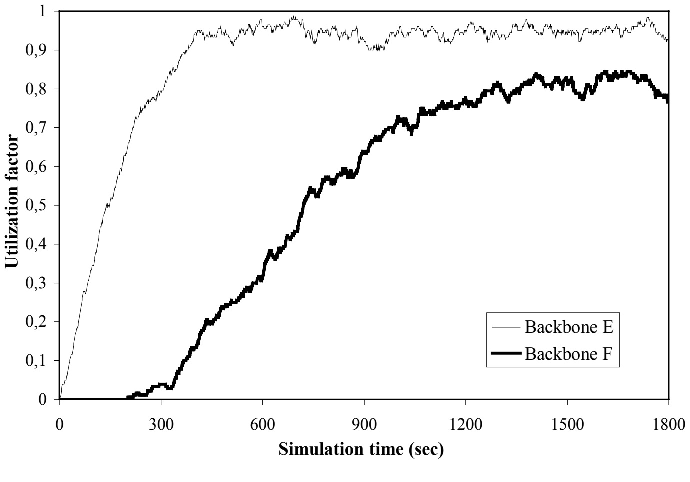

Fig. 12 - Utilization factor of backbone domains: heavily loaded case.

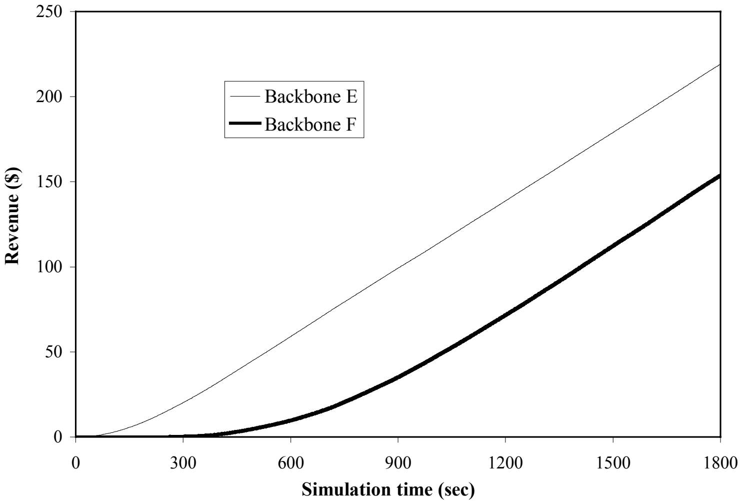

Fig. 13 - Total revenue of backbone domains: heavily loaded case.

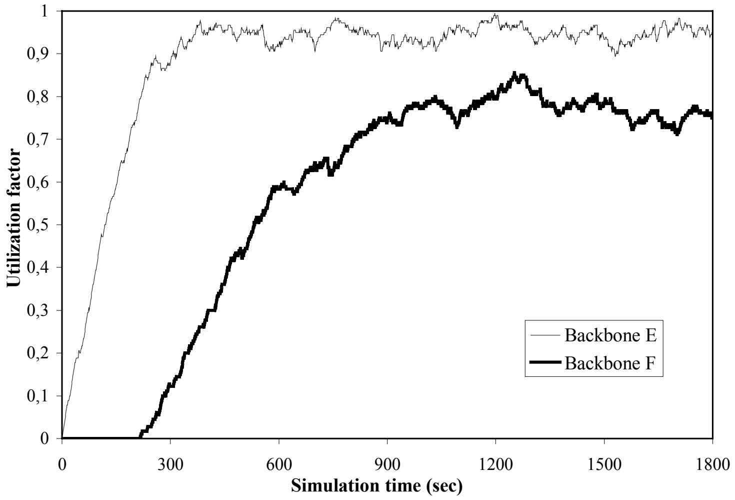

Fig. 14 - Utilization factor of backbone domains: very heavily loaded case.

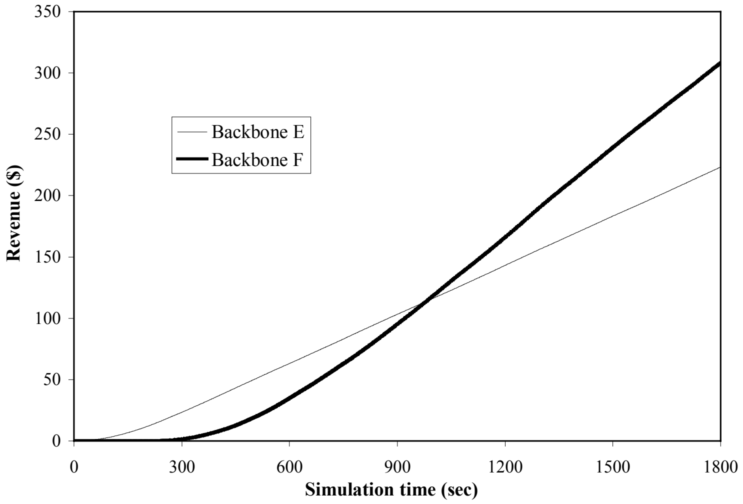

Fig. 15 - Total revenue of backbone domains: very heavily loaded case.

- Case study 2

We start considering that the per-second tariff charged by domain E is half the one of domain F, i.e., $f_E(d) = 2 - (1/14)d$ and $f_F(d) = 4 - (1/7)d$ (the other pricing parameters are chosen like in the previous case study). In addition, we consider that access domains apply the same tariff of domain E. Therefore, the tariff of domain A, B, C, D, and E varies in the range [0.044, 0.088] cent/s, whereas the tariff charged by domain F ranges in the interval [0.088, 0.176] cent/s. We assume that each network user provides the NRB with the demand curve shown in Fig. 16. We have identified three virtual delay intervals ([0, 4], [5, 9], and [10, 14]) associated with excellent, good and acceptable services, respectively. The willingness of users to pay is assumed to be: 0.0026 cent/s for an excellent service, 0.00246 cent/s for a good service, and 0.022 cent/s for an acceptable service. This means that users accept each level of service in the range [0, 14]; note that in this case $d_{\max}$=14.

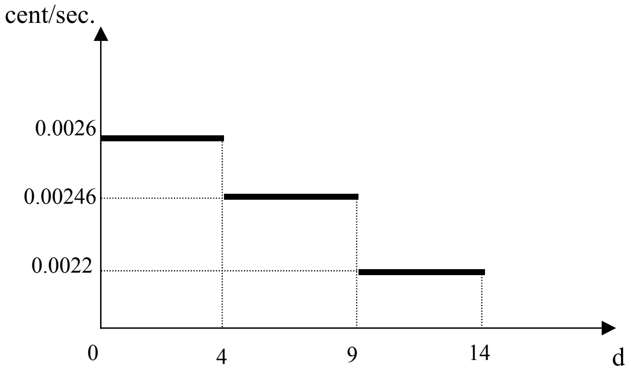

Fig. 16 - Demand curve.

We run the MB inter-domain routing algorithm under the hypothesis of very heavy load (i.e., $\lambda_{VH} = 0.8$ and $\tau = 0.002\,\mathrm{s}^{-1}$). The cost functions of paths through domain E and domain F are shown in Fig. 17. It is easy to verify that the cost function of the paths through the expensive backbone domain is higher than the users' demand curve. In other words, since domain F charges excessively high tariffs, then it is out of business. This would imply that it does not attract customers (see Fig. 18 and Fig. 19). It is worth noting that, while in the MP inter-domain routing algorithm the value of the virtual delay associated with the calls is always $d_{\max}$ (the cheapest one), the MB inter-domain routing algorithm selects the value of d that maximizes the customer surplus function. Since we have considered the demand curve as a simple step function, this value is the largest of the three intervals of virtual delay associated with a qualitative QoS level. In this case, it is equal to 9 for the paths through the cheap backbone domain; it corresponds to a surplus value of 0.00011 $/s.

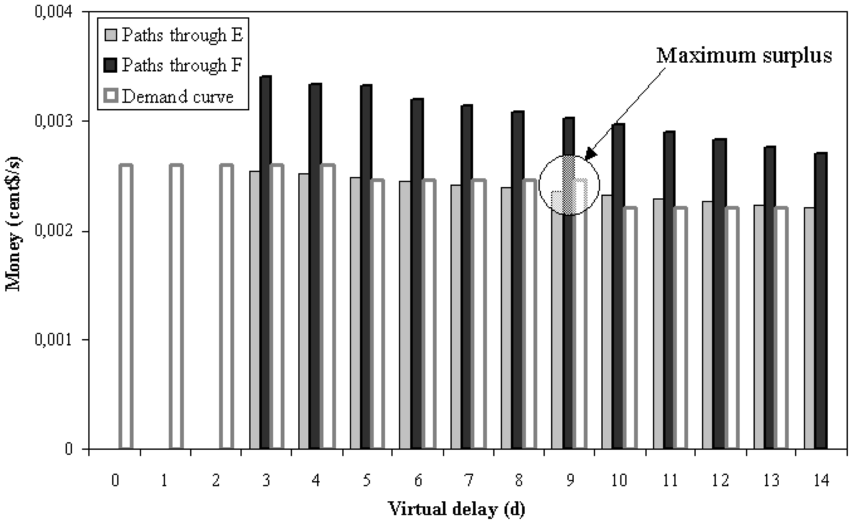

Fig. 17 - Cost functions of the paths through backbones E and F.

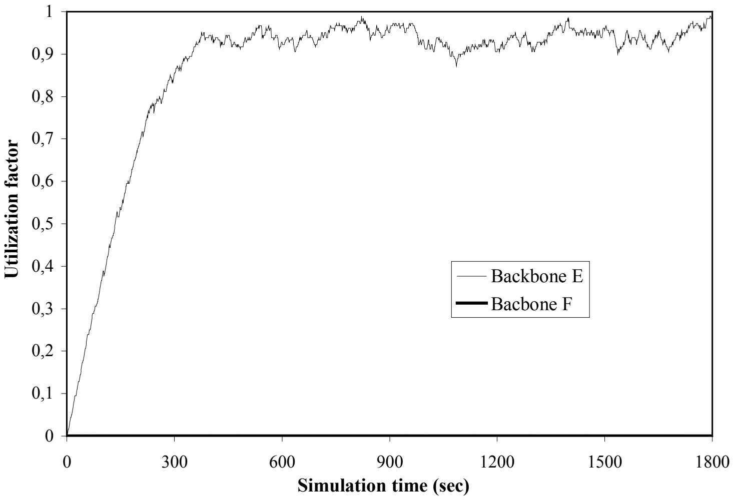

Fig. 18 - Utilization factor of backbone domains: very heavily loaded case and MB inter-domain routing algorithm.

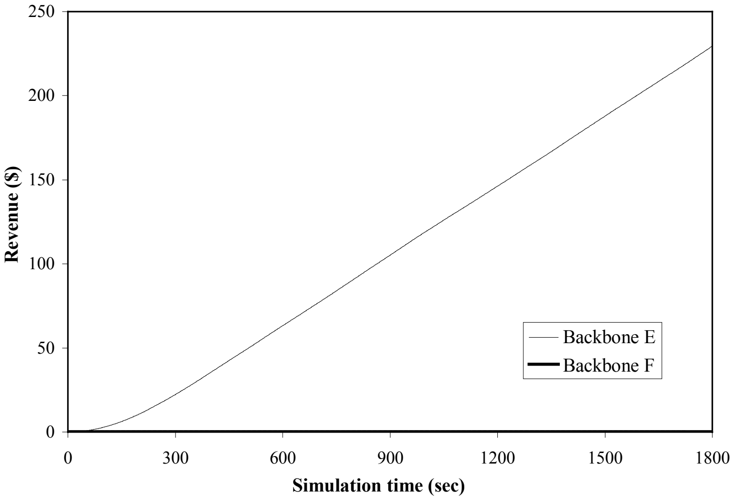

Fig. 19 - Total revenue of backbone domains: very heavily loaded case and MB inter-domain routing algorithm.

Since we are considering the very heavily loaded case, if we lower the tariff charged by domain F, it is possible to verify that domain F collects traffic and money when the network resources of the cheapest backbone are fully allocated. Let us consider $f_F(d) = 2.22 - (1.11/14)d$. The cost function of the paths through domain F is shown in Fig. 20. The calls obtain a QoS level corresponding to a value of *d*=9 until resources of domain E are available. Afterwards, they will obtain a QoS level corresponding to the value of *d* that maximizes the customer surplus function for the paths through domain F. Such a value remains equal to 9, also if the value of surplus decreases (0.00004 $/s). The utilization factor and the total revenue of backbone domains are shown in Fig. 21 and Fig. 22, respectively.

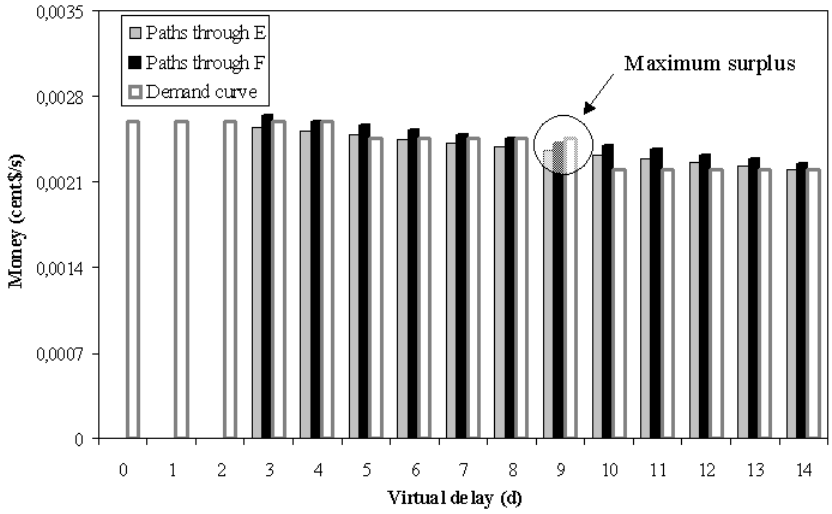

Fig. 20 - Cost function of the paths through domains E and F.

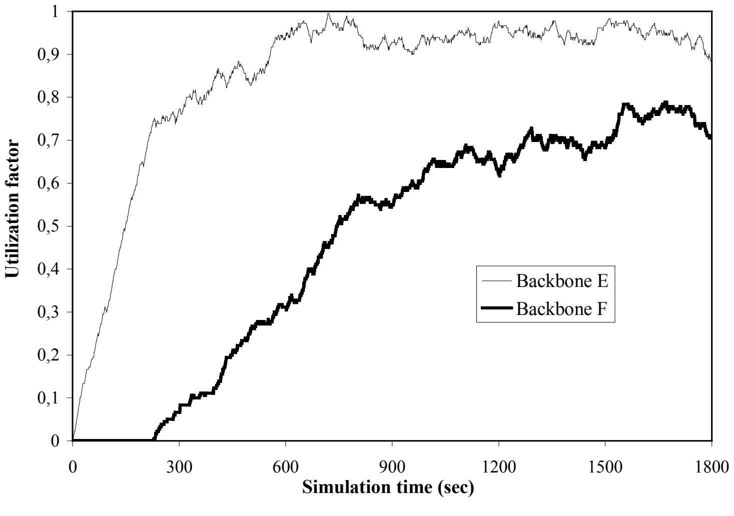

Fig. 21 - Utilization factor of backbone domains: very heavily loaded case and MB inter-domain routing algorithm.

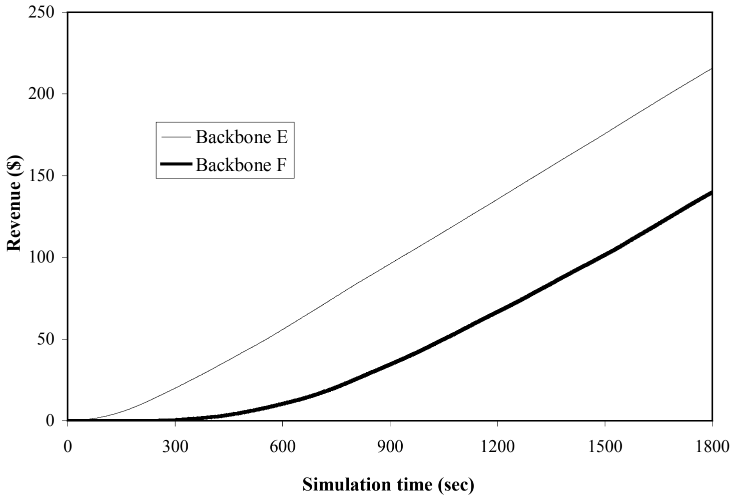

Fig. 22 - Total revenue of backbone domains: very heavily loaded case and MB inter-domain routing algorithm.

- *Case study 3*

So far, the price variation factor γ was set to 1, hence the resulting tariffs charged by domains are not time depending. In the following we assume that the price variation factor depends on the links congestion.

We assume that the network is very heavily loaded ($\lambda_{VH} = 0.8$ and $\tau = 0.002\,\text{s}^{-1}$). The function $f(d) = 4 - (1/7)d$, the price parameters $\alpha = \frac{11}{16}10^{-8}$ \$/commodity-unit and $B_{res} = 64000$ bit/s are common to all domains. As regards the price variation factor, γ, we consider that in backbone domains E and F it is equal to 1 until the busy bandwidth is lower than, or equal to, a given value; then the function increases linearly, as shown in Fig. 23. It is worth to note that the congestion pricing of domain F is more marked than in backbone E, since its γ value starts increasing from an allocated bandwidth equal to 40% (50% for domain E) and since the increasing rate is higher than the one of domain E.

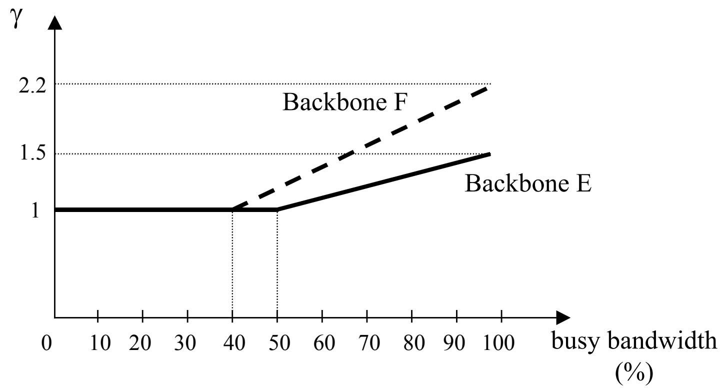

Fig. 23 - Price variation factors for backbone domains E and F.

The demand curve is assumed to be the same as in case study 3 (see Fig. 16). It is easy to verify that the tariff charged by domain E is out of business when the amount of available bandwidth is lower than 40%, and the tariff charged by domain F is out of business when the amount of available bandwidth is lower than 60%. The simulation confirms this expectation. In fact, the utilization factors of backbone E and F are always lower than 0.6 and 0.4, respectively (see Fig. 24). This clearly results in increased revenues for the cheapest backbone (Fig. 25).

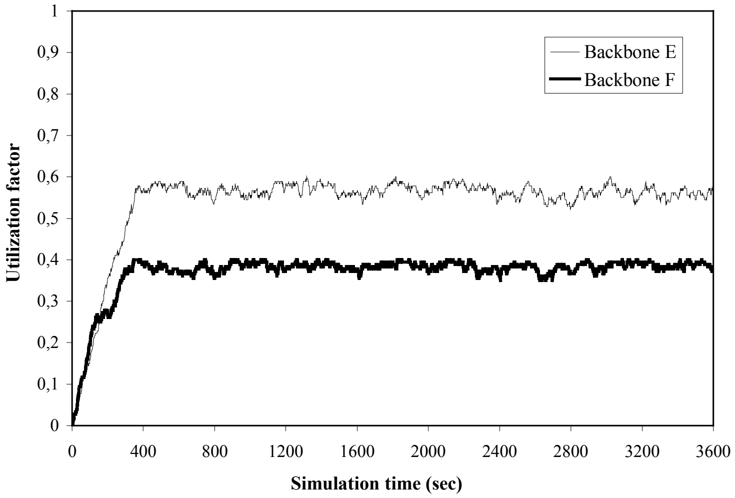

Fig. 24 - Utilization factor of backbone domains: MB case.

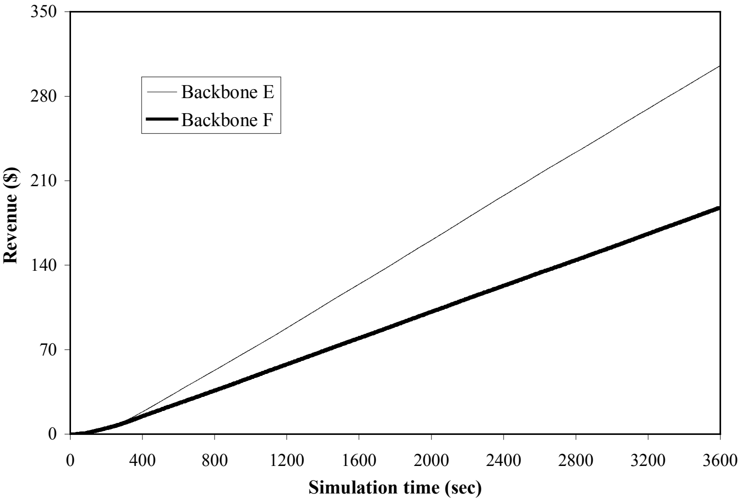

Fig. 25 - Total revenue of backbone domains: MB case.

6. End-to-End Service Establishment

In this section, we present the typical sequence of interactions that enable end-users to receive services from ASPs. The description refers to the analysis made so far. It represents a final summary of the architectures and procedures presented in this work.

The procedure evolves according to the following steps (refer to Fig. 26):

1. the end-user obtains physical connectivity from an NSP and IP connectivity from an ASP;
2. the end-user, if necessary with the help of an IB, identifies an ASP that can provide a specific service. This ASP can be either the same offering IP connectivity or another one;
3. the end-user (or the IB on his behalf) issues a request for the application service to the ASP and describes qualitatively the desired quality (e.g., acceptable, good, excellent) and his willingness to pay for it, and provides other information such as, for instance, the type of terminal used and the speed of the access link;
4. the ASP translates the customer's qualitative requirements into quantitative technical parameters, that is:
   - traffic descriptors, i.e., DLB parameters (r_S, P_S, B_{TS}) and, if necessary other information on the flow statistic (i.e., the functions $\lambda(B_{res})$ and $\theta(B_{res})$);
   - QoS constraints in terms of transfer delay, delay jitter and packet loss probability, i.e., $(\overline{D}_{transfer}, \overline{D}_{jitter}, \overline{P}_L)$, typical of the application service;
   - demand curve $U(d)$.

   This information set is sent to the NRB, with the request of finding the "best" path to deliver the application service to the end-user;
5. the NRB, which knows the topology of the network, identifies the domains which could potentially be involved in the network service and checks their offers. An administrative domain may send its own service offer on request, or may refresh them periodically according to its own policies. For each input-output port pair and the application considered, the offer is represented by the e-Table (see TABLE 1). The offer includes: (i) the range of QoS that the domain is able to guarantee ($[d_{\min}, d_{\max}]$); (ii) the available bandwidth for each service, B. According to this value, the NRB can decide if the flows of the requested application service can be allocated. For instance it verifies if the peak rate (or the average rate or a value of effective bandwidth) of the flow is lower than the available bandwidth; (iii) the value of the price parameters by which it is possible to determine the tariff applied by the domain for the specific flow. In particular, the per

volume-unit and per time-unit charges are respectively $a_V = \beta f(d)$ and $a_T = a_V B_{res}$, where $B_{res} = xP_S$ ($0 \le x \le 1$).

6. the NRB runs an inter-domain routing algorithm in order to find a path satisfying the QoS requirements and maximizing the customer surplus function;
7. the NRB communicates to the ASP the path found, the relevant value of end-to-end QoS, the tariffs charged, and the edge-to-edge QoS level relevant to the domains involved in the information transmission;
8. the ASP can offer the service with the relevant price to the customer. If he accepts, the agreement with the characteristics of the service requested and its price is drawn up. The price charged by the ASP depends on both application and network services and must be communicated to customers before the beginning of the transmission. The ASP may choose to charge the complete service according to the following options: (i) a tariff depending on the time length of the connection (inclusive of the costs of the network and application services); (ii) a fixed charge (inclusive of the costs of the network and application services). This option is clearly suitable for pre-determined time duration services; (iii) a tariff depending on the amount of traffic volume exchanged during the connection (inclusive of the costs of the network and application services); (iv) a per-time (or per-volume) tariff covering the cost of the network service plus a fixed charge for the application service.

   It is our opinion that, for their simplicity and clarity, the first two options represent the best choice for undetermined and pre-determined time duration services, respectively.
9. the ASP has to draw up one network service contract with each NSP involved in the end-to-end information transfer. A contract between a domain and an ASP for providing a guaranteed service must include, among other things, the following information: (i) the definition of the input port-output port pair; (ii) the definition of the so-called *compliant* traffic (i.e., DLB parameters); (iii) the definition of a rule for handling the *non-compliant* traffic; (iv) the edge-to-edge QoS performance of the compliant traffic; (v) the definition of the price of the service; (vi) the definition of the source address and destination address for accounting purposes.
10. each domain involved starts its own internal procedures to guarantee the promised level of QoS;

11. the application may start, and each domain involved implements its own accounting procedures to get information about the network resource consumption. The accounting record is used to charge the ASP, which, in turn, implement its accounting procedures to charge the end-user that enjoys the service.

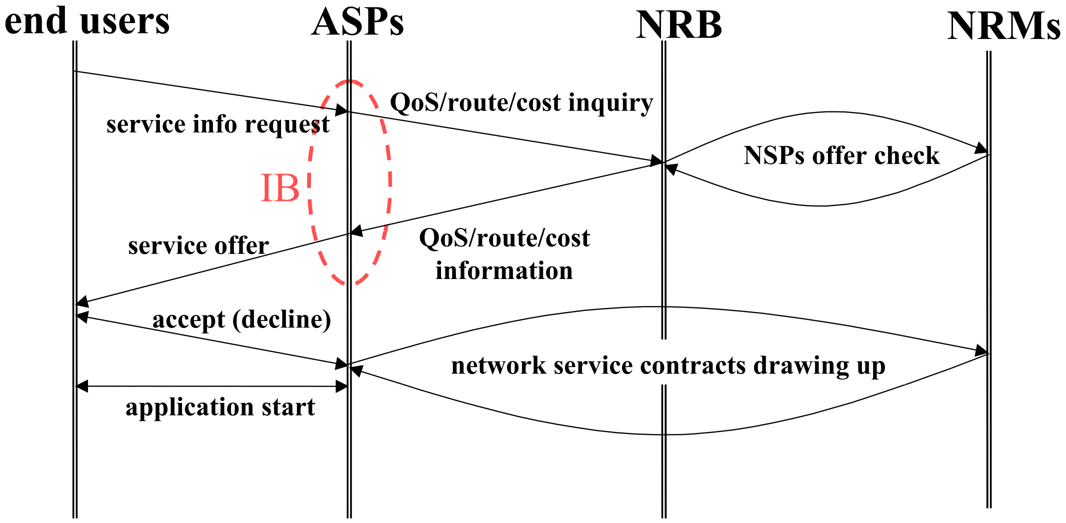

Fig. 26 - Temporal sequence of interactions.

TABLE 1 - E-TABLE

QoS class	**available bandwidth**	**price parameters**	
$[d_{\min}, d_{\max}]$	B	$x(d)$	$\beta, f(d)$

7. Conclusion

In this work, we have illustrated the set of rules that define a new information-trading paradigm, which can operate in parallel with legacy telecommunication systems. Operators, infrastructure owners, and service and content providers may freely decide to observe these rules, which could open new market space and attract customers. The business model is characterized by a clear separation of the communication infrastructure from the information content. We have defined interactions among service providers and customers, highlighting the importance of the brokerage activity, especially as regards the management of network resources, performed by the NRB. We have defined step by step the procedure to receive from an ASP an application service that needs a guaranteed network support to be satisfactorily delivered to customers. The key points of the approach are: (i) all technical aspects are completely transparent to end-users, and (ii) the business model would be actually customer-centered, in the sense that its goal is to improve service differentiation and offer, so enabling a barrier-reduced market.

Then, we have defined the concept of network commodity traded on the marketplace. The amount of the used commodity is directly related to the QoS parameters associated with the transmission of information units. We have shown how the number of commodity may be determined by using a function $f(d)$ of a specific service descriptor d, called *virtual delay*, summarizing the relevant QoS level. Such a function is used to define a per-volume and/or per-time tariff charging strategy for guaranteed network services, depending on both the actually used and the reserved network resources according to tunable weights.

Then, we have pointed out the role of the NRB and we have presented a novel inter-domain routing algorithm to maximize users' benefit in terms of perceived QoS and price

Finally, we have presented simulation results in some interesting case studies, which highlight the importance of the users' demand curve in a competitive market scenario.

Acknowledgements

The authors would like to thank the anonymous referees for their suggestions and comments that have improved the quality of the paper.

This work is co-supported by the European Union in the framework of the IST program (projects WHYLESS.COM [26] and SIMPLICITY [27]).

References

[1] G. Huston: “Next steps for the IP QoS architecture” IETF RFC 2990, November 2000

[2] Y. Bernet et al.: “A framework for Integrated Services operation over DiffServ networks”, IETF RFC 2998, November 2000.

[3] J. Altmann, K. Chu: “A proposal for a flexible service plan that is attractive to users and Internet service providers”, *IEEE INFOCOM 2001*, Anchorage, USA, April 2001.

[4] D. Di Sorte, M. Femminella, G. Reali, S. Zeisberg: “Network service provisioning in UWB open mobile access networks”, *IEEE Journal on Selected Areas in Communications (IEEE JSAC)*, 20(9), pp. 1745-1753, December 2002.

[5] D. Di Sorte, G. Reali: “Minimum price inter-domain routing algorithm”, *IEEE Communications Letters*, 6(4), April 2002.

[6] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss: “An architecture for Differentiated Services”, IETF RFC 2475, December 1998.

[7] R. Braden, D. Clark, S. Shenker: “Integrated services in the Internet architecture: an overview”, IETF RFC 1633, June 1994.

[8] A. Pras, B.J. van Beijnum, R. Sprenkels, R. Parhonyi: “Internet accounting”, *IEEE Communications Magazine*, May 2001.

[9] M. Falkner, M. Devetsikiotis, I. Lambadaris: “An overview of pricing concepts for broadband IP networks”, *IEEE Communications Surveys*, Second Quarter 2000.

[10] L.A. DaSilva: “Pricing for QoS-enabled networks: a survey”, *IEEE Communications Surveys*, Second Quarter 2000.

[11] N. Blefari-Melazzi, D. Di Sorte, G. Reali: “Accounting and pricing: a forecast of the scenario of the next generation Internet”, *Computer Communications*, Elsevier Science, 26(18), pp. 2037-2051, December 2003.

[12] X. Wang, H. Schulzrinne: “RNAP: a resource negotiation and pricing protocol”, *NOSSDAV 1999*, Basking Ridge, USA, June 1998.

[13] S. Shenker, D. Clark, D. Estrin, S. Herzog: “Pricing in computer networks: reshaping the research agenda”, *ACM Computer Communication Review*, 26(2), pp. 19-43, 1996.

[14] A.P. Markopoulou, F.A. Tobagi, M.J. Karam: "Assessment of VoIP quality over Internet backbones", *IEEE INFOCOM 2002*, New York, USA, June 2002.

[15] D. Di Sorte, M. Femminella, G. Reali: "A QoS index for IP services to effectively support usage-based charging”, to appear in IEEE Communications Letters.

[16] A.Elwalid, D. Mitra: “Traffic shaping at a network node: theory, optimum design, admission control”, *IEEE INFOCOM 1997*, pp. 445-455, Kobe, Japan, April 1997.

[17] G. Fankhauser, D. Schweikert, B. Plattner: "Service level agreement trading for the Differentiated Services architecture", Swiss Federal Institute of Technology, Computer Engineering and Networks Lab, Technical Report No. 59. November 1999.

[18] J. Hwang, J. Altman, H. Oliver, A. Suarez: “Enabling dynamic market-managed QoS interconnection in the next generation Internet by a modified BGP mechanism”, *IEEE/ICC2002*, New York, USA, April-May 2002.

[19] F. Baccelli, G. Cohen, G.J. Olsder, J.P. Quadrat: *Synchronization and linearity, an algebra for discrete event systems.* New York, Wiley, 1992.

[20] C.-S. Chang: “Deterministic traffic specification via projections under the min-plus algebra”, *IEEE INFOCOM 1999*, New York, USA, March 1999.

[21] J.Y. Le Boudec: “Application of Network Calculus to guaranteed service networks”, *IEEE Transactions on Information Theory*, 44(3), pp. 1087-1096, May 1998.

[22] K. Nichols, B. Carpenter: “Definition of Differentiated Services per-domain behaviors and rules for their specification,” IETF RFC 3086, April 2001.

[23] H.R. Varian: "Economics of Information Technology", Technical Report, March 2003, <http://www.sims.berkeley.edu/~hal/people/hal/papers.html>.

[24] H. Gruber, F. Verboven, “The evolution of markets under entry and standards regulation - The case of global mobile telecommunications”, *International Journal of Industrial Organization*, (19), p.1189-1212, 2001.

[25] Internet Economics, Edited by Lee W. McKnight and Joseph P. Bailey, Massachusetts Institute of Technology Press, 1998.

[26] Project IST-2000-25197 "WHYLESS.COM", <http:\\www.whyless.org>.

[27] Project IST-2004-507558 "SIMPLICITY", <http:\\www.ist-simplicity.org>.