

A Taxonomy and Design Considerations for Internet Accounting

Michel Kouadio*
mkouadio@acm.org
(* Corresponding Author)

Udo Pooch
pooch@cs.tamu.edu

Dept. of Computer Science, Texas A&M University,
501B Harvey Bright Bldg., College Station TX. 77843-3112.

ABSTRACT

Economic principles are increasingly being suggested for addressing some complex issues related to distributed resource allocation for QoS (Quality of Service) enhancement. Many proposals have been put forth, including various strategies from Pricing Theory and market-based insights for security in the Internet. A central need of these endeavors lies in the design of efficient accounting architecture for collecting, storing, processing and communicating relevant technical information to parties involved in transactions. This paper proceeds to a systematic classification of the major characteristics of accounting models with the aim of highlighting important features to optimize for building effective architectures.

Categories and Subject Descriptors:

C.2. Computer-Communication Networks. C.2.3. Network Operations: Network Management, Network Monitoring. C.2.4. Performance of Systems: Design Studies, Measurement Techniques. C.2.5. Local and Wide-Area Networks: Internet.

General Terms:

Measurement, Economics, Management.

Keywords:

Accounting Architecture, Pricing, Classification.

1. INTRODUCTION

The development of protocols for resource request and allocation is central to QoS provisioning efforts on the Internet. Two major architectures and protocols have been proposed to allow users to negotiate the availability and reservation of resources prior to initiating transactions.

The first was the Integrated Services Networks (IntServ) with the Reservation Setup Protocol (RSVP) [3], [36], and [43]. The

second, the Differentiated Services Networks (DiffServ), requires that users negotiate a Service Level Agreement (SLA) that provides upfront their traffic expected profile and terms of service provision with the network [2], [29]. In DiffServ, SLA is negotiated through a Bandwidth Broker (BB), which is an entity in a domain that keeps track of all resource allocated, and accepts or rejects new requests accordingly. BB is also designed to cooperate with its peers in other domains to confirm resource availability from source to destination.

In recent years, economic models of network management have emerged as complements to these architectures, or alternative frameworks purporting to contribute flexible solutions to distributed resource allocation issues and enhance QoS. Network economic approaches are fundamentally frameworks based on Game Theory and Pricing Theory that require resource owner (producer) to set "rules of the game" and leave to each user (consumer) the selection of an algorithmic behavior that optimizes his utility [9], [16], and [44]. Pricing studies generally posit that putting a price tag on network services will moderate and lead to self-regulation in user requests, thus ultimately resulting in the equilibrium of supply and demand of resources.

Application of economic principles to network operations introduces specific challenges encompassing modeling, computational complexities and user-network interactions. How to operate a network by auctioning resources? What signaling protocol may allow users and networks to exchange information on resource supply and demand to sustain tight real-time QoS constraints? Efficient accounting mechanisms should therefore play an important role in this picture by appropriately providing the necessary technical information to assess resource availability and support different charging policies. We define an accounting architecture as the set of components for monitoring resource availability, and for collecting, storing, processing and communicating various parameter data that define user QoS profile and consumption of network resources.

This paper investigates issues in the design of accounting architecture that may be used to support economic models of network management. In Section 2, we highlight the relationship between accounting and pricing. Section 3 provides a review of some of the main accounting studies found in the literature. Next, in Section 4, we propose six main criteria around which to organize and understand accounting models. Then we go on from Sections 5 through 10, to explore the implications of each of these characteristics from an overall performance and efficiency perspective. We wrap up our classification in Section 11, by stressing important design considerations and elements to optimize for building an effective accounting architecture.

2. ACCOUNTING AND PRICING

A great deal of research in economic models for the Internet has concentrated on pricing policies and strategies. Tutorials on pricing and the economics of networks may be found in [10] and [19]. General surveys of pricing models are conducted in [11], and recent trends in charging studies for Integrated Internet Services are reviewed in [40]. While pricing studies deal mainly with policy and charging strategies, they are closely dependent on accounting architectures and protocols for collecting, processing and communicating the supporting technical data. Hence, accounting architecture appears to be an enabling technology for economic-based resource allocation and management in networking.

Figure 1 presents the relationships between pricing model, accounting architecture and the network. As an example, pricing strategies such as auction and dynamic pricing require that the underlying accounting monitor the availability of resource and communicate relevant information to parties involved in the transaction (e.g. network and user agents). Pricing policies and accounting mechanisms have reciprocal influences. Accounting architecture and protocols determine the scope of feasible pricing models, and pricing models set out features and requirements for accounting. IETF (the Internet Engineering Task Force) has created a working group in Authentication, Authorization and Accounting (AAA) to investigate some of the issues involved [1], [5], [6], [31].

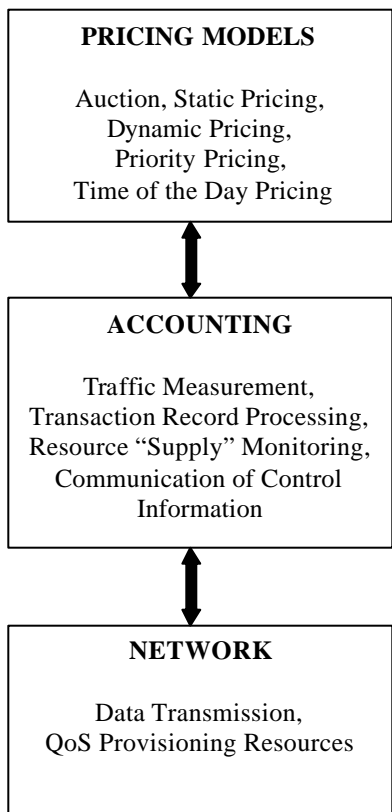


Figure 1: Layered Presentation of Pricing, Accounting and the Network.

3. REVIEW OF ACCOUNTING STUDIES

This review includes explicitly designed accounting models as well as some “implied” architectures put forth as part of pricing studies. Indeed, there is a natural interplay between pricing and accounting; and these fields being relatively new, their terminology is still evolving and is not standard in all studies.

3.1. Accounting for Optimal Pricing Models

Optimal pricing models aim to achieve maximal efficiency in the distribution of resources and maximal welfare in the use of network resources. They form the bulk of recent research in pricing and most are based on Game Theory [42]. The theoretical economical foundations of optimal pricing models is explained in [27] and [39]. Optimality refers to the perfect equilibrium of the supply of network resource and the demand by users, due to the regulating effect of the pricing strategy. Basically, the network is modeled as a supplier of QoS provisioning resources (bandwidth, buffer etc.), and users as consumers, are represented by agents, each endowed with a budget. The repeated interaction between the network advertising its resources (lightly loaded, overloaded etc.), and competing users adapting their strategies to optimize their QoS allocations is bound to reach an optimal point of convergence where each user is most satisfied with his gain and has no incentive to deviate for another QoS profile [16] and [24]. Example proposals are introduced in [9], [25], [30] and [44]. Typical strategies to reach optimality include dynamic pricing of congestion and auction [26], [28] and [34].

There is a diverse body of accounting mechanisms intertwined with these pricing models, to provide the architectural support. The architectures are generally not explicitly defined in these studies, but rather implied by the charging strategies. They share some common characteristics that follow from their features:

1. **Resource monitoring:** each network element (e.g. router, switch, and server) is endowed with mechanisms to collect data on aggregate resource consumption and availability (bandwidth, buffer, CPU), and sets conditions for release to users accordingly.
2. **User-network signaling:** there are typically two options that include interactive signaling like those used for auction models, or a one-way dynamic price advertising from the network. In the first option, a communication protocol allows network elements to advertise resource availability and current clearance conditions. Users submit resource requests and possibly their utility (valuation of transaction). The interaction continues until supply equals demands. The second option involves only the network setting clearance price and simply feeding it back to users who then may choose to adapt their QoS parameters accordingly. That was the option used in the network game simulator built at Microsoft Research and Cambridge [18] and [20]. The signaling load is often an issue with these models because of the frequent exchanges between users and networks, especially for auction-like strategies.
3. **Data collection:** typically data sent by users is measured in each network node and rated against the clearance price, based on the local resource availability conditions. At issue is the scalability of these mechanisms and their real-time performance when multiple nodes are involved from source to destination.

4. **Network models:** accounting architectures designed for these models may in principle work with most of QoS-enhanced protocols. Furthermore, they may be seen as alternatives to IntServ (RSVP) and DiffServ models of resource allocations. Reservation and classes of service are superseded by dynamic user appropriation of resources based on price level; therefore each application may be viewed as having their own class of service and continuously reserving the resources that they need. However, the communication model supported is unicast and does not extend to multicast.

3.2. Edge Pricing Architecture

In an attempt to deal with the practical scalability and flexibility issues of the accounting structures of optimal pricing models, Shenker suggested an architecture paradigm aimed at making accounting local to each network provider [35]. It is of interest to note that the induced architecture of optimal pricing models may require that uniform accounting mechanisms be implemented in different network domains. Such close inter-dependencies limits the flexibility that each network provider would have in the selection of different accounting schemes and charging models.

The main idea of the Edge Pricing architecture is to advocate that all accounting mechanisms be set up at the edge of each provider network, and appropriate cost approximations be performed to support different pricing models. Such approximations include the expected congestion conditions and traffic path from source to destination. That architecture fosters interconnection agreements between network providers, and has the advantage of breaking down the accounting complexity into local issues. The design approach is intended to provide flexibility in the scope of charging models that may be supported by each network domain.

Still, the Edge Pricing architecture is more a design framework than a detailed accounting model. It does not address either detailed data collection method, or data communication between parties. It left as open issues how to support models that charge receivers and multicast, which are not local issues to each provider.

3.3. Zone-based Cost Sharing Model

Clark is one of the first to consider cost sharing between sender and receiver [8]. The proposed accounting architecture recommends that the Internet be subdivided into zones where service would be provided at a uniform, distance-insensitive way. Zones might be site, city, region or country. Both sender and receiver would exchange signaling information on the zones for which they are willing to pay. That information is to be put into each packet header by the sending source when transmission starts. Each router is to be configured to examine the packet header to determine if any of the sender or the receiver is willing to pay for QoS resources in the zone where they are located, and to perform the relevant traffic metering. Therefore the accounting involves willingness-to-pay signaling messages, payment aware packet header and traffic measurement in routers along the path from source to destination.

In the case of multicast, in order to limit the load of signaling messages from receivers to the source, routers in the zone where the sender has expressed interest in paying, provide enhanced services and simply forward the packets. Then, as packets leave the sender-paying zone, routers in subsequent zones would mark packets to explicitly request from receivers whether they are willing to pay for enhanced services. Each receiver would be equipped with an agent "Receiver Traffic Meter" that makes decision on packet payment for zones of interest and would respond to such demands.

The significance of this architecture is best understood in relation with some popular activities on the Internet, such as web browsing, or emerging multimedia applications, such as videoconferencing and video-on-demand. As a matter of fact, in these cases, it is natural that pricing models target receivers of transactions, instead of charging the sender only. It is then important that accounting architectures provide mechanisms to support such cost sharing.

3.4. Internet Demand Experiment (INDEX)

Developed at UC Berkeley, INDEX is a field experiment that provides an architecture for charging and billing users for Internet traffic [12], [32]. The primary goal was to study user reaction to different QoS and usage sensitive pricing, as opposed to the currently flat pricing system used in the Internet. That study proposes exact data measurement over sampling for charging purposes to enhance credibility of collected data to users. The system is made up of many components. Users are equipped with a security agent that authenticates operations with the network and a purchasing agent that makes payment decisions. At the network entry point sits a billing gateway, a dedicated device such as a router that provides transaction metering and consolidates accounting records. An access controller coordinates user authentication procedures.

That architecture can support different static charging models, and provides a way to select which of the sender or the receiver to account for by explicitly asking confirmation from each end-user involved before allowing a transaction establishment. It deals with unicast traffic and does not support multicast. It is designed to work in the traditional Internet and does not extend to IntServ or DiffServ.

3.5. Generic Charging and Accounting (GenCA)

Departing from accounting schemes that are designed to support a few specific pricing models, Hartanto and Carle proposed a policy-based configurable accounting architecture for the European Information Infrastructure and Services Pilot, also known as SUSIE [21]. Their study developed a general architecture and billing system framework made of the following layers and functions:

- *Metering Layer:* tracks and records usage of resources.
- *Collecting Layer:* accesses data provided by meters and forwards them to accounting entities.
- *Accounting Layer:* consolidates the collected data into accounting records.
- *Charging Layer:* assesses costs of usage from the accounting records.
- *Billing Layer:* collects all charges for a customer over a time period and creates a bill.

In that framework, each new transaction by a user triggers the retrieval of his profile from a policy database to collect accounting and charging related parameters. These parameters are sent to metering entities to collect data relevant to the charging formula from the contract between the user and the provider. Typical parameters include volume, bandwidth, priority and duration. The data collected is sent to accounting consolidation entities to form records that later serve for charging.

GenCA accounting model is flexible and adaptable. Each user may have a custom-charging and service level agreement contract with a provider, and the accounting system is configured to measure, record and process traffic parameter data for the QoS resources in that contract. The system can support Integrated Services and Differentiated Services Networks. Provisions are made to support both unicast and multicast. That accounting model is bound for implementation in some European countries.

3.6. CATI (Charging and Accounting Technologies for the Internet)

CATI is a project developed in Switzerland as part of a charging program on Electronic Commerce transactions initiated by the Swiss National Science Foundation [17]. The accounting scheme encompasses many network protocols including Integrated and Differentiated Services Networks, ATM and MPLS. CATI accounting model uses pre-defined user profiles that contains default information on:

- Sender identity
- Default IP provider
- Payment method
- Invoice/billing information
- Default QoS
- Default cost-sharing scheme

Users have the opportunity to change some of the default parameters in their configuration file before initiating a transaction.

Two components are used in a network provider domain: Internal Service Broker (ISB) that manages service local to the provider and External Service Broker (ESB) that negotiates service with neighboring domains. Service Level Agreements (SLA) between a provider and his internal users or his peer providers are negotiated in form of volume of traffic, duration and price before the start of a session and can be dynamically updated.

Accounting signaling is performed through an extended RSVP protocol. The extension includes new objects on pricing scheme, identity of the payer of the transaction (e.g. sender or receiver), bidding fields where auction is applicable, fields for the type of service requested, and network provider information. The new RSVP objects extensions allow for example the implementation of auction pricing models by the use of bidding fields. That architecture was also used to implement dynamic pricing models between network providers, or Service Level Agreement Trading (SLAT), for an automatic negotiation of traffic volume for interconnection agreement.

3.7. Accounting and Charging Multicast

Herzog suggested a distributed architecture called “One Pass Accounting”, consisting of an accounting message that flows from the source through the multicast tree [22]. As the message traverses the tree, multicast nodes allocate costs to connected members. The accounting scheme of this approach requires the configuration of each network node in the multicast tree for running the costing algorithm. It is designed to support charging models based on link costs, and therefore this study appears to be more relevant to integrated network environments with the RSVP protocol that works with pre-determined paths. This architecture may not easily be extended to support other network models, such as DiffServ, where traffic path may vary. Overall, a link-by-link accounting mechanism may appear limited in flexibility and scalability, due to the complexity of node configuration.

Another multicast accounting architecture was proposed as an extension of the Generic Accounting Architecture [7]. This model sets up metering devices in edge nodes and all multicast routers of each domain. For a given multicast session, data collected in these meters is fed back to the traffic sender provider domain where the accounting server proceeds to a consolidation and compute the allocation share of each participant.

In both studies, the accounting architectures are designed to work across domains, on the whole multicast tree, which require somehow uniform mechanisms and inter-domain standards.

4. CHARACTERIZING ARCHITECTURES FOR INTERNET ECONOMICS

Accounting models for network economics differ by many characteristics. The design principles, parameters handled and components involved are crucial to the overall efficiency of each model. A taxonomy provides a general picture that helps highlight main model components, patterns and possible points that deserve more attention and efforts.

Accounting models may have many characteristics and this taxonomy, although broad is not meant to be exhaustive. Remember that accounting characteristics may be as broad as the set of architectural requirements of all conceivable pricing models. Still, we believe that the most recurrent characteristics can be classified in six major concepts:

- 1. Policy scope:** describes the different charging policies and economic principles that may be directly supported, or with minor adaptation, by the architecture. That often results from the design approach.
- 2. Source of data collection:** specifies where in the network data is collected to feed the accounting system.
- 3. Methods of data collection:** is a concept for understanding the granularity of data collection.
- 4. Collected elements:** refers to the different kind of data measured and provided by the accounting model; this provides insights in the kind of charging and pricing policies that may be supported.
- 5. Interactivity:** determines whether the accounting model has any provision to allow users to submit various parameters such as their valuation of utility and prices, or if all rules and parameters are set by network providers.
- 6. Supported network protocols:** indicates the scope of enhanced network service protocols and thus environments where the accounting system may operate.

Figure 2 synthesizes the classification of these characteristics, while Sections 5 through 10, provide further explanations.

5. POLICY

The scope of charging policies supported by an accounting architecture is typically determined by its design principles. Accounting models may thus be subdivided into two major groups: those that are driven by charging and pricing policies, and those that start from architectural grounds.

5.1. Close-Policy Accounting Models

Most of the works reviewed in the Section 3, start with investigating some specific pricing policy to eventually reach an accounting model that may support the proposed model. This may be considered a top-down approach or accounting driven by pricing-policy. Specific economic principles or pricing policy permeate the architecture and therefore the scope of policies, which it may directly support, is limited to those it embodies. Such accounting systems can readily support the pricing model and policy they relate to. However, they are generally not very flexible insofar as substantial effort may be needed to adapt them to support other policies. Whether they are centralized or distributed, their design entails the support of a limited set of policies. That may be particularly so, if the model implementation involves multiple network nodes. Obviously, adapting the accounting mechanisms in all nodes involved to support other policies may be demanding for an internetwork the size of the

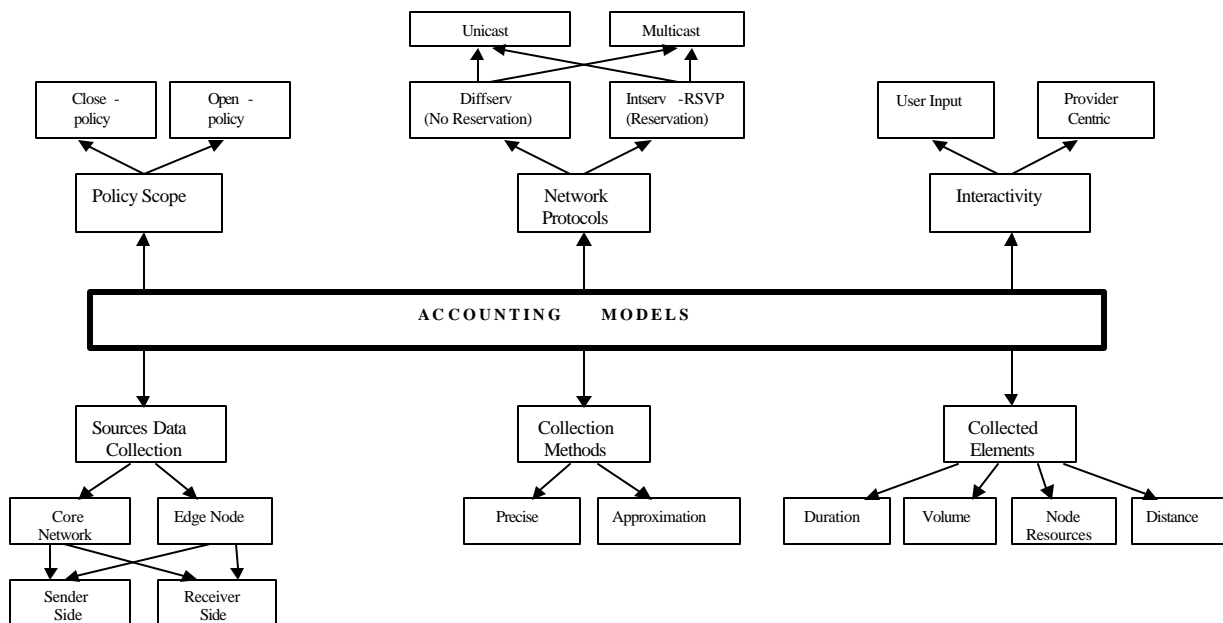


Figure 2: Taxonomy of Accounting Architecture Characteristics.

Internet. Hence, these models lean towards uniformity in accounting and policy across the Internet.

Such models may be called close-policy accounting models, because by virtue of their design, they are more or less explicitly geared to support one or a fixed set of policies. They may require, depending on the policies that they embody, minor or substantial change in the network architecture for their implementation. Examples include the induced accounting model of Smart-Market [28], One-Pass Accounting protocol for multicast suggested in [22], and some architectures for auctions and optimal pricing models [18] and [26].

5.2. Open-Policy Accounting Models

Accounting architecture may be independent of any specific pricing policy. Such model may be general so that a wide variety of economic models may be built on top of it. The scope of the pricing models supported, represents an important element that determines the flexibility of an accounting model. Accounting models whose design start by investigating architectural issues may be viewed as using a bottom-up approach. The advantage of such architectural-driven accounting models is that they fully integrate network architecture and protocol from the start. However, being independent of any specific policy, they may require systematic though limited adaptation work, in order to support any particular economic model. Since they generally provide a flexible foundation for supporting multiple pricing models, we will describe them as open-policy accounting models. Edge-Pricing architecture is an

example of this category [35]. It recommends an architectural design that would allow every provider to implement at the edge of his network a local accounting suitable to his economic policies and leaves inter-domain accounting to interconnection agreements.

Other examples that may be viewed as open-policy accounting include Arrow [15], CATI [17] and Generic Charging and Accounting (GenCA) [21].

6. SOURCES OF DATA COLLECTION

An important aspect of an accounting model lies in the source of data collection. That is typically carried out in the edge or the core of the network, and may target one side of the transaction.

6.1. Network Nodes

6.1.1. Network Core

This requires the configuration of multiple, if not all network nodes (routers/switches) in a domain or the Internet. This approach has the advantage of providing very detailed measurements at all stages of the network topology. It may provide a basis for dynamic pricing models that are based on the general network status, the actual path of traffic and multicast cost allocation schemes that assess imputations on a link basis. However, it does not scale well and it may bring more complexity in the engineering of the network. The overhead incurred may be very high and that may adversely affect network performance. Systems based on this approach may be involved as far as implementation, maintenance and upgrade are concerned. They are not very flexible, because any change in the

mode of operation requires change in each node. Examples of accounting models that require data collection in the core of the network include the induced accounting models of Smart-market [28], congestion-based optimal pricing models [33], and that of “Link Weighting” charging model [13].

6.1.2. Network Edge

Data is collected from some few nodes that serve as interfaces between users and the rest of the network, typically ingress or/and egress nodes. All accounting and pricing decisions may thus be managed at that local level. That leaves the internal structure of the network less complex with no additional overhead. However, this approach does not provide enough detailed elements on operations involving inner nodes such as multicast branching points and congestion bottlenecks. In general, it does not readily provide useful information for pricing and cost allocation to be applied to receivers. A typical example is Edge-Pricing architecture [35].

6.2. Side of Data Collection

The source of data collection may further be concerned with the side of transaction where elements are collected; that may lie with either the sender or the receiver. The implications of this architectural decision have a direct impact on pricing policies that may thus be supported.

6.2.1. Sender Side

Accounting on the sender side is the most popular form found in the literature. It has the advantage that information about the originator of the traffic is readily available to the network provider before any transactions take place. That also helps support pricing systems with built-in incentives for efficient network resource usage at the source of transmission, and auction models that require frequent reactions from the traffic sender. The main limitation of accounting on the sender side lies with the difficulty of dealing with charging models involving protocols such as multicast where costs are to be distributed among all participants (receivers), and require collecting information on both sides. Examples of this category include the accounting architectures of auction models [26], [34], that of INDEX [32], and Edge-Pricing architecture [35].

6.2.2. Receiver Side

Setting up accounting mechanisms on the side of the receiver of Internet transactions, is an option mostly used in multicast studies. Accounting on receiver sides is also used in unicast models that suggest ways to charge the traffic initiator. The difficulty of such systems is that in most cases, receiver information is only available in a remote region or domain, and it is not readily available for use where the traffic starts. Accounting mechanisms may therefore involve significant signaling and trusted relationships among Internet Service Providers etc. An example of such models may be found in [38]. Zone-based Cost Sharing architecture [8] and GenCA [21] also have provisions for accounting on the receiver side.

7. METHODS OF DATA COLLECTION

The methods used to collect data are important to accounting models, because different approaches determine the granularity of the collected elements and have implications on audit capacities, as well as the production of billing records. A good design must minimize the level of complexity and overhead that its implementation may incur.

7.1. Precise-Collection

The accuracy and fine-grained collection of cost elements is important for providing a full picture of resource consumption to users and help providers assess the availability of resource supply. Yet, the accuracy of the collection adds up to the complexity of the accounting system and increases the overhead in the computational load brought onto the system, especially if the data collected is the packet volume of traffic or congestion status [4], [21].

7.2. Approximation

Traffic measurements may be approximated to avoid costly accurate collection. Sampling may be used to that purpose, by exploiting statistical properties of traffic. Data is collected at various frequency intervals and extrapolated to determine the approximate resource consumption of each user. Sampling is typically used for volume of traffic, the packets sent over the network. The main limitation of sampling stems from its inability to provide a detailed figure of transactions, which may be necessary in case of disagreement between providers and users.

Approximation may also take the form of estimating user consumption from overall parameters of a negotiated SLA. The concept of effective bandwidth is a case in point [23] and [37].

8. COLLECTED COST OBJECTS

The collected cost objects provide a measurement of the resources to be charged for. There may actually be a great variety of elements to be collected based on the potential range of charging policies. In this sense, close-policy accounting models have the advantage of narrowing down these elements since they are designed to support specific policies. Open-policy models that are general and not designed to support explicit policies may collect more elements than necessary in many instances, and may need proper configuration to avoid needless overhead. Still, from an architectural viewpoint, there are commonly selected elements that are very representative of those found in most accounting models: duration, volume, distance, and node resources. A combination of several of these elements may be found in an accounting architecture proposal.

A. Duration: this refers to measurement for the length of time that traffic from a given source uses the network resources. It is more applicable to reservation-based transactions such as RSVP in IntServ, where it is easier to predict the amount of resources used per unit of time. In DiffServ, accounting for duration may also suit transactions in “Expedited Forwarding” class, which provides virtually no delay. Collecting traffic duration is easy to carry out with only a modest level of overhead.

B. Volume: the count of bytes or packets sent over the network during a session by a transaction is a commonly collected cost object. Contrary to duration measurement, volume metering involves dealing with the actual load put on the network. Measuring volume may require detailed metering, which may be costly from an overhead perspective. Volume-based accounting is used in NZGate [4]. In practice, a combination of volume and duration is quite frequent in the form of rate (i.e. $\text{Rate} = \text{Volume}/\text{Duration}$).

C. Distance: this generally represents the number of hops traversed by packets from source to destination. Distance may be evaluated from the network topology. Geographical information and IP address may also be used, although Internet address structure is

generally not directly expressive. Another method for distance evaluation involves the exploitation of IP packet TTL (Time To Live), which is decremented anytime that a packet crosses a router. Distance evaluation for accounting is used in [13].

D. Network Node Resources: packets need to use some amount of node resources (CPU and buffer) in order to meet QoS requirements such as low delay, low jitter and specific Per Hop Behavior (PHB) in DiffServ. There are accounting models that collect the usage of those elements, to typically support routing optimization policy, congestion control and distributed resource allocations [14], [33].

9. SUPPORTED NETWORK PROTOCOLS

Accounting models may be designed for specific network protocols and environments. The two main Internet QoS protocols that have major different accounting requirements are IntServ (Integrated Services Networks) and DiffServ (Differentiated Services Networks). Along each of these network environments, the accounting architecture may also target unicast or multicast.

9.1. QoS Protocols

A. IntServ: accounting for Integrated Services Networks with the RSVP protocol involves dealing with reserved and guaranteed resources, and pre-determined paths. In such an environment, all traffic parameters are well known in advance (e.g. duration, volume, and class of service). Data collection is hence made easier. Therefore, these QoS reservation parameters may be collected for accounting purposes as soon as resource availability and reservation are confirmed. Examples of accounting models for Integrated Services networks include Arrow [15], OCIT [41], One-Pass Accounting [22].

B. DiffServ: in Differentiated Services Networks, resources availability is checked and confirmed either by Bandwidth Brokers or light-weight RSVP across domains. There is no formal reservation in the sense that these resources are made available to other traffic if and when they are not used. Therefore, in DiffServ data must be collected on the fly and that brings about more performance constraints. Accounting needs to be multiform to cope with the heterogeneous systems and policy requirements of the various interconnected domains. Since there is no strict reservation per se, and no guarantee of QoS provisioning, it is important to use various adaptation strategies to account for packets whose class of service gets demoted from one domain to another. Interconnection may therefore be more relevant and accounting models need to be flexible enough to accommodate all these strategies and policies. Some representative studies in this category may be found in [7] and [17].

9.2. Communication Mode

The Internet has two major mode of transmission with very different accounting requirements.

A. Unicast: most pricing studies have implicit accounting mechanisms for unicast. In this case, communication is on a one-to-one basis and the identity of parties involved, is available in the packet header. Other information may be fairly easily collected.

B. Multicast: accounting methods for this mode require a distributed collection of information on many users. However, the identity of participants is not directly provided by standard multicast protocols. Heterogeneous QoS requests and traffic merging over shared trees, require that a method be devised to collect and consolidate accounting data in order to determine the cost to be attributed to each user. Multicast accounting models are proposed in [7] and [22].

10. INTERACTIVITY

Communication among users on one hand, and between users and the network on the other, may be indispensable in some approaches and therefore various parameters need to be exchanged. For instance, charges may be negotiated or determined solely on the basis of the cost elements collected. The interactivity component of an accounting system provides the opportunity to exchange custom information.

A. User input: users may submit their valuation of utility and reach an agreement through interaction with the network provider. In this case, accounting protocols or mechanisms provide the necessary elements for collecting such inputs, process them and send feedback to users. User input may take various forms: a bid to have a transaction completed, adaptive resource requests such as a strategy in Game Theory or utility characterization. INDEX architecture [32], and CATI [17] have provisions for user input. The level of interactivity may be a source for overhead in the accounting scheme and deserves careful design.

B. Provider-centric: accounting systems may operate solely from rules determined by providers, and standardized parameters are indiscriminately collected from traffic flows. The network provider is thus the only one involved in determining resource allocation policies and conditions. Such an approach has the potential advantage of reducing the level of signaling characteristic of interactive systems, and may have better performance. Yet, that approach limits user choices and self-optimization strategies. For example, the lack of interactive feature in accounting scheme may preclude implementation of auction like models.

11. DESIGN ISSUES

The design of accounting architecture and protocol needs to meet some important performance requirements. The taxonomy helps highlight some desirable characteristics. These may be very challenging to achieve in practice. Each proposal may show strength in some areas and limitations in some others. We further propose minimum engineering design requirements that should not be overlooked in order to optimize efficiency and effectiveness.

11.1. Scalability

This characteristic of accounting mechanisms refers to the ability for a model to be extended to multiple network nodes, domains, users etc., without incurring any significant additional load or overhead. This feature is all the more important in a large and heterogeneous network like the Internet, that accounting systems should be able to handle the high growth in the user population and the domains. For example, an architecture to support dynamic pricing and auction, should reliably provide price information from

and to users, most probably within some deadline. How to design a real-time, reliable communication mechanism to reach all users without much overhead? The characteristics of the taxonomy that are more relevant to scalability are the source of data collection and collected objects.

11.2. Flexibility

The flexibility of an accounting model refers to its ability to be easily used or adapted to support various network and computing environments, as well as different charging circumstances (e.g. pricing policies). An accounting system flexibility may be greatly enhanced if it can support a wide range of the characteristics surveyed in the taxonomy. The issue for a provider, is to have a single accounting architecture, on top of which different pricing policies can be implemented simultaneously or as needed over time.

11.3. Complexity-Overheads

The more sophisticated accounting models are, the more likely complex they are, and they increase the computational load in the network. For instance, detailed data collection requires a lot of CPU power and memory to store accounting records. An important feature of accounting systems that may add up to the overhead, is the communication and signaling mechanism to support the exchange of various type of information. The volume of information exchanged and its frequency may have major impact on the accounting system performance. For instance, in auction pricing models, it is important to provide a way to transfer user bids, which may be carried by dedicated bidding messages or it may be embedded in a portion of the packet. The former scenario may induce a lot of messages while the latter may use a significant portion of the header or the payload. A highly efficient and effective accounting scheme should leverage existing network architectures and protocols for a transparent integration.

11.4. Robustness

This feature refers to the ability of an accounting system to sustain operation under heavy loads. The system should reliably run and remain stable under various circumstances. For instance, an accounting system that has a very high frequency of signaling with poor communication control may have its various components severed from each other at time of congestion; the various components may not be able to communicate properly and the system may produce unreliable results.

11.5. Security

Accounting systems should include mechanisms to insure the integrity of data collected and prevent unauthorized users from having access and manipulating the records. Accounting ultimately leads to pricing and models of resource allocation for improved QoS performance. The absence of appropriate security mechanisms may not only cause prejudice to users, but also hurt the provider by allowing unauthorized access to the network infrastructure without proper data collection. Security includes authorization and authentication of users; as well as the protection of the integrity of the network system.

12. CONCLUSION AND FUTURE DIRECTIONS

In this paper, we investigated some of the main features of accounting architectures and how they contribute to overall efficiency. The features considered are not exhaustive; nevertheless they are very relevant in building effective accounting architecture. The insights provided by this taxonomy lay the ground for understanding and assessing models.

The effectiveness and the efficiency of accounting models are crucial in implementing viable pricing and economic based resource allocation models in the Internet. The bulk of research in network economics has focused on pricing models, and accounting architecture is still in need of further investigations. The design of efficient accounting is all the more difficult that some of the design issues described in this paper are hard to optimize altogether. Design approach must hence perform some trade-off, and optimization may ultimately depend on the class of pricing models that are likely to be implemented in a specific domain or by a set of cooperating providers.

We see multicast accounting as a fertile research area. As a matter of fact, multicast was designed to help improve network performance by having users share a single copy of QoS intensive applications. From an economic and game theoretic perspectives, users will cooperatively go for multicast, only if the cost allocation schemes allow clear saving over unicast. The support of equitable and incentive compatible cost allocation methods, requires detailed accounting of each user contribution to the global cost. The issues are that members are anonymous, have heterogeneous QoS demands; traffic is aggregated, the multicast tree topology has several splitting points where traffic diverges, and the tree crosses several domains under different administrative authorities. Hence, sophisticated accounting is important to capture the different cost relevant parameters in order to extract the share of each participant. It is still an open issue on whether, and how to make multicast accounting a local issue to each domain, so that each network provider may have more choice in the selection of their accounting model and cost allocation schemes.

Guided by the considerations made in this paper, we designed a new token-based accounting architecture (TOKENAC), with the aim of optimizing several of the components that we identified. It is designed to support different static and dynamic pricing models, as well as multicast. We leverage the principle of token-bucket, used for traffic regulation to perform QoS accounting, and thus avoid dedicated traffic metering. That architecture is currently being tested through simulation.

13. REFERENCES

- [1] Aboba B., Arkko J. and Harrington D. 2000. Introduction to accounting management. *Internet RFC 2975 (October)*, (<http://www.rfc-editor.org/>).
- [2] Blake S., Black D., Carlson M., Davies E., Wang Z. and Weiss W. 1998. An architecture for differentiated services. *IETF Internet RFC 2475 (Dec.)*, (URL: <http://www.rfc-editor.org/>).
- [3] Braden R., Zhang L., Berson S. and Jamin S. 1997. Resource Reservation Protocol (RSVP)- Version 1, Functional Specification. *Internet RFC 2205 (September)*, (URL: <http://www.rfc-editor.org/>).

- [4] Brownlee Nevil. 1996. New Zealand experiences with network traffic charging. *Journal of Electronic Publishing (JEP)*, vol. 2, no. 1, (May), (www.press.umich.edu/jep/econTOC.html).
- [5] Brownlee Nevil. 1997. Traffic flow measurement: architecture. *IETF RFC 2063* (January), ([URL:http://www.rfc-editor.org/](http://www.rfc-editor.org/)).
- [6] Calhoun P., Akhtar H., Arkko J., Guttman E., Rubens A. and Zorn G. 2001. DIAMETER base protocol. *IETF Internet Draft*, (www.ietf.org/).
- [7] Carle G., Hartanto F., Smirnov M. and Zseby T. 1999. Charging and accounting for QoS-enhanced IP multicast. *Proceedings of IFIP Six International Workshop on Protocols for High-Speed Networks*, Salem, MA, (August 25-27).
- [8] Clark David. 1996. Combining sender and receiver payments in the Internet. In *Interconnection and the Internet, Proceedings of Telecommunications Research Policy Conference 1996*, Gregory Rosston and Waterman (Eds.), Lawrence Erlbaum Pub.(1997), Mawah, New Jersey, pp. 95-112.
- [9] Cocchi R., Shenker S., Estrin D. and Zhang L. 1993. Pricing in computer networks: motivation, formulation and example. *IEEE/ACM Transactions on Networking*, vol. 1, no. 6 (December), pp. 614-627.
- [10] Courcoubetis Costas. 1998. Pricing and the economics of networks. *Infocom'98 Tutorial, San Francisco, USA, March 29-Apr.2*.
- [11] DaSilva Luiz A. 2000. Pricing for QoS-enabled networks: a survey. *IEEE Communications Surveys and Tutorials*, vol. 3, no. 2, Second Quarter 2000, pp. 2-8.
- [12] Edell R., Mckeown N. and Varaiya P.P. 1995. Billing users and pricing for TCP. *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 7 (September), pp. 1162-1175.
- [13] Einsiedler H. J. and Hurley P. 1998. Link weighting: an important basis for charging in the Internet. *Technical Report SSC/1998/017*, Communication Systems Division, EPFL (Swiss Polytechnic University, Lausanne), ([URL:http://lrcwww.epfl.ch/PS_files/publications.html](http://lrcwww.epfl.ch/PS_files/publications.html)).
- [14] Fankhauser G. and Plattner B. 1999. DiffServ bandwidth brokers as mini-markets. *International Workshop on Internet Services Quality Economics*, MIT, Cambridge, Massachusetts (December 2-3).
- [15] Fankhauser G., Stiller B. and Plattner B. 2000. Arrow: a flexible architecture for an accounting and charging infrastructure in the next generation Internet. *NETNOMICS: Economic Research and Electronic Networking*, vol. 1, no. 2, pp. 201-223, ([URL:www.baltzer.nl/netnomics/contents/1999/1-2.html](http://www.baltzer.nl/netnomics/contents/1999/1-2.html)).
- [16] Ferguson D.F, Nikolau C., Sairamesh J. and Yemini Y. 1996. Economic models for allocating resources in computer systems. In *Market Based Control of Distributed Systems*, Scott Clearwater (Ed.), World Scientific Press, River Edge, New Jersey, pp. 156-183.
- [17] Foukia N. and Billard D. 1999. Charging and accounting technology for the Internet: the CATI project. *6th OpenView University Association WS (HPOVUA'99)*, Bologna, Italy (June 12-15), (www.hpovua.org/PUBLICATIONS/PROCEEDINGS/6_HPOVUAWS/).
- [18] Ganesh A. and Laevens K. 2000. Dynamics of congestion pricing. *Technical Report MSR-TR-2000-70*, Microsoft, Research, Cambridge (June), (<http://research.microsoft.com/research/network/>).
- [19] Gibbens R. and Key P. 2000. Distributed control and resource pricing. *Tutorial ACM SIGCOMM 2000*, Stockholm (Aug. 29). (Also:http://research.microsoft.com/users/pbk/GKSig2kResPriceTutorial_files/frame.htm).
- [20] Gibbens R. and Key P. 2001. Distributed control and resource marking using best-effort routers. *IEEE Network*, vol. 15, no. 3 (May), pp. 54-59. (See also: <http://www.research.microsoft.com/network/>).
- [21] Hartanto Felix and Carle George. 1999. Policy based billing architecture for Internet differentiated services. *Proceeding of IFIP International Conference on Broadband Communications, BC'99* (November), Hong Kong.
- [22] Herzog S., Shenker S., and Estrin D. 1997. Sharing the cost of multicast trees: an axiomatic analysis. *IEEE/ACM Transactions on Networking*, vol. 5, no. 6 (December), pp. 847-860.
- [23] Kelly F.P. 1996. Notes on Effective Bandwidth. In *Stochastic Networks: Theory and Applications*, F. Kelly et al. (Eds.), pp. 141-168 Oxford University Press, 1996.
- [24] Kelly F.P. et al. 1998. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of Operational Research Society*, vol. 49, pp. 237-252.
- [25] Key Peter and Derek McAuley. 1999. Differential QoS and pricing in networks: where flow control meets game theory. *IEE Proceedings Software*, vol. 146, no. 2 (March), pp.39-43.
- [26] Lazar A.A. and Semret N. 1997. Auctions for network resource sharing. *CTR Technical Report No. 468-97-02*, Columbia University, New York, ([URL:http://comet.columbia.edu/publications/techreports.html](http://comet.columbia.edu/publications/techreports.html)).
- [27] Mackie-Mason and Varian H.. 1995. Pricing congestible network resources. *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 7 (Sept.), pp.1141-1149.
- [28] Mackie-Mason and Varian H. 1995. Pricing the Internet. In *Public Access to the Internet*, Brian Kahin and James Keller (Eds.), MIT Press 1995, Cambridge MA, pp. 269-314.
- [29] Nichols K., Jacobson V., and Zhang L. 1999. A two-bit differentiated services architecture for the Internet. *IETF Internet RFC 2638* (July), (<http://www.rfc-editor.org/>).

- [30] Orda A. and Shimkim N. 2000. Incentive pricing in multi-class systems. *Telecommunication Systems*, vol 13, no. 2/4, pp. 241-267.
- [31] Pras A. et al. 2001. Internet accounting. *IEEE Communications Magazine*, vol. 39, no. 5 (May), pp. 108-113.
- [32] Rupp B., Edell R., Chand and Varaiya P. 1998. INDEX: a platform for determining how people value the quality of their Internet access. *Sixth International Workshop on Quality of Service (IWQoS 98)*, Napa-California (May 18-20), pp. 85 -90.
- [33] Sairamesh J, Ferguson D. and Yemini Y. 1995. An approach to pricing, optimal allocations, and quality of service provisioning in high-speed networks. *Proceedings of the IEEE INFOCOM'9* (Apr. 2-6), Boston MA, vol. 1, pp.1111-1119.
- [34] Semret N., Campbell A. and Lazar A. 1999. Market pricing of differentiated Internet services. *IEEE Proceedings of 7th International Workshop on Quality of Service (IEEE/IFIP IWQOS'99)*, London,UK, pp. 184-193.
- [35] Shenker S., Clark D., Estrin D. and Herzog S. 1996. Pricing in computer networks: reshaping the research agenda. *ACM Computer Communication Review*, vol. 26, no. 2 (April 1996) pp. 19-43. Also in *Telecommunications Policy Research Conference*, vol. 20, no.3 (April), pp. 183-201.
- [36] Shenker S., Partridge C. and Guerin R. 1997. Specification of guaranteed quality of service. *IETF Internet RFC 2212 (September)*, (URL: <http://www.rfc-editor.org/>).
- [37] Songhurst D.J. (Ed.). *Charging Communication Networks: from Theory to Practice*. Songhurst D.J. (Eds.), Elsevier (1999), Amsterdam-Holland, pp-23-36.
- [38] Sprenkels R., Parhonyi R., Pras A. Beijnum B. and Goede L. 2000. An architecture for reverse charging in the Internet. *IEEE Workshop on IP-Oriented Operations and Management (IPOM 2000)*, Cracow, Poland (September).
- [39] Stahl Dale O. 1995. A critical survey of Internet pricing proposals. *Proceedings of OECD Workshop on the Economics of Information Society*, Istanbul, Turkey (December 14-15).
- [40] Stiller B., Fankhauser G., Plattner B. and Weiler N. 1998. Charging and accounting for integrated Internet services: state of the art, problems, and trends. *INET'99: The Internet Summit* (July 21-24), Geneva-Switzerland.
- [41] Stiller B., Fankhauser G., Joller G., Reichl P. and Weiler N. 1999. Open charging and QoS interfaces for IP telephony. *INET'99: The Internet Summit* (June 22-25), San Jose CA.
- [42] Von-Neuman John and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton-New Jersey.
- [43] Wroclawski J. 1997. Specification of the controlled load network element service. *IETF Internet RFC 2211 (September)*. (URL: <http://www.rfc-editor.org/>)
- [44] Yaiche H., Mazumdar R. and Rosenberg C. 2000. A Game theoretic framework for bandwidth allocation and pricing in broadband networks. *IEEE/ACM Transactions on Networking*, vol. 8, no. 5 (October), pp. 667-678.