A Longitudinal Analysis of Traffic Policing Across the Web

Tobias Flach, Luis Pedrosa, Ethan Katz-Bassett, and Ramesh Govindan

Technical Report 15-961
Department of Computer Science
University of Southern California

1 Introduction

Internet traffic has increased 5× in 5 years [4], much of it from high-volume services such as cloud storage and peer-to-peer sharing, and from the explosion of streaming video. YouTube and Netflix alone combine to contribute nearly half of the traffic to North American Internet users [8]. This is driven by the fact that these services deliver high-volume traffic, and by the vast popularity of some of these services—YouTube has one billion unique users per month [12] and more than 12% of the US population uses Netflix [6].

This high-volume traffic and its performance is important: users want a quality Internet experience; content providers rely on it for revenue; and Internet Service Providers (ISPs) must cope with delivering its volume. While content providers want to maximize the user quality of experience for their services, an ISP needs to accommodate traffic from a multitude of services and users, often through different service agreements such as tiered data plans. High-volume services like streaming video and bulk downloads (e.g., software updates) that require high goodput must coexist with smaller volume Web services that require low latency.

The question of how to manage high-volume traffic has generated both technical and policy discussions. Disputes arise over how to efficiently deliver it [9], and the question of how to handle the growing volumes of traffic has become important enough that even the President of the United States recently weighed in [7]. Content providers spend considerable effort optimizing their infrastructure to deliver data from a server to the client as well as possible [2, 5, 10].

This high-volume traffic has been subject to several forms of traffic management for several years now. A commonly-deployed traffic management mechanism, *policing*, limits a flow to a preconfigured throughput rate, for example to enforce a bandwidth corresponding to a data plan purchased by a user, with any traffic exceeding the rate being dropped immediately. However, a traffic *policer* (the logical entity in a router or middlebox that performs policing) can often be configured to accommodate short bursts that exceed the rate limit.

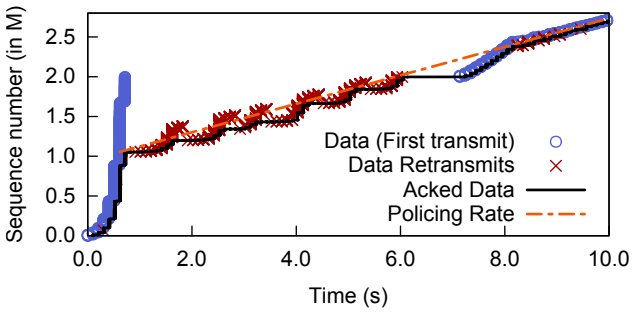

Figure 1: TCP sequence graph for a policed server-to-client flow.

To get a better grasp for the potential impact of policing on a flow, consider the trace in Figure 1. It shows the time-sequence plot of a policed flow collected in a lab experiment. The flow ramps up quickly to over 15 Mbps without any loss events. At $t \approx 1s$, the policer starts to throttle the connection to a rate of 1.5 Mbps. Since packets are transmitted at a rate exceeding the policed rate by an order of magnitude, most of them are dropped by the policer and retransmitted over a 5-second period (up to $t \approx 6.5s$). Following the delivery of the first 2MB, the sender remains idle for one second until more application data becomes available. Since the flow does not exhaust its allotted bandwidth in this time frame, the policer briefly allows the connection to exceed the policing rate once the sender resumes transmitting ($t \approx 7.5s$ till $t \approx 9s$). Overall, the flow observes a loss rate of over 30%.

Understanding Policing. Besides scarce anecdotal evidence [3, 11], little is known about how traffic policing is deployed in practice. Thus, we conduct an exhaustive study of traffic policing observed in the M-Lab NDT Dataset[1]. This work complements a larger study currently under submission where we analyzed policing in the context of traffic between CDN servers of a large video content provider and clients all over the world [1]. The details about the policing detection algorithm can be found in the joint study. We'll focus purely

[1] http://measurementlab.net/tools/ndt

on the analysis of the NDT dataset here.

2 Methodology

The NDT dataset comprises of millions of packet traces and metadata collected over the past seven years. Each trace is the result of a diagnostic task manually triggered between a client machine one of many M-Lab vantage points. While the dataset should not be seen as good representation of all Web traffic, it serves as a valuable addition to the CDN traffic that we analyzed in conjunction to this work [1]. Since the NDT dataset entails over 79TB of data we focused our work on a sample by only looking at the data from the first day of each month. In addition we filtered out traces with fewer than 100 packets for relevance.

We analyzed each of the remaining 7.5 million traces through our publicly available policing detection algorithm described in [1]. Each trace represents a single chunk of data, as NDT only performs a single transfer from the server to the client. As such, each trace becomes a datapoint with the following information:

- Timestamp
- The M-Lab vantage point that sent the data
- The client's geo-location as reported via the MaxMind GeoLite2 database[2], including country, continent, and autonomous system number (ASN)
- The result of our policing algorithm, including whether policing is detected and, if so, with what parameters (rate and burst size)
- The loss rate experienced in the trace, measured as the percentage of packets that are retransmitted
- The trace's overall goodput, as well as the goodput experienced until the first packet loss.

These datapoints form the base for our analysis. The collection of packet traces on a global scale over a time frame of more than six years enables us to reason about the prevalence and impact on policing across the Internet, and analyze longitudinal trends. For this we answer the following questions in subsequent sections.

1. How prevalent is policing across traces in the dataset? Are there differences depending on the region and/or client ISP? Did the results change over the past six years?
2. What impact does policing have on the performance of a TCP connection?
3. What policing configurations do we observe? Are there indicators for the underlying root cause leading to the deployment of these policers?

[2]http://www.maxmind.com

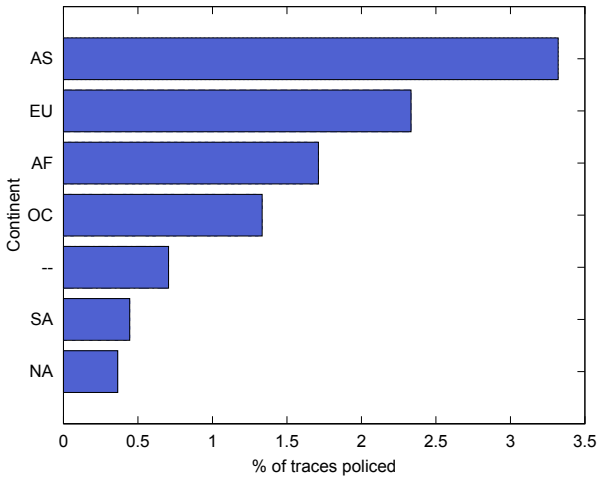

Figure 2: Prevalence of policing in the M-Lab NDT data-set across client continents. – represents clients that could not be geolocated with the MaxMind GeoLite2 Country database

Despite collecting data from clients on a global scale, the M-Lab dataset does have some limitations. NDT is an active measurement toolkit and as such requires the user to trigger the collection of data. While it is integrated into torrent clients like μtorrent and Vuze to reach a larger user base it potentially biases the data collection towards clients with connectivity problems who are more likely to run speed tests, etc. All results should therefore be taken with a grain of salt. Nevertheless we believe that the analysis presented in the following sections provides valuable insights about traffic policing in the wild and complements the passive measurement study conducted from the a content provider's vantage point, discussed in our parent paper [1].

3 Prevalence

We start our analysis with a quantitative look at how frequently we observed policing in our dataset. Overall, our algorithm marked 162,080 traces (2.2%), out of a total of 7.4 million relevant traces, as policed. As shown in Figure 2, the policing frequency varies widely across continents. While less than 0.4% of traces in North America are marked as policed, 3.4% of the traces from Asia are tagged. We see a similar disparity when clustering traces by the client's country as well (Figure 3). The bar diagram displays the policing frequency and the sample size for the countries with the largest fraction of their traces policed. Almost 25% of the traces matched to clients in Kyrgyzstan are policed. While most of the countries in the top-10 list contribute a small number of samples to the overall dataset, there are exceptions like India and Russia for which we have many data points and high policing frequencies (10 and 8%, respectively).

Since our dataset incorporates samples from a 6-year time

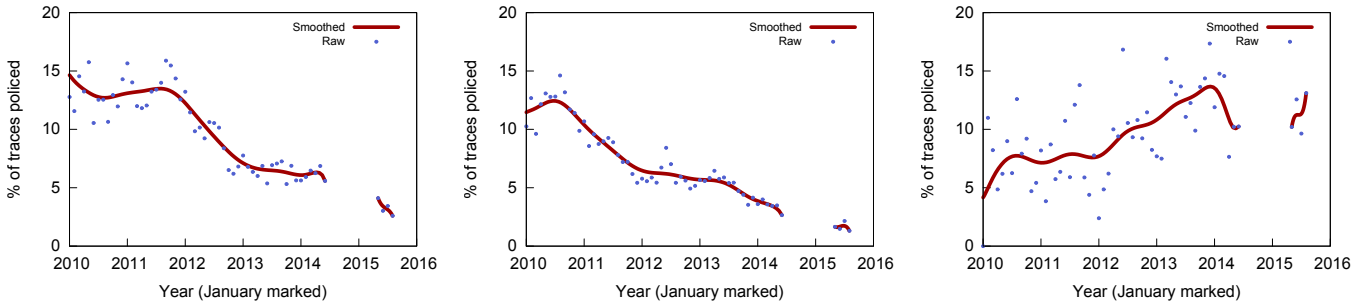

(a) India (b) Russia (c) Algeria

Figure 6: Prevalence of policing over time, depending on the country. The period of inactivity between July 2014 and May 2015 represents a period without traces due to a bug in the NDT software.

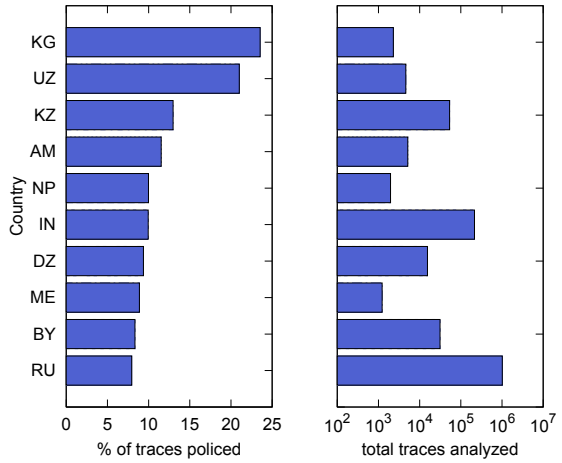

Figure 3: Top 10 countries with the most policing in the M-Lab NDT dataset over the whole measurement period. Countries with fewer than 1000 traces were excluded. The number of traces for each country is also represented on the right to give a measure of statistical relevance.

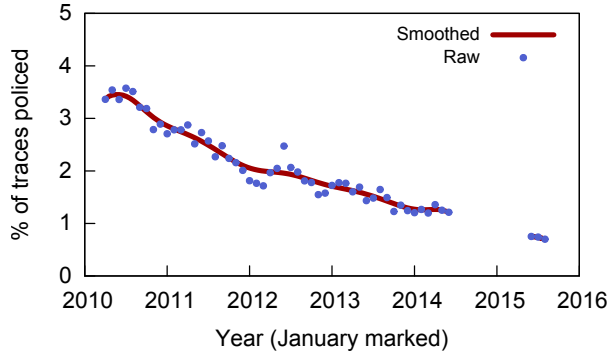

Figure 4: Prevalence of policing over time. The period of inactivity between July 2014 and May 2015 represents a period without traces due to a bug in the NDT software.

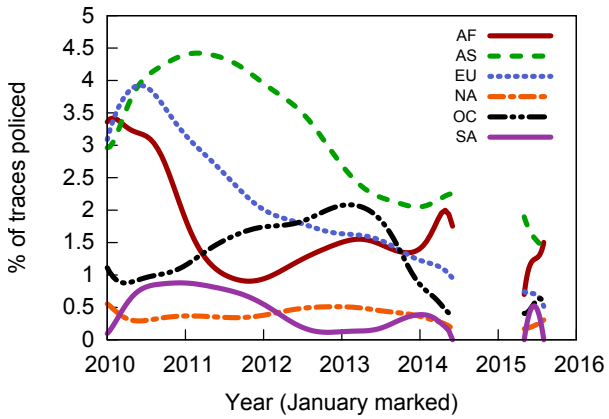

Figure 5: Prevalence of policing over time, depending on the continent. The period of inactivity between July 2014 and May 2015 represents a period without traces due to a bug in the NDT software.

frame (2010 to 2015), we also analyzed longitudinal trends. Figure 4 shows the global policing frequency seen in individual samples (the first day of every month) as well as the long-term trend over the past six years. In the oldest samples we analyzed (from early 2010), we detected policing in about 3.5% of the recorded traces. Over time, policing became less prevalent, with less than 1% of the traces policed in our latest samples. Again we broke down our dataset based on the client's continent (Figure 5) and country (Figure 6). For traces from Asia and Europe we see policing frequencies decline over time, whereas measurements from the remaining continents do not show a clear trend. For the per-country breakdown we selected three of the most policed countries with a substantial number of traces per sample to allow us to analyze long-term trends. For India, roughly 14% of the traces were policed in 2010, compared to less than 4% in late 2015. We observe a similar trend for traces tied to clients in Russia. However, this trend does not apply to all countries.

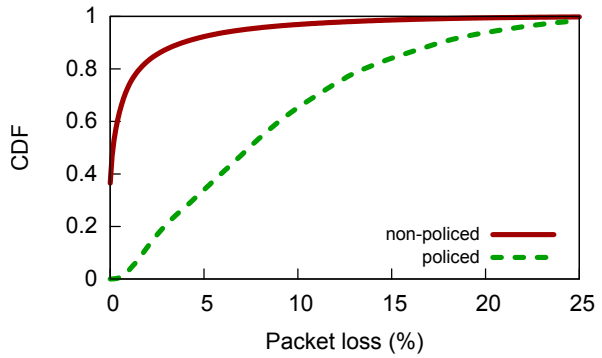

Figure 7: Distribution of packet loss rates seen for policed and unpoliced traces across the whole dataset.

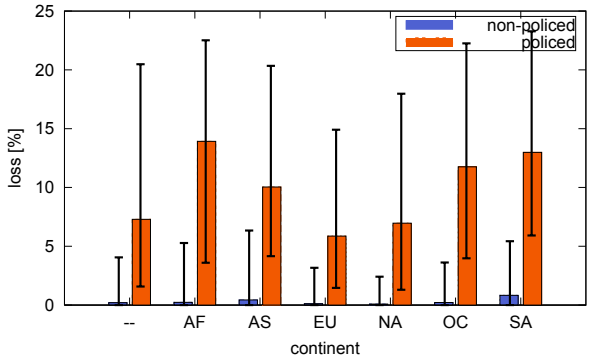

Figure 8: Packet loss rates seen for policed and unpoliced traces per continent (client location). Each bar shows the median, as well as the 10thand 90thpercentile using error bars.

For example, in Algeria we see the opposite trend with policing being more prevalent in the newer samples.

4 Loss Rates

The use of policing is not merely an academic discussion, as it has very real consequences to a connection, in particular with respect to packet loss. Every packet dropped by a policer must subsequently be retransmitted, further contributing to transit costs and network congestion at the content provider, the transit networks, and even the customer ISP which is doing the policing. Figure 7 compares the distribution of loss rates seen in policed vs. unpoliced traces. Generally, loss rates for policed traces are at least an order of magnitude larger compared to their unpoliced counterparts. In the median, we see a 7.4% packet loss when traces are policed vs. 0.14% for non-policed traces. The 90thpercentile further exacerbates this with 17.6% loss for policed traces, vs. 3.9% for unpoliced. Many of the locations with high policing rates also have high loss rates. To rule out the client location as a confounding factor we break down the results by regions with different granularities. We start with a breakdown by continent, as shown in Figure 8. For each continent, we compare the loss rates seen with policed vs. unpoliced traces. Clearly policed traces see much higher loss rates, often at least a magnitude larger compared to unpoliced traces from the same continent. For example, the median loss rate in Africa is 14% in policed traces vs. 0.5% in unpoliced traces. Next, we look at the loss rates when clustering by the client's country. Figure 9) shows results for the top-7 countries, based on the percentage of traces policed per country. Again, policed traces see much higher loss rates compared to unpoliced traces from the same country.

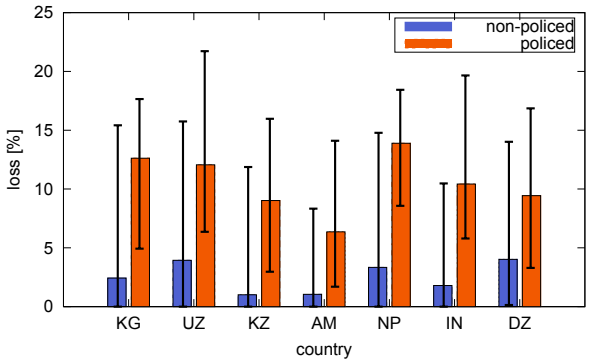

Figure 9: Packet loss rates seen for policed and unpoliced traces in the top-7 countries (based on the percentage of traces policed per country). Each bar shows the median, as well as the 10thand 90thpercentile using error bars.

Finally, we break down results based on the client's AS (Figure 10). A large number of the top ASes, based on the percentage of traces policed per AS, now belong to a large russian provider (Rostelecom). We therefore provide two plots, one for the top-7 ASes within Rostelecom and one for all other ASes. In comparison to the per-continent or per-country figures, the disparity between loss rates for policed and unpoliced traces is smaller. For most ASes the median loss rate is twice as high for policed traces. ASes 8443 and 42575 are an exception in this regard. Loss rates are particularly high for these ASes, even when policing is not detected, suggesting that non-policer induced loss is overshadowing the the effects of policing here.

5 Policing Rates

Having discussed the prevalence of policing and the impact it has on packet loss rates, we now describe the goodput rates enforced by policers when they are present, and how those rates break down for different clients. This also serves the purpose of further validating our policing detection algorithm since, as we will show, the policing rates found in many ASes coincide with the published rates for the contracts they offer to their customers.

Figure 11 shows the policing rate distributions broken

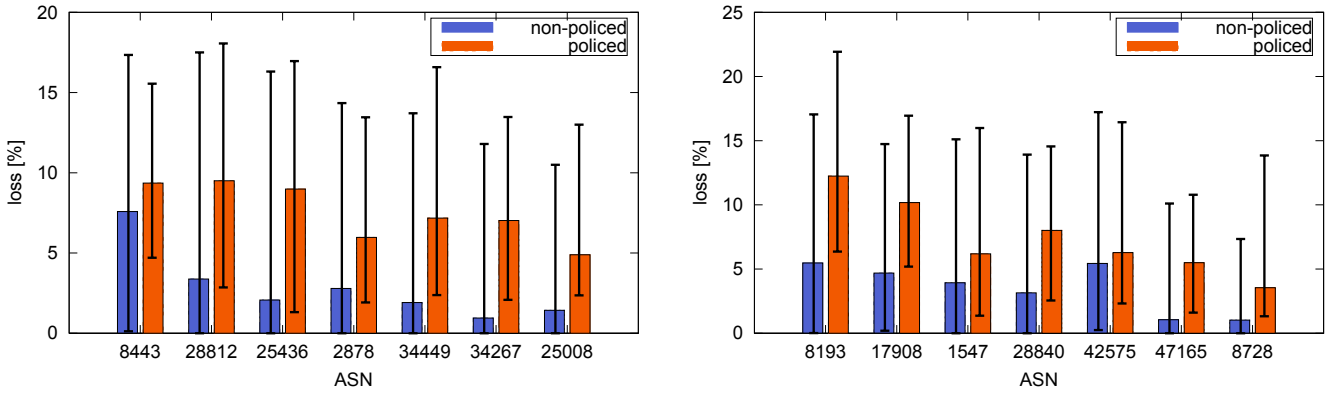

(a) Without Rostelecom

(b) Rostelecom only

Figure 10: Packet loss rates seen for policed and unpoliced traces in the top-7 ASes (separating Rostelecom and non-Rostelecom ASes; based on the percentage of traces policed per AS). Each bar shows the median, as well as the $10^{th}$and 90^{th}percentile using error bars.

ASN	Name	Country (TLD)	Matched to plan rates (Mbps)	Unmatched
8193	Uzbektelekom	Uzbekistan (UZ)	0.125, 0.25, 0.5, 1	None
28840	Tattelecom	Russia (RU)	0.5, 1	1.5
36947	Algerie Telecom	Algeria (DZ)	1, 2, 4	0.5
6697	Beltelecom	Belarus (BY)	1, 2, 3, 4	None
9829	BSNL	India (IN)	0.5, 1	None
6849	Ukrtelecom	Ukraine (UA)	0.5, 2	None
9198	Kazakhtelecom	Kazakhstan (KZ)	None	0.5, 1, 2

Table 1: Top-7 policing ASes using traces after May 2015 only.

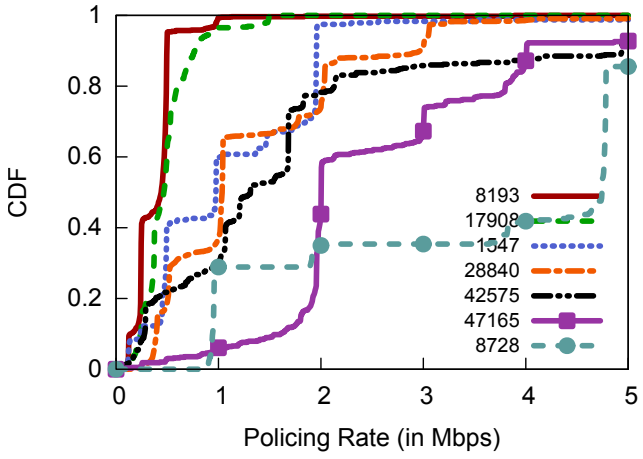

(a) Excluding Rostelecom

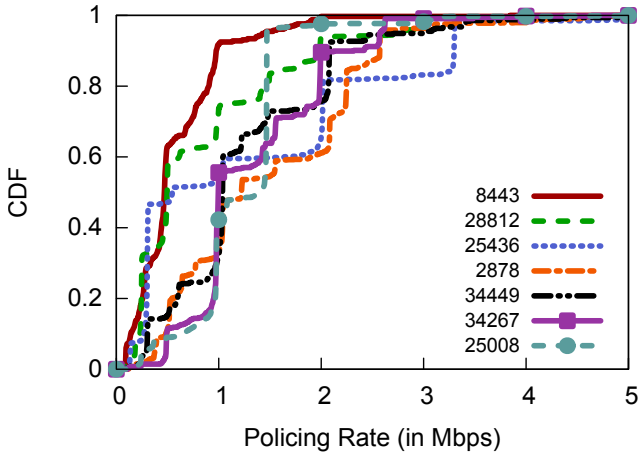

(b) Rostelecom only

Figure 11: Distribution of policing rates observed in the top-7 ASes (by prevalence of policing) excluding Rostelecom (left) and ASes that are now incorporated into Rostelecom (right), over the whole measurement period. Six of the top-10 ASes belong now to Rostelecom.

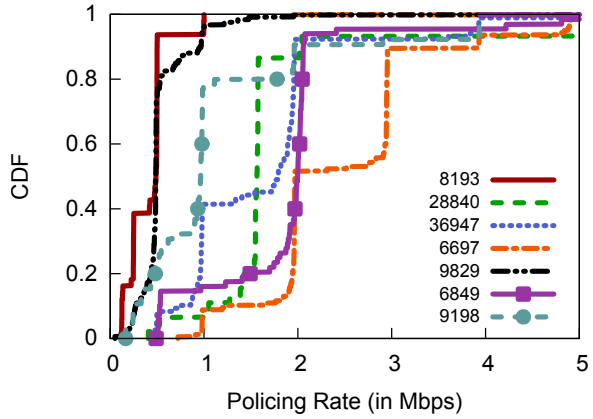

Figure 12: Distribution of policing rates observed in the top-7 policing ASes in 2015.

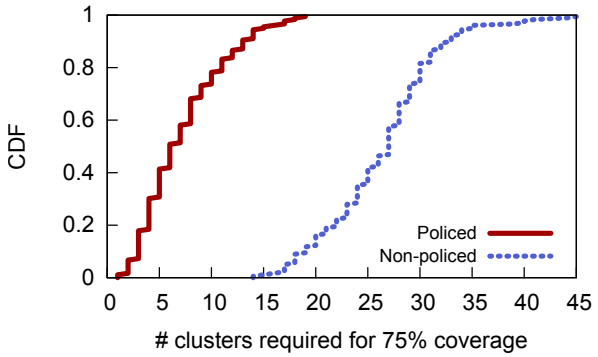

Figure 14: Number of rate clusters required to cover at least 75% of the samples within an AS.

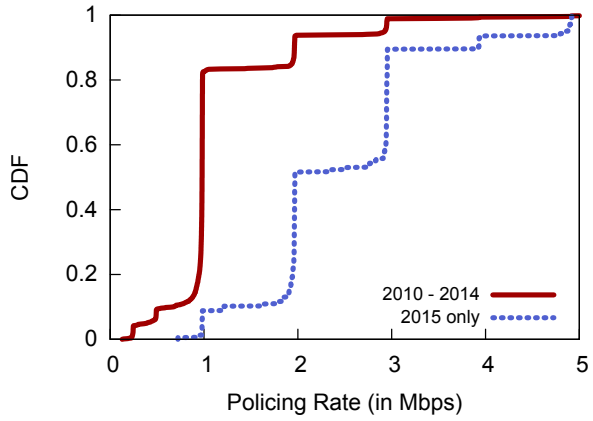

Figure 13: Distribution of policing rates observed in AS 6697, in two different time frames.

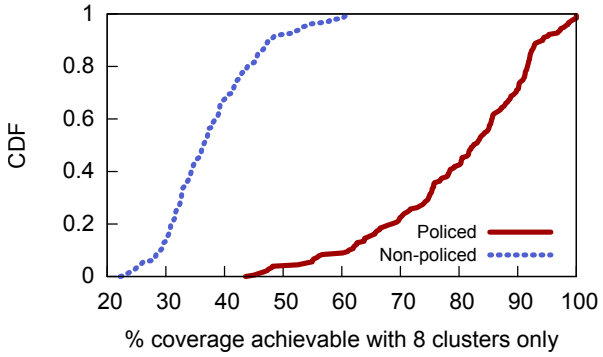

Figure 15: Maximum achievable coverage per AS when using only eight rate clusters.

down for the top ASes based on the fraction of their traces being policed. Rostelecom is a large Russian ISP which now owns six of the top-10 most policed ASes. To get a more representative view of policing rates used by different ISPs we plot the ASes now managed by Rostelecom separately. The distribution for each of the ASes shows a clear staircase pattern with few policing rates dominating per AS. For example, Uzbektelekom (ASN 8193) configures their policers primarily for throughput rates of 0.25 and 0.5 Mbps, with a few traces seeing rates of 0.125 and 1 Mbps. To get a more scoped view we narrow down our measurement window to traces from 2015 only. The results are graphed in Figure 12. Note that a different set of ASes are the top policers in this time frame. This is not surprising. As we mentioned earlier, policing became less prevalent over time. Thus, it is likely that some ASes decided to abandon policing as their method for traffic engineering over time.

For some ASes we see a larger spread of common policing rates for two reasons. First, we aggregate data from the whole six-year measurement period during which the configured policing rates can change. Figure 13 examplifies this for AS 6697. We generated the distribution for samples from 2010 to 2014, and for the samples from 2015 separately. While we observe rates of 1, 2, and 3 Mbps in both distributions (albeit in different quantities), the rates of 0.25, 0.5, 4, and 5 Mbps are only seen in one of the time frames. Second, even if an ISP currently offers a relatively small number of data plans, legacy data plans remain intact resulting in a variety of enforced bandwidths.

To quantify how well clustered the rate distributions are we devised a simple algorithm to identify the most prominent clusters (i.e. rates with a margin of error of $\pm 5\%$). Using this algorithm we can find how many clusters would be needed to represent a certain amount of the identified rates (Figure 14) or, conversely, we can find how many rates are represented by a fixed amount of clusters (Figure 15). The results show that policed traces are far more likely to fall in well defined clusters than non-policed ones. For most ASes, 75% of policed traces can fit in six clusters or less, whereas non-policed traces would need 27. On the flip-side, eight clusters can explain 82% of most policed traces, compared to just 36% for non-policed ones.

Finally, we note that the observed policing rates cluster around round numbers or fractions thereof that are commonly tied to data plans. For the top ASes based on the fraction of traffic policed in 2015, we looked up the data plans and bandwidth rates these ISPs offer to their customers and tried to find matches for the observed policing rates. The results are shown in Table 1. Generally, the policing rates do align with data plan rates with a few exceptions. For example, for Kazakhtelecom we could not find data plans that match any of the policing rates of 0.25, 0.5, and 1 Mbps that we see in our dataset. The current data plans that this ISP offers start at 4 Mbps. It is possible that the policing rates are tied to legacy plans. It is also possible that they reflect rates enforced for oversubscribers, i.e. when a customer exceeds a data limit. Our main paper discusses a case of this observation [1]. In conclusion, we conjecture that policers are indeed used to enforce data plans rather than just to mitigate network congestion.

6 Conclusions

Upon examination of the M-Lab NDT dataset, we find that traffic policing patterns mostly confirm what was seen in our larger and more representative dataset from the wild [1]. As before, policing is mainly prevalent in Asia, and relatively rare in the Americas. We also reconfirm the pattern of significant packet loss induced by policing, with policed traces suffering an order of magnitude more loss than non-policed traces. Finally, while the lack of ground-truth for the NDT data-set precludes a hard evaluation, we once again confirm that the detected policing rates are largely consistent with expected data-plan rates. As a result, we find that the goodput seen for policed traces much more closely fits within a few clusters than for non-policed traces. For the ASes where policing was more prevalent, we also confirm that these clusters largely coincide with the rates they publish for their data-plans.

While the other dataset represented a fairly short time-window, the NDT dataset spans six years, allowing for a more longitudinal study. We find that as the Internet becomes more widely deployed in developing nations, their use of policers to has been reducing over time. Up to 15% of traces were policed in India and Russia in 2010, compared to less than 5% in 2015. This feeds into a similar global trend, as the prevalence of policing reduced world-wide from 3.5% in 2010 to less than 1% in 2015.

We hope that the examination of this publicly available data, in addition to the private data from a large video content provider, will further help inform network operators of the disadvantages of policing over other forms of traffic management. To further help the academic community in studying these patterns, we have made the raw data behind this report publicly available on the accompanying website[3].

7 Acknowledgements

We appreciate M-Lab and NDT making datasets available for analysis. We thank Google for supporting our research. This work was supported in part by the National Science Foundation grants CNS-1564242.

[3] https://usc-nsl.github.io/policing-detection/

References

[1] Anonymous. An Analysis of Traffic Policing in the Wild (under submission), 2016.

[2] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan. Mapping the Expansion of Google's Serving Infrastructure. In *Proc. of the ACM Internet Measurement Conference (IMC '13)*, 2013.

[3] Cisco. Comparing Traffic Policing and Traffic Shaping for Bandwidth Limiting. `http://www.cisco.com/c/en/us/support/docs/quality-of-service-qos/qos-policing/19645-policevsshape.html#traffic`.

[4] Cisco. The Zettabyte Era – Trends and Analysis. White Paper, 2014.

[5] A. Ganjam, F. Siddiqui, J. Zhan, X. Liu, I. Stoica, J. Jiang, V. Sekar, and H. Zhang. C3: Internet-Scale Control Plane for Video Quality Optimization. In *Proc. of the USENIX Symposium on Networked Systems Design and Implementation (NSDI '15)*, 2015.

[6] Netflix. Letter to Shareholders (Q4 2014). `http://ir.netflix.com/results.cfm`.

[7] B. Obama. Net Neutrality: President Obama's Plan for a Free and Open Internet. `http://www.whitehouse.gov/net-neutrality`.

[8] Sandvine. Global Internet Phenomena Report 1H 2014. 2014.

[9] M. Taylor. Verizon's Accidental Mea Culpa. `http://blog.level3.com/open-internet/verizons-accidental-mea-culpa/`.

[10] Te-Yuan Huang and Ramesh Johari and Nick McKeown and Matthew Trunnell and Mark Watson. A Buffer-Based Approach to Rate Adaptation: Evidence from a Large Video Streaming Service. In *Proc. of the ACM Conference of the Special Interest Group on Data Communication (SIGCOMM '14)*, 2014.

[11] C. Wittbrodt. CAR Talk: Configuration Considerations for Cisco's Committed Access Rate. `https://www.nanog.org/meetings/abstract?id=1290`, 1998.

[12] YouTube Statistics. `http://www.youtube.com/yt/press/statistics.html`.