Rating Autonomous Systems

Laurent Zimmerli
Open Systems AG
Zurich, Switzerland
Email: laz@open.ch

Bernhard Tellenbach
ETH Zurich
Zurich, Switzerland
Email: tellenbach@tik.ee.ethz.ch

Arno Wagner
Independent Consultant
Zurich, Switzerland
Email: arno@wagner.name

Bernhard Plattner
ETH Zurich
Zurich, Switzerland
Email: plattner@tik.ee.ethz.ch

***Abstract*—Rating Autonomous Systems helps in establishing and maintaining mission critical Internet communication paths. We elaborate performance metrics, tools, and quality indicators to rate Autonomous Systems. An initial rating approach based on traceroute measurements led to the discovery of the frequent effect of non-increasing round-trip times in traceroute measurements. Our improved outlier-based rating approach addresses this issue and allows the real-time detection of Autonomous Systems causing poor Internet connection performance as well as the comparison of Autonomous Systems against each other over an extended period of time.**

I. Introduction

Data sent over the Internet typically traverses more than one Autonomous System (AS). The performance of each traversed AS is crucial for the performance of the end-to-end transfer. Small disruptions in an AS, such as congestion of network links, configuration errors, or router failures, can lead to significant performance degradations. Furthermore, incidents of the past, for example earthquakes, have demonstrated that disruptions within one or multiple AS's degrade the performance of data transfers significantly [10] and for a considerable amount of time (the routing infrastructure requires time to adapt to such an abrupt degradation). Detecting and, if possible, avoiding AS's exposing poor performance more often than others, is therefore critical to providing reliable and fast Internet based services. Avoiding one or multiple AS's can be done by e.g., dynamically switching to an upstream provider which routes traffic on a different path, or by setting up a path that bypasses the low-quality AS using an overlay network.

In this paper, we present a simple but efficient way to measure the quality of AS's. Our contributions are the following:

- We propose and analyze an initial approach to rate the quality of AS's using paris-traceroute [3] and tcptraceroute [11]. We report on the shortcomings of traceroute measurements (i.e., non-increasing round-trip times).
- We propose replacements for the problematic metrics and explore them using a simple yet beneficial AS quality rating.
- We report on the results of extensive experiments monitoring the full mesh connections of 159 hosts located in 68 different countries during 38 days.

The remainder of this paper is organized as follows: The following section introduces performance metrics, tools and quality indicators for AS's. Section III explains the initial rating approach, discusses traceroute issues and presents improvements to cope with them. The evaluation of the rating and the discussion of the results are presented in Section IV. Section V describes related work, and Section VI concludes the paper.

II. AS Quality

Assessing the overall quality of an AS is difficult since it is almost impossible to e.g., collect and include information about all of the paths in an AS. Because of this, we limit the scope of our rating to a specific *rating domain*. The rating domain is a set of end-hosts of particular interest communicating with each other on a regular basis. Examples are the VPN overlays operated by Open Systems AG. One of those VPNs forms the basis of our measurements. It interconnects 159 gateways spread among 68 countries world-wide. Therefore, measurements can only be performed from end-hosts (VPN gateway) to end-hosts (VPN gateway) and no sensors can be placed on the connections between them. Another important restriction arising from the fact that our measurements must not interfere with the processes running on the end-hosts (VPN gateways) or the services provided by them.

A. Performance Metrics

The IP Performance Metrics Working Group of the IETF recommends the following set of metrics [9] for Internet data transfer quality description:

- connectivity
- one-way and round-trip delay and loss
- delay jitter
- loss patterns
- packet reordering
- bulk transport and link bandwidth capacity

Because we want to rate the sub paths within every traversed AS's in our rating domain, and not just specific end-to-end paths, we used a restricted set of metrics and evaluated their individual impact: round-trip packet delay and loss, delay variation and *AS internal hop count*. In addition, we use the characteristics of the *outgoing AS-path*, i.e. the sequence of AS's between the current AS and the destination, as part of a quality metric.

A possible extension is to do more precise measurements of the links within each AS wherever access to a host in the AS is available. This is planned as a future extension.

In this work, all measurements were performed on a full mesh of the rating domain, i.e. on every possible end-to-end

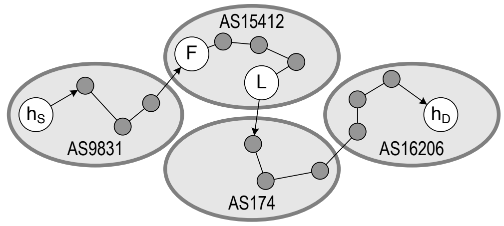

Fig. 1. The AS-RTT of AS15412 is RTT to the last host **L** in AS15412 minus the RTT to the first host **F** in AS15412

connection of the rating domain. One measurement on an end-to-end connection is referred to as a *measurement connection*.

B. Traceroute Measurements

The best tool for our measurements was a traceroute-based software, combined with AS lookup services. We used paris-traceroute as well as tcptraceroute to perform the traceroute measurements. Both tools do not suffer from traceroute anomalies caused by routers that employ load balancing on packet header fields, as described by Augustin et al. [2]. Our measurements consist of traceroute runs between all end-hosts of the rating domain. The measurements were done uniformly randomly distributed with an inter-measurement time between 10 and 20 minutes. According to Baccelli et al. [4] Poisson Sampling does not have advantages over a uniform random distribution when doing *rare probing*.

C. Quality Indicators

The following set of quality indicators form the basis of our rating approach described in Section III:

a) AS-RTT: The AS-RTT represents the time a packet spends within an AS. It is calculated by subtracting the RTT to the first host within an AS on the chosen path, from the RTT to the last host in the AS. For example, in Fig. 1 the AS-RTT of AS15412 is the RTT to *F* minus the RTT to *L*.

b) Packet Loss: The round-trip packet loss is the number of packets lost within an AS. There are two different types of loss: Packets which are lost within an AS are called *internal loss*. Packets which are lost between hosts in neighboring AS's are called *border loss*. Border losses are added to the packet loss measurement for both involved AS's.

c) Internal Hop Count: The internal hop count is the number of hops within the AS on the measurement connection.

d) Outgoing AS-Path: The outgoing AS-path of an AS is the sequence of AS's traversed on the measurement connection following the AS until the destination is reached. In Fig. 1 the outgoing AS-path of AS15412 is *174,16206*.

e) Outgoing RTT: The RTT from the AS to the destination of the measurement.

f) Uplinks and Peers: There are two possibilities to determine uplink and peers of an AS. One is querying a public database holding information about peerings. An alternative is recording the AS's to which an AS forwards the packets on the measurement connections. The first method gives the complete set of neighbors of an AS. The second one determines only the relevant neighbors of the AS within the rating domain, which is sufficient for our purpose.

D. AS Measurements

We refer to the parameters measured for an AS using one measurement connection as an *AS measurement*. For example, the measurement connection in Fig. 1 delivers AS measurements for AS9831, AS15412, AS174, and AS16206.

III. AS Rating

As a result of the measurements each AS receives an AS rating. This rating consists of two values related to the two main tasks of an AS: Routing packets through its own network and maintaining connections to other (high quality) AS's:

- *Intra-AS quality:* Reflects the network performance within an AS
- *Inter-AS quality:* Reflects the quality of its neighbors.

A. Initial Rating Approach

Our initial approach to rate the quality of Autonomous Systems exposed important shortcomings of traceroute measurements. Nevertheless, we discuss the initial approach for two reasons: **(1)** We propose and explore replacements for the problematic measurements for fixing the initial approach. However, we do not yet reintegrate them but explore their potential using a simpler AS quality rating first. **(2)** It serves as an entry point to the discussion of traceroute issues with regard to measurements at AS granularity.

1) Delay Stability: A smaller variation of the AS-RTT indicates higher stability and hence better quality of the AS (with regard to this indicator). The variation of the AS-RTT differs for different paths through an AS and different sets of entry- and exit points. We group the AS measurements into subsets which are *comparable*. The RTT variation is then computed within each set of *comparable* measurements. We call two AS measurements *comparable* if the measurement paths enter and leave the AS at the same points.

The top part of Fig. 2 shows two non-comparable measurement connections through AS15412. The three measurement paths on the bottom of the figure are comparable for AS15412.

There are different ways to capture the variation of a set of values. One approach is by computing the coefficient of variation (COV) which is only meaningful for unimodal and symmetrical distributions. Another is by using the interquartile range (IQR) [12] which does not require a specific data distribution. The IQR is defined as $IQR = Q_3 - Q_1$, with Q_i denoting the i-th quartile.

Using the IQR, we can compute the delay stability of each set of comparable AS measurements as follows: Normalize the IQR values by the median AS-RTT of the comparable subset and take the median of the normalized IQRs. This stability value can then be used for the comparison of different AS's.

2) Internal Loss: We normalize each internal loss count for each set of comparable measurements by dividing it by the internal hop count. The quality measure for internal loss is then the median of all normalized values.

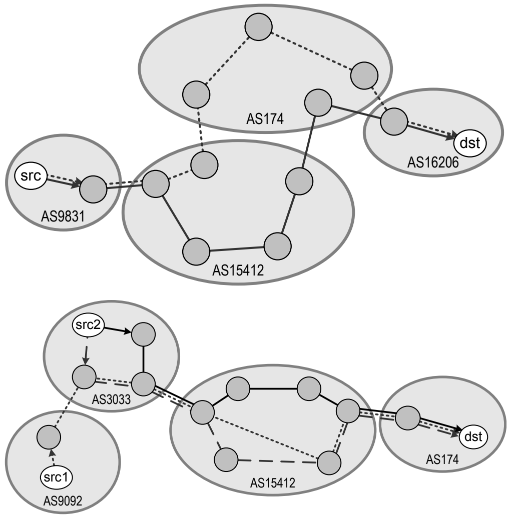

Fig. 2. **Top:** Two non-comparable measurement paths for AS15412. **Bottom:** Three comparable measurement paths for AS15412.

3) Border Loss: Border loss is similar to internal loss. However, border loss does not have to be normalized since there is only one hop on a connection between two different AS's (on the IP layer, which is the relevant layer for our work).

4) Internal Hop Count Stability: We correlate the stability of the internal hop count with the delay stability (cf. Section III-A1). For example unstable internal hop count combined with high delay stability is good, since it indicates good load balancing. And stable internal hop count combined with low delay stability is bad, since it indicates bad or no load balancing. To calculate the stability of the internal hop count, we follow the same approach as for the AS-RTT.

5) Outgoing AS-Path Adaptation: The outgoing AS-path stability is similar to the internal hop count stability. However, instead of correlating the internal hop count with the internal delay stability, the outgoing AS-path is correlated with the outgoing delay. An unstable AS-path correlated with stable outgoing delay indicates good adaptation of forwarding paths, whereas a stable AS-path with unstable outgoing delay indicates no or bad adaptations.

B. Traceroute Issue

Delay stability, internal hop count stability, and the outgoing AS-path adaptation measure are based on the AS-RTT. The AS-RTT results from the difference of two round trip times to intermediate hosts of the traceroute output.

After starting the measurements, we analyzed a first set of results. We discovered that the AS-RTT was often negative, i.e. the RTT of consecutive hops non-increasing. The top of Fig. 3 shows all AS-RTT measurement results of a specific AS for a total of 2600 measurements on one specific measurement connection. The bottom of Fig. 3 shows the plot of the RTTs of a single traceroute measurement versus the hop count to the target host.

In 100'000 traceroute measurements, we found 98% with at least one hop with decreasing RTT. In 92% of the measurements there was at least one hop where the RTT decreased by at least 10% of the total RTT. 83% contained at least one hop where the RTT decreased by 20% of the total RTT or more. Altogether, on 29% of the measured hops the RTT decreased.

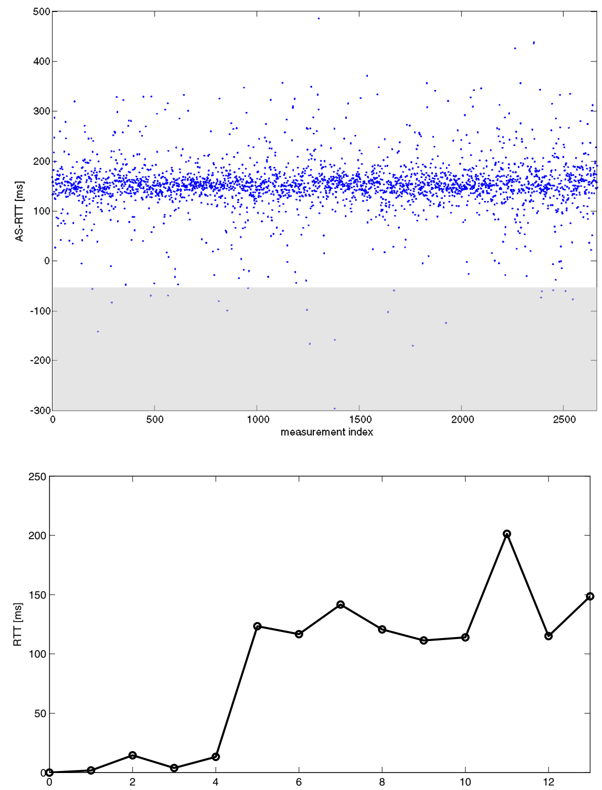

Fig. 3. **Top:** Sample AS-RTTs **Bottom:** Sample traceroute RTTs

The most likely explanation of this effect is that notifications for expired TTLs are a low-priority task in routers. If, on an end-to-end connection, a router which is further away from the source does this faster than one closer to the source, then the measurable RTT can decrease with distance.

As a result, AS-RTT based quality measures suffer from large measurement errors and are not representative for the quality of an AS.

In [14] where the authors describe a model for building a comprehensive network map, this problem is solved by applying a filter to get a single representative hop-by-hop latency value. However, this is not an option if we want to see short term changes, especially when restricted to rare probing. We might simply filter out the interesting outliers.

Overall, our initial AS rating approach turned out to be problematic enough that we needed to replace the AS-RTT and Outgoing RTT by more reliable measures. We explore a promising replacement using a simple yet beneficial AS's quality rating.

C. Improved Rating Approach

The basic idea of the refined approach is the following: If the end-to-end RTT of a measurement is significantly higher than usual, then the AS which causes the outlying RTT is determined and recorded. The real-time measurements are compared to historic data and the AS causing the current poor network performance is detected and its identity recorded.

The actual detection is done as follows: The AS-route is detected for every measurement connection. Then the RTT to the last host of each AS (called the *AS-end-RTT*) is determined. Once an outlying end-to-end RTT is observed, each AS-end-RTT is compared to the historic values to find out whether it is also an outlier. An AS which has an outlying AS-end-RTT and whose following AS's on the AS-route all have outlying AS-end-RTTs, is identified as the one causing the performance degradation.

If the AS-end-RTT of an AS is outlying but one of the following AS on the route does not have an outlying AS-end-RTT, then we can assume that the outlier is caused by a router that sends messages about exceeded TTLs slowly. In this case the observation is discarded to avoid misinterpretations as discussed in Section III-B.

Observing stable AS-end-RTTs gives strong evidence of an AS handling traffic well. Based on our approach, an AS identified to cause a lot of outlying end-to-end RTTs is likely to suffer from frequent temporary bad performance, degrading network performance on connections traversing it. Hence, the initial approach might be fixed by using the stability of the AS-end-RTT and its outliers instead of the *AS-RTT* [1].

D. Outlier Definition

The analysis of our RTT measurement data shows that RTT data has a heavy-tailed statistical distribution. An outlier definition suitable for heavy-tailed and other distributions can be found in [12]. The definition is based on the interquartile range as follows: An outlier is any value which is more than three IQRs below the first or above the third quartile. That is, an outlier is any value outside of the range:

$$[Q_1 - 3 \times IQR, Q_3 + 3 \times IQR] \text{ with } IQR := Q_3 - Q_1 \quad (1)$$

where Q_i denotes the i-th quartile. This outlier definition does not assume any specific statistical distribution and is therefore applicable to heavy-tailed distributions. Figure 4 illustrates the outlier definition for a sample series.

E. Quality Measures

In order to explore if the AS-end-RTT and its outliers are well suited to characterize the quality of an AS, we define a simple and intuitive quality measures and leave the reintegration with our initial approach to future work.

Our primary measure for intra-AS quality is the ratio of the number of times the AS caused an outlier to the total number of end-to-end measurements including the AS:

$$q_{intra}(AS_x) = \frac{times\ AS_x\ caused\ outlier}{total\ measurements\ in\ AS_x} \quad (2)$$

The lower the intra-AS quality (called *outlier-ratio*) the better the quality of the AS. Note that the reliability of the outlier-ratio depends on the number of measurements including the AS. When comparing two AS's, it is important to not only compare the outlier-ratio, but also to take the number of measurements for each AS into account.

[1]And combining the same information for the remaining AS's on the path to the destination could replace the *outgoing RTT*

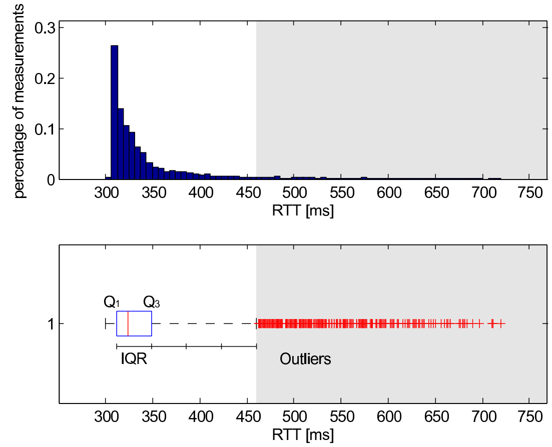

Fig. 4. Histogram and corresponding boxplot illustrating the outliers of a sample measurement series

Our primary measure for inter-AS quality is the weighted sum of the qualities of all neighbors of an AS (see II-C):

$$q_{inter}(AS) = \sum_{i \in neighbor(AS)} weight(i) \times q_{intra}(i) \quad (3)$$

The weight of a neighbor is the fraction of all measurement connections through the AS which had this specific neighbor as the next AS in the outgoing AS path.

The given quality measures represent the long-term quality of the AS's. Hence, the improved rating approach does not give real-time quality measurements, but only quality measurements for a certain observation time interval. Nevertheless, it is still possible to detect the AS responsible for a connection performance degradation in real-time, by comparing a real-time measurement to recorded data the same way it is done during measurement data collection.

IV. Evaluation

The measurements were performed on a VPN overlay network operated by Open Systems, spanning 159 hosts located in 68 countries on five continents. 223 different Autonomous Systems were involved. The duration of the measurements was 38 days between April, 22nd and May, 29th in 2008. The measurement data includes:

Total measurements:	15’991’277
Measurements not reaching the destination:	1’588’900
Valid measurements used for AS ratings:	10’095’928
Measurements with outliers:	406’084
Outliers with detected responsible AS:	234’137

A. Intra-AS Rating

The intra-AS ranking rates the AS's according to their outlier-ratio. The best and worst AS's found in the 38 days ratings are shown in Table I.

TABLE I
MEASUREMENT RESULT: BEST AND WORST INTRA-AS QUALITIES

Rank	ASN	# Measurements	# Outliers	Outlier-Ratio
1	16260	37691	0	0
2	24748	12931	0	0
⋮	⋮	⋮	⋮	⋮
39	39010	99743	1	0.00001
40	8402	93004	1	0.00001
⋮	⋮	⋮	⋮	⋮
220	5384	100337	4997	0.05237
221	19262	100416	5368	0.05645

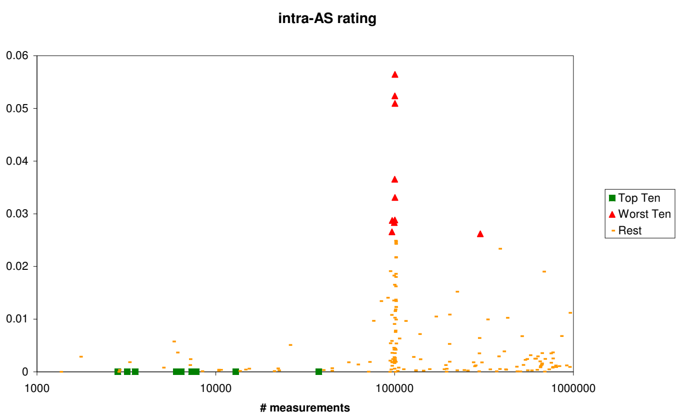

Fig. 5. Scatter plot of the intra-AS quality

The column *Measurements* in Table I shows how many times the AS was traversed by a valid measurement. Every measurement which reached its destination and contains all AS-end-RTTs is valid. The column *Outliers* shows how often the AS caused an outlying end-to-end RTT. Finally, the last column lists the outlier-ratio, which is the actual intra-AS quality measure.

As can be seen, some AS's were traversed much more frequently than others. As a consequence, the rating is more reliable for some AS's than for others. To show the reliability level of the rating, we give a scatter plot of the results. Figure 5 shows the number of measurements an AS was involved in versus its outlier-ratio for the intra-AS rating.

The scatter plot can be interpreted as follows: The lower an AS on the plot, the better its outlier-ratio, i.e., the better its intra-AS quality. The more an AS to the right of the plot, the more reliable the statement about its quality.

The red triangles indicate the worst ten AS's whereas the green squares show the best ten AS's of our rating.

B. Inter-AS Rating

In order to determine the inter-AS rating, for each AS, we determined to which other AS's it routed connections during the measurements as well as how often that happened for each neighbor. Table II shows the results of one example AS.

The inter-AS quality of the top and worst Autonomous Systems are shown in Table III, Fig. 6 shows the scatter plot.

The scatter plot in Fig. 6 shows the number of times an AS routed a connection to another AS versus the weighted sum

TABLE II
NEIGHBORS OF AS1120

ASN	neigh. ASN	count	total count	weight	neigh. quality
1120	5588	1644	3343	0.4918	0.0017
1120	8447	1654	3343	0.4948	0.0288
1120	8514	2	3343	0.0006	0.0123
1120	12558	43	3343	0.0129	0.0054

TABLE III
MEASUREMENT RESULT: BEST AND WORST INTER-AS QUALITIES

Rank	ASN	Routing Count	Inter-AS Quality
1	8501	3264	0
2	13284	8	0
3	41398	3264	0
4	3786	314902	0.0001
⋮	⋮	⋮	⋮
156	27699	52881	0.0287
157	8961	74139	0.0524

of the intra-AS qualities of all its neighbors. The lower an AS is on the vertical axis, the better the qualities of its neighbors. The more an AS is located to the right of the plot, the more reliable the statement about its inter-AS quality. Again, the red triangles indicate the worst ten AS's whereas the green squares show the best ten AS's of our rating.

C. Discussion

Using an outlier-based rating approach, we were able to compare AS's against each other, and thereby to examine the network quality of different ISPs and intermediate networks over an extended period of time. Furthermore, using the collected measurement data, the generated rating data allows us to determine Autonomous Systems with emerging poor performance in real-time. The long term rating supports the selection of local ISPs by choosing the available ISP whose AS performs best. One example for the comparison of ISPs are the two AS's 3786 and 4766. They are operated by two different ISPs in South Korea. Both AS's were part of our rating. Figure 7 illustrates their considerably differing intra-AS qualities on the scatter plot. As can be seen, AS3786 clearly outperforms AS4766 in terms of intra-AS quality. The real-time rating can e.g., be used in fail-over scenarios where (temporary) re-routing over an overlay network is an option or where upstream providers with different routes toward the destination exist. This is especially useful and efficient if the majority of the Internet connections of a host have the same destination. Remote locations of global businesses communicating with application servers or an Internet proxy located at the major location of the company are an example for such a scenario. Note that while real-time detection is possible, our evaluation makes use of data collected beforehand.

V. RELATED WORK

There are rating approaches for Autonomous Systems according to different criteria. E.g. CAIDA[2] performs AS rankings according to topology criteria and customer cone size [5].

[2] Cooperative Association for Internet Data Analysis (CAIDA). http://www.caida.org/, last visit August 2008

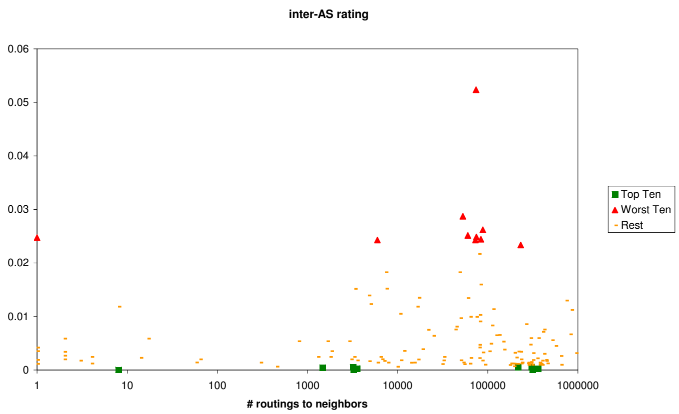

Fig. 6. Scatter plot of the inter-AS quality rating

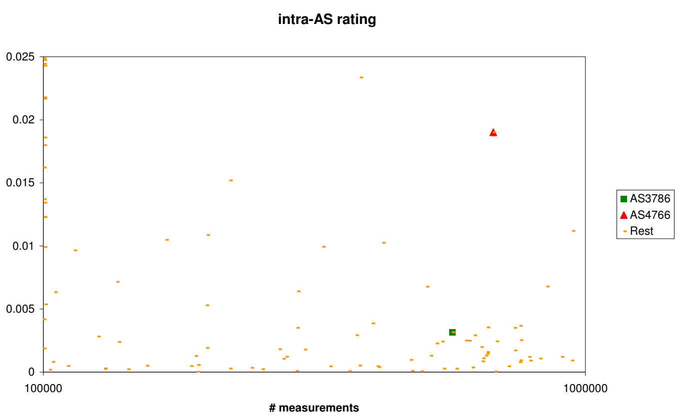

Fig. 7. AS3786 is represented by the green square whereas AS4766 is shown as red triangle

However, we are not aware of any work that actually does a quality-rating of Autonomous Systems or defines AS-centric quality measures according to network performance. Any known Internet performance measurements relate to end-to-end connections, none of them focuses on AS related sub paths like we do. The work that is probably most related to our work is [14] where the authors describe a model for building a comprehensive network map with more than just connectivity information. They combine connectivity, latency, and routing information into an interface-level network map based on a large collection of traceroute data.

Furthermore there are several approaches for measuring Internet performance. E.g. Chen et al. [6] surveyed different projects and tools for Internet performance monitoring. However, no work has been found focussing the preformance monitoring on traversed AS's rather then measuring the performance of end-to-end connections.

Alves et al. [1] characterized the Internet connections evolution over the last 10 years at the Autonomous Systems level analyzing BGP data. Dimitropoulos et al. [8] introduced heuristics for inferring AS relationships—which delivered better results than from BGP data—and offer a data repository at CAIDA containing up-to-date data about AS relationships. Dimitropoulos et al. [7] also classified AS's and calculated relevant characteristics for all current AS's from that time. Wang et al. ranked Autonomous Systems according to their routing contribution to the Internet using their IDAV model [13].

All of these results relate to characteristics of AS's, but do not study their quality.

VI. Conclusion

We propose an approach to rate the quality of Autonomous Systems based on traceroute measurements. It was investigated using a test environment consisting of 159 hosts located in 68 different countries on all five continents. Data collection has shown that traceroute measurements often suffer from anomalies expressed by non-increasing RTTs of consecutive hops. These anomalies have been explored using an outlier-based rating approach. The approach allows comparison of AS's over extended periods of time as well as real-time detection of bad performing AS's.

Rating Autonomous Systems supports the maintenance and operation of mission critical Internet connections. Businesses, which rely on the quality of Internet links, can benefit from the rating to improve their service quality. Rating Autonomous systems also helps in real-time network debugging by detecting weak AS's on poor performing Internet connections and it supports the selection of best performing local ISPs.

References

[1] Alves, N., de Albuquerque, M. P., Albuquerque, M. P., and de Assis, J. T. *Topology and Shortest Path Length Evolution of The Internet Autonomous Systems Interconnectivity*. PhD thesis, Universidade do Estado do Rio de Janeiro, 2008.

[2] Augustin, B., Cuvellier, X., Orgogozo, B., Viger, F., Friedman, T., Latapy, M., Magnien, C., and Teixeira, R. Avoiding traceroute anomalies with paris traceroute. In *IMC '06: Proceedings of the 6th ACM SIGCOMM conference on Internet measurement* (New York, NY, USA, 2006), ACM, pp. 153–158.

[3] Augustin, B. et al. Paris traceroute. http://www.paris-traceroute.net.

[4] Baccelli, F., Machiraju, S., Veitch, D., and Bolot, J. C. The role of pasta in network measurement. *CCR 36*, 4 (2006), 231–242.

[5] CAIDA. As ranking. http://as-rank.caida.org/. last visit February 2009.

[6] Chen, T., and Hu, L. Internet performance monitoring. In *Proceedings of the IEEE* (September 2002).

[7] Dimitropoulos, X., Krioukov, D., Riley, G. F., and Claffy, K. Classifying the types of autonomous systems in the internet. In *SIGCOMM Poster* (2005), vol. 151.

[8] Dimitropoulos, X. A., Krioukov, D. V., Fomenkov, M., Huffaker, B., Hyun, Y., Claffy, K. C., and Riley, G. F. As relationships: inference and validation. *CCR 37*, 1 (2007), 29–40.

[9] IP Performance Metrics Working Group. Ip performance metrics (ippm). http://www.ietf.org/html.charters/ippm-charter.html.

[10] Lemon, S. Earthquake disrupts internet access in asia. *www.computerworld.com* (27.12.2006).

[11] Toren, M. C. Tcp traceroute. http://michael.toren.net/code/tcptraceroute/. last visit November 2008.

[12] Tukey, J. W. *Exploratory data analysis*. Addison-Wesley, 1977.

[13] Wang, Y., Wang, Y., Chen, M., and Li, X. Inter-domain access volume model: Ranking autonomous systems. In *The International Conference on Information Networking (ICOIN)* (January 2007).

[14] Wang, Z., Jin, C., and Jamin, S. Network maps beyond connectivity. *IEEE GLOBECOM* (Dec. 2005).