

Providing Public Intradomain Traffic Matrices to the Research Community

Steve Uhlig¹, Bruno Quoitin
Computing Science and Engineering
Department
University of Louvain-la-Neuve, Belgium
{suh,bqu}@info.ucl.ac.be

Jean Leprore², Simon Balon²
Research Unit in Networking
University of Liege, Belgium
{leprore,balon}@run.montefiore.ulg.ac.be

ABSTRACT

This paper presents a new publicly available dataset from GÉANT, the European Research and Educational Network. This dataset consists of traffic matrices built using full IGP routing information, sampled Netflow data and BGP routing information of the GÉANT network, one per 15 minutes interval for several months. Potential benefits of publicly available traffic matrices comprise improving our understanding of real traffic matrices, their dynamics, and to make possible the benchmarking of intradomain traffic engineering methods.

Categories and Subject Descriptors

C.2.3 [Computer-communication Networks]: Network Operations—*Network Monitoring*; C.2.5 [Computer-communication Networks]: Local and Wide-Area Networks—*Internet traffic*

Keywords

traffic statistics, intra-domain traffic matrices

1. MOTIVATION

A lot of effort has been put the last few years on trying to infer traffic matrices based on SNMP link counts [1, 2, 3, 4]. The approach of relying on the raw traffic demand [5, 6] is rarely used as the burden of the measurement and storage infrastructure is significant [7]. Still, recent works [8, 9] indicate that obtaining precise traffic matrices is not out of reach.

Contrary to single capture points traffic traces [10, 11, 12] or BGP routing data [13, 14] for which numerous publicly available datasets exist, publicly available traffic matrices coming from a real network are rare. The only publicly available set of traffic matrices to our knowledge is at <http://www.cs.utexas.edu/~yzhang/research/AbileneTM/> based on data from the Abilene network. Developing intradomain traffic engineering tools or traffic matrix modeling require real datasets to validate the tools or the models. Without publicly available datasets, no comparisons with alternative techniques or models can be performed. To contribute to filling this lack in the networking community, this paper presents a publicly available dataset consisting of intradomain

¹Corresponding author. Steve Uhlig is Postdoctoral fellow of the Belgian National Fund for Scientific Research (F.N.R.S).

²Jean Leprore and Simon Balon are Research fellows of the Belgian National Fund for Scientific Research (F.N.R.S).

traffic matrices coming from the GÉANT network, the European Research and Educational Network.

2. THE GÉANT NETWORK

GÉANT is the pan-European research network and it is operated by DANTE. It carries research traffic from the European National Research and Education Networks (NRENs) connecting universities and research institutions. GÉANT has a PoP in each European country³. All the routers of GÉANT are border routers. GÉANT is composed of 23 routers interconnected using 38 links. In addition, GÉANT has 53 links with other domains. GÉANT uses ISIS to compute its intradomain routes. The IGP weights of GÉANT are mainly based on the inverse of the link capacities with some manual tunings. We obtained a libpcap trace of ISIS for the purpose of building a model of the GÉANT topology.

In order to build an accurate model of GÉANT suitable for the computation of its intradomain traffic matrices, we also obtained from DANTE the interdomain routes known by GÉANT as well as a trace of the traffic transiting across GÉANT [15]. The interdomain routes are obtained from BGP and the traffic trace is collected using Netflow. We describe these two datasets in the following paragraphs.

BGP Routing data

In GÉANT, the BGP routes are collected using a dedicated workstation running GNU Zebra [16], a software implementation of different routing protocols including BGP. The workstation has an iBGP session with all the border routers of the network. Using this technique, it was possible to collect all the best BGP routes selected by the border routers of the AS.

Figure 1 provides the evolution over time of the number of prefixes for each 15 minutes interval during the 4 months. The number of prefixes rises slowly from about 160,000 at the beginning of the considered period, and more than 170,000 at the end. We do not have neither routing data nor Netflow for a 6 days period.

Netflow data

To build an accurate picture of the traffic, we obtained the Netflow statistics collected on all the edge links of the GÉANT network. In order to limit the volume of the Netflow traces, a Netflow sampling rate of 1/1000 is used. This still generates in the order of 150 GB of gzipped traces per month of traffic. Hence, we further summarized the Netflow information by aggregating the raw Netflow flows on

³An overview map of the GÉANT network is publicly available from http://www.geant.net/upload/pdf/Topology_Oct_2004.pdf.




Figure 1: Evolution of prefixes

a \langle source prefix,destination prefix \rangle basis, keeping only the byte volume for each \langle source prefix,destination prefix \rangle pair for each Netflow file. We call the result of the summarization process the "aggregated netflow". We used a single BGP Routing Information Base (RIB) to aggregate the raw Netflow flows, as the purpose is only to limit the size of the traces. Note that this summarization introduces a loss of information concerning the time resolution of a given Netflow flow, as we attribute all flows present in a given Netflow file to the 5 minutes time interval that the file covers.




Figure 2: Evolution of traffic

Figure 2 shows the evolution of the traffic over the four months inside the GÉANT network. We show on the y-axis the amount of bytes seen in the sampled Netflow during each 15 minutes interval, multiplied by 1000 to reverse the sampling performed by Netflow. The total traffic per 15 minutes interval ranges typically between 600 GBytes and 1.4 TBytes. Both the weekly and the daily trends are visible from Figure 2. The period without Netflow statistics and routing data appears as a hole in the traffic evolution.

3. GENERATING TRAFFIC MATRICES

In this section we explain our methodology to generate traffic matrices from the routing and traffic data of the GÉANT network. We rely on the TOTEM toolbox [17] to perform the computation. The TOTEM toolbox provides an open-source framework for integrating various Traffic Engineering algorithms. It targets researchers

as well as network operators that want to implement or test traffic engineering (TE) algorithms. The methods provided within the toolbox cover intra-domain and inter-domain TE, IP-based and MPLS-based TE.

The TOTEM toolbox is a suitable environment for implementing our traffic matrices generation method for the following reasons. First, it defines flexible input and output file formats based on XML to represent traffic matrices (both intra- and interdomain) as well as the topology of networks, including the configuration of an IGP protocol. Second, it contains C-BGP [18], a routing solver that we use to compute both the intra- and interdomain routes known by each router. Finally, by integrating our method in the TOTEM toolbox and generating traffic matrices in the TOTEM format, we allow other researchers to directly make use of our dataset. It is also directly possible to test on these traffic matrices all the intra-domain TE methods that are integrated in the TOTEM toolbox.




Figure 3: Architecture of the TM generation setup.

The architecture of the traffic matrix (TM) generation setup is illustrated on Figure 3. The initial step of the methodology is to provide the TOTEM toolbox with the network topology of GÉANT in XML format. To build the network topology of GÉANT, we relied on a one-day IS-IS trace. We cross-checked the obtained topology with a map of the network provided by DANTE.

The next steps of the methodology are as follows. We compute the routes known by each router to reach each destination prefix. This is done by the C-BGP routing solver. Based on these routes, for each destination prefix, we route the traffic received at each ingress router of the network until it reaches the egress router for this destination. We build a matrix with the sum of traffic going from any ingress router to any egress router. We output this matrix in TOTEM XML format. We do this for every 15 minutes period of our routing and traffic data.

In the publicly available version of the traffic matrices, all IP addresses were anonymized using a random permutation of the IP addresses and replacing the IP addresses by the index of the IP address in the permutation. The network topology in XML format is also anonymized consistently with the traffic matrices in the publicly available dataset. The time information inside the traffic matrices has also been anonymized, by making it begin on January 1 2005 00:00 CET. Only the sequence of 15 minutes time intervals is indicated in the traffic matrices, no absolute time information. For more information and access to the non-anonymized version of the topology and the traffic matrices, please contact the corresponding author.

3.1 Reproducing the routing of the GÉANT network

To route the traffic received at any ingress router towards any egress router, we need to reproduce the routing of the GÉANT net-

work. We use C-BGP, the routing solver contained in the TOTEM toolbox for this purpose. C-BGP computes for each router the next-hop that would have been selected to reach each destination prefix. C-BGP computes the routes inside the GÉANT network by using steady-state models of the IGP and BGP protocols. The IGP protocol is modeled using Dijkstra's SPF algorithm. In contrast, the BGP protocol is modeled using a global linear queue to propagate the BGP messages between routers as well as an accurate implementation of the BGP decision process [19].

C-BGP relies on the network topology and the BGP routing tables of GÉANT. However, we do not have the BGP routing tables for every 15 minutes period since GÉANT only captures its BGP routing tables (RIB) once a day. Fortunately, GÉANT also captures all the BGP updates that occur between two routing table dumps. Hence, we start from a given RIB, then we apply BGP updates till the end of the considered time interval to get the state of the routing at the time of the considered time interval.

The computational complexity of computing the best routes inside C-BGP is directly proportional to the number of prefixes in the routing tables. A full BGP routing table of GÉANT can contain more than 170,000 prefixes. However, when considering the routes announced by all the neighbors of GÉANT, it appears that a lot of prefixes are learned from the same neighbors, with the same BGP quality. The BGP quality of a route is measured based on its attributes (Local-Preference, length of AS-Path, MED, next-hop) [20]. The outcome of the decision process will be the same for these prefixes. Therefore, we group together the prefixes announced with the same attributes by the same neighbor routers in order to decrease the time required to recompute during each time interval the state of the best routes inside the AS.

The number of clusters varies between 350 and 1350, representing between 0.2% to 0.8% of the total number of prefixes. The gain using clustering is thus significant for the scalability of the traffic matrix generation methodology with respect to the number of prefixes present in the BGP routing tables.

3.2 Computing the traffic matrices

One interdomain traffic matrix of GÉANT is a set of triples (ingress router, destination prefix, traffic volume). Each triple represents the traffic volume that has been received by an ingress router and to be sent towards the destination prefix over a 15 minutes period of time.

To compute the intradomain traffic matrix, we replay the flow of the traffic across GÉANT. For this purpose, we take each triple from the interdomain traffic matrix, one at a time. Then, we perform a longest-matching in the routing table computed by the routing solver for the considered ingress router, in order to find the prefix that contains the destination. We then use the route associated with this prefix to route the traffic. We repeat this step on a hop-by-hop basis until the egress router is reached. The intradomain traffic matrix is obtained by keeping the volume of traffic exchanged between all the pairs of ingress and egress routers for the period of time.

4. TRAFFIC MATRICES FROM THE GÉANT NETWORK

In this section we analyze the available traffic matrices to give some of their characteristics. We will begin by providing statistics about the dynamics of the actual (i.e., using actual GÉANT weights) link utilization over the network during the considered period. Then we will conduct some experiments on the dataset to present a subset of the capabilities of the TOTEM toolbox.

4.1 Statistics about the dataset

Figure 4 provides the evolution over the four months of the maximum link utilization inside the GÉANT network, as well as the mean and the percentile 90, for each 15 minutes time interval. We used the actual IGP weights of GÉANT to produce this figure. Note that we have no data for a 7 days period.




Figure 4: Evolution of link load over time (actual IGP weights)

In the GÉANT network, due to the wide range of link capacities (from 155Mbps to 10Gbps), the link with the maximal utilization is always the same low capacity link. This is why we present also the percentile 90 and the mean link utilization because these values reflect more the state of all the links of the network.

Periodic trends are apparent on Figure 4, with weekly and daily patterns. The dataset provides traffic matrices for different days in a week and for different hours in a day. It is interesting to have a closer look at a typical week of our dataset. The Figure 5 shows maximum link load, mean and percentile 90 of the network for a typical week. We can clearly see peak load during work time and a decrease of the load when it is a day off.

4.2 Experiments with TOTEM

Once all the traffic matrices are generated, we can use the TOTEM toolbox to run some experiments on the dataset. With TOTEM, it is easy to load the GÉANT topology and one traffic matrix, select a specific routing algorithm and analyze the resulting network state




Figure 5: Evolution of link load for a typical week

(see [21] for a deep analysis). We can choose different intra-domain routing algorithms (IP-based or MPLS-based), and see the effect of these algorithms on all the generated traffic matrices.

This can be very useful to investigate the effects of a new routing scheme. Indeed, operators are often afraid of changing a link weight because it is difficult to be sure that this weight change will not affect undesired paths. With the TOTEM toolbox, it is possible to investigate how the load inside the network might change if some intradomain change is performed on the network.




Figure 6: Evolution of link load over time (inverse capacity weights)

One very simple simulation we can easily obtain with the toolbox is presented on figure 6. This graph shows the simulated link loads when relying on the Cisco default $\frac{1}{\text{link capacity}}$ weights (referred to as inverse capacity weights). We can see on this figure that by using the Cisco default IGP weights the maximum load inside the network might be smaller than the actual one, provided that the change in the IGP weights does not result in a change of the traffic demand. This is simply the result of a better load balancing on the most loaded links. However, we have to point out that GÉANT takes into account delays to engineer their weights and we do not consider them in our routing scheme.

Note that many complicated simulations involving more complicated algorithms or combination of several algorithms can be performed in the TOTEM toolbox. The goal of this section was only to give a very brief overview of the TOTEM features.

5. WEBSITE AND SOURCE CODE

The anonymized traffic matrices are available on the TOTEM website [17] in the datasets section (<http://totem.info.ucl.ac.be/dataset.html>). More details about the scripts used to build the traffic matrices or the non-anonymized versions of the traffic matrices can be obtained on request from the corresponding author.

6. ACKNOWLEDGMENTS

This work has been partly supported by the European Union under the E-Next Project FP6-506869 and by the Walloon Region (TOTEM project). We are grateful to DANTE for the information given about GÉANT, especially Nicolas Simar. We are also grateful to everyone involved in the Intel-DANTE monitoring project [15] that made possible the collection of the routing and Netflow data from GÉANT, and in particular Richard Gass from Intel research Cambridge.

7. REFERENCES

- [1] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot, "Traffic matrix estimation: existing techniques and new directions," in *Proc. of ACM SIGCOMM*, 2002.
- [2] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg, "Fast accurate computation of large-scale IP traffic matrices from link loads," in *Proc. of ACM SIGMETRICS*, 2003.
- [3] Y. Zhang, M. Roughan, C. Lund, and D. Donoho, "An information-theoretic approach to traffic matrix estimation," in *Proc. of ACM SIGCOMM*, 2003.
- [4] A. Soule, A. Nucci, R. Cruz, E. Leonardi, and N. Taft, "How to identify and estimate the largest traffic matrix elements in a dynamic environment," in *Proc. of ACM SIGMETRICS*, 2004.
- [5] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving traffic demands for operational IP networks: methodology and experience," in *Proc. of ACM SIGCOMM2000*, September 2000.
- [6] A. Gunnar, M. Johansson, and T. Telkamp, "Traffic matrix estimation on a large IP backbone: a comparison on real data," in *Proc. of ACM SIGCOMM IMC*, 2004.
- [7] G. Varghese and C. Estan, "The measurement manifesto," *Comput. Commun. Rev.*, vol. 34, no. 1, pp. 9–14, 2004.
- [8] K. Papagiannaki, N. Taft, and A. Lakhina, "A distributed approach to measure IP traffic matrices," in *Proc. of ACM SIGCOMM Internet Measurement Conference*, 2004.
- [9] A. Soule, A. Lakhina, N. Taft, K. Papagiannaki, K. Salamatian, A. Nucci, M. Crovella, and C. Diot, "Traffic matrices: balancing measurements, inference and modeling," in *Proc. of ACM SIGMETRICS*, 2005.
- [10] "Internet Traffic Archive," <http://www.acm.org/sigcomm/ITA/>.
- [11] University of Waikato, "The DAG project," <http://dag.cs.waikato.ac.nz/>.
- [12] T. McGregor, H.-W. Braun, and J. Brown, "The NLANR Network Analysis Infrastructure," *IEEE Communications Magazine*, May 2000.
- [13] "University of Oregon Route Views Project," <http://www.routeviews.org/>.
- [14] "Nanog Looking Glass Sites," <http://www.nanog.org/lookingglass.html>.
- [15] Intel-DANTE, "Intel-DANTE monitoring project," <http://www.cambridge.intel-research.net/monitoring/dante/>.
- [16] Kunihiro Ishiguro, "GNU Zebra: free routing software," Available from <http://www.zebra.org>, 1996.
- [17] TOTEM, "A TOolbox for Traffic Engineering Methods," <http://totem.info.ucl.ac.be>, February 2005.
- [18] B. Quoitin, "C-BGP, an efficient BGP simulator," <http://cbgp.info.ucl.ac.be/>, September 2003.
- [19] B. Quoitin and S. Uhlig, "Modeling the Routing of an Autonomous System with C-BGP," *IEEE Network*, vol. 19, no. 6, November 2005.
- [20] B. Halabi and D. Mc Pherson, *Internet Routing Architectures (2nd Edition)*, Cisco Press, January 2001.
- [21] S. Balon, O. Delcourt, J. Lepropre, F. Skivée, and G. Leduc, "A traffic engineering toolbox and its application to the geant network," *Submitted to IEEE eTransactions on Network and Service Management*, 2005.