Title

Connectivity in the South American Internet

Authors
Junqueira, Flavio
Teixeira, Renata

Publication Date
2003-01-13

Peer reviewed

Connectivity in the South American Internet

Flavio Junqueira   Renata Teixeira

1 Introduction

Understanding the reliability and performance of wide-area communications is crucial to the evolution of the Internet. Unfortunately, despite the explosive growth of the Internet in other countries, research in this area has focused predominantly on the US. Most recent research efforts have focused on four aspects: network stability, routing, traffic patterns, and topological analysis.

Labovitz et al. present a study on Internet network stability both inter- and intra-domain traffic [1]. They observed a high level of path fluctuation and a low MTBF for individual paths. Chandra et al. defined a network failure model and derived values for the parameters of this model [2]. Based on these values, they evaluate several techniques for overcoming network failures. In terms of routing, previous works assess the efficiency of the mechanisms currently implemented on the Internet. Savage et al. showed that the path taken between a given source-destination pair is many times suboptimal [3]. Estrin et al. [4] investigate the impact of routing policies on Internet paths. According to their study, many paths are inflated by five router-level hops or more. Other works address the problem of estimating traffic demands. Fang et al. focus on traffic patterns among ASs. The main conclusion of their work is that a small percentage of the flows is responsible for most of the traffic [5]. Feldmann et al. describe a methodology for deriving traffic demands based on flow-level measurements on ingress links [6]. Finally, some research efforts focus on observing topology characteristics. Gao proposes heuristics for inferring AS relationships [7]. Huffaker et al., as part of the CAIDA effort on providing information about the Internet structure, are the only group we are aware of that measures the Internet outside the US. They present a study of network connectivity in the Asia-Pacific region [8].

In this paper, we address connectivity in South America. Our experimental framework is similar to that adopted by Huffaker et al. We analyze routing traces from sources inside and outside South America in order to investigate the behavior of communication both within and from outside the continent. In particular, we examine the round-trip time (RTT) and the routing paths between sources and destinations.

This paper proceeds as follows: Section 2 describes our approach for collecting and processing connectivity data. Section 3 analyzes RTTs and routing paths for the different datasets. Finally, we present our conclusions and suggestions for future work in Section 4.

2 Methodology

In order to evaluate connectivity we observed traces to a set of destinations in South America from different sources. We obtained data that specify nodes in the path from source to destination and round-trip times (RTTs). This data is then extended to contain geographic and autonomous system (AS) information. Based on this information we determine ASs that connect countries in South America, the RTTs between source and destination, and the location of nodes in the routing paths.

2.1 Data Collection

We used data gathered from inside and outside South America. We selected one source in Rio de Janeiro, Brazil (GTA) and another in Buenos Aires, Argentina (HAB). The Brazilian source is a restricted-use machine which belongs to the Federal University of Rio de Janeiro, whereas the Argentinian one is a traceroute server available through a CGI interface. The third source is in Oregon, US. The Oregon source is a *skitter* box maintained by the CAIDA organization. This box probes daily several destinations around the world, including sites in South America.

Table 1: Number of destination hosts per country.

Country	*Domain Suffix*	*Number of Destinations*
Argentina	ar	26
Bolivia	bo	22
Brazil	br	198
Chile	cl	32
Colombia	co	18
Ecuador	ec	12
Guyana	gy	10
Peru	pe	18
Paraguay	py	21
Suriname	sr	6
Uruguay	uy	16
Venezuela	ve	14
Overall		393

Table 2: Percentage of destination hosts per domain type.

Domain Type	*Percentage*
com	40.04 %
gov	12.47 %
edu	9.2 %
net	3.72 %
others	34.57 %

First, we needed to build the list of destinations to probe from the two South American sources. To construct the list, we used the *google* [9] search engine specifying domain names ending with a South American suffix, resulting a list of 456 destinations. A close analysis of the original list of destinations determined that 55 of those were not in South America. One destination was hosted in the Netherlands, and the rest in the US despite their South American domain suffixes. In addition to the destinations hosted outside South America, there were 8 destinations that were not reachable during the experiments. Therefore, we discarded 14 % of the destinations in our original list. Table 1 presents the number of destinations per country in the final list. Table 2 shows the distribution of destination in the different domain types. It is important to stress that a considerable number of universities in South America do not use the suffix `.edu` and hence fall into the *others* category.

Initially, the list is randomized to avoid concentrating the measurements to particular countries at certain times of the day. For each destination on the randomized list, the source in Rio de Janeiro performed a traceroute every 30 seconds. The source in Buenos Aires is a traceroute server, to which we issued a request every 30 seconds from the same randomized list. For each pair of source and destination, five traces were obtained. Our first set of measurements from Rio, to which we refer as GTA-01, were taken from 02/23 to 02/24. Unfortunately, these days correspond to Carnaval, an important holiday in Brazil which we suspect may have affected traffic patterns. Thus, we conducted a new set of measurement from Rio between 02/28 and 03/02 called GTA-02. The dataset from Buenos Aires is called HAB-01 and was collected from 03/02 to 03/05.

The CAIDA machine at Oregon uses a different list of destinations. In the data provided to us, there are 4038 destinations in South America and only one trace per destination. The traces were executed on 02/06 and 02/07 and call this dataset CAIDA-01.

2.2 Data Processing

The datasets have information on both RTTs and intermediate nodes in the routing path. To compute the distance between source and destination, and between intermediate nodes we need to have the latitude and longitude of each node in the path. Using a set of tools from CAIDA, it is possible to determine geographic and AS information about the routers on a path. NetGeo is a system implemented using a client-server architecture which returns latitude and longitude of a router given its IP address. AS numbers are determined by a function in the CoralReef perl library. Given a dump of a BGP routing table, `get_as` returns the AS number for an IP address. The BGP routing table information utilized was downloaded from the National Laboratory for Applied Network Research (NLANR) [10].

The geographical information provided by the NetGeo server is based on entries returned by WHOIS servers. The NetGeo server parses these entries looking for any piece of information that indicates the origin of the IP address. This approach, however, is not entirely accurate. WHOIS entries often only identify the contact address of the company to which an address block is assigned. This does not necessarily correspond to the physical location of specific routers. Moreover, the granularity of the latitude and longitude is not fine-grained. In the best case it refers to a city and in the worst case to an entire country. Whenever it was possible to infer location data from other sources, such as web pages or router names, we modified the data manually.

Customer-provider relationships are also a source of problems for the NetGeo data. If a block of addresses is assigned to a provider and a subset of those is assigned by the provider to a customer, it is not possible for NetGeo to identify this relationship if there is no entry for the customer in the WHOIS servers. We attempt to infer this information from other means, including browsing web pages and inspecting router names.

A problem shared by CoralReef and NetGeo is related to 1918 addresses [11]. There are three blocks of addresses assigned for private use, over which the Internet Assigned Numbers Authority (IANA) has no control. These addresses are to be used in machines inside organizations and are not supposed to appear in any host connected directly to the Internet. It is impossible to determine geographic information or AS number for these routers since these addresses are not assigned to any specific organization. There are paths, however, that include routers with 1918 addresses. Usually this 1918 address appear between two routers of the same organization. Therefore, it is possible to infer the data for a router with an 1918 address from previous and next hops.

Finally, multi-homed networks introduce multiple entries in the routing table for the same address prefix. For addresses in multi-homed networks, the `get_as` function returns an error indicating multiple origins. In most cases, it is possible to infer the AS number from the previous hop. However, this is not always accurate because the AS of the previous hop may not be in the set of possible origins. In this case, we choose a random AS from the set of possible origins.

In summary, the tools used to determine geographic and AS information are not perfect due to limitations of the knowledge-base used by them. Consequently, it was necessary to manually modify the datasets in order to increase the robustness of our results.

3 Results

In this section we investigate connectivity from both outside and inside South America based on the datasets GTA-01, GTA-02, HAB-01, and CAIDA-01. First, we present the round-trip time exhibited in those datasets and compare the countries in South America according to this measurement. Then, we discuss the paths traversed by packets to South America and their most important ASs.

3.1 Round-trip Time

Figure 1 shows the variability of round-trip time as a function of geographic distance for each dataset. The three graphs show a relatively high degree of RTT variability. Note that traces originating in Rio (Figure 1(a)) have RTTs varying from 3 ms to 3000 ms, and 90 % of the traces have RTTs less than 762 ms. In the traces from Buenos Aires (Figure 1(b)) RTTs vary from 14 ms to 2800 ms, but only 41 % of them have RTTs shorter than 762 ms. Figure 1(c) shows the dataset with the source outside South America. RTTs from Oregon are more variable (the y-axis is in log scale), ranging from 97 ms to 16900 ms. One interesting observation is that approximately 83 % of the traces from outside South America have RTTs less than 762 ms, more than double the percentage of traces from Buenos Aires.

Nevertheless, comparing absolute RTTs for the three different sources is not accurate for a number of reasons. First, datasets GTA-01, GTA-02, and HAB-01 probe a different set of sites than those probed in CAIDA-01. Second, the measurements were taken in different days, and even if taken in the same day it is hard to synchronize all the sources to probe the same sites at the same time. Further, RTT is inherently related to distance, hence traces from outside South America should present higher RTTs. Another factor that influences the measurements is that the number of sites probed in Brazil is much larger than those probed for other countries. Thus, measurements from Brazil present lower RTTs.

In order to overcome some of this factors, we randomly discarded traces to Brazil in GTA-01, GTA-02, and HAB-01 datasets leaving only 26 destinations in Brazil. Since all the traces were collected during weekdays, we assume that the traffic is similar. Instead of comparing absolute values of RTT, we divide the shortest geographic distance between source and destination by the RTT between them. This measurement is the speed with which packets are transferred. Figure 2 shows the cumulative distribution of speed in miles/ms for each of the datasets. We also show the speed of light (186 miles/ms) that is the theoretical maximum speed. We observe that CAIDA-01 has the highest speed, with more than 65 % of the traces exhibiting speeds exceeding 10 miles/ms or more. GTA-01 and GTA-02 are the second best with approximately 70 % of the traces with speed more than 2.5 miles/ms. HAB-01 has the lowest speed, 75 % of the traces present speed less than 2.5 miles/ms. We examine the reasons for this behavior in Section 3.2.

Figure 3 compares the speed among the different countries in South America for traces with the source in Rio. Since the source is inside Brazil, some of the traces within

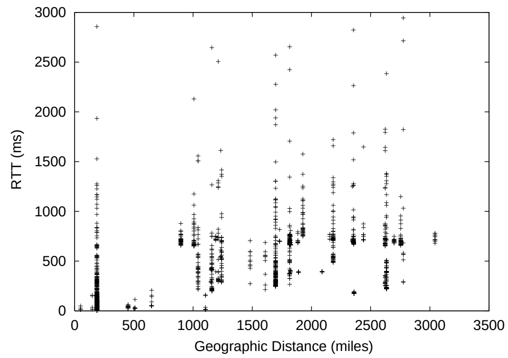

(a) GTA-02

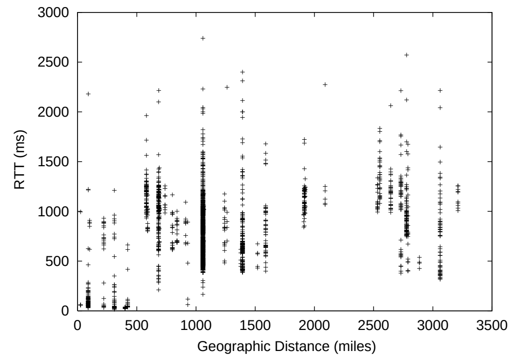

(b) HAB-01

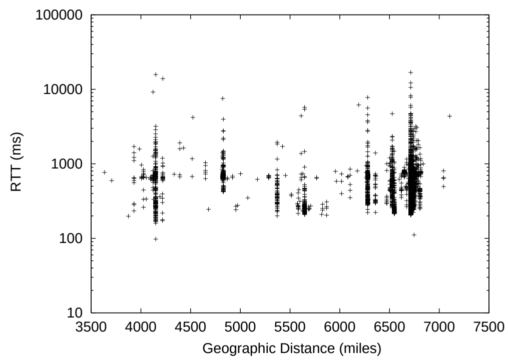

(c) CAIDA-01

Figure 1: Round-trip time by distance for each traces.

Brazil present small RTTs and hence high speeds: 33 % of those have speed higher than 10 miles/ms. The traces to Brazil also present the largest speed range – from 0.07 miles/ms to 97.74 miles/ms. Note that for the majority of the other countries the speed range is small. This means that from Brazil to each one of these countries the ratio between geographic distance and RTT is approximately the same. The same behavior is observed in the traces from Argentina as shown in Figure 4.

As seen in Figure 3 and Figure 4 the country that has lowest speed is Paraguay and the one with highest speed is Venezuela. Paraguay and Ecuador exhibit delays above 650 ms from Brazil and above 900 ms from Argentina. Suriname also has a minimum RTT of 750 ms from Brazil and 1000 ms from Argentina. This minimum delay suggests that these traces traverse a satellite link. AS data presented in the next section confirm that Paraguay, Ecuador, and Suriname are connected to Brazil and Argentina via satellite links.

3.2 Routing Paths

In this section we discuss the paths traversed by the traces from both inside and outside South America. In order to determine whether the paths to South America follow the shortest geographic path, we compared the sum of the distance between each hop (*idist*) and the shortest distance between source and destination (*gdist*). Figure 5 shows the ratio

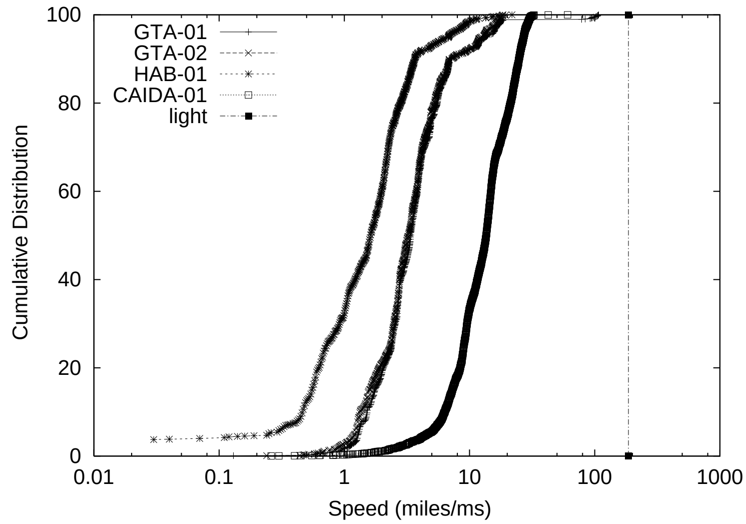

Figure 2: Cumulative distribution of speed (overall measurements).

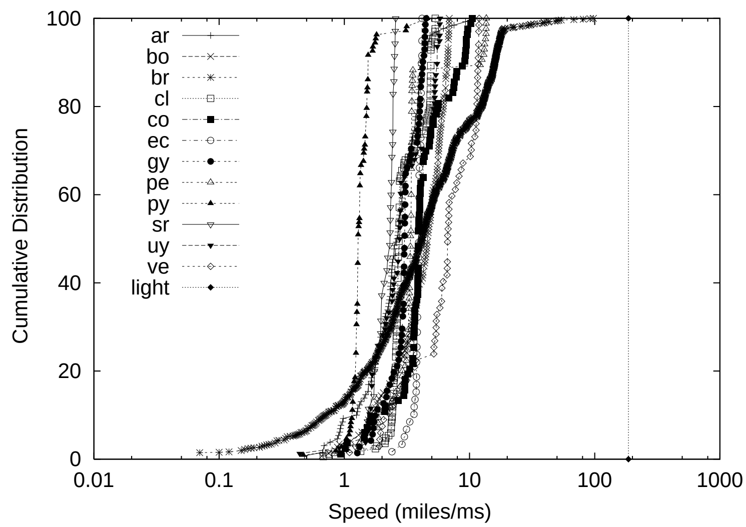

Figure 3: Cumulative distribution of speed (GTA-02).

idist/*gdist* for each datasets. For the reasons discussed in Section 3.1 the datasets GTA-01, GTA-02, and HAB-01 were modified to contain only 26 destinations in Brazil. A ratio of 1 means that *idist* equals *gdist*, i.e., the routing path used by the trace corresponds to the shortest geographic path between source and destination. A ratio of x means that the path used in the trace is x times longer than the shortest geographic path.

Figure 5 shows that for 99.8 % of traces from Oregon *idist* is at most 4 times longer than *gdist*. Paths originating in Rio, both GTA-01 and GTA-02, are at most 15 times longer than the shortest path for 99 % of traces while paths from Buenos Aires are 55 times longer for the same percentage of traces. Note that some traces from HAB-01 are more than 128 times longer than the shortest path.

Figures 6 and 7 present the cumulative distribution of *idist*/*gdist* for each country. The majority (95 %) of traces from Rio to destinations in Brazil follow the shortest path. This

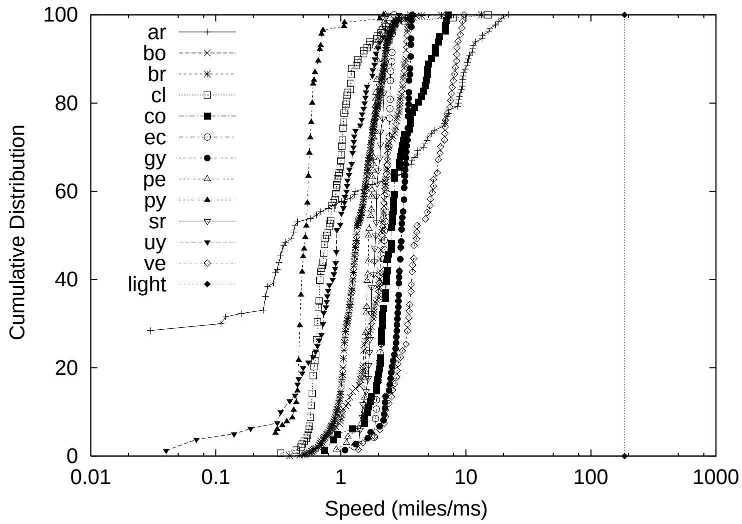

Figure 4: Cumulative distribution of speed (HAB-01).

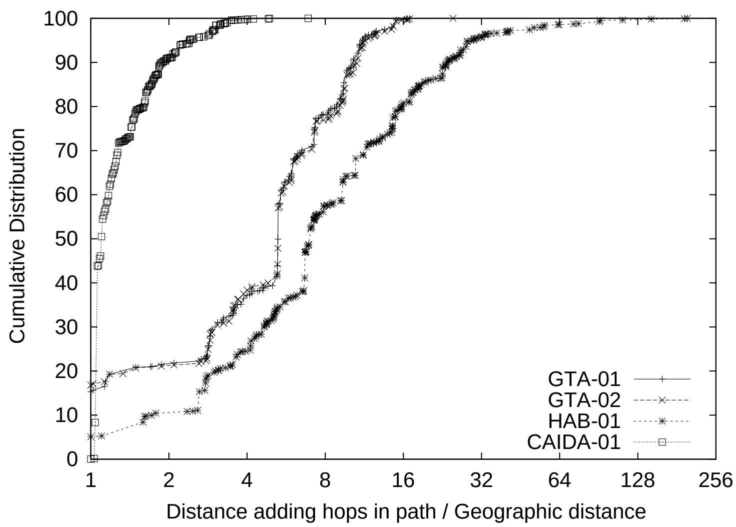

Figure 5: Cumulative distribution of distance adding up the geographic distance between individual hops in the path versus the geographic distance from source to destination (overall measurements).

shows that inside Brazil there is good connectivity. Only 45 % of the traces from Buenos Aires to destinations in Argentina have a ratio of 1. As in 3.1, Paraguay exhibits the worst *idist*/*gdist* ratio. The longer paths experienced by Paraguay are one reason for the lower speeds. In Figure 6 we observe a ratio equal to 1 in 80 % of the traces for Uruguay. By examining the AS paths from Rio to Uruguay we discovered that around 80 % of the traces go via Embratel, a Brazilian provider, straight to ANTEL1, the most important AS in Uruguay (see Table 5).

The high ratio of the distance traversed by the paths to the shortest geographic distance led us to investigate the paths with sources and destinations in South America. Tables 3 and 4 present the percentage of paths from GTA-02 and HAB-01 respectively that go to

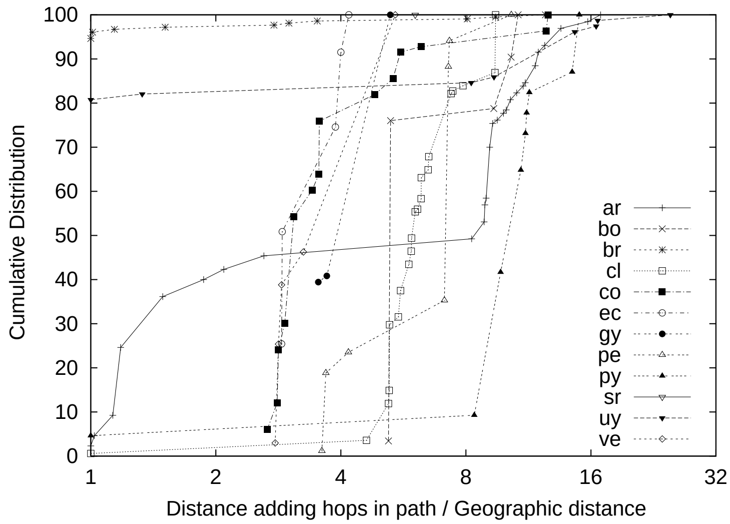

Figure 6: Cumulative distribution of distance adding up the geographic distance between individual hops in the path versus the geographic distance from source to destination (GTA-02).

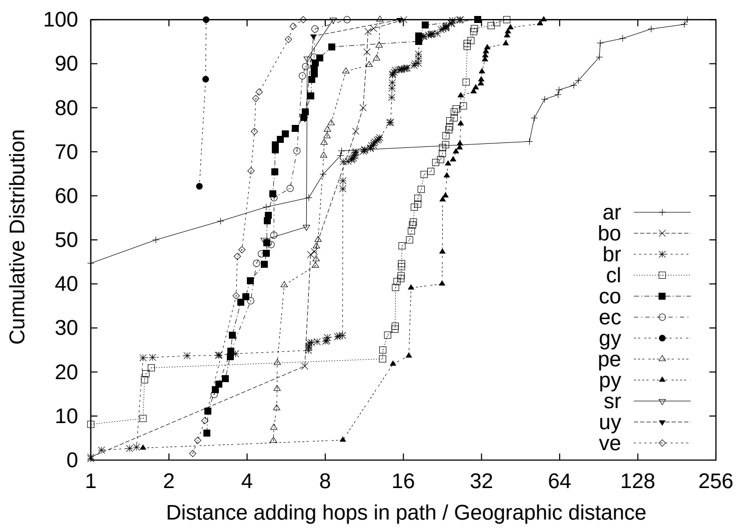

Figure 7: Cumulative distribution of distance adding up the geographic distance between individual hops in the path versus the geographic distance from source to destination (HAB-01).

North America and Europe. Note that for traces originating in Rio 7 out of 12 countries have 100 % of the paths going to North America and 100 % of traces from Rio to Peru and to Suriname go to both North America and Europe. Traces with source in Buenos Aires exhibit similar behavior. We observed a sample trace from Buenos Aires to Porto Alegre, in the south of Brazil. The path taken is: Buenos Aires, Paris, Frankfurt, Amsterdam, Washington D.C., New York, São Paulo, and finally Porto Alegre.

Table 5 presents the stub ASs and the last transit ASs for each country based on data from GTA-02. The results for the other datasets are similar. We obtained this information

Table 3: Percentage of paths that go outside South America (GTA-02).

Country	*North America*	*Europe*
Argentina	54.62 %	11.54 %
Bolivia	100 %	27.4 %
Brazil	0.46 %	0.46 %
Chile	99.4 %	30.36 %
Colombia	100 %	22.89 %
Ecuador	100 %	33.9 %
Guyana	100 %	0 %
Peru	100 %	100 %
Paraguay	95.37 %	41.67 %
Suriname	100 %	100 %
Uruguay	17.95 %	10.26 %
Venezuela	100 %	53.73 %
Overall	42.9 %	17 %

Table 4: Percentage of paths that go outside South America (HAB-01).

Country	*North America*	*Europe*
Argentina	20 %	16.15 %
Bolivia	89.33 %	65.33 %
Brazil	73.87 %	52.64 %
Chile	79.05 %	59.46 %
Colombia	100 %	58.02 %
Ecuador	100 %	87.23 %
Guyana	100 %	13.51 %
Peru	100 %	70.59 %
Paraguay	96.52 %	82.61 %
Suriname	100 %	97.06 %
Uruguay	0 %	0 %
Venezuela	100 %	80.6 %
Overall	75.07 %	53.21 %

by considering the last AS in the AS path to be the stub AS and the second to last to be the last transit AS. This assumption, however, does not always hold because of customer-provider relationships and because a stub in an AS path may also be a transit in another AS path. Brazil has the largest number of stub ASs (37) before Argentina (14). Bolivia, Guyana, and Suriname are serviced by only one provider: BolNet, Cable & Wireless, and Interpacket respectively. Paraguay, Guyana, and Suriname have no local provider. Even though Paraguay has 8 different stub ASs, approximately 90 % of them are registered in the US and the remaining are registered in Brazil or Argentina. The main stub ASs in Paraguay are Panamsat and SAT-TEL and in Suriname is Interpacket. These three ASs provide Internet connectivity via satellite. This explains the RTTs presented in Section 3.1. Note that the majority of last transit ASs are in the US. The ASs outside South America that appear in most paths to South America are UUNet (USA), Cable & Wireless (USA), France Telecom (France), and Teleglobe (Canada).

Table 5: ASs per country (GTA-02).

country	Last Transit AS		Stub AS	
	#ASs	*Main ASs*	*#ASs*	*Main ASs*
ar	8	TASF/AR (35.16 %) IMPSAT-AR (23.44 %) Telintar/AR (19.53 %)	14	Advance/AR (19.53 %)
bo	3	AT&T/USA (72.41 %) Telecom Italia (24.14 %)	1	BolNet/BO (100 %)
br	10	RNP/BR (47.88 %) Embratel/BR (18.30 %)	37	Embratel/BR (42.87 %)
cl	10	CWUSA (29.17 %) Americatel/USA (16.67 %)	9	PROVDESERV/CL (26.19 %) RdC Internet/CL (19.64 %) Entel/CL (16.67 %)
co	10	UUNet/USA (21.69 %) IMPSAT-BR (18.07 %) Global One/CO (16.87 %)	11	IMPSAT-AR (18.07 %)
ec	3	Interpacket/USA (49.15 %) UUNet/USA (25.42 %) IMPSAT-BR (25.42 %)	4	Cyberweb/EC (32.2 %) Publicom/USA (25.42 %) IMPSAT-AR (25.42 %) ECUANET/EC (16.95 %)
gy	2	CWUSA (59.15 %) UUNet/USA (40.85 %)	1	CWUSA-2BLK (100 %)
pe	4	12956 (70.59 %) UUNet/USA (18.82 %)	4	Telefonica del Peru (76.47 %)
py	6	UUNet (46.73 %) CWUSA (23.36 %)	8	Panamsat/USA (37.38 %) SAT-TEL/USA (23.36 %)
sr	1	UUNet/USA (100 %)	1	Interpacket/USA (100 %)
uy	4	Embratel/BR (82.05 %)	2	ANTEL1/UY (83.33 %) AS1797/UY (16.67 %)
ve	4	Sprint Inter./USA (53.73 %) Sprint/USA (22.39 %) Publicom/USA (16.41 %)	5	Global One/VE (53.73 %) 2027(16.42 %)

4 Conclusion

In this paper we aimed to study connectivity in the South American Internet. We have presented an analysis of four datasets with traces to destinations in South America. The study characterizes connectivity in terms of RTTs and both AS and routing paths. Based on this information we have determined that all communication from Brazil and Argentina to the majority of the other countries in South America traverses North American ASs. Countries like Bolivia, Guyana, and Suriname have only one stub AS. Hence all the communication to those countries depend on these ASs. We observed that Paraguay perform worst in both speed and *idist*/*gdist* ratio. The reason for the worst performance is that Paraguay is close to both our sources, but it is only connected to other countries via satellite links, which present high RTTs, and all its paths go through the US, thereby having a high value of *idist*. Suriname and Ecuador, which are also connected via satellite, are located farther from both our sources, thereby performing better according to those metrics.

Our study would benefit from increasing the number of sources. In particular, having at least one source per country would enable further analysis of communication between countries and within each country. Another important next step is to add other performance metrics like loss rate, congestion, and bandwidth. Finally, a study of peering relationships among South American providers and with outside providers would further the understanding of path selection.

5 Acknowledgments

We would like to thank Bradley Huffaker for providing us data from the *skitter* box, Otto Duarte and José Rezende for allowing us to gather data using their machines, Stefan Savage and Geoff Voelker for the helpful discussions, and Alan Su and Henri Casanova for the comments on the paper.

References

[1] C. Labovitz, A. Ahuja, and F. Jahanian, "Experimental Study of Internet Stability and Wide-Area Backbone Failures," in *The 29th International Symposium on Fault-Tolerant Computing*, (Madison, Wisconsin), June 1999.

[2] B. Chandra, M. Dahlin, L. Gao, and A. Nayate, "End-to-end WAN Service Availability," in *USENIX USITS*, (San Francisco, CA), Mar. 2001.

[3] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson, "The End-to-end Effects of Internet Path Selection," in *ACM SIGCOMM*, (Cambridge, MA), Sept. 1999.

[4] H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin, "The Impact of Routing Policy on Internet Paths," in *INFOCOM*, (Anchorage, AK), Apr. 2001.

[5] W. Fang and L. Peterson, "Inter-AS Traffic Patterns and Their Implications," in *IEEE GLOBECOM*, (Rio de Janeiro, Brazil), Dec. 1999.

[6] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving Traffic Demands for Operational IP Networks: Methodology and Experience," in *ACM SIGCOMM*, (Stockholm, Sweden), Aug. 2000.

[7] L. Gao., "On Inferring Autonomous System Relationships in the Internet," in *IEEE GLOBECOM*, (San Francisco, CA), Nov. 2000.

[8] B. Huffaker, M. Fomenkov, D. Moore, E. Nemeth, and K. Claffy, "Measurements of the Internet topology in the Asia-Pacific Region," in *INET*, (Yokohama, Japan), July 2000.

[9] "Google wepage at `http://www.google.com/`."

[10] "NLANR webpage at `http://moat.nlanr.net/AS/`."

[11] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, *Address Allocation for Private Internets.* RFC-1918, Feb. 1996.