# Design Considerations for Usage Accounting and Feedback in Internetworks

**Deborah Estrin and Lixia Zhang**[1]
estrin@usc.edu, lixia@parc.xerox.com

## Abstract

This paper investigates the design of resource usage feedback mechanisms for packet switched internetworks. After a discussion of the *motivations* for feedback mechanisms, *feedback channels* and *policies* are described. We then outline issues raised by the design of *mechanisms* to realize these policies, including: network service disciplines, accounting granularity, metrics, authentication, and coordination among transit carriers.

Usage-based charging is only one means of feedback. Our purpose is to begin a systematic discussion of the technical issues associated with a range of usage feedback alternatives. Therefore the paper should not be read as a policy statement promoting usage sensitive charging in internets. In fact, one of the goals of the feedback mechanisms explored in this paper is to allow network service providers and users to *avoid* the introduction of usage sensitive charges if they so wish; while still realizing the benefits of statistical resource sharing offered by packet switching and the benefits of efficient resource utilization offered by usage feedback.

**Keywords: Network Accounting, Inter-Enterprise Networking.**

## 1 Introduction

This paper concerns resource usage feedback for interconnected, packet-switched, computer communications networks; hereafter referred to as internetworks, or internets. The global internetwork has developed through the interconnection of thousands of commercial and private networks.[2] As the technology matures, the role of commercial service providers is expected to grow, along with the demand for accounting mechanisms. At the same time, increasing connectivity brings with it the need for mechanisms that motivate efficient behavior on the part of the larger and more heterogeneous user population.

The traditional circuit-switched telephone network provides a possible model for resource usage accounting and feedback. However, many of the mechanisms do not translate directly into a packet switched environment. In this paper we investigate the design space for resource usage accounting and feedback mechanisms in a large scale, packet-switched internetwork. In particular, unless stated otherwise, most of our discussions below assume a **connectionless** internet that provides datagram services.

Although we do not address issues of cost recovery specifically, charging is one form of feedback and therefore this discussion is of relevance to cost recovery as well. Cost recovery entails additional tasks such as setting prices based on a careful assessment of both fixed and incremental cost factors; further discussion is beyond the scope of this paper.

### 1.1 Internet model and terminology

Internet technology has developed primarily within private and consortium networks. Commercial carriers have participated mostly through leasing of lines used to connect network nodes within the private networks. More recently there is increased interest in commercial offerings of datagram delivery services, e.g., SMDS [14]. The advent of commercial offerings introduces new incentives, and in some cases a necessity, for resource usage feedback mechanisms; and the accounting necessary to collect information for the feedback channel.

We refer to the different administrative entities and their associated network resources as Administrative Domains (ADs). As described in [10, 13, 12], an AD is a set of resources (network links, routers, bridges and end systems) under the control of a single administrative authority. In this context, a *stub* AD is one that does not carry transit traffic for other ADs, e.g., private customers/consumers of communications services. That is, all traffic entering a stub AD is destined for end systems within that AD, and all traffic exiting a stub AD originated within the AD. Most campus and corporate networks are examples of stub ADs. *Transit* AD refers to an AD whose primary function is to provide transit services for other ADs. Long haul backbone and regional networks are examples of transit ADs. In addition, some private networks that are connected to more than one

[2] Commercial networks refer to those that offer services to anyone and for any purpose, so long as they pay the established fees (e.g., AT&T, GTE Telenet, MCI, PSI Inc.). Private networks refer to those that are operated and used by a restricted set (often one) of organizations (and/or for a restricted set of uses) based upon administrative, instead of (or in addition to) monetary, arrangements (e.g., NSFnet, Xerox Corporation's internal network).

transit or stub network offer limited transit services to select ADs. We note the existence of bypass links, along side the more common hierarchical structure. The term *end user* refers to the human beings who make use of the communication resources via the end-systems that lie within the ADs. These distinctions are relevant to our discussion because we must identify which entities provide the feedback, and likewise to which entities the feedback is provided.

### 1.2 Accounting in Packet Switched Internets

The effort required to account for traffic depends upon the network architecture. Circuit switched networks reserve resources for each user call, and therefore feedback and accounting can be performed along with call setup and teardown. Connection-oriented, packet-switched, networks maintain state per connection inside the network and successive packets in a connection typically travel via a fixed route (although some architectures allow the connection to switch routes in midstream). If the connection protocol reserves resources then the accounting and feedback needs are analagous to the circuit switched case. If there is no reservation, then connection state and switch function must be augmented with accounting related information and packet-counting, respectively. In a pure datagram network there is neither resource reservation nor per-user state maintained within the network. Packets from the same end-to-end connection (i.e., source-destination, transport level association) are forwarded independently and may travel through different routers.

There is also a further interaction among application types, network architecture, and accounting. In computer communications, the range of application behavior and desired services is much greater than in voice telephony. Human to human voice communication represents a single type of application, and the entire telephone network has been built to optimize the service quality and pricing mechanisms of the application.[3] Moreover, whereas voice traffic is handled relatively efficiently with circuit switching, computer communications are often bursty. The more varied and bursty the traffic patterns, the more important it is to avoid inefficient forms of resource reservation.[4] The diversity of traffic patterns presented by computer communication applications implies that the network should distinguish between classes, or types, of service (e.g. delay sensitivity). When different *types of service* (TOS) are provided, the accounting function will need a more complex mechanism than a simple packet meter. The extra packet processing involved in supporting TOS specific performance guarantees may offer some opportunity for supporting accounting related functions *if* the TOS is implemented on a connection basis. Otherwise if TOS is offered on per-packet basis, additional work would be required to account for usage on a per packet, per TOS basis.

Another difficulty with respect to accounting in a packet switched, computer communications context is the unit of accounting. The units of accounting in packet switching potentially are much smaller than in circuit switching (i.e., a packet instead of a call) so the overhead of accounting could be much higher. This small unit is also problematic for the end user. A user can easily estimate the cost of a telephone call based on the call duration. In the current computer communications environment, however, it would be difficult for a user to predict the network usage implied by his or her application-level actions, if the network accounting is based on the unit of packet. The packet is too low level of an abstraction for the user; and today's applications and transport protocols are not instrumented to translate packet counts into units that are meaningful to the end user.

Despite these difficulties, usage accounting and feedback have some particular benefits in the context of packet switching. Computers served by packet switched internets differ in many ways from human users served by telephone networks. Real time, voice communication exhibits rigid requirements for stable transmission delay and rate. Many computer applications, on the other hand, exhibit "softer", more elastic, service requirements. For example, a circuit switched phone call must have a 3 Khz channel allocated, otherwise the call cannot start. A packetized voice session, however, can tolerate some degree of packet loss and still support intelligible communication.[5] Due to their asynchronous characteristics, certain applications can even accept temporary postponement of services; electronic mail and background file transfers are such examples. Therefore it is possible to regulate packet traffic by usage feedback, and thereby enable a service provider to offer better service, at lower prices, to all end users. For this reason, mechanisms for usage feedback could

[3] Although today's telephone networks also carry non-voice applications, such as FAX and dialup terminal-to-computer connections (through the use of a modem), voice remains the dominant load in the system. FAX traffic makes efficient use of the communication circuit. Terminal to computer connections, however, make relatively inefficient use because of their bursty nature.

[4] However reservation may be necessary whenever you need to guarantee a service and it is possible to implement efficient reservation for bursty traffic; this is the subject of ongoing research.

[5] One could argue that the telephone systems do not exploit the complete market. There is a potential of multi-TOS for voice communication as well. For example, calls can be sorted to interruptable and non-interruptable ones, so that the former can be cut off during peak hours but also receive a lower charging. However, unlike the electrical power market and data communications, the market size and network efficiency gains do not appear to warrant the overhead associated with differentiating between the traffic types.

benefit both service providers and consumers, if appropriately designed and implemented.

In summary, there are several interesting technical issues raised by the question of accounting and feedback in packet switched internets. We discuss motivations and models for usage feedback in Sections 2 and 3. Section 4 outlines several issues associated with the design of supporting mechanisms.

# 2 Motivations

There may be multiple purposes served by accounting and feedback for resource usage. One goal may be to recover costs. Another may be to motivate users to behave more efficiently from the perspective of the shared resources (i.e., the network). In the latter case, feedback signals should be different when the network is lightly loaded than when it is heavily loaded. Although both cost recovery and efficient network usage can be achieved using accounting and feedback, accomplishing one does not necessarily accomplish the other. Moreover, a usage-sensitive charging mechanism in one part of the internet may introduce the need for a feedback scheme in another part (e.g., a transit carrier's charging mechanism may motivate a stub AD to introduce usage sensitive feedback in order to motivate efficient use of the communications budget). This paper focuses on design considerations for usage feedback mechanisms. However, because of the potential interaction and frequent confusion, we begin with a brief discussion of cost recovery.

## 2.1 Cost Recovery

The most basic cost recovery goal is to generate revenues that are adequate to pay for physical facilities (links, routers, etc.), operation, maintenance, software development, personnel, etc. This model is complicated somewhat by the need to generate enough revenue to fund improvement and expansion.[6]

A more unique problem in the context of data networking is an environment in which additional capacity can be called up on demand (at greater expense than had it been planned for and installed privately). As traffic load increases, decisions must be made concerning a) whether to dial up additional resources, b) how long and under what conditions to maintain them, c) how to distribute this additional cost among users, d) whether to redistribute existing capacity, and e) at what point to invest in permanent facilities instead.

[6] This problem has been studied extensively by economists in the areas of telephony and utility company capacity planning and tariffs.

Cost is recovered by charging users for their network usage. Therefore the charging itself is one means of feedback. As such the charging policies may have a great impact on users' behavior. For example, the most common form of cost recovery today in packet switched networks is a fixed-fee per physical connection, where the fee is often a function of the bandwidth of the leased lines utilized by the connection. Neighboring transit ADs agree upon procedures for carrying each others traffic. The mechanisms for supporting various settlement and allocation procedures among the transit ADs is an intersting issue beyond the scope of this paper; it has been addressed extensivly in the case of telephony.

This approach provides no feedback to the end user regarding the actual resource usage and so does little to encourage efficient network usage. The feedback only provides a signal to the organization as to what bandwidth connection to select, or whether to connect at all. In the absence of any other feedback, connected users would have little incentive either to upgrade a poor protocol implementation to the best available one (which may cost both effort and money), or to carefully plan their network usage to avoid congesting the network unnecessarily.

Another concern is the desirability, from a policy perspective, of exposing all users and usage to usage-sensitive billing. It may be preferable in some environments to decouple cost recovery and usage feedback in order to encourage communication among all, or some special subset, of users (e.g. promoting communication among members of the research community). In other words, global efficiency is very hard to measure when one takes into account the externalities (goods and costs) associated with communication. Therefore, it is not appropriate to simply minimize network usage to the exclusion of other factors. For this reason, we discuss alternative feedback models below.

## 2.2 Feedback

Feedback is needed in any service system to motivate users to make globally-efficient use out of existing resources. ¿From the systems' perspective, when the system is lightly loaded feedback should encourage (or at least not discourage) usage to maximize system throughput. When the system is heavily loaded (i.e., demand approaches or exceeds the finite capacity) feedback should motivate deferable users to delay submitted traffic or expendible users to back off altogether. That is, an ideal feedback system would encourage intelligent usage while preventing the system from being overloaded. In the context of internets, there are two particular types of efficiency that we want to motivate: efficient implementations and efficient end-user behavior. For example, a good transport protocol implementation that eliminates superfluous retransmissions should help reduce the probability of network congestion.

An example of motivating efficient user behavior is feedback that encourages users to shift time-insensitive traffic to off peak hours. The current Internet, for example, may be considered as providing a very crude form of such feedback, e.g, during peak hours the network performance degrades so that performance-sensitive users are forced to shift their work to less crowded times of day. However, users less sensitive to performance degradation, might even have an incentive to transmit more to compensate for the losses caused by congestion. The inefficient users are not penalized adequately by the total queueing delay increase or packet losses that is caused by their action. The current Internet provides a first-come-first-serve (FCFS) datagram service, therefore the increased delay and losses are shared among all users. When demand exceeds capacity, the result is a network that is overly congested during peak hours and consequently provides poor performance to all users. In other words, what is locally efficient behavior for some users results in globally inefficient resource usage from the perspective of the network resources. It illustrates the tragedy-of-the-commons phenomenon [9].

Before describing feedback models in Section 3 we address the interaction and distinction between different internet participants.

### 2.3 Transit Carrier, Stub AD, and End User Goals

Motivations for usage accounting and feedback differ for transit and stub ADs, and for end systems.[7] At the same time, the motivations of the three types of internet entities interact with one another in some predictable ways.

- Transit carriers are concerned with cost recovery through collection of user fees or third party subsidies. In a competitive internet environment, cost recovery increases in importance. Carriers compete by offering attractive services at relatively low prices while still covering expenses and expansion.

  To keep the price low, transit carriers are concerned with efficient usage of their resources. If users behave efficiently then the same service can be provided to the same number of users at lower cost than if if users behave inefficiently.

- Stub ADs want to minimize, or at least contain, costs in the presence of whatever feedback scheme transit carriers implement. In a flat rate environment, stub ADs may be concerned with recovering costs of network attachment charges, and/or with promoting efficient use of a limited capacity connection. Where transit carriers introduce usage feedback, some stub ADs may want to pass such signals back to some or all end systems or users in order to encourage their more efficient behavior. In addition, as transit carriers introduce usage sensitive pricing, stub ADs will be increasingly concerned with verifying that their bills are accurate, i.e., they will want to take measures to prevent fraud. Stub ADs will be concerned with developing accurate models of usage in order to anticipate, plan for, and detect anomalies in, usage and charging.

- End systems and users will similarly want to minimize, or at least contain, costs in the presence of transit and stub AD feedback mechanisms. Some end system administrators may wish to avoid the overhead and inhibited communication that can result from too fine-grain accounting (while still controlling costs), while others will want to propagate feedback signals all the way to the end user. In addition, some users may be considered billable and others not. For similar reasons as stub ADs, end systems and users will require (better) tools with which to predict, assess, and verify the communication costs associated with their transactions.

We will elaborate on stub AD and end system requirements for implementing usage accounting and feedback in Section 4.8. For now we return to our discussion of feedback in more detail.

## 3 Feedback Models

Feedback schemes can be characterized by the *feedback channel* used and the *policies* implemented.

### 3.1 Feedback Channel

Usage sensitive *charging* implies billing for services, by definition. But feedback to end systems or users regarding resource usage can also be achieved in terms of network signaling, service quality degradation (e.g. delay), or even administrative means; as an alternative, or in addition, to actual monetary feedback. Each of these can be thought of as a different *feedback channel*. The feedback is usually sent to the traffic source, but in some cases may be sent to the destination or some third party. Below we discuss the features and merits of different feedback channels.

[7] Of course the motivations for feedback are not identical for all stub ADs, or for all transit ADs. However, there is more commonality among entities of a particular type.

- **Monetary feedback** has very explicit impact on user behavior. Individuals and groups have limited budget resources, and therefore are motivated to to economize on their usage (i.e., communication expenses).

  However, explicit, direct impact does not necessarily mean that this channel is always optimal or desirable. Consider the research community as an example. Externalities such as inhibiting communication based on price-elasticity may well be *undesirable* from the perspective of the social good of "research communication, productivity, and technology transfer", for example. From the perspective of global efficiency, individuals may make suboptimal decisions to underinvest in communications. In other words, some individuals will experience all the cost but not all of the benefit of their expenditure, when the benefit of their communication is partially (or largely) to other members of the community. Consequently, if each individual optimizes his or her own behavior based on local costs and benefits, a social optimum may *not* be achieved.

  Allocation or quota schemes can act as a proxy for monetary billing. Traffic sources (which may be end systems or stub ADs, depending upon the accounting granularity) are encouraged to behave efficiently because they have a limited resource, their quota. Various quota schemes have been used in computer systems for usage accounting (e.g. MIT Multics). Such allocation schemes do have drawbacks. For example, users may overly constrain their communication early in the quota period and over utilize at the end or vice versa (i.e., a user could flood the network with traffic at the beginning of the period and then starve for the duration). Unlike real money, the quota is not exchangable for other goods or services and is more likely to result in this sort of inefficient usage.

- **Performance feedback** can take different forms: an explicit message from the network warning of overload condition (e.g. ICMP source-quench [16]), or an implicit increase in delay or packet-loss rate. This type of feedback has no relation to cost recovery. Its function is to influence user behavior (directly, or indirectly through intermediate protocol layers). For example, upon receiving an ICMP source-quench message requesting a slow down in data transmission, users who find the service inadequate may decide to shift to a less congested time of day, or adjust their usage in some other way.

  However, in the absence of other mechanisms, applications or users who are insensitive to the performance parameters may not modify their behavior. For example, electronic mail runs in the background and the end user would not notice whether transfer of some message incurred 50% retransmissions. This can lead to a situation in which performance-sensitive users under-utilize the system (because they find it of less value) and performance-insensitive applications over-utilize the system, from a global efficiency perspective. Therefore performance feedback is most effective when TOS support mechanisms are in place, so that performance-sensitive users can be given priority in utilizing network resources.

- **Administrative feedback**, such as monthly usage reports or allocation schemes, may be used alone, or in combination with performance feedback. Administrative feedback can be effective to the extent users are sensitive to administrative (or peer) pressures. Usage levels can be posted or broadcasted at regular intervals; the performance-insensitive users described above might then be discouraged via administrative pressure from overutilizing the resources. The result would be a more attractive network for performance-sensitive users, and relatively little degradation for the performance-insensitive user who could shift usage to uncongested times of day.

Depending on the feedback channel in use, the receiver of the feedback signal can be different. For example, performance feedback will be received directly by the end user. Administrative feedback may target the stub AD, which may then redistribute the signal internally through whatever channel it deems appropriate. Regardless of the channel type, in order for feedback to be most effective end users should be the ultimate receiver of *some* form of the feedback signal. But how the feedback is provided internally, and whether to associate it with internal accounting and billing actions, is the stub AD's decision.

In summary, the granularity of the feedback recipient is tightly coupled to the intended objective. We suggest that the first objective is to carry the collection of users' traffic in an efficient manner, e.g., introduce delay for deferrable traffic such as asynchronous mail when the network is heavily loaded. This may be achieved, at least in part, through relatively coarse-grained feedback. A second objective may be, in some cases, to provide feedback to finer grain traffic sources (human users) in order to alter users' demand, i.e., offered load, most directly.

## 3.2 Feedback Policies

In this section we describe four alternative usage accounting and feedback policies: flat per-packet fee, TOS based, peak load, and priority based. We are interested in the potential impact each policy may have on the user's behavior, and thus its effectiveness in regulating network usage. These policies typically are described in the context of monetary feedback, i.e., billing. However, schemes can be devised using the other types of feedback channels.

Unless otherwise explicitly stated, we assume the network serves each datagram on a first-come-first-serve (FCFS) basis. (We explore the impact of alternative network service disciplines in Section 4.1.) Moreover, we assume that some form of feedback signal is passed to end users, directly or indirectly, in order to influence their behavior.

#### 3.2.1 Flat Per-Packet Charging

To the extent carriers' costs are related to usage, flat per-packet charging schemes provide a means for distributing costs among users (e.g. SMDS) [14]. Moreover, this approach provides fine grain feedback to the user to promote efficient use of network resources. However, because the feedback is based on a flat per-packet fee, which is independent of current system load and service quality received, it does not encourage users to delay non-time critical usage and may lead to underutilization when the network is not loaded. The network provides all users with either a best-effort service (e.g., IP) which may be inadequate for real-time applications, or with a guaranteed high-quality service (e.g., SMDS), which may not be cost effective for less demanding applications. The uniform service type provides no incentive (or support) for users to sort their applications into different categories in order to share the resources most efficiently.

#### 3.2.2 TOS Based

If internets offer different types of services, the accounting should be based on the TOS service quality provided. When the network is fully loaded, however, additional traffic that requires a high quality TOS will have to be rejected in order to guarantee service qualities to the current users. In this case, users can predict the cost for a required level of service. They either get the requested level of service or nothing.

If the accounting in a TOS-guarantee network is independent of the current or expected system load, and the network simply applies a FCFS policy to resource requests, late comers during peak hours will be forced to shift to different usage times. However, the network would prefer to encourage demanding-TOS users when the network is under-utilized, and discourage them when the network is loaded, by having load-sensitive TOS accounting. This can be achieved by one of the load-sensitive accounting policies discussed next.

#### 3.2.3 Peak load

Peak load pricing provides different feedback (e.g., charges different rates) depending upon the aggregate demands placed on the system [2]. If there are regular, predictable times of day at which the network will be heavily loaded, then the charge for transmission during those hours can be raised significantly to shift flexible users off the peak. The charge may be in terms of real money, monthly-report-units, or allocated credit-units; corresponding to the different types of feedback channels. The accounting procedure may be activated at connection setup time in a connection-oriented internet, or on a per packet basis in a datagram internet.

This scheme is most effective when peak periods are predictable so that users can plan and behave accordingly. Network traffic measurements from different sources have shown consistent gross patterns of network busy hours [1, 8]. If traffic patterns are not so predictable, peak load rates could be varied dynamically with network load. However, traffic sources would not be able to predict their charges accurately, thereby undermining the utility of the feedback channel for budget planning purposes.

Network facilities may be expanded to meet demand on a dynamic basis, i.e., the network provider may dial up additional facilities to meet peak hour demand.[8] There is a symbiotic interaction between peak load policies and dynamic network costs. By setting a higher charge for peak hours, the revenue may be used to cover the extra cost of dialup lines. More generally, if traffic load variations are predictable, the need for dialing up additional capacity can be predicted and the situation can be made to resemble the fixed capacity case.

#### 3.2.4 Priority based

An alternative to peak load rates is priority pricing.[9] Under this policy, the network will serve users in the order of their priority levels, and the rate charged for carrying traffic will be computed accordingly. These rates are slow to change and are advertised to all traffic sources. This scheme is more adaptive than peak load schemes because the priority labels provide a basis for the network to delay lower priority traffic in favor of higher priority when loaded.

[8] Even a fixed-facility network is not faced with static costs, i.e., capacity planning decisions are made continually regarding installation of additional facilities. But in the dynamic case, dial-up circuits interact directly with real-time performance and monetary feedback channels.

[9] Scott Shenker of Xerox PARC originally proposed this approach for use in datagram internets.

Priority pricing has been implemented by utility companies and appears quite promising for network resources as well [17]. For example, in electrical power systems, at peak load the service provider may not be able to meet the peak demand from all users. The priority pricing implementation charges less to customers who are willing to have their service curtailed/cut-off when demand rises above capacity. Inflexible users pay more to receive a guaranteed continuous service. The scheme is relatively static because users vary their priority level infrequently and slowly (i.e., they put in a request and expect it to take some time to go into effect).[10]

In the data network context, performance feedback and priority adaptation could be more dynamic. For example, a user first sets a certain priority level; if the experienced delay is too great (or some other quality metric is too low), and if the users' demand is relatively elastic to performance but inelastic to price (or the administrative equivalent), they may increase the priority levels until acceptable performance is achieved. This means that the actual cost of a particular transaction will depend on the network conditions at that time. To the extent network load is predictable, users will distribute their usage more evenly. The net result is more efficient use of network resources. However, a concern here is the potential inefficiency of highly dynamic, real-time, tuning of priorities to optimize end-user service and cost.

If there is no accounting system associated with a priority scheme, however, all users have incentive to set high priority on all traffic, and the scheme will not be effective. Consequently, whether through administrative means, or using an actual or proxy (quota system) monetary channel, users' priority setting must be regulated.

In the subsequent section we investigate the supporting mechanisms required, and the design issues raised, by consideration of usage based feedback in internets, with a particular interest in load-sensitive, TOS feedback.

# 4 Design Issues

In this section we enumerate several essential choices that must be made in designing usage-based feedback mechanisms for transit and stub ADs, i.e., network service discipline, accounting granularity, feedback frequency, cost metrics, dynamic capacity issue, authentication, and coordination required among transit ADs. Based on the very early state of work in this area, we raise more questions than we answer. Much more work is needed to analyze design choices and tradeoffs in detail.

## 4.1 Network Service Discipline and TOS Implementation

The network service discipline employed influences the feedback signals directly. Thus far we have assumed FCFS packet handling. We now consider the interaction of network service and feedback channels in more detail.

To make feedback messages meaningful to the end users, the network must have a clearly defined service discipline in addition to accounting policies; especially when the messages are through performance or monetary channels. The most common discipline in today's networks is FCFS, best effort service. The primary merit of FCFS is simplicity in implementation. Under FCFS, routers have no need to identify or discriminate among users; they merely forward each packet as quickly as possible. However, due to lack of user identifications, FCFS networks do not provide any load-sensitive feedback to individual users other than through the aggregated, implicit, performance channel; which has the unfortunate tragedy-of-the-commons inefficiencies described earlier. Alternatively, if the network provides a fair-queueing service to all end systems [4], when an end system detects an increased loss rate, it will understand that its current data transmission is going faster than its fair share of the network resources. If the end system chooses to ignore the signal it will harm only itself.

Recently, there has been an increased interest for some transit ADs to control the usage of their resource by different user groups, and to provide insulation among users to minimize traffic interference. Providing this functionality will require identification of users, and an appropriate service discipline, for each user or user group. The same mechanisms may then be used to support usage accounting and feedback, although possibly at a different granularity.

An internet may provide multiple levels of resource control through the use of multiple service disciplines. One level may implement a fairness mechanism that simply insulates all users from one another. A second level may provide a resource guarantee to particular users (or user groups).[11] A third level may implement complete TOS support mechanisms to fully exploit the benefit of statistical sharing in packet switching and allow each user to pay the minimal possible while receiving adequate service.

[10] Another analogy for priority pricing is the airline industry, in which you pay more for flexibility (i.e, making reservations with short notice, without restrictions for cancellation, and without restrictions on time of day and day of week) and certainty (standby pays less than reservations). To be efficient the airline scheme also requires some predictability—certain spaces are reserved for the higher-cost, last minute reservations. If predictions are not accurate, the seats will go unused or will be sold at lower standby rates.

[11] Suggested by D. Clark, private communication.

## 4.2 Accounting Granularity

There is a cost tradeoff associated with fine grain accounting. In general, finer granularity offers more accurate control at the expense of greater overhead to the system. Granularity decisions must be made regarding both the unit of traffic and user monitored.

### 4.2.1 Traffic Granularity

Some form of packet or byte counting is required to support the usage feedback policies described above. Counting individual packets and associating them with particular users/subscribers in very high speed networks presents technical challenges. However, high performance gateways often process and maintain state information in terms of source destination pairs for the purpose of route caching and possible queueing practices, in addition to possible access control [11]. Consequently, maintaining packet counts or statistics may only require a minimal incremental action.[12] Nevertheless, the packet count cannot be regenerated easily if lost, and additional mechanisms are needed to make the counter resilient to gateway crashes (e.g., written to disk or sent out over the network to some collection agent).

An alternative to exact packet counts is statistical accounting, i.e., taking samples of high speed traffic sources. However, the length of the sampling or averaging period must be matched to the dynamics of the network traffic.

### 4.2.2 User granularity

A related issue is the granularity with which traffic records will be built, i.e., the granularity of users. Transit ADs can identify traffic sources and sinks at different levels of granularity. In particular, the transit system can track individual end systems or may aggregate traffic counts for an entire AD. The telephone network tracks the particular end systems (i.e., one's telephone number). On the one hand, the overhead of accounting in transit networks could be reduced significantly if this were *not* the case. On the other hand, the accounting granularity and feedback are tightly coupled, in the sense that the feedback system cannot identify users in more detail than the accounting record. As stated earlier, some form of feedback signal must be provided to the traffic sources if behavior is to be affected. Nevertheless it is possible to provide performance or administrative feedback to end users and restrict monetary feedback to the AD level.[13]

In very large internet environments the cost of fine user granularity accounting may be prohibitive, and may be considered undesirable for policy reasons; for example, to prevent usage sensitive accounting data from being used for traffic analysis purpose (e.g. NSFnet policy).

## 4.3 Frequency of feedback

Another dimension of all feedback schemes is the frequency with which the information is collected and returned to the traffic source. Network management protocols can be used to collect aggregated statistics and return them to the traffic sources on a regular, but infrequent basis. In contrast, some feedback mechanisms are based on real time (minimal delay) information akin to congestion and flow control feedback.

Performance feedback channels provide feedback signals to the end user in real time. Whereas, administrative feedback is usually provided at a much lower frequency. Monetary feedback channels can be implemented at either rate but with significantly different implications for end-user behavior, e.g., a real time "meter" vs. a monthly bill. Feedback rate presents a clear tradeoff between the rapidity of user adaptation and the cost of realizing the scheme itself.

## 4.4 Cost Metrics

Whether the feedback channel is monetary or administrative, there is the question of cost metrics, i.e., the appropriate measure or metric for network cost recovery and feedback. The simplest metric is a function of the number of packets. If packets are not of equal size, however, some function of packets and bytes may be preferred. In addition, the number of gateway hops traveled may be a factor in the actual marginal cost of delivering a packet. In a best-effort-service internet (the IP Internet, for example), the metric used in the implicit performance feedback channel is only packet-count and distance related. If we introduce a TOS and/or priority mechanisms, the metric should be a function of packets, bytes, hops, TOS, and priority.

There are additional issues to consider when the feedback channel is monetary. In addition to the pricing problem (i.e., how to set prices appropriately), accounting rates and procedures must be agreed between the transit ADs that carry each others traffic (e.g., in order to allocate costs). Moreover, the interaction of end user feedback and inter-transit AD accounting can result in undesired externalities such as have occurred in the telephone network. For example, in telephony the long-haul carriers collect from end users and pay local telephone companies for use of their resources. In some

[12] Personal communications, David D. Clark.

[13] Another motivation for monetary feedback to the end systems themselves is the greater opportunity for transit carriers to dynamically compete for end-user traffic. Otherwise, stub ADs impose a single decision upon the aggregate traffic instead of allowing individual users within the AD to select.

cases this has contributed to bypass that is inefficient from a global perspective.

## 4.5 Dynamic Capacity

If a service provider routinely dials up additional capacity when the load exceeds a certain threshold, the service provider's costs may rise significantly during crowded periods and may need to be reflected back in user charges. In addition, the users will not experience sustained reduced performance at overload because the service provider will compensate by dialing up additional circuits. Therefore both the performance feedback and cost recovery assumptions are different in the dynamic resource-cost case. Since the natural performance feedback that one gets with a fixed capacity system is now absent (or diminished), monetary or administrative channels must be used.

It may not be fair to charge more only to the users whose demand instigates the dialing up of additional facilities (unless the desire is to implement a FCFS policy in which case late comers have a choice of paying more or waiting). Instead, one might like to inform all users of the network load, so that they can decide whether the additional cost (of dialing/switching in additional lines) is worthwhile. However, each user may face increased difficulty in accurately predicting the cost of a performance-sensitive transaction given the dependency on other users behavior. To avoid this, it may be preferable for the service provider to estimate the network load and the cost for dynamically added resources and to set a relatively stable price.

## 4.6 Identifying and Authenticating the Collection Points and Billable Entities

Another issue in a multi-transit AD internet (e.g. backbones and regionals are transits) is how transit ADs and end users will be identified, and how they will be authenticated.

Carriers may either feedback information to traffic sources directly (whatever the granularity, stub AD or end system), or charge the previous or next hop AD. In either case, once monetary or administrative mechanisms are in place, the incentive for fraud and the desire of traffic sources and carriers alike to detect and prevent fraud is great. It may be difficult to impersonate a neighbor AD because of the physical realization of AD connections. However, detecting fraudulent AD identifiers for ADs other than one's direct neighbors presents a problem.

Feedback schemes discussed rely on identification of the endpoint (end user or AD) or charge code. Applying cryptographic signatures and checks on a per packet basis to protect the integrity of these identifiers may introduce excessive performance overhead for many environments. One alternative is to use statistical mechanisms, e.g., to sign every Nth packet only [5]. A feedback channel could also be used to inform the charged entity of the (exact or statistical) usage value on a frequent basis so that the subscriber can check whether the charges are within range of their expectations (or the local measures). The communication of usage values could travel out of band of the data flow itself, but over the same network.

## 4.7 Coordination among Transit Carriers

Coordination is required among transit carriers with respect to both billing arrangements (in the case of monetary feedback channels) and TOS.

In the case of billing, the various transit hops along a source to destination path must agree to use mechanisms that, if not equivalent, are compatible. In addition to deciding the billing scheme for network end points, neighbor ADs must reconcile their accounts on a regular basis; similar to inter-carrier telephony accounting.

Billing schemes can vary in several dimensions:

- Who is billed and who is paid: source AD, destination AD, previous-hop AD, next-hop AD, or a third party.
- What unit of traffic is billed for: per packet, per byte, per connection, etc. (i.e., depending upon the monitoring granularity in use).
- The nature of the payment: dollars, funny-money, exchange of resources, etc. (i.e., depending upon the feedback channel in use).

TOS guarantees are useful to the source and destination only if they are supported by all the transit ADs along a path. For instance, congestion at some points on the path can lead to clumping or spreading of traffic that makes it difficult for ADs farther down the path to live up to their guarantees. In such cases, should the user still pay the full rate at those places that were not responsible for the congestion or degraded service? Billing may raise a similar issue; i.e., if a packet is dropped half way through to the destination, should the user be billed by the first AD hops in the path?

In summary, when an end-to-end service is carried out by joint effort of several ADs, the service provided by one AD will have an impact on the service interface and quality provided by the others. The nature of such interactions, and the interaction with usage accounting or charging, require further investigation.

### 4.8 Additional Stub-AD and End-System Requirements

Stub ADs must manage their connections to transit carriers. Consequently, they face several additional requirements. For example, in order to control their communications budgets, stub ADs must either be able to predict or bound variable costs, or they must be able to recover over-expenditures from end systems. It is desirable that stub ADs have available to them mechanisms with which to bound their variable costs as an alternative to passing back usage sensitive charges to end systems on the one hand, or restricting internet connectivity on the other. As discussed earlier, to the extent the goal is modified end user behavior, it is almost always desirable for signals to travel all the way to the individual traffic source. However, stub ADs should be able to select the particular type of feedback channel or mechanisms used internally.

Another requirement faced by stub ADs is the need for prevention and/or detection of fraud. For example authentication mechanisms can be employed to prevent unauthorized usage. Similarly, stub ADs will want to verify the actual service quality delivered if they incur TOS-based charges.

End systems require similar cost control and verification capabilities. In addition, end systems require instrumented applications that can assist users in developing communication cost expectations and that can translate low-level usage feedback signals into higher level units that are meaningful to the end user (e.g., cost per electronic mail message distribution instead of per packet).

## 5 Conclusions

At its best, resource usage accounting and feedback presents an opportunity to promote efficient usage of network resources, and to reduce end-user communication costs by setting charges that reflect the statistical resource sharing possible with packet switching architectures. Design, simulation, and experimental research is needed to develop appropriate technical mechanisms to realize these benefits, and avoid the many negative behavioral and technical consequences of poorly designed approaches. This paper represents a first attempt to articulate the design space of usage feedback mechanisms.

Before concluding we reiterate a few caveats and recommendations with respect to resource usage feedback in internets. First, network administrators should avoid charging end users on an usage-basis without understanding users' demand elasticity, the impacts of the charging policy chosen, and the technical overhead of doing so. Secondly, stub ADs and end users should be urged to develop tools necessary to manage their communication budgets before usage sensitive charges to stub ADs are introduced widely. Moreover, effective TOS support mechanisms should be employed in the network to fully exploit the benefits of statistical resource sharing. Finally, in a network environment that supports multiple TOS, it may prove most beneficial to introduce usage feedback first for the most demanding applications (performance sensitive) only – it is likely to have a significant impact on network usage efficiency and at the same time will more likely impact users that can afford the expense.

We conclude with a brief enumeration of recommended research and development tasks.

### 5.1 Future Research and Development Work

Below we enumerate some concrete development and research tasks that could be pursued to gain a better understanding of the issues discussed in this paper:

- Packet counting experiments
  Instrument internet gateways to count packets and associate them with host/network pairs, as an investigation into performance overhead.

- Instrument communication oriented applications
  As described earlier there is a need to better model user behavior and the communication costs of user level applications. File transfer and remote login are relatively straightforward applications in this respect. On the other hand, electronic mail and conferencing are less straight forward to instrument because of their group/multicast nature.

- Transport protocol meters
  These meters can be one of the tools that help the end users to understand the network usage of various applications as well as the efficiency of the protocol implementations. Meters can also be used to verify the usage counts received from networks.

- Survey needs
  Informal discussion with network administrators to define needs for stub AD procedures and mechanisms needed to protect liability/costs without adversely inhibiting end user connectivity.

- Priority pricing/TOS Simulations
  Explore priority pricing through simulation, with

a comparison to flat per packet charging. Investigate the interaction with existing and proposed TOS and network service disciplines (e.g., congestion control mechanisms). Investigate granularity, feedback rate, and feedback channel design tradeoffs and interactions.

- Fraud detection protocol
  Investigate alternatives for fraud detection protocols. Investigate use of existing communication channels between end point and transit nodes (e.g., congestion control signals, policy routing related communication [12]).

- Multicast
  Many applications are of a multicast nature, e.g., teleconferencing and information distribution lists. To the extent explicit multicast support leads to more efficient delivery of multi-destination messages, the introduction of usage accounting and feedback could motivate additional multicast support and use. However, one characteristic of multicast is that the sender rarely knows how many destinations will receive the message, nor where those destinations are located, nor the intermediate path. Consequently, users might face unknown charges when sending multicast traffic if packets are accounted for on a per-packet basis. Because multicast implementation is a globally more efficient approach than multiple, directed unicasts, it would be very unfortunate if its use was discouraged because users did not know what to expect. Much work is needed to investigate alternative accounting and feedback mechanisms in the light of multicast.[14]

# 6 Acknowledgments

We are very grateful to Bob Braden, Vint Cerf, Alex McKenzie, Yakov Rekhter, and Scott Shenker for comments provided on an earlier draft.

# References

[1] Amer, P., et. al. *Local Area Broadcast Network Measurement: Traffic Characterization,* **University of Delaware Tech report 86-12**, January 1986.

[2] Boiteux, M. *Peak Load Pricing,* **Journal of Business**, 1980, volume 33, p 157-179.

[3] D. Clark, *Policy Routing in Internet Protocols,* **RFC 1102, SRI Network Information Center**, May 1989.

[4] A. Demers et al, *Analysis and Simulation of a Fair Queueing Algorithm,* **Proceedings of ACM SIGCOMM**, September, 1989.

[5] D. Estrin and G. Tsudik, *Security Issues in Policy Routing,* **Proceedings of 1989 IEEE Symposium on Security and Privacy**, May 1989.

[6] D. Estrin, J. Mogul, G. Tsudik, *Visa Protocols for Controlling Inter-Organizational Datagram Flow,* **IEEE Journal on Selected Areas in Communications**, May 1989.

[7] D. Estrin, *Policy Requirements for Inter Administrative Domain Routing,* **RFC 1125, SRI Network Information Center**, November 1989.

[8] D. Feldmeier, *Empirical Analysis of a Token Ring Network,* **Massachusetts Institute of Technology Technical Report MIT-LCS-TM254**, January 1984

[9] *The Tragedy of the Commons,* **Science**, December 13, 1968, volume 162, p 1243-1248.

[10] S. Hares, D. Katz *Administrative Domains and Routing Domains, a Model for Routing in the Internet,* **RFC 1136, SRI Network Information Center**, December 1989.

[11] S. Heimlich, *Traffic Characterization of the NSFNET National Backbone,* **Proceedings, USENIX, January 1990.**

[12] M. Lepp and M. Steenstrup. *An Architecture for Inter-Domain Policy Routing* **DRAFT RFC**.

[13] ISO *OSI Routeing Framework,* **ISO/TF 9575**, 1989.

[14] D. Piscitello and M. Kramer, *Internetworking Using Switched Multi-Megabit Data Service in TCP/IP Environments,* **Computer Communication Review**, Vol. 20, No.3, July 1990.

[15] J. Postel, *Internet Protocol,* **RFC 791, SRI Network Information Center**, September 1981.

[16] J. Postel, *Internet Control Message Protocol,* **RFC 792, SRI Network Information Center**, September 1981.

[17] R. Wilson, *Efficient and Competitive Rationing,* **Econometrica**, January 1989, volume 57, no. 1, p 1-40.

[14] Private communications with Steve Deering, Xerox PARC.