

Network Analysis in Support of Internet Policy Requirements ¹

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

In this paper we describe the importance of network analysis in support of policy considerations and evaluate a number of examples. We propound and support our hypothesis that, in the face of today's critical era in the evolution of global information infrastructure, Internet policy considerations and network analysis must begin to interact in ways not previously recognized or implemented.

I. Introduction

While initially conceived as a demonstration project of a then new networking technology 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, 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, fair 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 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 backbone², 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, pro-

¹ This research is supported by the National Science Foundation as an extension to the cooperative agreement to the San Diego Supercomputer Center, and also via NSF grant NCR-9119473.

² The "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.

vide unrestricted transmission service for any legal traffic from any paying customer.

In this paper we describe the importance of network analysis in support of certain 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.

II. 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.

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., a user transmitting

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 or even several network numbers. 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 certain flows, especially for real-time needs, easily exceeds the capabilities of available Internet technology.

III. Network analysis for accounting

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

Prerequisite to fair cost allocation and accounting is a secure mechanism for attribution of resource consumption, an historically difficult task in globally shared 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-term evolution. As a result, operationally collected statistics are generally useful for 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.

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, prevents such attribution of resource consumption to a user, much less a user in conjunction with an executing network application. An obvious alternative is to step up to the next coarsest layer of

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

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




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

aggregation: individual networked end-systems or IP network numbers.

Since most IP networks receive connectivity to the Internet via intermediate service providers, we could also consider 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. Attribution to service providers involves an aggregation of multiple IP network numbers into *Administrative Domains* (which in routing protocols are typically identified by one or more uniquely assigned Autonomous System numbers).

Some special cases of provider/client accounting may be amenable to perhaps the simplest accounting model: attribution of resource consumption according to aggregated packet/byte flows at service interfaces. This model assumes that a simple volume expression is a sufficient definition of traffic exchange, and typically will require that the client perform sub-accounting within its own area.

However most situations are not so simple. 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 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 bandwidth for NSF via its rich network infrastructure⁵ As IP providers continue to expand and leverage existing infrastructure, it will be imperative to find mechanisms to differentiate service components and performances and to assure clients that they are receiving contracted network services.

⁵ NSF funds Sprint, via the International Connections Manager (ICM) cooperative agreement, for components of its international connectivity to NSF clientele in other nations.

III.A. 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).⁶

Unfortunately, currently collected data does not allow such simultaneous attribution of traffic type and geographic distribution. Furthermore, the method used for the NSFNET backbone to attribute traffic by type demonstrates a decomposition of flows more reflective of traditional applications: electronic mail; interactive access; bulk file transfer; name/address translation services; and aggregated other TCP/UDP applications. Table 1⁷ shows, for example, the distribution of traffic by port on the NSFNET backbone for the month of March 1993.

Some newer applications, such as 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 greatest growth, in fact tripling in traffic volume between November 1992 and March 1993, and during that month of April constituted in excess of @@update to May 1.3% of overall NSFNET backbone traffic

⁶ The 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.

⁷ Source: Merit, Inc., March 1993.

⁸ These tools provide for distributed document search and retrieval.

Table 1: 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

volume.⁹

In addition to resource directory services, other applications are also gaining gains a greater proportion of network bandwidth: MUD¹⁰; X11¹¹; 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.

With the increasing diversity in applications, it will be even more critical to develop effective categories of transmission, perhaps along the lines of: information retrieval; real-time video; conferencing; multicasting; non-real-time messaging; low-priority bulk transfer; distributed computation; etc. Customer profiles must incorporate characteristic proportions from these descriptive categories, but should also account for the impact of time of day and time zone differences on network contention.

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 at-

tribute 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.

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 supercomputer before locally displaying resulting data. Performance profiles and resulting accounting characteristics for such applications will differ from those used for more conventional networking applications.

IV. Assessment of international flows

Higher level political goals may also require attention. National or international policy may necessitate attribution of resource consumption to individual countries. Figure 2 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. Figure 3 presents the matrix for non-U.S. countries for the same time period.

The operationally collected data sets also allow one to explore aspects of the data such as those in Table 2, which shows for February 1993 the directional asymmetries in traffic volume; average packet size by country; and skewness of distributions through time. The seventh column in Table 2 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 back-

⁹ WWW has also grown considerably, from .002% of the packets in January to .066% in April 1993. These statistics assume that a given port maps to only one service@@hw how do i say. There are exceptions to this behavior, such as when gopher serves as a front end to wais; this table will classify such "wais" traffic as gopher traffic. It may be better to aggregate all the various information retrieval tools into one group. Such end-system behavior is another example of the difficulties with interpreting this kind of data.

¹⁰ MUD (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.

¹¹ X11, or X-windows, can provide remote graphical displays across the network

Table 2: 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

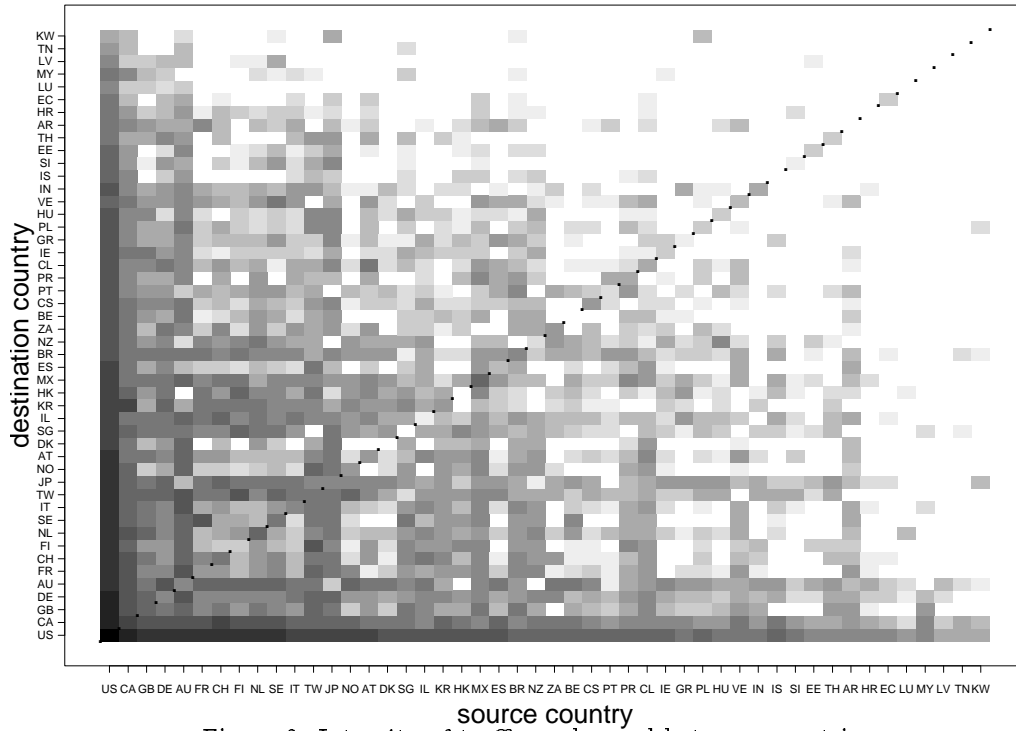


Figure 2: Intensity of traffic exchanged between countries




Figure 3: Intensity of traffic exchanged between non-U.S. countries

bone to the number of bytes that country sent into the backbone. Figure 4 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 paper also reflect the same one week time window.

Table 2 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)! These European countries seem to be characterized by only a few IP network numbers (Luxembourg has only four IP network numbers.) except for Germany, where there is a far greater of networks with very efficient outbound NSFNET traffic. There are also a few countries who are sending traffic *to* the backbone via very large (i.e., efficient) 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 5 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 6 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 then 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 7 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.

V. Capacity planning

A more general need for flow determination relates to the accounting considerations discussed above, but extends to a wider variety of applications, most notably large scale capacity planning and flow policy consideration, both possibly including international environments. A range of service providers, from local companies or campuses to international backbone service providers, will 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 designers 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, including perhaps service categories based on traffic priorities, service quality, and/or application distribution.

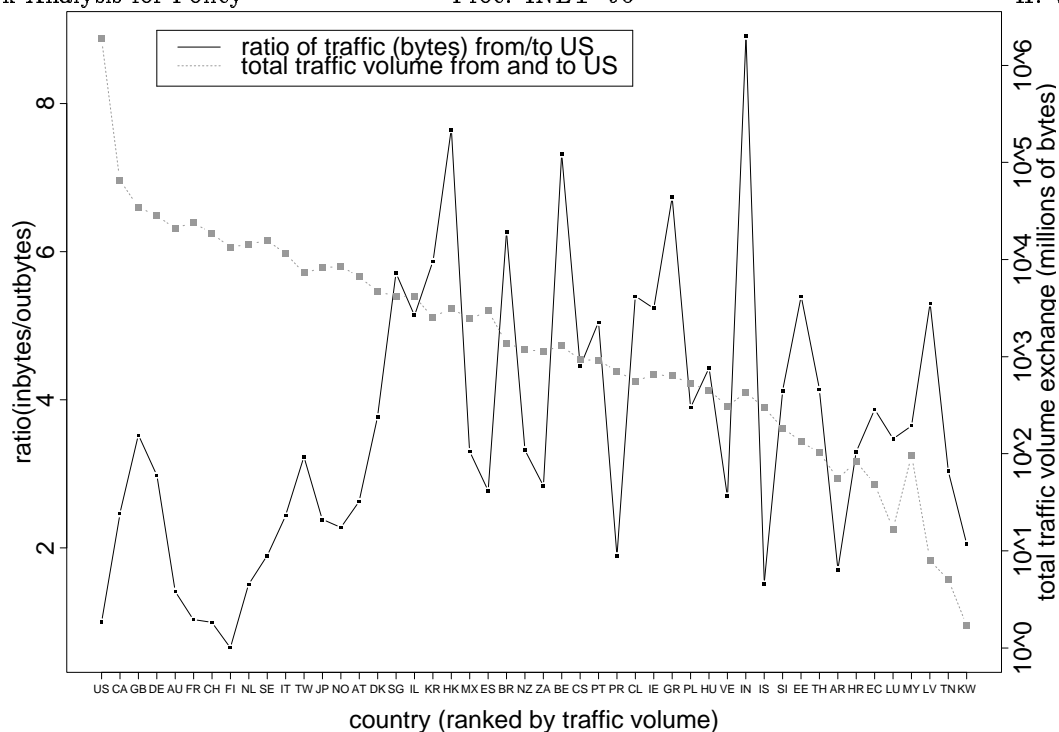


Figure 4: Intensity of traffic exchanged between countries

One example of the implementation of the traffic service classes is the priority queueing prototyped in the Fuzzball-based 56kbps NSFNET backbone in 1986. The backbone then queued traffic based on both the IP precedence field as well as the interactiveness and thus required responsiveness of the protocol.¹² 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. 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 queueing for end-user traffic. The new infrastructure used multiple queues only to differentiate between user traffic and network management traffic, based on NSFNET backbone IP addresses. 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.

While performance optimization and accounting considerations are the dominating factors moti-

vating 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. 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 length histograms and burstiness profiles. The integration of networked video and other multimedia capabilities as standard services on modern workstations will drive these requirements further.

VI. 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 of local as well as global scope. Results of traffic matrices by

¹² NSF decided the categories based on experiences and user feedback during the course of the NSFNET backbone project.




Figure 5: Traffic exchanged from Japan, Mexico, and Great Britain to United States

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 intimidate 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.

VII. Acknowledgements and Support

We would like to express our appreciation for valuable discussions we had with Steve Goldstein of the National Science Foundation.

Disclaimer:

Any opinions, conclusions, or recommendations in this report are those of the authors and do not necessarily reflect the views of the National Science Foundation, other supporting organizations, General Atomics, SDSC, UCSD or the SDSC Consortium members.




Figure 6: Ratio of traffic volume from/to the U.S. through the NSFNET backbone




Figure 7: Traffic from countries to themselves through the NSFNET backbone