

Network analysis issues for a public Internet

Hans-Werner Braun and K. Claffy
 hwb@sdsdsc.edu, kc@sdsdsc.edu
 Applied Network Research Group
 San Diego Supercomputer Center
 San Diego, CA 92186-9784

Abstract

While initially conceived as a dedicated communications facility for the United States federal government, today's Internet aggregates traffic from among a far wider set of constituencies. Pooling resources of so many constituents into a massively interconnected environment raises the issue of resource and cost allocation. In this paper we describe the importance of network analysis in support of resource attribution and evaluate a number of examples. We offer evidence to support our hypothesis that, in the face of the current evolution of global information infrastructure, vastly expanding both in ubiquity and sophistication of applications, Internet policy considerations and network analysis must begin to interact in ways not previously recognized or implemented. In particular, as the scale of, access to, commercialization of the Internet broadens, cost allocation among (even globally) shared resources will require the development of new accounting and billing models to accommodate the wide range of players and services.

1 Introduction

While initially conceived as a dedicated communications research facility for the United States federal government, today's Internet aggregates traffic from a far wider set of constituencies. As the number of client networks of the Internet heads into the tens of thousands, with millions of users world-wide, the image of a ubiquitous network, relying on globally shared resources, has already become a reality.

A key characteristic of the Internet is the role of the constituent networks. These networks are not simply clients which pay for a service from a transit provider, but rather integrated entities which actively contribute network resources. These resources range from vast national and international backbones to regional transmission services and even local network service within individual campuses and companies, many of which are themselves multi-million dollar institutions.

Pooling resources of so many constituents into a massively interconnected environment raises the issue of resource and cost allocation. In the early days of the Internet when one or a few US government agencies assumed the financial burden of building and maintaining the infrastructure, there was little controversy over proportioning of costs. However as the number of constituencies, including federal, academic, and commercial entities, increases on a global scale, equitable resource allocation dominates many discussions of Internet development. Usage policy considerations complicate the discussions further.

Cost allocation and policy considerations in the Internet require models different from those used by phone companies in the past, where end-users pay their service provider directly, and service providers use among themselves a settlement process that is largely transparent to the end-user. Impediments to using such a model, for example in the U.S. portion of the Internet, include the current funding framework, where major government agencies fund significant fractions of the infrastructure based on often abstract goals, such as fostering scientific research. Many times these goals in turn impose specific criteria for transmitted traffic, resulting in Acceptable Usage Policies (AUPs) for the network. An example is the NSFNET backbone¹, a major core switching fabric that aggregates traffic from a vast set of clients. The United States National Science Foundation (NSF) pays for this network, in line with its objective to foster research and education. In turn the NSF requests that traffic crossing the backbone conform to its AUP, which essentially restricts the network to traffic in support of NSF programmatic requirements.

Other U.S. federal agencies provide even more restricted network services, e.g., NASA, DOE and DoD all run their own dedicated agency networks in direct support of their individual missions. Other organizations, such as commercial entities within the US or the pan-European EBONE network, provide unrestricted transmission ser-

¹The "NSFNET backbone" now refers to a virtual backbone service, i.e., a set of services provided across the ANSnet physical backbone. In this paper we refer to the "T3 NSFNET backbone" with the understanding that we are referring to a service provided to NSF, not a dedicated NSFNET infrastructure.

vice for any legal traffic from any paying customer.

With today's large number of service providers, who in fact compete with one another, the ad hoc interconnection approach used thus far² has begun to break down. For example, during the recent establishment of a major multi-service-provider interconnection facility, or Network Access Point (NAP), some service providers, trying to protect their own assets and marketing opportunities, have refused to connect without more clearly articulated interconnection policies. The problems arising out of the NAP environment hinder the conceptualization and implementation of even more extensive international interconnection points (Global Interexchanges (GIX)).

The *Reach-As-Far-and-As-Fast-As-You-Can* paradigm sufficed for the initial phase of the Internet, but as the network matures into a community of strictly operational and often commercial service providers, we must consider how the Internet differs from the telecommunications industry: individual bandwidth demand is constantly and rapidly increasing; an increasing number of service providers must cooperate to aggregate resources; and finally, those using the network as end-users are often also the ones developing the multi-protocol technologies to advance it, and they want to see those newly developed technologies deployed far sooner than traditional telecommunications carriers have ever had to imagine. Combined with the demand for ever increasing bandwidth, predictability, and ubiquity, the resulting environment requires rapid adaptation to new technologies and user needs, and must compensate for an ever-increasing base of constituents. The developing complexity of the Internet system renders imperative the clear definition of network policies in crisp, implementable, and verifiable terms, if there is to be any chance for their applicability to today's environment.

In this paper we describe the importance of network analysis in support of these policy considerations and evaluate a number of examples. We offer evidence to support our hypothesis that, in the face of the current evolution of global information infrastructure, vastly expanding both in ubiquity and sophistication of applications, Internet policy considerations and network analysis must begin to interact in ways not previously recognized or implemented. In particular, as the scale of, access to, commercialization of the Internet broadens, cost allocation among (even globally) shared resources will require the development of new models to accommodate the wide range of players and services.

2 The policy space of the current Internet environment

The United States component of the Internet currently consists of a three-level hierarchical architecture of national agency backbones, attached mid-level networks,³ and connected local sites. Similar architectures have evolved in other areas of the globe, perhaps most visibly in Europe, where the EBONE pan-European backbone supports communication among participating countries.

Figure 1 depicts several logical levels of interest to the U.S. portion of the Internet community.⁴ Components at each layer are typically operated and managed by autonomous organizations, each with their own rules and policies for the usage of their network. The collection of these autonomous entities within the structure of the global networking environment defines a policy space for the Internet, with policy boundaries typically at the interfaces between component networks on the same or different layers. While Figure 1 constitutes an abstract illustration of the interconnectivity, the actual implementation of all the connections forms a much more complex framework.

3 Aggregation granularity

Since a core focus of any network policy is the flow of traffic, it is critical to develop a common model of flow definitions. At one extreme, such a model may describe a flow matrix among countries participating in the Internet, and the impact of such flows on major constituent networks such as the NSFNET backbone. At the other extreme, one may attribute network usage to individual users, applications of the user, or even some more abstract context definition (e.g., transmission of a high volume packet video stream). Other granularities of service aggregation between those extremes include traffic flows by multibackbone environment (e.g., of different agencies), single backbone at large, backbone node, external interface of a backbone node, backbone client service provider, Administrative Domain, IP network number, and individual hosts. These granularities do not have an inherent order, as a single user or application might straddle several hosts, network numbers, or other aggregation mechanisms. There is no inherently best granularity to use for network analysis; the appropriate selection depends on the question of interest. However, as the complexity of such possible questions continues to grow, the ability to account for fine-grained flows, especially for real-time

³Mid-level networks have also been called "regionals", reflecting their geographical span, but we will use the term "mid-level" to reflect its hierarchical position in the architecture.

⁴Internet interconnectivity is evolving in different ways in different areas of the world.

²Mandelbaum, in [7] also refers to this as the "throw me a line" approach.




Figure 1: Model of U.S. Internet interconnectivity architecture

needs, easily exceeds the capabilities of available Internet technology.

The issue of granularity plays perhaps its most critical role with respect to implementing mechanisms for cost allocation and accounting. As accounting matures, it may eventually be used for billing purposes, at which time the developed accounting models must even more accurately collate network usage at whatever level of aggregation the billing mechanisms use.

Prerequisite to equitable cost allocation and accounting is a secure mechanism for attribution of resource consumption, an historically difficult task in globally shared datagram infrastructures. Wide area network infrastructures are typically strongly focused on the real time operational and near term engineering requirements to keep the fabric alive, while ensuring short to medium term evolution. As a result, operationally collected statistics are generally geared toward day-to-day operations and management, such as indicators of real-time utilization and outages. Collected statistics also often allow near term network engineering based on network capacity and utilization. However, as the Internet grows in geographic and functional scope, the requirements for statistics reporting grow more complex, and the Internet community must assume a proactive role in defining an appropriate

structure for information pertaining to resource consumption.

For example, from the perspective of a service provider, attributing Internet usage to individual users is not feasible with current technology. The underlying datagram service, as well as the heavy aggregation of many users via multiple service providers, would hinder most service providers from being able to attribute resource consumption to a user, much less a user in conjunction with an executing network application.

Since most IP networks receive connectivity to the Internet via intermediate service providers, an obvious alternative is a hierarchical model of attribution, where higher level providers can attribute resources to intermediate providers, who can in turn re-attribute resource consumption among their clients. Some special cases of provider/client accounting may be amenable to perhaps the simplest accounting model: aggregated packet/byte flow counters at service interfaces. This model assumes that a simple volume expression is a sufficient definition of traffic exchange, and typically requires that the client perform sub-accounting within its own area.

However most situations are not so simple. Often attribution to service providers will require measuring traf-

fic volume not just as total packets traversing an interface, but according to the source and/or destination, as well as type, of each packet. Since current network instrumentation for collecting such traffic matrices on the NSFNET backbone supports only the granularity of individual clients (identified by IP network numbers), one is limited to this or a coarser granularity. For example one could group multiple IP network numbers into their associated *Administrative Domain*.

Furthermore, for performance reasons many wide area network infrastructures must rely on sampling mechanisms to determine traffic flows. The NSFNET T3 backbone is an example of such an infrastructure; in Section 4 we discuss the impact of sampling on their flow assessment capabilities.

4 Effect of Sampling on Assessing Accuracy of Traffic Locality

The National Science Foundation requires that its backbone service provider furnish an account of monthly source-destination matrices based on IP network numbers using the NSFNET. These matrices describe how much traffic, in bytes and packets, each IP network sends to other IP network destinations.

However, for performance reasons the T3 backbone routers only support this statistics collection by sampling every 50th packet. Thus these net-to-net matrices, collected and stored in fifteen minute increments, will be incomplete. Achieving accurate network number matrices via sampling poses some difficulty, since the dispersion is high: the number of networks exchanging traffic via NSFNET was as of May 1993 over 12,000 and rapidly growing. Particularly for a fifteen minute interval, this sampling has a detrimental effect on the integrity of a typical net-to-net matrix. In this section we present a few statistics on the impact of sampling for a source-destination matrix.

The sampled net-to-net matrix for December 1992 accounted for communication between 1378065 site pairs. For 281680 (19.8%) of these net pairs, the sampling mechanism only captured 1 packet. There is no way to know whether that pair exchanged only 1 packet, which just happened to be a in the 2% that were sampled, or 1 million packets, of which only 1 was sampled.

It is impossible to compare this sampled matrix to the actual traffic flows to verify its integrity, since the network does not support complete collection. In order to assess the impact of sampling, we therefore had to collect a dedicated packet trace in a similar environment. We chose a single interface into one of the T3 E-NSS backbone nodes, specifically the one located in San Diego. We were moti-

vated for this selection by three reasons. First, we selected a site which would be representative of a wide-area environment, in terms of traffic intensity and aggregation, in order to have a realistic sampling scenario.

Second, an E-NSS interface is a natural location for accounting by other means, i.e., non-sampled SNMP packet/byte counters, if these metrics were deemed sufficient for billing purposes. Unfortunately, accounting and billing according to traffic source and/or destination will likely require more sophisticated statistics objects, i.e., those which in the current NSFNET infrastructure are only possible via sampling. Therefore it would be useful to assess the accuracy of objects built at these interface points.

The third reason is somewhat pragmatic. The E-NSS locations are convenient since they constitute the backbone perimeter and are attached to LANs where statistics collection is more straightforward.

4.1 Time granularity: day vs. month

At our selected interface into the T3 backbone cloud we collected a 24-hour trace on 22 March 1993, and simulated a variety of sampling granularities on this trace. The trace resulted in more than 650MB of data for one 24-hour period, making the collection method suitable for this research investigation, but not for operational day-to-day measurements. However, the disk space limitation prevents us from collecting longer traces, which would be more useful if desired billing periods were monthly rather than daily.

4.2 Space granularity: end entities versus pairs

The NSF requirement for net pair matrices was based on prior experiences with the 56kbps NSFNET backbone, when network planners thought the resulting objects might be useful for general insight into network behavior. Since there was no guarantee for accuracy of the statistics collectors, the backbone cooperative agreement also included a clause that the net-net matrix not be used for accounting purposes. From these matrices the NSFNET project makes publicly available a collapsed data set which attributes traffic by source and destination network numbers, i.e., the marginal totals of the net-to-net matrix.

For some accounting and billing models such entry and exit traffic attribution may be sufficient. One would only need the pair-wise attribution if both the source and destination were required to charge one of the endpoints, e.g., based on distance, similar to long-distance phone networks. Since the sampling inaccuracies dis-

cussed above are greater for data objects that spread out over many buckets, the net pair matrix with 144 million (the square of the number of communicating networks) entries is much harder to accurately assess than for example sampling on a linear vector of just sources and destinations for IP traffic. The number of source and destination buckets grows only linearly, rather than quadratically, with newly observed network numbers.

4.3 Metrics indicating integrity of sampling

In this section we offer metrics comparing the sampled matrix to that of the larger population.

4.3.1 “Hit” pairs

One metric indicating the gap between the flow matrices of the sampled vs. parent population of packets is a simple counter of the number of communicating site pairs accounted for during that interval. During a fifteen minute interval, thousands of packets might typically traverse a typical backbone NSS. Every fiftieth of these packets will provide only 0.2% of that data from which to build a net-to-net matrix.

As mentioned previously, aggregation in time, e.g., over a month or so, or in space, e.g., by country or autonomous system, may mitigate the effect of this sampling. For example, Figure 2 plots the number of net pairs “hit” by sampling as a function of elapsed time of the sampling mechanism. The number of net pairs seems to level off after around 30 minutes of sampling, at around 50% of the total communicating net pairs for the day. This graph also includes the “hit” ratios for the sources and destinations networks. As one would expect, the performance is better for the granularity of network numbers than network pairs. Whereas the captured net pairs leveled off after around 30 minutes of sampling, at around 50% of the total communicating net pairs for the day, the number of captured networks seems to reach almost 70% of the total communicating networks before leveling off.

4.3.2 Error metric

The two graphs in Figure 3 plot one example error metric for sampling net-to-net traffic. The top graph plots for the top 15% net-pairs (in terms of traffic volume exchanged for the day) the ratio of the number of sampled packets multiplied by 50 over the true number of packets, exchanged for the 24-hour interval. The x-axis is the volume of packets exchanged for the day. There are 2500 points in this plot, for the 15% most communicative of the approximately 17000 net pairs who exchanged traffic that day.

The lower graph plots for the same 2500 net pairs the same metric for bytes; the x-axis in this graph is the volume of bytes exchanged for the day. These graphs indicate that the net pairs who exchange more than, for example, 20,000 packets during the 24-hour period, are sampled with less than 5% inaccuracy. Bytes are somewhat less accurately assessed, consistent with the greater possible number of possibilities (1500: 1 to 1500 bytes) than there are with packets (2: packet or no packet).

The two graphs in Figure 4 plot the same error metric plotted in Figure 3 for source and destination networks rather than net pairs. The top graph plots for the top 15% source or destination networks (in terms of traffic volume sourced or received for the day) the ratio of the number of sampled packets multiplied by 50 over the true number of packets, exchanged for the 24-hour interval. The x-axis is the volume of packets sourced or received for the day. There are 440 points in this plot, for the 15% most communicative of the approximately 3000 networks who either sent or received a packet that was sampled for that 24-hour period. The lower graph plots the same metric for the same networks for bytes; the x-axis in this graph is the volume of bytes sourced or destined for the day. These graphs indicate that the networks who send or receive more than, for example, 20,000 packets during the 24-hour period, are sampled with less than 5% inaccuracy.

Note that these two pairs of graphs account for the same fraction of data across the backbone (for both networks and network pairs, the top 15% account for approximately 97% of the byte volume). However since the lower graphs have fewer points (buckets) among which to distribute the data, they depict more accurate assessments.

4.4 Implications of sampling investigation

We have discussed some difficulties of wide-area network infrastructures that may have to rely on sampling methods in both the short term, for billing, and in the long-term, for determination of which network locations require upgrade. If applied to billing, sampling methodologies must consider the tradeoff between granularity of accounting entity and the accuracy of traffic attribution. Accurate accounting by source or destination is more feasible than by network pair. Capacity planning objectives may require knowledge only of major network flows imposed on the infrastructure (hotspots, or “hotflows”) for future upgrades or improved design. For this purpose only the high volume entries of the network pair matrices are relevant, and thus network pair matrices are still important objects to maintain.

This discussion constitutes only a preliminary investigation into the effects of sampling, but provides a beginning to wide-area network administrators on how to best sup-




Figure 2: Number of net pairs “hit” by sampling (22 March 1993 24-hour data trace into SD E-NSS)





Figure 3: Sampling error metrics for top 15% (2,000) of net-net pairs (22 March 1993 data trace into SD E-NSS)





Figure 4: Sampling error metrics for top 15% of networks (22 March 1993 data trace into SD E-NSS)

port, in this case, sampled statistical objects which track network usage.

5 Assessment of international flows

As mentioned in the Section 4, aggregation in time, e.g., over a month of so, or in space, e.g., by country or autonomous system, may mitigate the effect of sampled statistics. We present in this section some examples of statistics on international traffic flows across the NSFNET backbone for a week in February 1993. These statistics are of particular interests to national or international policy makers who want to attribute of resource consumption to individual countries for evaluation of cost-sharing models. Figure 5 presents a matrix of traffic volume exchanged by country during the first week of February 1993. We use the operationally collected data sets for the NSFNET backbone, which include source-destination matrices by network numbers, to create this matrix. We exclude the United States from this graph, as those values dwarf the values of the other countries. Table 1 provides figures for relative proportions of traffic by country.

The operationally collected data sets also allow one to explore aspects of the data such as those in Table 1, which shows for February 1993 the directional asymmetries in traffic volume; average packet size by country; and skewness of distributions through time. The sixth column in Table 1 provides an indication of the asymmetry with which countries utilize the backbone; this column measures for each country the ratio of bytes received from the backbone to the number of bytes that country sent into the backbone. Figure 6 plots these ratios, along with the traffic volume each country sources into the backbone, for this first week of February 1993. The other graphs in this section also reflect the same one week time window.

Table 1 also provides one example performance characteristic related to the asymmetry in traffic volume discussed above: the distribution of packet sizes among countries, which provides a measure of indication of the payload per packet each country is getting from the network. The last three columns in this table show the average packet size (in bytes) used by each country into and out of the backbone, and the ratio of the two values, for the month of February 1993. Most countries have an average packet size into the backbone of under 90 bytes, while the average sizes of packets from the backbone to non-U.S. destination countries is substantially larger. We interpret this to mean that these countries are likely requesting bulk traffic from U.S. sites.

Another point of interest is the significantly higher payload which some European countries are receiving from their NSFNET outbound traffic. In particular, Luxembourg's average packet size into the backbone is 41 bytes

and its average packet size out of the backbone is 514 bytes (almost twice the number two country of Korea)! Of course the traffic volume of Luxembourg, and many other countries, is relatively low, as is the number of IP network numbers (Luxembourg has only four IP network numbers). Germany is notable for having a relatively large number of networks with rather larger outbound NSFNET packets; these packets are more efficient in terms of per packet payload. There are also a few countries who are sending traffic to the backbone via large packets; we assume the top networks in that category are major FTP data sources.

We can also use currently collected data to explore traffic shifts between the U.S. and specific countries via the NSFNET backbone. NSF already had repeated occasions where they needed such analyses of traffic volume exchanged among countries, often to address policy and funding related questions relative to global interconnectivity. Using the same one-week window in February, Figure 7 shows the bidirectional flow of traffic between the U.S. and three countries in different time zones. The impact of the time zones, in this case in Japan, Mexico and Great Britain, is quite visible in relationship to the flows of traffic volume, where the traffic peaks tend to coincide with the business hours of the particular country.

Figure 8 depicts the directional ratios of traffic volume with other countries, as seen relative to the NSFNET backbone. Over the seven day period almost all countries receive more bytes from the United States than vice versa, though the discrepancies vary dramatically by individual country. The data indicate that the this asymmetry tendency is a long term effect; at shorter time periods, for example by two-hour intervals, the data demonstrated periods where the traffic flow into the U.S. is higher.

Figure 9 is an NSFNET backbone centric illustration of countries using the U.S. for their own domestic communications, both in terms of absolute volume, as well as in relationship to the overall traffic those countries exchanged with the NSFNET. This effect typically derives from multiple connections between some country and the U.S., and is at times being addressed on a case-by-case basis by the constituents of the connections.

Such attribution of international traffic flows is rapidly becoming an important issue, as the mechanism of splitting the costs evenly between the two end-point countries of a connection breaks down. Several recent international connection scenarios have required the reevaluation of this current model of interconnection. Since all international networking resources contribute to the quality of the global Internet, including the emergence of major international data base servers, better instrumentation will be necessary to assess the service qualities and network impact of such resources.

Table 1: Traffic to and from NSFNET backbone per country for February 1993

country	country code	existing networks	% of total bytes into NSFNET	% of total bytes from NSFNET	bytes ratio from/to bb	mean pkt sz inbound	mean pkt sz out-bound	ratio to/from NSFNET
United States	US	4170	90.89	80.93	0.89	195	178	0.91
Canada	CA	289	1.64	4.51	2.76	110	276	2.51
United Kingdom	GB	214	0.64	2.01	3.12	112	254	2.27
Australia	AU	171	0.88	1.19	1.35	172	238	1.38
Germany	DE	297	0.71	1.89	2.68	151	324	2.15
Sweden	SE	67	0.60	1.02	1.69	153	193	1.26
Switzerland	CH	58	0.77	0.75	0.97	201	190	0.95
France	FR	291	0.73	1.17	1.59	230	276	1.20
Finland	FI	59	0.79	0.50	0.63	257	138	0.54
Netherlands	NL	96	0.54	0.70	1.31	180	258	1.43
Taiwan	TW	73	0.23	0.58	2.49	121	250	2.06
Norway	NO	38	0.20	0.53	2.65	105	221	2.10
Italy	IT	116	0.18	0.67	3.73	96	309	3.20
Japan	JP	189	0.24	0.46	1.92	145	262	1.81
Austria	AT	59	0.13	0.41	3.24	103	279	2.72
Mexico	MX	19	0.07	0.21	2.77	78	196	2.51
Denmark	DK	7	0.28	0.27	0.93	313	213	0.68
Singapore	SG	16	0.06	0.33	5.42	75	329	4.38
Israel	IL	22	0.07	0.30	4.51	96	303	3.15
Hong Kong	HK	8	0.04	0.29	7.99	60	349	5.83
Korea	KR	30	0.04	0.24	5.83	84	355	4.22
Spain	ES	29	0.03	0.13	4.47	84	322	3.84
New Zealand	NZ	38	0.02	0.10	4.29	76	304	4.00
Brazil	BR	38	0.02	0.10	5.27	70	290	4.15
Belgium	BE	11	0.03	0.11	3.64	116	313	2.70
South Africa	ZA	32	0.03	0.11	3.53	123	320	2.61
Czechoslovakia	CS	35	0.02	0.09	4.50	78	341	4.36
Chile	CL	9	0.02	0.06	2.62	103	253	2.46
Puerto Rico	PR	3	0.02	0.03	1.94	80	171	2.15
Ireland	IE	16	0.01	0.06	5.11	78	273	3.51
Poland	PL	19	0.01	0.04	4.81	67	244	3.62
Portugal	PT	26	0.02	0.04	2.48	152	284	1.87
Greece	GR	11	0.01	0.04	5.88	71	188	2.64
Hungary	HU	8	0.01	0.02	3.17	80	262	3.26
Venezuela	VE	5	0.00	0.02	3.43	73	194	2.65
Iceland	IS	5	0.01	0.01	2.11	109	184	1.69
Slovenia	SI	6	0.00	0.01	5.77	56	305	5.46
India	IN	2	0.00	0.01	5.27	57	112	1.96
Thailand	TH	3	0.00	0.01	2.65	62	178	2.85
Luxembourg	LX	4	0.00	0.02	17.62	42	514	12.31
Argentina	AR	1	0.00	0.00	2.13	77	131	1.70
Estonia	EE	3	0.00	0.01	7.39	63	294	4.69
Malaysia	MY	3	0.00	0.00	5.42	66	318	4.81
Ecuador	EC	10	0.00	0.00	3.96	66	203	3.08
Croatia	HR	2	0.00	0.00	2.70	70	141	2.03
Tunisia	TN	1	0.00	0.00	2.78	74	196	2.64
Latvia	LV	1	0.00	0.00	5.13	55	194	3.56
Cyprus	CY	6	0	0	5.03	61	184	3.00
Kuwait	KW	1	0	0	3.15	53	114	2.15
Costa Rica	CR	1	0	0	12.62	58	90	1.56
Turkey	TR	5	0	0	3.46	276	159	0.58
Cameroon	CM	1	0	0	NA	NA	40	NA

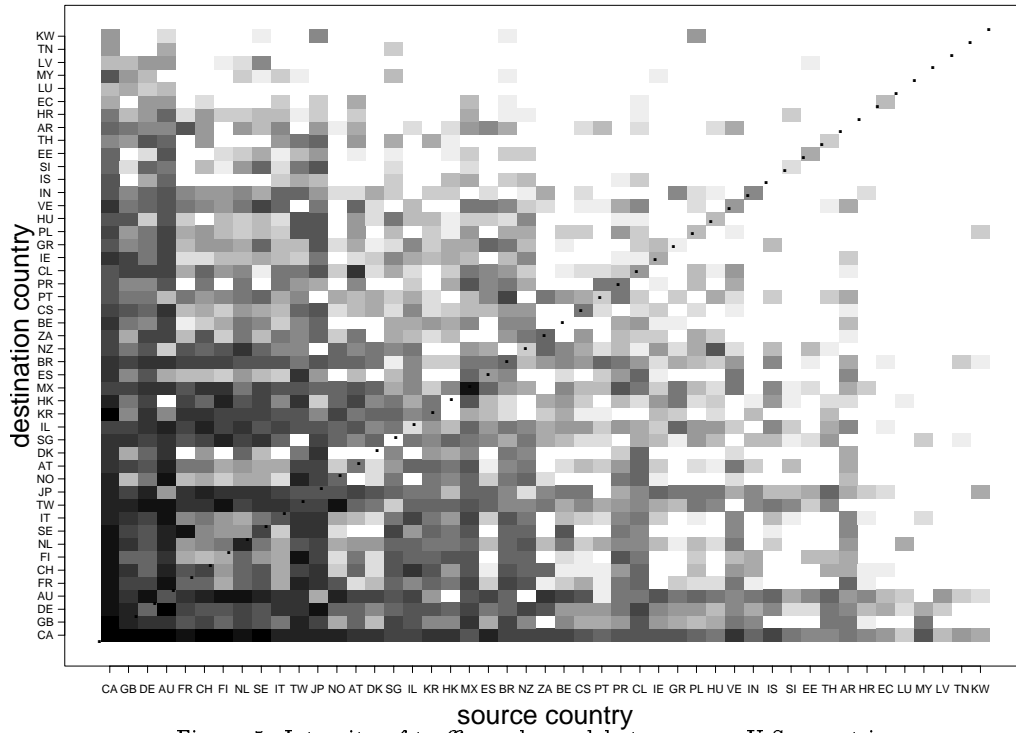


Figure 5: Intensity of traffic exchanged between non-U.S. countries

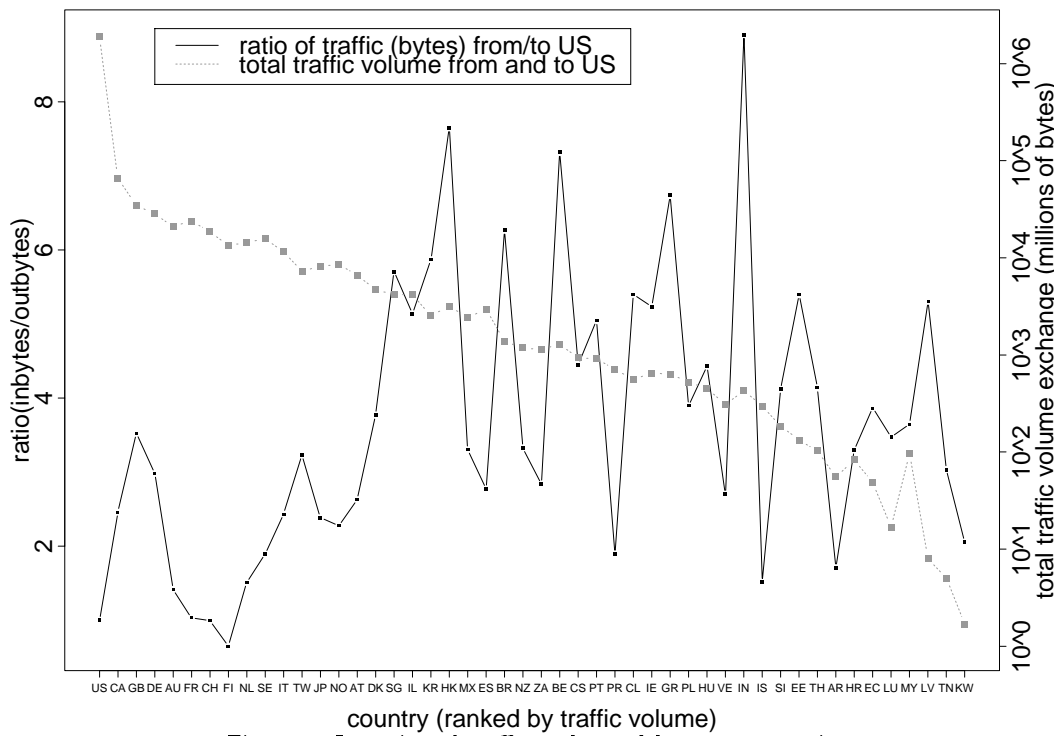


Figure 6: Intensity of traffic exchanged between countries




Figure 7: Traffic exchanged from Japan, Mexico, and Great Britain to United States

6 Application diversity

A further complication of flow attribution involves the increasing variety of network applications. A reasonable model of flow attribution among specific sites must transcend gross flows, conditioning the attribution on the nature of the service carried. One may want to assign (financial, political, etc.) responsibility for file transfer traffic volume to the destination site, while at the same time assign responsibility for electronic mail to the source site (analogous to the postal service).⁵ Unfortunately, currently collected data does not allow such simultaneous attribution of traffic type and geographic distribution. Claffy *et al.* [1] provides a description of the limitations of the current methodology used to track the traffic cross-section. In this section we provide only a brief summary.

The majority of applications on the NSFNET are built on top of the Transmission Control Protocol (TCP), and a few on top of the User Datagram Protocol (UDP). Both TCP and UDP packets use *port numbers* to identify the Internet application that each packet supports. Each TCP or UDP header has two fields for the 16 bit values identifying the source and destination ports of the

⁵The U.S. infrastructure provides an interesting example in the flow attribution context, where especially the NSFNET backbone network functions as a switching hub among many countries. While the reachable countries typically exchange the bulk of their NSFNET traffic with the U.S., a large fraction often goes to other countries, via the U.S., as well.

packet. Originally, ISI (Information Sciences Institute, University of Southern California), on behalf of DARPA, administered a space of 1 to 255 as the group of “Well Known Port” numbers, reserved for specific applications. For example, Telnet received port assignment 23. To open a Telnet connection to a remote machine, the packet carries the destination IP address of that machine in its destination IP address field, and the value of 23 in the destination port field. (In the case of Telnet, the packet uses some arbitrarily assigned source port that has significance only to the originating host. Often these “return address ports” have values greater than 1000.)

Although ISI administers the number range for Well Known Port numbers, at some point Unix developers injected a bit of anarchy into the system by unilaterally assuming that numbers below 1024 identify specific applications. They then began to use that numbering space as they deployed applications, such as port 513 for rlogin. Eventually network users began to use numbers even above 1024 to specify further services, extending the lack of coordination further. Examples include XDR/NFS (port 2049), and X-Windows (port 6000+), and port 4444 for (some) MBONE video multicasts.

Port numbers are the only mechanism via which the NSFNET can monitor statistics on the distribution of applications on the backbone. Thus the proliferation of uncoordinated number assignments imposes ambiguity into this categorization of packets by application.

For NSFNET statistics gathering on port distribution for




Figure 9: Traffic from countries to themselves through the NSFNET backbone




Figure 8: Ratio of traffic volume from/to the U.S. through the NSFNET backbone

the backbone, Merit (and now ANS) specifically monitors ports in the ranges 0-1023, 2049 (for NFS) and 6000-6003 (for X-window traffic). Merit/ANS categorizes packets into these ports if either the source or destination port in a given packet matched one of these numbers. However even within this range not all ports have a generally known assignment, so packets using such undefined ports go into an *unknown* port category. The Merit monthly statistics report for February 1993 states:

3) Services with TCP/UDP port numbers 0-1023, 2049, 6000-6003 and 6667 are tracked individually. Service names are given, as documented in RFC 1340, but otherwise are labeled "(unknown)". All other TCP/UDP ports are grouped into the single category "(other_tcp/udp_ports)".

We can derive the following four categories from this measurement methodology:

1. source or destination port < 1023 (or 2049 or 6000-6003) and known
2. source or destination port < 1023 (or 2049 or 6000-6003) but unknown
3. neither source nor destination in 0...1023, 2048 or 6000...6003
4. Non TCP or UDP protocol (i.e., no port number in use)

The first category results in a defined assignment; the second in a defined numerical port number, but no port definition; the third in unknown ports; and the fourth in unknown protocols.

Figures 10 and 11 uses this collected data to categorize the proportion of traffic on the network by category since August 1989. The categories in these figures correspond to:⁶

- File exchange: ftp data and control (tcp ports 20, 21)
- Mail: smtp, nntp, vmnet, uucp (tcp ports 25, 119, 175, 540)
- Interactive: telnet, finger, who, login (tcp ports 23, 79, 513, udp port 513)

⁶Note that Merit began to use sampling for this collection on the backbone in September 1991. In November 1991 traffic migration to the T3 backbone began; the majority had migrated by May 1992 and in November 1992 the T1 backbone was dismantled. For June to October 1992 no data was available for either the T1 or T3 backbones.

- Name lookup/dns: (udp port 53, tcp port 53)
- Other TCP/UDP services all tcp/udp ports not included above (e.g. irc, talk, X-windows, appletalk)
- Non-TCP/UDP services Internet protocols other than tcp or udp (e.g. icmp, igmp, egp, hmp, ax.25, etc.)

Figure 10 illustrates the difficulty of tracking changes in the cross-section of traffic on the backbone. The decomposition of flows in Figures 10 and 11 reflect these traditional applications used over the last several years: electronic mail; interactive access; bulk file transfer; name/address translation services; and aggregated other TCP/UDP applications. The "other protocol" category corresponds to applications using a transport protocol other than TCP or UDP; the "other port" category to non-standard or not well-defined ports. Both of these categories have grown much larger over the years⁷, reflecting in the first case an increasingly multi-protocol environment, and in the second the diminishing ability to track individual new applications which often use non-standard or not well-defined ports. In fact, the "other port" category is, as of November 1992, the largest single category of traffic on the backbone, exposing the trend of application developers arbitrarily choosing their own port numbers for applications that collectively utilize much bandwidth. Since these port numbers are undefined to anyone but the end site using them, the growth of traffic volume for such applications is difficult to track; most statistics collection mechanisms can only attribute traffic to well-known port numbers, leaving other traffic in a large "unknown" category. In particular, the statistics collection process for the NSFNET backbone only classifies port numbers lower than 1024, plus a few select ports above 1024, and they are thus unable to attribute the traffic of the growing number of applications with port numbers above 1024. They must bundle such traffic into an "other protocols" category, making attribution of more than the base services (telnet, ftp, etc.) close to impossible.

Table 2 shows a more detailed distribution of traffic by port on the NSFNET backbone for the month of March 1993, and shows some indication of the growing range of applications. For example, several Internet resource discovery services (WAIS, WWW, gopher, prospero)⁸, have experienced tremendous growth in volume since their deployment, filling a significant void in network services. Of these applications, the available NSFNET backbone statistics indicate that the *gopher* service has exhibited the

⁷Claffy *et. al* [2] presents similar statistics for port usage on the T1 backbone, before the T3 backbone was fully deployed.

⁸These freely available tools provide for distributed document search and retrieval aimed to enhance the comfort and productivity of average network users. See Danzig *et. al.* [5] for a more detailed description of these and other resource discovery services.

Table 2: Traffic on NSFNET backbone by port for March 1993

Packet Total: 34,874,064,400				Byte Total: 6,502,203,065,800			
Service Name	Port	Rank	Packet Count	% Pkts	Rank	Byte Count	% Bytes
ftp-data	20	1	8279042350	23.740	1	2933157697150	45.110
telnet	23	2	5265928200	15.100	4	361378044900	5.558
nntp	119	3	2926178750	8.391	2	609322233900	9.371
smtp	25	4	2443215200	7.006	3	396478596800	6.098
domain	53	5	1731471000	4.965	6	157806711950	2.427
ftp	21	6	730566100	2.095	9	64429501750	0.991
irc	6667	7	703252650	2.017	8	69347837550	1.067
icmp	-1	8	634413950	1.819	10	50857619650	0.782
vmnet	175	9	454947500	1.305	5	165006133800	2.538
gopher	70	10	327717650	0.940	7	79023945150	1.215
X0	6000	11	279602550	0.802	11	48300762100	0.743
cmd/syslog	514	12	271915300	0.780	12	35153809700	0.541
login/who	513	13	223685900	0.641	13	22262183800	0.342
talk	517	14	212462050	0.609	14	21820335300	0.336
(unknown)	1023	15	172610350	0.495	16	16767055550	0.258
finger	79	16	166695800	0.478	17	15385492150	0.237
snmp	161	17	164575050	0.472	15	18249319150	0.281
ntp	123	18	125367100	0.359	25	9544144250	0.147
(unknown)	1022	19	86481600	0.248	19	14542602850	0.224
uucp	540	20	63177700	0.181	21	12344993750	0.190
(unknown)	1020	21	58279550	0.167	20	13987812450	0.215
(unknown)	1021	22	48658900	0.140	26	8956301150	0.138
ip	-4	23	43916400	0.126	22	12148087450	0.187
ntalk	518	24	38390450	0.110	31	3940355450	0.061
unidata-ldm	388	25	37887200	0.109	18	15213706250	0.234
efs/router	520	26	33235450	0.095	24	9694732350	0.149
bgp	179	27	27590100	0.079	44	1920440300	0.030
(unknown)	703	28	19975600	0.057	28	6197171350	0.095
z39.50	210	29	19506350	0.056	29	5415741150	0.083
(unknown)	700	30	18819800	0.054	30	4147485950	0.064
www	80	35	11294550	0.032	32	3613584700	0.056
shilp/sun-nfs	2049	57	5071450	0.015	63	709518550	0.011
shilp/sun-nfs	2049	57	5071450	0.015	63	709518550	0.011
X1	6001	72	2636100	0.008	83	281638250	0.004
iso-ip	-80	97	1131650	0.003	69	563371500	0.009
X2	6002	386	87100	0.000	346	17533600	0.000
X3	6003	567	36600	0.000	462	8546500	0.000
prospero	191	700	13950	0.000	432	10205800	0.000




Figure 10: Distribution of packets offered into NSFNET backbone by protocol




Figure 11: Distribution of bytes offered into NSFNET backbone by protocol

greatest growth, in fact tripling in traffic volume between November 1992 and March 1993, and during that month of March constituted in excess of 1.2% of overall NSFNET backbone traffic volume.

In addition to resource directory services, other applications are also gaining a greater proportion of network bandwidth: MUD⁹; X11¹⁰; and more recently and still only in its infancy, packet video and audio. Many of these application use inconsistent TDP/UDP port numbers, or port numbers unknown to the anyone but the end site using it. The growth of traffic volume for such applications is therefore difficult to track, since most statistics collection mechanisms can only attribute traffic to well-known port numbers, leaving other traffic in a large “unknown” category.

7 Impact on accounting and pricing

The problems outlined in the above discussion have obvious implications for the task of accounting in the Internet. Most instrumentation for “accounting” in the Internet reflects its status as bulk-funded rather than free market datagram environment. It is not at all clear how to implement accounting and billing in such an environment. In this section we discuss several problems related to accounting and pricing as the network evolves from a research environment with relatively narrow scope to a more commercialized environment that will eventually render data networks more of a utility, similar to the water, electric power and telephony systems.

Comparison to dedicated voice or data circuits may illuminate the difficulty of network usage accounting in a datagram environment which aggregates many end users and their applications. When providing dedicated circuits or services to a single customer, verifying the delivery of the promised product is relatively straightforward. In contrast, a network provider in the multiplexed datagram-based Internet environment promises a customer a probability of service resources rather than a dedicated and constantly verifiable physical pipe. In this scenario it is far more difficult to verify the promised level of service to any given customer. The evaluation of network performance and integrity of services becomes even more complicated when a virtual network service is mapped into a larger physical infrastructure, such as ANS’s provision of the virtual NSFNET backbone via its larger physical infrastructure, or Sprint’s provision of international band-

⁹MUD (Multi-User Dungeon) is a distributed electronic role playing game. What MUD enhances is beyond the scope of this study. MUDs have also been commonly used for a purpose similar to that of the Internet Resource Chat (IRC) protocol.

¹⁰X11, or X-windows, can provide remote graphical displays across the network

width for NSF via its rich network infrastructure¹¹ As IP providers continue to expand and leverage across existing infrastructures, it will be imperative to find mechanisms to differentiate service components and performances and to assure clients that they are receiving contracted network services.

A further complication arises even within certain service categories, when charging by the bit per source does not take into account the true beneficiary of a service. Shaping charging policies thus demands consensus on accounting conventions, and the distribution of benefits not only across transactions but also within the transactions themselves, such as the relative costs and benefits to the end points of the transactions. Unfortunately, statistics collection mechanisms, especially at service interfaces, inhibit the attribution of traffic to the transaction-requesting country; one can only attribute the traffic volume according to its physical source and destination countries. This distinction is important in the Internet: the generator of a TCP network connection request may not be the entity benefiting from the transaction. For example, charging for File Transfer Protocol (FTP) services based on the specific flow of IP packets from source ports to destination ports would be unacceptable to most sites sponsoring FTP-servers, which respond to requests for data with requests of their own to transmit the data. End-point accounting was not a goal in the initial design of the FTP protocol, and retrofitting a market-based environment to such underlying protocols will be challenging at best.

8 Proposed/relevant pricing models

Research in network pricing for both computer as well as other network infrastructures has led to several possible models of billing in networks. We discuss some of the models recently proposed in the literature, and then discuss how current Internet infrastructure and its instrumentation constrain the viability of proposed models.

The telephone network offers several models for billing, and even offers customers options depending on their self-assessed profile. Possible billing options include a flat service charge (typically for unlimited service within a local area), or a base price for a certain number of local calls plus an incremental charge for any calls above that limit. Long-distance telephone service billing is typically completely measured according to individual calls.

Applying telephone service billing models to the Internet imposes several difficulties since the Internet has far greater, and ever-increasing, functionality and diversity.

¹¹NSF funds Sprint, via the International Connections Manager (ICM) cooperative agreement, for components of its international connectivity to NSF clientele in other nations.

Currently, some larger institutions lease bandwidth from a network service provider in the form of dedicated circuits for Internet services, e.g., paying for a T1 line from Altnet and using any amount up to that limit. This model of bulk bandwidth distribution is not conducive to the vast majority of the Internet community, who have traffic profiles which could not justify the expensive of a leased line.

Billing in an environment with varying qualities of service will require effective categories of transmission, reflecting the required levels of service. Examples of services using such categories may include: information retrieval; real-time video; conferencing; multicasting; non-real-time messaging; low-priority bulk transfer; distributed computation; etc. The classification of traffic will include priority versus standard versus deferrable traffic flows, as described above, but may also extend to distributions of low-level traffic characteristics such as packet length histograms and burstiness profiles. The impact of time of day and time zone differences on network contention will also require considerations.

Cocchi [3] [4] presents a scheme for pricing in a computer network with multiple priorities. Cocchi provides evidence by computer simulations to confirm his thesis that in a network with multiple service priorities it is possible to set prices so that users of every application type are more satisfied with the combined cost and performance of a network with service-class sensitive prices than they would be with flat pricing. Thus a priority pricing scheme is always achievable which will enhance total community utility.

Parris *et. al* offer a scheme for real-time pricing in computer networks which can support reservation of resources. Their scheme is based on charging per real-time channel based on the resources reserved, including the type of service, time of day, and channel lifetime. For computational feasibility, the authors assume a homogeneous network, and reduce their analysis to a single node. We are not aware of how to scale their scheme to the multiple, very heterogeneous nodes, of the Internet. A bigger obstacle is the lack of capability for resource reservation in the current Internet.

MacKie-Mason and Varian [6] offer an alternative model of real-time Internet pricing. Their model addresses the inability to predict in advance the optimal price of network service based on internal congestion, which is a short-term, bursty, and hence unpredictable phenomenon in most components of the Internet. They propose a two-component pricing scheme: a flat connection charge, based on characteristics such as the type of customer or size of bandwidth, and a per-packet congestion charge assessed during times of network congestion.

The congestion charge would occur via a

“smart market”: a price for packet access to the net that varies minute-by-minute to reflect the current state of the network congestion....

Each packet would have a “bid” field in the header that would indicate the willingness-to-pay for that packet... The network would then admit all packets whose bid exceeded some cutoff amount. The cutoff amount is determined by the condition that the marginal willingness-to-pay for an additional packet has to equal the marginal congestion costs imposed by that packet.

Similar to the Parris scheme, this scheme would require a separate pricing and queuing “auction” to occur at each network router. The authors claim that their model, drawn from other economic applications of network pricing and applied to the Internet, has the long-term advantage that properly set congestion prices are the appropriate prices for valuing capacity expansion. In other words, efficient pricing of network congestion in the short run with investment of resulting revenue in capacity expansion provides the optimal investment in future capacity.

8.1 Current instrumentation

In this section we present how current instrumentation will constrain the feasibility of implementing aspects of several proposed pricing models.

Unfortunately, schemes which require each router to assess congestion and recompute and attribute charges would likely unacceptably interfere with packet-switching performance. Traffic flows contributing to network congestion often modulate within sub-seconds, although in many cases congestion may sustain, and block network resources, for minutes or even longer.¹² It is unlikely that pricing recomputation could keep up with the frequency of these changes in congestion.

Proposals of dynamic adaptation of pricing strategies to the existing network situation are attractive, but several obstacles will render such schemes difficult to implement in the existing architecture:

- overhead in the router
- difficulty with broad acceptance by users due to:

- inability of independent billing verification;
- unpredictability of the actual networking cost;
- cost increase due to other clients using the network

Also, one viewpoint is that the end user should not need to worry about congestion in real time. Real-time contention for network services is the responsibility of the service provider issue, who must base available resources on long term planning. Long term planning may yield fluctuations that may suggest graded pricing schemes through the day similar to telephone service rates. The user can predict and, perhaps more importantly, independently verify such payment schemas. Real time pricing adjustments based on network resource contention would result in network providers being attracted to a semi-congested state of their network, as it would drive up the price of network access to the network customer. Furthermore the price a client has to pay would depend on the demand of networking resources by other network clients.

Thus in our opinion allocation of the task of pricing to external systems will have far more auspicious effects on overall network stability. Routers may be able to supply statistics, and reallocate bandwidth among multiple prioritized queues, but it is critical to save router cycles for switching. Pricing schemes based on longer-term [hourly or daily] fluctuations in utilization are more feasible, predictable, and verifiable, than the rapid price recomputation required in a smart market.

Regardless of the pricing scheme, billing models amenable to implementation will have to deal with constraints of current network instrumentation. NSFNET, which is typical of most wide-area Internet service providers, can currently keep measurements of: raw volume (bytes); transactions (packets); or both (packets and payload). The mechanism for these measurements is via the SNMP protocol at network interfaces; it cannot attribute the traffic to packet type or geographic source or destination. The NSFNET does collect a flow matrix by network number for traffic crossing the T3 NSFNET backbone, but the overhead of computing the matrix prevents the complete collation of traffic; the computation must rely on only sampled (currently every fiftieth) packets. Attributing individual network numbers to other granularities, such as network external interfaces, Administrative Domains, or sub-service providers may allow for a greater accuracy of the collected statistics information. Such aggregation would allow the backbone provider to assess separate charges to these administrative entities, who would then have to perform their own accounting to redistribute the cost.

The NSFNET backbone project also provides a distribution of ports as described in Section 6. There are two

¹²Congestion manifests itself via queue growth in routers that do not have the local CPU capacity or external bandwidth to handle all received network traffic. As a router could starve more and more for resources under congestion, the contention may intensify sufficiently to consume all available buffer storage for additional packets. Alternatively, packet contention for processing via the CPU could consume all CPU resources, leaving the processor incapable of handling additional traffic.

clear problems with the current implementation of this mechanism. First, the service classes include only conventional Internet application categories such as interactive; file transfer; and transactions. It is difficult to track other service categories, including the new and foreboding continuous video/voice traffic type. A second difficulty relative to the NSFNET data is that the collation of this distribution however, is decoupled from the traffic flow matrix described above, and thus one cannot attribute a packet of a particular service category to a specific end network. Currently there is a recognized need within the NSFNET as well as other service providers to address both of these issues, but solutions in the actual infrastructure will require community resolution on standards of accounting for various service qualities, and in general what should be used for accounting analysis. Port information may be one criterion, but precedence and geographic source and destinations may suffice.

8.2 Capacity planning

The issue of unknown applications is not by itself necessarily as disturbing as the dramatically changing nature of the newly introduced traffic. The recent deployment of prototype packet video and audio applications bodes ominously for an infrastructure not able to preferentially deal with certain traffic. In this section we describe the dangers of our increasing inability to monitor traffic type in a “high-end” Internet.

Today’s Internet is inherently based on a datagram architecture, typically with no admission control in packet forwarders. Most entrance points into transit networks can not provide back pressure to peer points of the network that deliver more traffic than the network can handle. End systems can thus unfairly monopolize available bandwidth and cause significant congestion in the larger network.

During the life of the 56kbps NSFNET backbone in the mid-80s, this state of congestion developed to a dangerous degree, and in response the NSFNET engineers deployed an emergency measure to provide certain interactive network applications, specifically Telnet, preferential treatment over other traffic. The priority service prototyped in the Fuzzball-based 56kbps backbone in 1986 queued traffic based on both the IP precedence field as well as the interactivenss and thus required responsiveness of the protocol.¹³ The objective of this classification of applications into service types, and priority queueing of traffic based on type and IP precedence value, was to address real-time service contention under heavy congestion situations; The priority transit allowed interactive

users requiring better network responsiveness to continue working under highly network congested circumstances.

When the NSFNET was upgraded to T1 capacity, offering a 24-fold bandwidth increase and a richer topology, the designers did not re-introduce the priority queuing for end-user traffic. The new infrastructure used multiple queues only to differentiate between user traffic and network management traffic. An overabundance of bandwidth, with flat rather than per-volume payment scheme, rendered superfluous the use of multiple queues. In the case of the NSFNET backbone, the project partners bore all the costs of maintaining this bandwidth ahead of demand. The subsequent upgrade to the T3 network exemplified further this method of coping with network congestion: increase network capacity,

However today software developers continue to build advanced network applications which can consume as much bandwidth as network engineers provide. In particular, applications using packet voice and video do not exhibit the same “burstiness” characteristics of more conventional applications such as file transfer and electronic mail, but rather require continuous delivery of large amounts of traffic in “real-time”, and thus continuously consume significant fractions of the available bandwidth. Usage of such applications will not scale in the current Internet architecture, which may potentially need to support many such continuous point-to-point connections simultaneously.

It is difficult to overestimate the dramatic impact which such digital continuous media applications will have on the Internet fabric. No other phenomenon could more strongly drive the research community to instrument the network for admission control, as well as accounting and billing. Prerequisite to accounting and billing instrumentation will be a more accurate model for the attribution of resource consumption, derived from how particular applications impact network performance. Such a model may have to reliably attribute applications, or traffic profiles, to the clients, if multiple levels of services exist.

While performance optimization and accounting considerations are the dominating factors motivating the establishment of various traffic priorities/types, network engineers must incorporate the burden of this additional complexity into a longer term horizon. It will be a challenge for an inter-provider infrastructure to remain robust to, or even take advantage of, a greater number of possible traffic profiles based on an increasing range of diversity in service quality categories. A range of service providers, from local companies or campuses to international backbone service providers, will thus find it critical to stay aware of both short and longer term fluctuations in flows within the increasingly dynamic infrastructure. Longer term trends in flows can enable network providers and de-

¹³NSF based these categories on experiences and user feedback during the course of the NSFNET backbone project.

signers to plan or improve various aspects of the network, including topologies, application profiles, and underlying transmission technologies. Consideration of such flows requires the definition of a granularity model, as with the accounting case, but will also require greater focus on the traffic type and characteristics, including perhaps service categories based on traffic priorities, service quality, and/or application distribution.

A final consideration is accommodation of the diverse interests of network funding agencies, such as the NSF, that aim to encourage the development, deployment, and use of advanced, network-transparent applications on the network. An accurate assessment of traffic profiles could demonstrate conclusively the extent to which the overall infrastructure supports advanced applications, which could thus motivate planning for a higher performance network. An example might be a high-volume image rendering software package that routinely and invisibly to the user executes some software module on a remote super-computer before locally displaying resulting data. Performance profiles and resulting accounting characteristics for such applications will differ from those used for more conventional networking applications.

9 Summary

High level goals often qualify if not define the relationship between network analysis and network policy. We have offered evidence to support the hypothesis that in the face of today's critical point in the evolution of global information infrastructure, Internet policy considerations and network analysis must interact and support each other.

In particular, network analysis can offer insight into service categories relevant to accounting and policy considerations in network environments ranging from local to global scope. Results of traffic matrices by country have already proven useful to the U.S. NSF to illustrate international exchange of traffic among its constituents. In addition to quantifying network flows by various granularity, it will also be important to quantify and validate performance. As the threshold of high performance continues to expand into high volume real-time applications and advanced distributed computing paradigms, mechanisms to verify performance over shared infrastructures will be essential to clients as well as funding agencies.

Network analysis methodologies will also have obvious value for the integration of Internet accounting and billing mechanisms. As the functional and geographic scope of network performance continues to diversify, so does the financial structure of the Internet. Currently a transitional and somewhat confusing blend of public vs. private funding sources, some of which impose usage policies on critical pieces of the infrastructure, this structure can in-

timidate potential service providers as well as end-users. Creative and innovative developments in network analysis, with feedback to the developers of network policy, may dispel fears that a concerted effort between public and private networking efforts is not possible. On the contrary, such collaboration can enhance rather than retard Internet evolution.

References

- [1] T. Asaba, K. Claffy, O. Nakamura, and J. Murai. An analysis of international academic research network traffic between Japan and other nations. In *Inet '92*, June 1992.
- [2] K. Claffy, G. C. Polyzos, and H.-W. Braun. Traffic characteristics of the T1 NSFNET backbone. In *Proc. INFOCOM '93, San Francisco, CA*, April 1993.
- [3] R. Cocchi. *Pricing in multiple service class computer communication networks*. PhD thesis, U.C. Berkeley, 1992.
- [4] R. Cocchi, D. Estrin, S. Shenker, and L. Zhang. Pricing in computer networks: Motivation, formulation, and example. Technical report, U.C. Berkeley, 1992.
- [5] P. B. Danzig, K. Obraczka, and S.-H. Li. Internet resource discovery services. Technical report, USC, March 1993.
- [6] J. Mackie-Mason and H. Varian. Pricing the Internet. Technical report, U. Michigan, May 1993.
- [7] R. and P. Mandelbaum. The strategic future of mid-level networks. In Brian Kahin, editor, *Building Information Infrastructure*, 1992.