

An Empirical Approach to Modeling Inter-AS Traffic Matrices

Hyunseok Chang

Department of EECS
University of Michigan
Ann Arbor, MI 48109-2122
hschang@eecs.umich.edu

Sugih Jamin*

Department of EECS
University of Michigan
Ann Arbor, MI 48109-2122
jamin@eecs.umich.edu

Z. Morley Mao†

Department of EECS
University of Michigan
Ann Arbor, MI 48109-2122
zmao@eecs.umich.edu

Walter Willinger

AT&T Labs-Research
180 Park Ave.
Florham Park, NJ 07932-0971
walter@research.att.com

Abstract

Recently developed techniques have been very successful in accurately estimating *intra-Autonomous System (AS)* traffic matrices. These techniques rely on link measurements, flow measurements, or routing-related data to infer traffic demand between every pair of ingress-egress points of an AS. They also illustrate an inherent mismatch between data needed (e.g., ingress-egress demand) and data most readily available (e.g., link measurements). This mismatch is exacerbated when we try to estimate *inter-AS* traffic matrices, i.e., snapshots of Internet-wide traffic behavior over coarse time scale (a week or longer) between ASs. We present a method for modeling inter-AS traffic demand that relies exclusively on publicly available/obtainable measurements. We first perform extensive Internet-wide measurement experiments to infer the “business rationale” of individual ASs. We then use these business profiles to characterize individual ASs, classifying them by their “utility” into ASs providing Web hosting, residential access, and business access. We rank ASs by their utilities which drive our gravity-model based approach for generating inter-AS traffic demand. In a first attempt to validate our methodology, we test our inter-AS traffic generation method on an inferred Internet AS graph and present some preliminary findings about the resulting inter-AS traffic matrices.

1 Introduction

Motivated by recent successes of *intra-domain* traffic matrix estimation techniques [8, 23, 16], we attempt to obtain accurate estimates of traffic volume exchanged between individual ASs. Knowledge of such a global but coarse-scale spatio-temporal picture of Internet traffic is vital for a number of practical networking problems, including evaluating

the impact of emerging technologies such as “intelligent routing control” (e.g., multi-homing, overlay routing) [3]; identifying or forecasting potential network bottlenecks or investigating the effectiveness of proposed remedies aimed at alleviating growing congestion [2]; assessing the performance of new protocols or Internet-wide applications such as new overlay networks; or improving the current understanding of how intra-domain traffic engineering (TE) impacts inter-domain TE and vice versa [1, 15, 14]. In the spirit of the work by Feldman et al. [9], we can use these inter-domain traffic matrices to partially answer such questions as “How stable are AS traffic volumes over time?” or “Which ASs carry most of the traffic, and how much?”

Given the highly competitive nature of today’s Internet Service Provider (ISP) market, ISPs do not make public such sensitive data as traffic volume statistics. The task of capturing the global behavior of AS-level Internet over long time scale (e.g., days) has been left to experimentalists. Due to the immense difficulties in collecting high volume of traffic data on an Internet-wide scale, research efforts to model and estimate inter-domain traffic demand are still in their infancy. With some exceptions [6, 22, 9], most studies that require knowledge of inter-domain traffic demand typically employ extremely simple (and untested) demand models, often assuming uniform traffic demand between every pair of ASs [2, 21]. Studies that rely on traces collected from a single vantage point (typically located in some stub network) are inherently constrained in their ability to provide a global view of inter-domain traffic. Nevertheless, an analysis of these traces revealed that while any given AS may exchange traffic with most of the Internet, only a small number of ASs are responsible for a large fraction of inter-domain traffic. In contrast to [6, 22], Feldmann et al. [9] also use server logs from a large CDN and develop a methodology for estimating inter-domain demand matrices for Web traffic.

Our model of inter-domain traffic demand encompasses the overall traffic, with Web traffic representing just one (though quite significant) component. Considering the

*This project is funded in part by NSF grant number ANI-0082287 and by ONR grant number N000140110617. Sugih Jamin is further supported by the NSF CAREER Award ANI-9734145, and the Presidential Early Career Award for Scientists and Engineers 1998.

†Z. Morley Mao is supported by the NSF Award CNS-0430204.

highly restrictive nature of access to proprietary data (such as server logs of large CDNs), we rely only on publicly available or obtainable data in our modeling of inter-domain traffic demand. However, the proposed approach is flexible enough to incorporate proprietary data, should it become available and should it be relevant to the problem at hand. In addition, we argue that a useful inter-domain traffic demand model should possess the following two characteristics. First, given the challenging task of obtaining actual inter-domain traffic matrices, the model should be *flexible* enough to allow for a systematic exploration of different traffic scenarios and traffic engineering strategies. At the same time, the model should be *parsimonious* enough to provide an intuitive understanding of the generated traffic demand and allow for a network-grounded interpretation of its parameters.

Our approach to develop such a model is partly empirical and partly analytical. We combine information gained from performing extensive Internet-wide measurement experiments with a modeling framework known as the “general gravity model.” The general gravity model has recently been used in estimating *intra*-domain traffic matrices [23] and has a long history in the social sciences, where it has been applied to describe the movement of people, goods, and information between geographic regions [17, 7]. To apply the gravity model to inter-domain traffic modeling, we first need to define the concepts of “mass” and “distance” within the context of inter-domain traffic exchange. To that end, we start by discussing the *business model* (or *operational characteristics*) of individual ASs.

In Section 2 we make the null hypothesis that to a first approximation, the traffic volume exchanged between two ASs necessarily reflects the business model of their operators. For example, an AS in the business of hosting various web and multimedia content will exhibit a very lopsided traffic profile (i.e., disproportionately heavy outbound traffic volumes). For another example, if two ASs are mainly in the business of providing access to residential customers, with comparable customer bases, traffic demand between the two networks can be expected to be more symmetric. By exploiting a range of publicly available data sets and by relying on information collected from our own Internet-wide experiments, we develop in Section 3 a combined measurement and business “profiling” methodology to infer the “utility” of an AS’s physical network. We identify the utility of an AS as providing Web hosting, residential access, or business access services. Depending on a simple high/low classification of these utility values, we infer seven natural AS business models. In Section 3, we classify ASs into one of these models, and rank the ASs within each class by their combined utility. These rankings then constitute the key input data to our general gravity model presented in Section 5 and determine the generation of inter-domain traffic volumes exchanged between individual ASs.

In particular, we illustrate that our model, besides being flexible and parsimonious, is capable of generating realistic traffic demand with different characteristics. As a final contribution, we focus in Section 6 on model validation and attempt to partially address this issue, even though inter-domain traffic matrix estimation is a case where even the most basic “ground truth” appears elusive. We conclude in Section 7 with a discussion of unresolved issues and open problems.

2 AS Business Models

We broadly define an AS’s business model as the *utility* of its physical networks, i.e., the primary reason(s) behind the design, operation, and management of its physical infrastructure. We associate a “business model” with each AS, not with each ISP or company, mainly because some large ISPs maintain multiple ASs (or domains) and typically assign these (sub)domains to separate business divisions, each with its own business characteristics or utility. We avoid the daunting task of identifying and enumerating the business purposes and operating strategies of all ASs by restricting our attention to the most *generic* utilities of existing networks that clearly affect their resulting traffic demand. In the following, we identify three such utilities (i.e., Web hosting, residential access, and business access) and describe how they can impact inter-domain traffic demand.

Web Hosting (Web). The success of the World Wide Web has led to an explosion of web sites that host various web content and streaming media. Powered by sophisticated content distribution technologies, a number of Web hosting companies have also emerged to host content outsourced by popular web sites. So much so that Web hosting is now a common service included in an AS’s service portfolio. An AS that hosts popular web content or e-commerce engines and distributes this content to the global Internet can be expected to carry voluminous outbound traffic and relatively little inbound traffic.

Residential Access (RA). Retail Internet business that directly deals with residential customers has existed since the inception of the commercial Internet. With advances in Internet access technologies, along with access speed, the number of residential users equipped with Internet access has risen steadily.¹With the proliferation of high-speed Internet users, the influence of end-user applications on the global traffic pattern becomes increasingly pronounced. A prime example are the bandwidth-demanding peer-to-peer (P2P) file-sharing applications. ASs populated by a large pool of residential users can exchange nontrivial amount of traffic among themselves as well as receive a large amount of web download traffic from Web content networks.

Business Access (BA). The low barrier of entry into the ISP market creates an environment whereby wholesaling and reselling of Internet access is actively pursued by both

Table 1: AS business profiles and models

Utility Profile			AS Business Model
Web	RA	BA	
H	H	H	Tier-1
H	H	L	Retail service
H	L	H	Business service
L	H	H	Network access
H	L	L	Web hosting
L	H	L	Residential access
L	L	H	Business access

incumbent carriers and new market entrants. Customers of large ISPs with nation- or continent-wide footprints are often themselves ISPs that resell the purchased Internet access to their own customers. For the purposes of this paper, “business access providers” are ISPs that resell purchased Internet access. Traffic demand of ASs that are business access providers can be estimated by the quality of service these ASs provide. Transit ASs guaranteeing good quality of service are likely to attract and retain a high number of customers. In turn, customers with reliable Internet connections can rely on the Internet for a large part of their business transactions, resulting in high traffic demand.

While the first two of these three utilities are traditionally attributed to stub networks, the third one is typically associated with transit networks. In today’s ISP market the service portfolio of large ISPs typically reflect multiple concurrent utilities (e.g., a business access provider may also be in the business of providing residential access or Web hosting services). We therefore do not follow the traditional classification of ASs into stub and transit networks. Instead, we attempt to determine a given AS’s business model by inferring from relevant data which of the three utilities dominate the AS’s operation, and what combination of utilities best characterizes the AS’s business. Table 1 lists seven AS business profiles based on the three identified utilities. Whether a given utility is primary or secondary to an AS’s business profile is denoted by “H” (high) and “L” (low). We tag each of the seven business profiles with an appropriately named business model listed in the right column.

3 Method for Inferring AS Business Model

In the highly competitive ISP market, the business plans of existing ISPs are generally confidential and cannot be inspected. Commercial ISP market research studies, e.g., Prisma (telegeography.com), are often dated, cover only a handful of well-known ISPs, or are based on some very narrowly-defined criteria. They are in general ill-suited for inferring business profiles of ASs, as defined in Table 1, in a comprehensive and coherent manner. As a viable alternative, we propose a methodology for inferring an AS’s business model that involves performing extensive Internet-wide measurement experiments and also involves collecting data indicative of individual ASs’ utilities. We

Table 2: Computation of $U_{web}(\cdot)$

initialize $U_{web}(X)$ to 0 for every AS X .
For each URL u ,
let $size(u)$ be the size of a file referred to by u .
extract a web server name N from u .
find the IP address set \mathcal{S} that is resolved to N .
For each IP address $I \in \mathcal{S}$,
find AS X that I belongs to.
$U_{web}(X) = U_{web}(X) + \frac{size(u)}{ \mathcal{S} }$

rely exclusively on publicly available/obtainable data and assume that we have no access to any proprietary data. We discuss the limitations imposed by this restriction and comment on how they could be alleviated if different types of proprietary data (e.g., server logs from a large CDN as used by Feldman et al. [9]) should become available.

3.1 Web Hosting

Our approach to quantifying web service utility is based on locating popular content on the Internet. ASs that host a large amount of popular content are considered to have high utility as a web service provider. To determine popular web content, we first consulted the web site Aleksika (www.skyart.org) and obtained a list of the top 10,000 search keywords most frequently submitted to search engines in the years 2003-2004. For each keyword, we queried the Google search engine, using the Google Web Services Application Programmer Interface (API), to retrieve a set of most closely matched URLs. For each submitted query, the Google Web API returned the top-10 matched results (URLs). We collected about 85,000 distinct URLs from all Google responses. By extracting web server addresses from these URLs, we inferred the ASs hosting widely accessed web content. To prevent bias towards discovering English-only content, we repeated the above experiment six more times, directing the Google Web Services API to return URLs in Chinese, French, German, Japanese, Korean, and Spanish respectively. The same set of English keywords was used in each experiment. For each language, we obtained between 80,000 and 90,000 URLs for these keywords. Merging all these results yielded close to 650,000 distinct URLs.

One source of web traffic not captured by the above measurement experiments is *embedded* web content, which includes media files delivered by dedicated media servers secondary to a web server, private CDNs, or third-party online advertisement objects (e.g., doubleclick.com). To uncover such web traffic, we crawled the above keyword-retrieved URLs individually, and extracted from the crawled pages URLs associated with embedded objects. Combining keyword-retrieved URLs with embedded-object URLs increased the total number of our collected URLs by a factor of three.

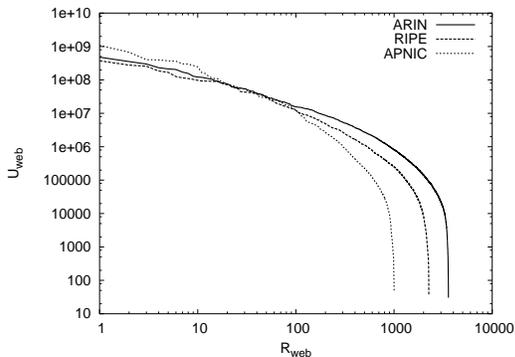


Figure 1: U_{web} distribution

Next we extracted a web server IP address from each URL and mapped it to its corresponding AS. A number of issues complicated this seemingly straightforward step. For one, some frequently-accessed web sites employ DNS-based load balancing, whereby their domain names are resolved to multiple addresses in a round-robin manner. In more sophisticated cases, a given domain name is resolved to a set of addresses depending on the querying client’s geographic location, web server availability, and network condition. CNN and AOL are examples of two content providers employing such DNS-based web request routing, and their domain names are typically resolved to hundreds of addresses spread out geographically and administratively over the Internet. To obtain all IP addresses associated with a given web server’s domain name, we performed a reverse-DNS lookup of each domain name from 96 geographically dispersed PlanetLab nodes (`planet-lab.org`), and collected all resulting addresses. Using BGP routing tables, we then mapped each resulting IP address to its corresponding AS. In Table 2, we summarize the steps taken to obtain the web-hosting utility $U_{web}(X)$ for every AS X , where $U_{web}(X)$ can be viewed as an estimate of the byte counts of popular web content hosted by AS X .

Finally, we sort the ASs by their utility $U_{web}(\cdot)$ in decreasing order and assign them ranks, denoted $R_{web}(\cdot)$. Table 3 lists the top-10 Web hosting ASs by rank, in three different geographic regions: North America (ARIN), Europe (RIPE) and Asia-Pacific (APNIC). As expected, the top-ranking Web hosting ASs include e-commerce companies (e.g., Amazon and eBay), telecom companies (e.g., AT&T, Deutsche Telekom, Korea Telecom), and well-known portal sites (e.g., Yahoo). One notable observation is the dominance of telecom companies in the Web hosting business in the Asia-Pacific region.

Fig. 1 shows the entire R_{web} vs. U_{web} distribution. Consistent with previous findings [6, 22, 9], $U_{web}(\cdot)$ associated with high-ranking ASs (e.g., up to rank 100 – 400) in all three geographic regions are characterized by a Zipf-type law (i.e., $U_{web} \sim (R_{web})^c$, where $c \approx -0.9$ for ARIN and

RIPE, and $c \approx -1.1$ for APNIC). The steep fall-off of the curves in the region of low-ranking ASs is likely due to the limited coverage our keyword-based crawling process has of their web content.

The key underlying assumption in our empirical method is that to infer the popularity of content networks, a feasible alternative to using actual web traffic measurement is to rely on data that measures the appearance of web content in search results. Strictly speaking, this alternative method can only account for traffic from users actively seeking specific information, not traffic from users visiting bookmarked pages or links on web pages. However, due to our decision to use the top 10,000 most popular keywords, we expect the resulting bias to be small. It is well-known that Google’s PageRank weighting algorithm carefully considers web link structures (e.g., links that a page receives) in calculating link values. As a result, we believe that our method does implicitly account for some aspects of actual link traffic.

Ideally, characterizing web service utility should make use of measured inter-domain web traffic, as is done for example by Feldman et al. [9]. However, this requires access to server logs of widely-deployed private CDNs. Such data is not publicly available. These data sets are also limited in their coverage: they capture only web content served by the CDN, and they do not capture web traffic emanating from content providers who are not clients of the CDNs. These difficulties illustrate the technical challenges associated with accurate estimation of inter-domain web traffic. Viewing our method as one of many viable approaches to making progress in this area, it exemplifies how a combination of publicly obtainable and publicly available data sets can be used to infer inter-domain traffic volume. At the same time, the method is flexible enough to incorporate web server logs from CDNs should such logs become publicly available.

3.2 Residential Access

To infer an AS’s utility in providing residential Internet access, we estimate it by the number of P2P file sharing users of the AS. Besides web browsing, P2P file-sharing is currently one of the most popular applications on the Internet (`cachelogic.com`). To estimate the number of users per AS, we perform measurement experiments involving three different file sharing systems: BitTorrent (`bittorrent.com`), eDonkey (`edonkey2000.com`), and Gnutella (`gnutella.com`). At the time of our study, these were among the most popular file sharing systems in use on the Internet [4].²

Most P2P file sharing systems have built-in mechanisms to discover existing users, which makes estimation of file sharing population relatively straight forward. The Gnutella system employs a decentralized approach to file

Table 3: Top-10 web service ASs (As of Sept. 2004)

R_{web}	ARIN		RIPE		APNIC	
	AS#	Name	AS#	Name	AS#	Name
1	3561	Savvis	8560	Schlund	3786	Dacom
2	2914	Verio	8220	Colt	4766	Korea Telecom
3	16509	Amazon	16276	Ovh	9304	Hutchison
4	21844	ThePlanet	3320	Deutsche Telekom	9318	Hanaro Telecom
5	11643	eBay	559	SWITCH	4808	Chinanet
6	13749	Everyones Internet	680	DFN	4134	China Telecom
7	7018	AT&T WorldNet	1273	C&W	2514	NTT
8	209	Qwest	702	MCI Europe	9848	GNG
9	701	UUNet	12312	Tiscali	4812	China Telecom
10	14134	Navisite	12322	Proxad	23880	Yahoo-KR

searching: individual Gnutella peers form an overlay network to propagate search messages. The eDonkey systems, on the other hand, relies on dedicated, centralized servers which peers must contact to search for a file. Similarly, BitTorrent relies on centralized *Trackers* from which a peer can obtain a list of other peers serving a particular file.

To estimate the population of BitTorrent, we downloaded about 2,800 *torrent files* from a well-known BitTorrent web site (torrentspy.com). A torrent file contains the meta data of a shared file, among which is the address of the BitTorrent *tracker* cognizant of peers sharing the file. Using BTtools (bagley.org/~doug/project/bttools), we obtained from each tracker a list of peers in possession of complete copies of the file (*seeds* in BitTorrent parlance). Over a period of four days, we collected about 634,000 distinct IP addresses of BitTorrent peers.

To estimate the population of eDonkey, we ran a toy eDonkey server and recorded the IP addresses of all the peers that contacted our server. eDonkey servers run a gossip protocol among themselves to maintain an up-to-date list of the server population. Each eDonkey peer maintains a list of servers and sends them periodic ping messages. In one day, our eDonkey server collected about 1,014,000 distinct IP addresses of eDonkey peers from these ping messages.

Finally, for the purpose of estimating the population of Gnutella, we ran a Gnutella client application on 20 PlanetLab nodes (10 in North America, 5 in Asia-Pacific, and 5 in Europe) and recorded the IP addresses of all Gnutella peers which exchanged traffic with us. Over a period of one week, we collected about 542,000 distinct IP addresses of Gnutella peers.

In total, we collected about 2.19 million distinct P2P IP addresses, which we subsequently mapped to their corresponding AS using BGP routing tables. Denoting by $U_{RA}(X)$ the utility of AS X as a residential Internet access provider, we computed this quantity for every AS X by counting the number of distinct P2P IP addresses associated with AS X . Of course, some P2P users perform more active downloads/uploads than others. By aggregating a sufficient number of IP addresses, we try to minimize any error that may be caused by ignoring such fine-

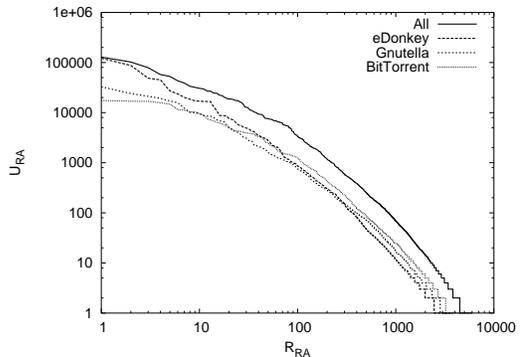


Figure 2: U_{RA} distribution

grained file sharing activity. We sort ASs by their utility $U_{RA}(\cdot)$ in decreasing order and assign them ranks, denoted $R_{RA}(\cdot)$. Table 4 lists the top-10 residential access ASs by region. In the European and Asia-Pacific regions, most of the high-ranking residential access providers are associated with telecom companies. In the North American region, retail Internet access business is more diversely distributed among telecom carriers and cable companies. Fig. 2 shows the entire R_{RA} vs. U_{RA} distribution, for all three P2P file-sharing applications individually as well as for their aggregate. In agreement with a recent study based on proprietary data set [20], $U_{RA}(\cdot)$ associated with the top 100 or so highest-ranking ASs can be characterized as a Zipf-type law with parameter -0.9.

One caveat in measuring P2P network usage is that the user base of different P2P systems is not uniformly distributed across residential networks [19]. Since several P2P applications with distinct features and evolving popularity coexist, relying on a single application may introduce sampling bias in capturing residential population. Our use of three popular file-sharing systems provides a reasonable coverage of the current file-sharing population. Our utility measurements can also be extended to not only cover other emerging P2P systems as they become popular, but also to include proprietary data such as per-AS statistics on residential subscriptions (including the amount of traffic generated by its connected residential customers), should the latter become publicly available.

Table 4: Top-10 residential access ASs (Oct. 2004–April 2005)

U_{RA}	ARIN		RIPE		APNIC	
	AS#	Name	AS#	Name	AS#	Name
1	1668	AOL	3320	Deutsche Telekom	4134	China Telecom
2	7132	SBC	3352	Telefonica España	4837	China Network
3	6478	AT&T WorldNet	3215	France Telecom	3462	HiNet
4	22909	Comcast Cable	12322	Proxad	4788	TMnet
5	577	Bell Canada	5617	Polish Telecom	1221	Telstra
6	22773	Cox	3269	Telecom Italia	4804	Microplex
7	812	Rogers Cable	5089	NTL	10091	SCV
8	7843	Adelphia	2856	BTnet	4812	Chinanet
9	6327	Shaw	6739	Cableuropa	9506	Magix
10	6128	Cablevision	3209	Arcor	7545	TPG

3.3 Business Access

Our approach to infer an AS’s utility in providing business access relies on publicly available BGP routing tables to estimate the AS’s bandwidth distribution. From a BGP routing table, one can infer provider-customer relationship among different ASs [10], and a naive estimate of an AS’s bandwidth distribution would be the number of its customer ASs. However, this lump-sum measure does not distinguish between customer ASs of different sizes. A more meaningful measure of an AS’s bandwidth distribution is the number of *downstream* ASs that are reachable from the AS, following the provider-customer relationship chains. A large transit customer AS with a high bandwidth requirement will then be properly weighted by its number of downstream customers. If a customer AS is multi-homed, i.e., obtaining its Internet access from several providers, it would typically impose lower bandwidth requirements on each provider than if it were single-homed. To infer $U_{BA}(X)$, the utility of AS X of providing business access, we assume that every AS has a unit bandwidth requirement. We then percolate each AS’s bandwidth requirement up the provider-customer relationship hierarchy. When an AS is multi-homed, its per-provider bandwidth requirement gets divided by the number of its providers. We estimate $U_{BA}(X)$ in terms of the bandwidth distribution of AS X , computed as shown in Table 5.

Table 5: Computation of $U_{BA}(\cdot)$

for every AS X ,
 unmark X .
 $U_{BA}(X) = 0$.
 $C(X) = \#$ of X ’s customer ASs.
 $P(X) = \#$ of X ’s provider ASs.
 while (1)
 for each unmarked X with $C(X) = 0$,
 mark X .
 for each provider Y of X ,
 $U_{BA}(Y) = U_{BA}(Y) + \frac{U_{BA}(X)+1.0}{P(X)}$
 decrement $C(Y)$ by one.
 if no AS has $C(\cdot) > 0$, exit.

We then sort those ASs with $U_{BA}(\cdot) > 0$ in decreasing order and assign them ranks, denoted $R_{BA}(\cdot)$. When different ASs have the same $U_{BA}(\cdot)$ value, we break ties by

the size of the ASs’ BGP-advertised address space. Table 6 lists the top-10 bandwidth reseller ASs. Most of them are associated with well-known tier-1 ISPs that operate continent-wide backbone networks. Note that in the European and Asia-Pacific regions, many top-ranking ASs are telecom companies.

We caution that our assumption of unit bandwidth requirement per AS may be too simplistic. For example, business customers which are not assigned public AS numbers are ignored in the computation of the $U_{BA}(\cdot)$ value. A more precise estimate of such intra-AS business customers could be obtained by examining intra-AS router-level connectivity, which, unfortunately, is not easy to discover from passive measurements. We rely here on an AS’s address space size to partially account for the presence of such “hidden” customer ASs. Having access to proprietary information on intra-AS business customers would simplify this problem considerably.

3.4 Discussion

In the absence of readily available information about AS business models, our proposed methodology for inferring an AS’s business model is based on the following assumptions: (1) ASs’ web service utilities can be gleaned from the usage patterns of a popular web search engine, (2) utilizing widely adopted file sharing applications, an AS’s residential access utility can be inferred from the size of its file sharing population, and (3) an AS’s business access utility, as measured by its bandwidth distribution, can be estimated by counting its downstream AS customers. A careful study of the robustness of the proposed methodology to violations of these assumptions (e.g., a more URL-dependent network usage, non P2P-based residential access, a more cost-driven approach to providing business access) is necessary, but is left for future work.

Our measurement method identified about 40% of all BGP-advertised ASs as providing some form of Web hosting service, about 30% as providing residential access, and about 15% as providing business access.³ The union of all identified ASs covers about 56% of all BGP-advertised ASs. Although close to 50% of all ASs are not categorized by our method, these are typically small ASs generating negligible traffic volume. For example, according to Net-

Table 6: Top-10 business access ASs (As of Sept. 2004)

R_{BA}	ARIN		RIPE		APNIC	
	AS#	Name	AS#	Name	AS#	Name
1	701	UUNet	1299	TeliaNet	4637	Reach
2	1239	SprintLink	702	MCI Europe	10026	ANC
3	3356	Level3	3320	Deutsche Telekom	2516	KDDI
4	7018	AT&T WorldNet	8220	Colt	3786	Dacom
5	209	Qwest	5511	France Telecom	703	UUNet AP
6	2914	Verio	1273	C&W	4766	Korea Telecom
7	3549	Global Crossing	6762	Telecom Italia	7474	Optus
8	3561	Savvis	3292	TDC	9225	Level3 AP
9	174	Cogent	6849	UKR Telecom	2764	Connect
10	3491	Beyond The Network	5400	BT Europe	7473	SingTel

Flow statistics obtained from a regional ISP,⁴ the 56% identified ASs were responsible for 99% of all traffic observed by this ISP. This suggests that our methodology works well for the set of ASs that are responsible for the bulk of Internet traffic, but is of limited use for the many small ASs that contribute little to the overall traffic volume.

As the Internet continues to evolve, so does an AS’s business model. However, since our proposed methodology relies mainly on being able to (1) identify the most “generic” business elements shared by existing ASs, and (2) infer each of these business elements by relying on appropriate “surrogate” measurements (or direct measurements, if available), we argue that our approach will remain applicable under changing Internet conditions (e.g., emergence of new “killer” applications), at least as long as generic business elements can be defined and viable surrogate measurements can be identified.

4 AS Business Characterization

From Tables 3, 4, and 6, one can see that some ASs rank very high with respect to more than one utility (e.g., Deutsche Telekom appears in all three), whereas other ASs rank high in only one category (e.g., Amazon in Web hosting, and Comcast in residential access). To compare these multi-variate AS utility profiles, and to associate each AS with one of the seven profiles listed in Table 1, we introduce the following quantitative metric. We first convert the ranks $R_{web}(X)$, $R_{RA}(X)$, and $R_{BA}(X)$ of an AS into their normalized counterparts, denoted by $r_{web}(X)$, $r_{RA}(X)$, and $r_{BA}(X)$, respectively. More specifically, we set $r_{web}(X) = R_{web}(X)/\max\{R_{web}(i), i \in \text{set of all ASs}\}$, so that $0 \leq r_{web}(X) \leq 1.0$. If $R_{web}(X)$ is unknown, we set it to $\max\{R_{web}(i)\}$, reflecting our intuition that X has negligible web utility, so that in this case, $r_{web}(X) = 1.0$. Likewise for $r_{RA}(X)$ and $r_{BA}(X)$. We then define the *rank vector* $\mathbb{R}(X)$ corresponding to AS X as $\mathbb{R}(X) \triangleq (r_{web}(X), r_{RA}(X), r_{BA}(X))$. Note that $\mathbb{R}(X)$ can be interpreted as a point in the 3-dimensional hypercube, with the seven business models listed in Table 1 representing the extreme or corner points (i.e., $(0, 0, 0)$, $(0, 0, 1)$, $(0, 1, 0)$, \dots , $(1, 1, 0)$) of this hypercube, where 0 and 1 corresponds to “H” and “L” respectively.

Intuitively, the business model of an AS X is determined by the minimal distance between its rank vector $\mathbb{R}(X)$ and the seven corner points. For example, as the rank vector $\mathbb{R}(X)$ gets closer to $(1, 0, 1)$, the business model of AS X is considered to be increasingly that of a residential access provider.

In Table 7, we list the top 10 North American ASs for each of the seven business models presented in Table 1. The top-10 ASs in the “Tier-1” category are those whose rank vectors are closest to the $(0, 0, 0)$ point in the 3D hypercube. Likewise for the remaining six categories. We observe that, first, ASs that are dominant in all three utility categories are indeed well-known tier-1 ASs. Second, in the “Network access” category where the primary utilities of the ASs are to provide Internet access to both business and residential customers, several of the high-ranking ASs are telecom companies. Third, the business profile of educational institutions falls under the “Retail service” category, where the primary utilities of an AS are to provide both Web hosting and residential access. Networks belonging to educational institutions usually host various academic web sites and at the same time provide Internet access to students living in university-owned housing. Fourth, several educational and research ASs are categorized under “Business access.” These ASs serve purely as backbone networks connecting other smaller institutions’ networks.

Next, we examine the correlation between the three utilities of an AS, for example, is an AS hosting a large volume of popular web content likely to serve a large number of residential customers as well? We use Kendall’s rank correlation coefficient [12] to quantify these pairwise correlations. For samples $(X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n)$ from a bivariate distribution, Kendall’s (sample) τ coefficient is defined as $\sum_i \sum_{i < j} \binom{n}{2}^{-1} [\text{sign}(X_j - X_i) \cdot \text{sign}(Y_j - Y_i)]$, where $\text{sign}(x) = 1$ if $x > 0$, and -1 if $x < 0$. Kendall’s τ provides a distribution-free measure of the strength of the association between two variables (i.e., monotonicity between two variables). The traditional Pearson product-moment correlation coefficient is less useful if Gaussian assumptions do not hold for the random variables at hand, as is the case for U_{web} , U_{RA} , and U_{BA} . To visualize the relationships between the three inferred utilities, we show in

Table 7: Top-10 ASs in different business categories (North American region)

Tier-1		Retail service		Business service		Network access	
AS#	Name	AS#	Name	AS#	Name	AS#	Name
3356	Level3	87	Indiana Univ.	14742	Internap	6383	BellSouth
7018	AT&T WorldNet	18566	Covad	297	NASA	6385	BellSouth
7132	SBC	1249	Univ. of Massachusetts	6922	Texas Backbone	13675	Verizon
209	Qwest	23504	Speakeasy	19782	Indiana Univ.	19158	USCarrier
1239	Sprint	2637	Georgia Tech.	5663	EDCnet	19752	Hydro One
3561	Savvis	14	Columbia Univ.	18695	Arbinet	22573	Northwestel
701	UUNet	25	UC Berkeley	5693	InteleNet	7776	Mebtel
852	Telus	4130	PSC	12179	N/A	25899	NOAnet
577	Bell Canada	18	Univ. of Texas	11588	El Dorado	7843	Adelphia
5650	ELI	20001	RoadRunner	4436	nLayer	10796	RoadRunner
Web hosting		Residential access		Business access			
AS#	Name	AS#	Name	AS#	Name		
16509	Amazon	7757	Comcast	11537	UCAID		
11643	eBay	20231	RoadRunner	6347	Savvis		
14134	Navisite	13367	RoadRunner	1784	Global NAPs		
8070	Microsoft	29737	WideOpenWest	6020	DCInet		
7224	Amazon	27699	TSP	19151	IBIS7		
11305	Interland	19115	Charter	2548	DIGEX		
6432	Doubleclick	22269	Charter	3643	Sprint		
26101	Yahoo	21508	Comcast	6509	Canarie		
7859	Pair	11683	Earthlink	293	Energy Science Net		
11443	OLM	4999	Sprint	2153	CSU		

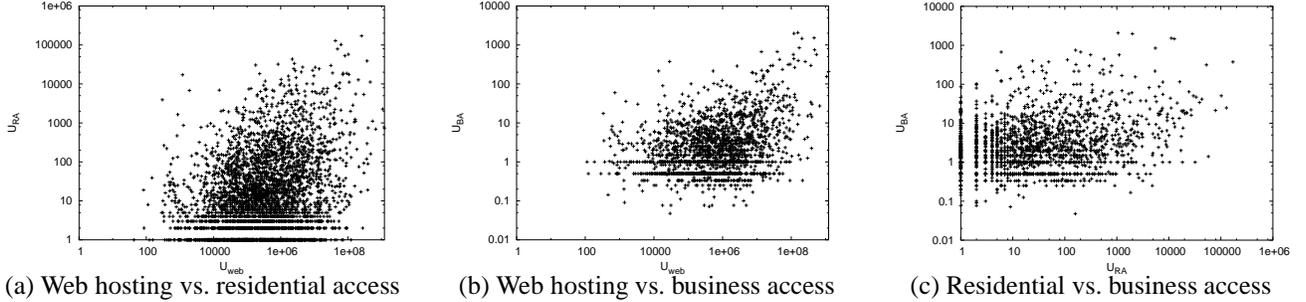


Figure 3: Pairwise utility correlation

Table 8: Pairwise Kendall's τ for $(U_{web}, U_{RA}, U_{BA})$

Kendall's τ	ARIN	RIPE	APNIC	All
U_{web} vs. U_{RA}	0.1562	0.1540	0.1820	0.1540
U_{web} vs. U_{BA}	0.2332	0.1439	0.2032	0.1856
U_{RA} vs. U_{BA}	0.1864	0.2088	0.2483	0.2068

Fig. 3 pairwise scatterplots for U_{web} , U_{RA} , and U_{BA} . The graphs show that there exists non-negligible correlation between each pair of utilities.

Table 8 lists the resulting Kendall's (sample) τ values for pairwise correlations between U_{web} , U_{RA} , and U_{BA} . By calculating the pairwise correlation for the three geographic regions separately, we find that in the European and Asia-Pacific regions, the correlation between "residential access" and "business access" is higher than the other two pairwise correlations. As hypothesized earlier, this higher correlation may be due to dominance of the ISP market in these regions by incumbent telecom carriers. Since the ISP market in the Asia-Pacific countries, in particular, is highly regulated by the government, it is to be expected that the entire ISP business in this region is dominated by a few large telecom carriers. In the North American region, on the other hand, the correlation between "residential access" and "business access" is relatively low, reflecting a less regulated environment.

5 Inter-AS Traffic Demand Model

The proposed traffic demand model builds on our AS business characterization by using as key input the ASs' inferred utility profiles (i.e., U_{web} , U_{RA} and U_{BA}). We show how the new model can be used in conjunction with a given AS graph to generate realistic inter-domain traffic demand.

5.1 Modeling Framework

For describing inter-domain traffic demands, we postulate a general gravity model (see for example [23] and references therein), where traffic flow from body i to body j (denoted X_{ij}) is assumed to satisfy

$$X_{ij} \propto \frac{S_i \times T_j}{F_{ij}}, \quad (1)$$

where:

- S_i : repulsive factor associated with "leaving" i ,
- T_j : attractive factor, "approaching" j ,
- F_{ij} : friction factor from i to j .

S_i , T_j , and F_{ij} are defined appropriately for each environment under study. For example, in applying the model to

urban transportation networks, S_i and T_j typically represent the populations in areas i and j , with F_{ij} being in general a function of the distance between the two areas. When applying the model in the context of modeling intra-domain traffic demand, Zhang et al. [23] take S_i as the traffic volume entering at location i and T_j as the amount of traffic exiting at location j , and assume a common friction factor F_{ij} that does not depend on i and j [23]. They show that this model works remarkably well and yields traffic demand that is consistent with measured intra-domain traffic volumes. To put the gravity model to good use for describing inter-domain traffic demand, we need to define S_i , T_j , and F_{ij} within our context so as to reflect the specifics the Internet’s AS environment.

Repulsive, attractive factors. We assume that a majority of inter-domain traffic in today’s Internet can be attributed to two kinds of interactions: (i) communication between web servers and clients (called “web” traffic), and (ii) communication between two clients (called “inter-residential” traffic). Web surfing and media streaming belong to the first category, Email and file sharing belong to the second category.

A typical web transaction is an asymmetric two-way communication: client’s request for web resource, and the corresponding response from the server. Thus, web traffic from AS i to AS j can be attributed to either a server in AS i returning web content requested by a client in AS j , or a client in AS i sending a request for web content served by a server in AS j . In the context of the gravity model, we model the volume of web traffic as a function of an AS’s client population size and web content population size. Let $P_{RA}(X)$ be the client population size of AS X and $P_{web}(X)$ the web content population size of the AS. The volume of “web” traffic from AS i to AS j is then quantitatively expressed as the weighted sum of two products: $P_{web}(i) \cdot P_{RA}(j) + \kappa_w \cdot P_{RA}(i) \cdot P_{web}(j)$. The first term corresponds to the “response” traffic from AS i (of population size $P_{web}(i)$) to clients in AS j (of population size $P_{RA}(j)$). The second term corresponds to the “request” traffic from clients in AS i for web content in AS j . The parameter κ_w is the ratio of request traffic over response traffic, usually significantly less than 1.

The symmetric inter-residential traffic from AS i to AS j is modeled as $\kappa_r \cdot P_{RA}(i) \cdot P_{RA}(j)$. The parameter κ_r is a normalization factor that determines the relative weight of web traffic and inter-residential traffic. Combining web and inter-residential traffic, the total traffic volume from AS i to AS j can be estimated as: $P_{web}(i) \cdot P_{RA}(j) + \kappa_w \cdot P_{RA}(i) \cdot P_{web}(j) + \kappa_r \cdot P_{RA}(i) \cdot P_{RA}(j)$.

A key factor in specifying our traffic demand model concerns the issue of modeling the quantities $P_{web}(\cdot)$ and $P_{RA}(\cdot)$. However, as seen earlier in Section 4, the empirically derived quantities $U_{web}(\cdot)$ and $U_{RA}(\cdot)$ are in fact estimates of the population of web content and the population

of clients, respectively, and are thus natural drivers of our gravity model. More generally, $P_{web}(i)$ and $P_{RA}(i)$ can be modeled as $f_1(R_{web}(i))$ and $f_2(R_{RA}(i))$, where $f_1(\cdot)$ and $f_2(\cdot)$ are monotonically decreasing rank-size functions; for example, in the case where $U_{web}(\cdot)$ follows roughly a type-1 Pareto distribution, $f_1(x) = x^{-\omega}$.

Friction factor. In urban transportation studies, the friction factor of the gravity model is typically a function of the distance between two regions. In estimating intra-domain traffic demand using the gravity model, Zhang et al. [23] assume a common, constant friction factor. In the inter-AS environment, such an assumption may not be realistic. For example, an over-provisioned path between two ASs may increase traffic flow between them, whereas an under-provisioned path is likely to decrease traffic flow between them.

Based on this observation, we define the friction factor between AS i and AS j as $R_{BA}(i, j)^\beta$, where $R_{BA}(i, j) = \max\{R_{BA}(X) \mid X \in \text{path}(i, j), X \neq i, j\}$ and $\text{path}(i, j)$ denotes the set of transit ASs in the path between AS i and AS j . As defined in the previous section, $R_{BA}(X)$ is the rank of AS X among all business access providers. $R_{BA}(i, j)$ is thus the maximum rank of a transit AS between AS i and AS j . Assuming that an AS with higher transit rank is more likely to maintain a well-provisioned network, this definition of the friction factor captures the transit quality of the *bottleneck* AS of a given path. By tuning the parameter β , we can study the sensitivity of traffic demand as a function of the transit network quality (a smaller β means lower variability in transit quality).

Remarks: Note that the original gravity model postulates that interactions between nodes are *independent*. On the one hand, this assumption seems reasonable for modeling highly aggregated quantities such as inter-domain traffic flows, where the latter are, in general, sufficiently aggregated so that possible dependencies among finer-grained traffic flows can be safely ignored. On the other hand, by defining the friction factor in terms of $R_{BA}(\cdot)$, we may introduce subtle dependencies among inter-domain traffic flows, as $R_{BA}(X)$ is not independent of $R_{BA}(Y)$ if X is a downstream customer of Y , or vice versa. In this sense, the proposed “general gravity model” is not a gravity model in the strict sense, but allows for dependencies that may be genuine at the Internet’s AS level, where inter-dependent traffic engineering is not uncommon.

5.2 Generation of Inter-AS Traffic Demand

Given an AS graph with N nodes and using the above gravity model, we can express the traffic demand from i to j as

$$X_{ij} \sim \frac{T_w(i, j) + \kappa_r \cdot T_r(i, j)}{R_{BA}(i, j)^\beta}, \quad (2)$$

where

- $T_w(i, j) = f_1(R_{web}(i)) \cdot f_2(R_{RA}(j)) + \kappa_w \cdot f_2(R_{RA}(i)) \cdot f_1(R_{web}(j))$, and
- $T_r(i, j) = f_2(R_{RA}(i)) \cdot f_2(R_{RA}(j))$.

To produce the resulting inter-domain traffic demand matrix, we first *generate* for each node (AS X) in the graph its rankings in terms of the three utilities we identified ($\hat{R}_{web}(X)$, $\hat{R}_{RA}(X)$, $\hat{R}_{BA}(X)$). When generating these rank vectors, we must account for the pairwise correlation between the rankings as reported in Section 4. By definition, the ranking $\hat{R}_{BA}(\cdot)$ is determined solely by the topology of a given graph. Given a graph, $\hat{R}_{BA}(\cdot)$ can be computed independent of the other two rankings. Using $\hat{R}_{BA}(\cdot)$ as an *anchor*, we next generate $\hat{R}_{web}(\cdot)$ and $\hat{R}_{RA}(\cdot)$ based on a well-known method for generating multi-variate normal random numbers [18]. Our rank generation algorithm is described in Table 9. The input parameters to our algorithm are the AS graph and a 3×3 rank correlation matrix $\Sigma_\tau = \{\tau_{ij}\}$.

Table 9: Generation of $\hat{R}_{web}(\cdot)$, $\hat{R}_{RA}(\cdot)$, and $\hat{R}_{BA}(\cdot)$

<p>Input: AS graph with N nodes, $\Sigma_\tau = \{\tau_{ij}\}$</p> <p>Algorithm:</p> <p>// generation of $\hat{R}_{BA}(\cdot)$</p> <p>Compute $U_{BA}(\cdot)$ by the method shown in Table 5.</p> <p>Assign $\hat{R}_{BA}(\cdot)$ to ASs in a decreasing order of $U_{BA}(\cdot)$.</p> <p>// generation of $\hat{R}_{web}(\cdot)$ and $\hat{R}_{RA}(\cdot)$</p> <p>Convert $\Sigma_\tau = \{\tau_{ij}\}$ into $\Sigma_r = \{r_{ij}\}$, where Σ_r is product moment correlation matrix with $r_{ij} = \sin(\frac{\pi}{2} \cdot \tau_{ij})$ (due to Kruskal [13]).</p> <p>Compute a lower-triangular matrix L, such that $\Sigma_r = L \cdot L^T$.</p> <p>Generate $(\hat{x}_i, \hat{y}_i, \hat{z}_i)$ for each AS i, where $\hat{x}_i = \frac{\hat{R}_{BA}(i)}{N}$, and \hat{y}_i and $\hat{z}_i =$ uniform random numbers $\in [0,1]$</p> <p>Obtain $(x_i, y_i, z_i)^T = L \cdot (\hat{x}_i, \hat{y}_i, \hat{z}_i)^T$.</p> <p>Assign $\hat{R}_{web}(\cdot)$ to ASs in a decreasing order of y_i.</p> <p>Assign $\hat{R}_{RA}(\cdot)$ to ASs in a decreasing order of z_i.</p> <p>Output: ($\hat{R}_{web}(X)$, $\hat{R}_{RA}(X)$, $\hat{R}_{BA}(X)$) of all ASs</p>

6 Toward Model Validations

Recall that our empirical approach to determining the input data (i.e., utility profile-based AS business models) that drives the gravity model proposed in Section 5 is based exclusively on publicly obtainable/available data sets and does not use any actual traffic volume measurements. Thus a natural starting point for attempting to validate many aspects of our inter-domain traffic demand model is to gain access to traffic volume-related AS-specific data sets which

are in general *not* publicly available. We follow this strategy by relying on a week’s worth of (sampled) NetFlow measurements from a regional ISP. The data sets were collected from one of its access routers around the same time when we performed our own measurement experiments described in Section 3 (i.e., Oct. 2004). The captured traffic originates from or is destined to the ISP’s networks, and the data sets contain, among other information, source/destination IP address prefixes of length at most 24 (due to anonymization), source/destination port numbers, and size of each traffic flow. We use these actual traffic measurements (i) to check a basic assumption underlying our gravity model, namely that the inter-domain traffic demand is determined by “web” and “inter-residential” traffic, (ii) to explore the adequacy of our key decision to use “surrogate” traffic measurements (e.g., data measuring the appearance of web content in search results, estimates for the number of P2P file-sharing users) instead of actual traffic volume estimates, and (iii) to provide a preliminary comparison between actual inter-domain traffic demand and those generated by our model.

6.1 Traffic Classification

To check the assumption explicit in Equation (2) that inter-AS traffic demand consists of the two components, “web” traffic and “inter-residential” traffic, we classify the flows in our NetFlow data set into “web” traffic and “inter-residential” traffic, using an up-to-date list of well-known port numbers. Given a traffic flow, if either source or destination port number is assigned to well-known web service (e.g., http, nntp, streaming), the flow is marked as “web” traffic. If either source or destination port number is associated with a well-known P2P file sharing application, the flow is marked as “inter-residential” traffic. However, an increasing number of applications do not use well-known port numbers, which makes it difficult to fully identify network traffic type based on port numbers alone. To improve upon the above naive traffic classification, for flows we cannot identify by source/destination port pair, we examine their source and destination address prefixes to heuristically infer its application type. More specifically, we randomly choose two IP addresses, one from each the source address prefix and destination address prefix, and perform a reverse DNS lookup. If either one of them has web-service related domain names (e.g., www* or web*), then we mark the flow as “web” traffic. If both IP addresses are resolved to well-known residential network domains (e.g., reshall.umich.edu or comcast.net), we mark the flow as “inter-residential” traffic.

Table 10 reports our traffic classification result. The reported percentages are based on total volumes of traffic. Although applying our heuristic reduces the amount of unknown traffic by 8%, uncategorized traffic still accounts for

Table 10: Inter-AS traffic classification

Classification	Web	Inter-residential	Unknown
Port-based	30.2%	28.8%	40.0%
Port-based + Heuristic	31.6%	36.0%	32.4%

one-third of all traffic, consistent with other available numbers [4]. We thus observe that with currently available traffic classification methods, our model appears to capture *at least* two thirds of actual inter-AS traffic. Improvements of state-of-the-art traffic classification techniques (e.g., [11]) can be expected to show a more accurate coverage of inter-AS traffic by our model.

6.2 Measurement Methodologies

Our methodology described in Section 3 for inferring an AS’s utility profile and determining in turn its business model avoids on purpose actual traffic volume-related measurements. Instead, we rely on “surrogate” traffic measurements such as appearances of web content in search results or estimates of an AS’s P2P file sharing population and assume that the latter are viable substitutes for the largely inaccessible actual traffic data. To check this assumption, we rely again on our NetFlow data sets and extract from them three distinct traffic volume measurements for each AS X : $T_{web}(X)$, $T_{RA}(X)$, and $T_{BA}(X)$, which correspond to “web-hosting” traffic, “residential access” traffic, and “business access” traffic, respectively. Using our classification of traffic in Section 6.1 into “web” and “inter-residential” traffic, we compute $T_{web}(\cdot)$ and $T_{RA}(\cdot)$, as described in Table 11.

Table 11: Computation of $T_{web}(\cdot)$ and $T_{RA}(\cdot)$

$T_{web}(\cdot) = T_{RA}(\cdot) = 0$ for every AS. for each traffic flow f , if f is web traffic, let $X =$ web hosting AS for f let $Y =$ client AS for f $T_{web}(X) = T_{web}(X) + size(f)$ $T_{RA}(Y) = T_{RA}(Y) + size(f)$ else if f is inter-residential traffic, let $X =$ client AS 1 for f let $Y =$ client AS 2 for f $T_{RA}(X) = T_{RA}(X) + size(f)$ $T_{RA}(Y) = T_{RA}(Y) + size(f)$

The “business-access” traffic $T_{BA}(X)$ captures the volume of traffic going through AS X . To compute $T_{BA}(X)$, we use Gao’s heuristics [10] to construct an AS-level routing path for each source-destination pair. We then increment $T_{BA}(\cdot)$ of every transit AS between the source-destination pair by the size of each flow.

Table 12: Pairwise Kendall’s τ for $(T_{web}, T_{RA}, T_{BA})$

Kendall’s τ	ARIN	RIPE	APNIC	All
T_{web} vs. T_{RA}	0.2490	0.1816	0.2752	0.2410
T_{web} vs. T_{BA}	0.1970	0.2467	0.2826	0.2440
T_{RA} vs. T_{BA}	0.1973	0.2489	0.2157	0.2371

Fig. 4 shows how well our inferred utilities (i.e., U_{web} , U_{RA} and U_{BA}) compare to their actual traffic-derived counterparts (i.e., T_{web} , T_{RA} and T_{BA}). We observe that in all three cases, ASs with high inferred utilities also have high measured traffic volumes. In Table 12, we quantify the pairwise correlation of T_{web} , T_{RA} , T_{BA} , just as we did earlier in Table 8 with the inferred utilities U_{web} , U_{RA} , and U_{BA} . Comparing Tables 8 and 12, we note that the pairwise correlation values are slightly underestimated by our methodologies, more so in the North American region than in the other regions. Overall, while the actual values differ, the generally low degree of pairwise correlations in T_{web} , T_{RA} , and T_{BA} is consistent with what we observed earlier for the inferred utilities U_{web} , U_{RA} , and U_{BA} , which suggests that our assumption of using appropriate “surrogate” traffic measurements instead of actual traffic measurements is not obviously unreasonable. However, there is clearly room for significant improvements.

6.3 Traffic Demand Model

Ultimately, any inter-domain traffic demand model will be judged by how well the results compare to actual Internet data. In our attempt to provide such an initial comparison, we check whether our model, when combined with an inferred Internet AS graph and when appropriately parameterized, is capable of generating realistic traffic demand, consistent with actual demand measured in the Internet. To this end, we use an inferred Internet AS graph consisting of 18,221 nodes and 39,558 edges, and since our NetFlow data sets were collected from a single vantage point on the Internet, we compare model-generated traffic demand of a single source node against NetFlow traffic information, where the source node is chosen to be an AS S that has a business model comparable to that of the AS from which the NetFlow measurements were collected. We also present some preliminary result concerning the sensitivity of the generated inter-domain traffic matrix to the choice of model parameters.

Given the Internet-like pairwise rank correlation (i.e., τ_{ij} taken from the “All” column of Table 8), we assign ranks to each node based the method described in Section 5.2, considering the special case where the quantities $P_{web}(\cdot)$ and $P_{RA}(\cdot)$ are given by type-1 Pareto distributions, i.e., $P_{web}(X) = R_{web}(X)^{-\omega}$ and $P_{RA}(X) = R_{RA}(X)^{-\rho}$, with $\omega, \rho > 0$. Using this model, we then generate the traffic demand T_X that each node X maintains with our

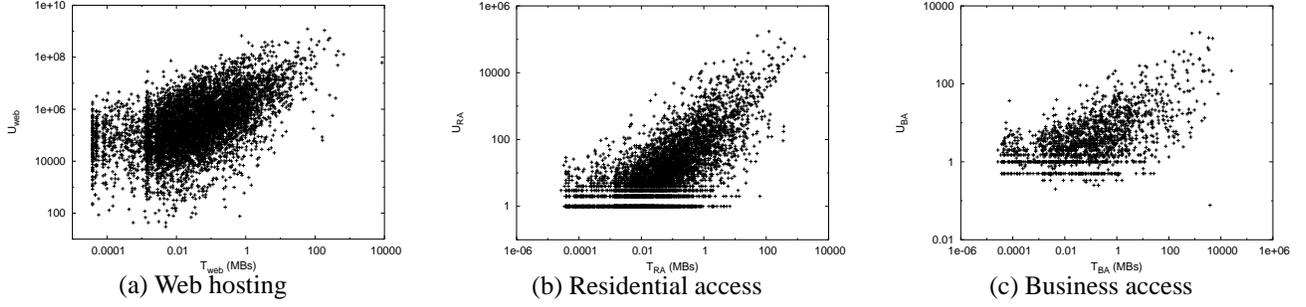


Figure 4: Measured utilities vs. traffic

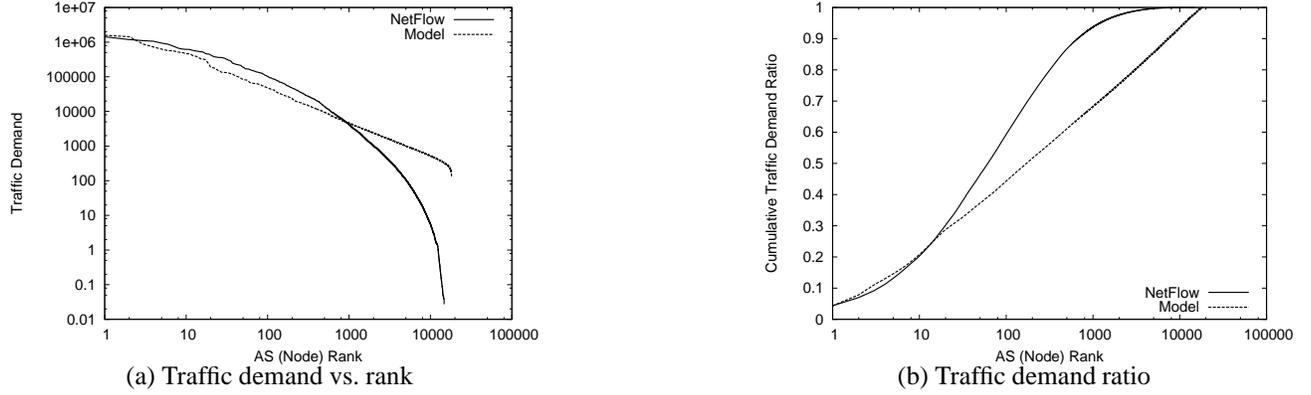


Figure 5: Single-source traffic demand comparison ($\omega = \rho = 1.0$, $\beta = 0.1$, $\kappa_w \in [0, 0.5]$, $\kappa_r \in [0, 1.0]$)

chosen source node S (i.e., $T_X = T_{SX} + T_{XS}$). At the same time, relying on the NetFlow data sets, we also obtain the measured traffic demand between each AS and the NetFlow collector AS, the regional ISP from which the NetFlow data was obtained.

Fig. 5(a) compares model-generated traffic volumes with NetFlow-derived traffic volumes by showing volume vs. rank distributions.⁵ In the case of the NetFlow-derived volume measurements, the traffic demand starts to fall steeply after rank 1000 or so, while the model-generated demand does not. This deviation is likely due to a deficiency of the type-1 Pareto distribution (used to model P_{web} and P_{RA}) as suggested by Figs. 1 and 2. In Fig. 5(b), we plot the cumulative traffic demand ratio; i.e., we examine what percentage of the total traffic demand the top- x percent ASs are responsible for. The actual traffic demand ratios of the top-ranking ASs (up to rank 30 or so) are matched well by our model, but for the same reason as before, the generated demand ratios start to deviate considerably from their actual counterparts when we include the lower-ranking ASs.

Finally, to compare model-generated and actual demand in terms of outbound and inbound traffic profiles, we plot in Fig. 6 for each node or AS X on the x -axis the volume

of traffic from X to S , and on the y -axis the volume of traffic from S to X . Since we already know that our model is inadequate for predicting small ASs' traffic demand, we focus in the figure on the large ASs (i.e., top-1000 ASs in terms of their total traffic demand) to get an idea about how well our model predicts the inbound/outbound traffic for the critical large ASs. Fig. 6(a) shows the profile obtained from using NetFlow measurements, while Fig. 6(b) depicts the profile resulting from our model-generated inter-domain demand. As the traffic profile of an AS moves further below the diagonal line, its business profile becomes increasingly that of a web service provider. Conversely, an AS whose traffic profile is located above the diagonal line is a typical residential access provider. Comparing the two profiles in Fig. 6, we see that the NetFlow-derived values are comparable to their model-generated counterparts, with less concentration around the diagonal, though.

To illustrate the effect of parameterization of our gravity model (2), we consider the parameter β , which determines how variable the transit quality of different networks is. A high value of β means that the transit quality of the higher-tier ASs is significantly better than that of the lower-tier ASs, and a low value of β implies a more uniform transit quality. In an extreme case where $\beta = 0$, the transit quality

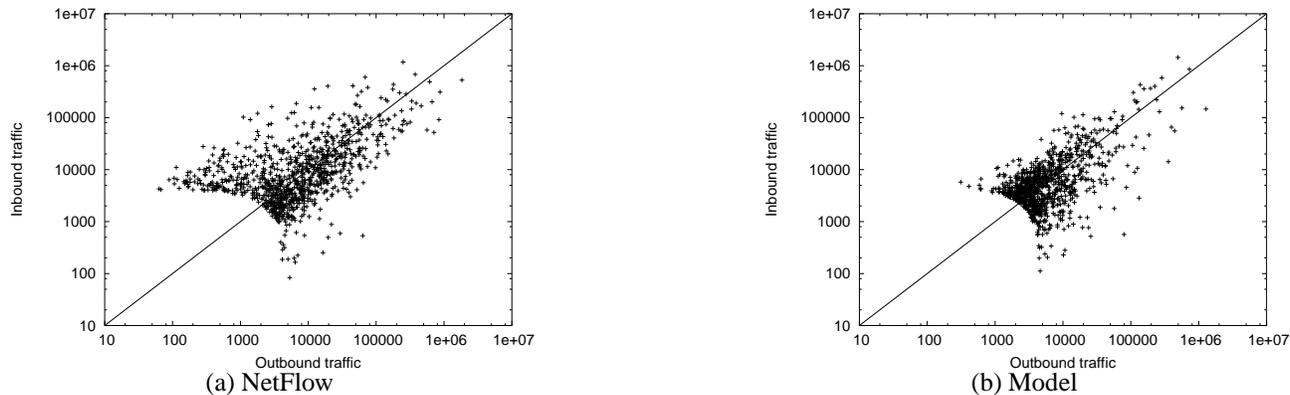


Figure 6: Single-source outbound vs. inbound traffic distribution ($\omega = \rho = 1.0$, $\beta = 0.1$, $\kappa_w \in [0, 0.5]$, $\kappa_r \in [0, 1.0]$)

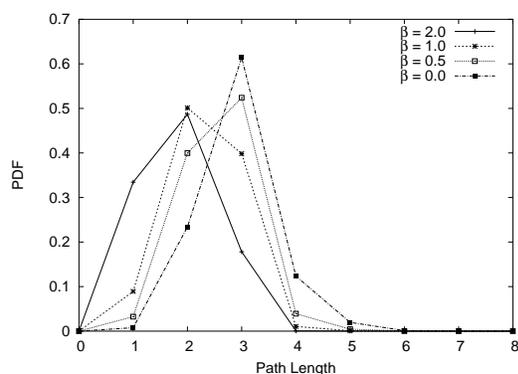


Figure 7: Topological distribution of traffic ($\omega = \rho = 1.0$)

of all transit networks is considered the same. As described in Section 5, the friction factor F_{ij} , which is a function of β , reflects the transit quality of the path bottleneck and is thus a topological metric that depends on graph connectivity. To examine the effect of β in terms of its impact on the *topological distribution* of traffic demand, for a given path length l , let S_l denote the total sum of traffic exchanged between every pair of ASs that are distance l apart. We are interested in understanding how, depending on the parameter β , S_l changes with different path length l .

Fig. 7 shows the probability density functions of S_l for different values of β . As β increases, the entire density function is shifted to the left. Intuitively, in a high- β setting, the total traffic demand is more likely to be dominated by the demand between close-by source-destination pairs; demand between far-away pairs becomes negligible, (long paths are more likely to encounter a bottleneck than short paths and will be avoided). In short, a high- β setting imitates an Internet environment where traffic demand tends to be highly localized [22]. In contrast, in a low- β setting, traffic demand tends to be less sensitive to the bottleneck quality between source and destination, resulting in distant source-destination pairs exchanging non-negligible amount of traffic.

Besides the parameter β , we also examine the pairwise rank correlation matrix $\Sigma_\tau = \{\tau_{ij}\}$ and study the traffic profiles resulting from different degrees of rank correlation. As expected, increasing pairwise rank correlation results in higher correlation between outbound traffic associated with web hosting utility and inbound traffic attributed to residential access utility. Due to space limitation, we do not include the detailed results.

7 Conclusions

The Internet AS environment is a setting where establishing “ground truth” is notoriously difficult. For example, while it is relatively easy to infer AS maps of the Internet from publicly available BGP-derived data, the underlying measurements are known to provide only a very incomplete picture of Internet connectivity at the AS-level [5]. In turn, this creates significant challenges for accurately modeling the Internet’s AS topology and large unresolved problems as far as validating the resulting models is concerned.

In this paper, we are concerned with an even more elusive aspect of the Internet’s AS environment, namely the AS-level traffic matrix giving the traffic demand between any pair of connected ASs. For one, there exists no equivalent of the publicly available BGP-derived data, and this has led researchers to pursue a mostly model-based approach. Even worse, for fear of losing competitive advantage, ASs are very reluctant to provide any AS-related data. As a result, AS-specific traffic data is by and large not publicly available, causing researchers to look for “surrogate” measurements that are publicly available or obtainable (i.e., via measurement experiments that can be performed by anyone connected to the Internet) and that may shed some light on the nature of the actual inter-AS traffic demand. As far as model validation is concerned, this situation causes nightmares, because on top of examining the validity of a proposed model, it first requires checking that the considered surrogate measurements are indeed suitable and relevant as substitutes for the largely unavailable data.

By developing a flexible approach to generating inter-AS traffic matrices, we make four specific contributions:

1. Identification of relevant surrogate measurements that are publicly obtainable/available;
2. Derivation of AS-specific statistics from the measurements in 1 that are the key inputs to a general gravity model for inter-domain traffic demand;
3. Generation of inter-domain traffic demand from the gravity model in 2 that are not obviously inconsistent with actual demand;
4. Methodology for validating AS-level traffic demand models that puts to good use the few and rare proprietary data sets that some ISP have been willing to share with the networking research community.

While many of the specific details of our approach can be questioned and much room for improvements exist, we have demonstrated that overall, it is not only feasible, but also generates realistic inter-AS traffic demand with Internet-like characteristics. The chosen parameterization makes our model an attractive object for exploring “what-if” scenarios; by relating it to recent successful attempts at modeling intra-AS traffic demand, our model provides an initial framework within which one can start exploring the impact of intra-AS traffic engineering on inter-AS traffic engineering and vice versa.

References

- [1] AGARWAL, S., NUCCI, A., AND BHATTACHARYYA, S. Towards Internet-Wide Network Management. In *Proc. of IEEE Infocom 2005 Poster Session* (March 2005).
- [2] AKELLA, A., CHAWLA, S., KANNAN, A., AND SESHAN, S. On the Scaling of Congestion in the Internet Graph. *ACM SIGCOMM Computer Communication Review* 34, 3 (July 2004).
- [3] AKELLA, A., MAGGS, B., SESHAN, S., SHAIKH, A., AND SITARAMAN, R. A Measurement-Based Analysis of Multihoming. In *Proc. of ACM SIGCOMM* (Aug 2003).
- [4] CACHELOGIC, INC. The True Picture of Peer-to-Peer Filesharing. <http://www.cachelogic.com>.
- [5] CHANG, H., GOVINDAN, R., JAMIN, S., SHENKER, S. J., AND WILLINGER, W. Towards Capturing Representative AS-level Internet Topologies. *Computer Networks* 44, 6 (2004).
- [6] FANG, W., AND PETERSON, L. Inter-AS Traffic Patterns and Their Implications. In *Proc. of IEEE Global Internet Symposium* (1999).
- [7] FEDERAL HIGHWAY ADMINISTRATION. Calibrating and Testing a Gravity Model for Any Size Urban Area, August 1983. U.S. Department of Transportation.
- [8] FELDMANN, A., GREENBERG, A., LUND, C., REINGOLD, N., REXFORD, J., AND TRUE, F. Deriving Traffic Demands for Operational IP Networks: Methodology and Experience. *IEEE/ACM Transactions on Networking* 9, 3 (June 2001).
- [9] FELDMANN, A., KAMMENHUBER, N., MAENNEL, O., MAGGS, B., PRISCO, R. D., AND SUNDARAM, R. A Methodology for Estimating Interdomain Web Traffic Demand. In *Proc. of ACM Internet Measurement Conference* (2004).
- [10] GAO, L. On inferring autonomous system relationships on the Internet. In *Proc. IEEE Global Internet Symposium* (Nov 2000).
- [11] KARAGIANNIS, T., PAPAGIANNAKI, D., AND FALOUTSOS, M. BLINC: Multilevel Traffic Classification in the Dark. In *Proceedings of ACM SIGCOMM '05* (2005).
- [12] KENDALL, M. G. *Rank Correlation Methods*. Charles Griffin & Company, 1962.
- [13] KRUSKAL, W. Ordinal Measures of Association. *Journal of the American Statistical Association* 53 (1958).
- [14] MACHIRAJU, S., AND KATZ, R. H. Verifying Global Invariants in Multi-Provider Distributed Systems. In *Proc. of HotNets-III* (2004).
- [15] MAHAJAN, R., WETHERALL, D., AND ANDERSON, T. Towards Coordinated Interdomain Traffic Engineering. In *Proc. of HotNets-III* (2004).
- [16] MEDINA, A., TAFT, N., SALAMATIAN, K., BHATTACHARYYA, S., AND DIOT, C. Traffic Matrix Estimation: Existing Techniques and New Directions. In *Proc. of ACM SIGCOMM* (2002).
- [17] PÖYHÖNEN, P. A tentative model for the volume of trade between countries. *Weltwirtschaftliches Archive* 90 (1963).
- [18] RUBINSTEIN, R. Y. *Simulation and the Monte Carlo Method*. John Wiley & Sons, 1981.
- [19] SANDVINE, INC. Regional characteristics of P2P - File sharing as a multi-application, multi-national phenomenon, 2003. White paper.
- [20] SEN, S., AND WANG, J. Analyzing Peer-To-Peer Traffic Across Large Networks. *IEEE/ACM Transactions on Networking* 12, 2 (April 2004).
- [21] TANGMUNARUNKIT, H., GOVINDAN, R., JAMIN, S., SHENKER, S., AND WILLINGER, W. Network Topology Generators: Degree-Based vs. Structural. In *Proc. of ACM SIGCOMM* (2002).
- [22] UHLIG, S., BONAVENTURE, O., MAGNIN, V., RAPIER, C., AND DERI, L. Implications of the Topological Properties of Internet Traffic on Traffic Engineering. In *Proc. of ACM Symposium on Applied Computing* (2004).
- [23] ZHANG, Y., ROUGHAN, M., DUFFIELD, N., AND GREENBERG, A. Fast Accurate Computation of Large-Scale IP Traffic Matrices from Link Loads. In *Proc. of ACM SIGMETRICS* (2003).

Notes

¹Recent years have seen accelerated migration of dialup modem users to broadband subscription and, according to some recent statistics, DSL subscription in the US nearly doubled in 2003 (point-topic.com).

²Since traffic on the FastTrack/KaZaA network has declined sharply in the past few years and continues to decline, we did not include it in our study.

³The collected measurement data sets are available at <http://topology.eecs.umich.edu/traffic/>.

⁴We thank Manish Karir for making the NetFlow data available to us.

⁵The random parameterization of κ_w and κ_r reflect AS-dependent variations of traffic components.