

Capturing Internet traffic dynamics through graph distances

Steve Uhlig¹, Bingjie Fu², and Almerima Jamakovic³

¹ TU Berlin/Deutsche Telekom labs, Berlin, Germany,
steve@net.t-labs.tu-berlin.de

² Delft University of Technology, Delft, Netherlands,
B.Fu@tudelft.nl

³ TNO ICT, Delft, Netherlands,
almerima.jamakovic@tno.nl

Abstract. Studies of the Internet have typically focused either on the routing system, i.e. the paths chosen to reach a given destination, or on the evolution of traffic on a physical link. In this paper, we combine routing and traffic, and study for the first time the evolution of the traffic on the Internet topology. We rely on the traffic and routing data of a large transit provider, spanning almost a month.

We compute distances between the traffic graph over small and large timescales. We find that the global traffic distribution on the AS graph largely differs from traffic observed at small timescales. However, variations between consecutive time periods are relatively limited, i.e. the topology spanned by the traffic from one time period to the next is small. This difference between local and global traffic distribution is found in the timescales at which traffic dynamics occurs on AS-level links. Small timescales, i.e. less than a few hours, do not account for a significant fraction of the traffic dynamics. Most of the traffic variability is concentrated at timescales of days. Models of Internet traffic on its topology should thus focus on capturing the long-term changes in the global traffic pattern.

Key words: Internet traffic, AS topology, graph distance, multi-resolution analysis

1 Introduction

Most of the studies on traffic dynamics focus on a single link [6, 7, 11, 17, 9, 24]. In reality, Internet traffic is the outcome of end-hosts exchanging data, not through a single link, but over paths¹. The Internet is composed of more than 30,000 autonomous systems (AS). An AS is a network under a single administrative authority. Each AS chooses independently its paths to reach destinations, among the paths that its neighboring ASs advertise. Typical examples of ASs are In-

¹ Paths in the Internet are typically asymmetric [16, 4], so that packets exchanged between two hosts follow different paths in the two directions.

ternet Service Provider networks, or university campuses. In this paper, we use the abstraction of the Internet topology at the AS-level.

When an AS receives traffic that has to be sent towards a destination, it relies on the interdomain routing protocol, BGP [14], to find the next AS on the path to reach the destination. Each AS knows the full AS path that will be followed by its packets to reach a destination. Each path is made of AS-level links or edges. As topological failures happen within ASs or on the links between two ASs, and as ASs change their path preferences over time, AS paths may change. Tracking the actual dynamics in traffic on the AS topology requires to model the routing state of the considered AS over time [13], as explained in Section 2.

In this paper, we study the dynamics in the set of AS-level edges used for forwarding traffic, as well as the dynamics of the amount of traffic carried by each AS-level edge over time, for a large transit AS. This knowledge of the flow of the traffic in the Internet is important not only for operational purposes like traffic engineering [12, 21, 23], but also to understand the Internet as a complex system [10]. For the first time, we study in this paper the global dynamics of the traffic on the Internet topology, as seen from a large transit AS. More specifically, we try to understand the dynamics of the AS-level topology spanned by the traffic. We find that this topology at small timescales differs considerably from the global traffic distribution over a long time period. This indicates that modeling Internet traffic requires models that capture the small timescales behavior of the topological traffic distribution. This small timescales topological traffic distribution is highly dependent on the traffic dynamics observed by individual AS-level edges.

We present the data used in this paper and how the traffic is mapped to the AS-level connectivity in Section 2. In Section 3, we define the distance between two AS-level graphs, and the distance between two traffic distributions on the corresponding AS-level graphs. We first study the distance between individual time intervals and the global traffic topology in Section 4. Then, we analyze changes of traffic distribution between consecutive time intervals in Section 5. We rely on multi-resolution analysis to study the variance of traffic on each AS-level edge across different timescales in Section 6. Finally, Section 7 concludes this paper.

2 Data and methodology

We obtained traffic and routing information from the GÉANT network. GÉANT is the pan-European research network. It carries research traffic from the European National Research and Education Networks (NRENs) connecting universities and research institutions. GÉANT has a point of presence in each European country.

To properly reconstruct paths followed by the traffic, a model of the routing of GÉANT must be built [13]. To compute paths between routers inside its network, GÉANT uses the ISIS routing protocol. We obtained a trace of its ISIS messages. With these messages, we keep an up-to-date view of the internal state of GÉANT and compute the paths from any router to any other router inside

the GÉANT network during the whole time of the study. Once we know the internal path followed by the traffic inside the GÉANT network, we can find out the exit router of GÉANT that forwarded traffic outside the network.

Then, we rely on information from the BGP routing protocol to determine the global AS-level paths taken by traffic observed by GÉANT to reach its destinations. BGP [14] is the current routing protocol used between ASs. With BGP, each AS learns the paths to reach each destination in the Internet. In GÉANT, the BGP routes are collected using a dedicated workstation running GNU Zebra [25], 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 is possible to collect all the BGP routes selected by the border routers of GÉANT and thus find out the global AS-level path followed by traffic entering GÉANT towards any destination in the Internet. With this, we know the set of ASs crossed by traffic entering GÉANT towards any destination, at any time instant of the study [13].

We also obtained Netflow [8] traces collected from all external links of the GÉANT network, i.e. all the traffic entering the network was recorded. Netflow provides the aggregated information of the layer-4 flows, by recording the starting time, the ending time and the total volume in bytes for each unidirectional TCP and UDP flow. Netflow was configured with a 1/1000 packet sampling rate. With this sampling, only one out of 1000 is considered by Netflow. In a large network such as GÉANT, the amount of traffic prohibits to use low sampling rates as it is unsafe for the proper operation of the routers. Given that the aim of this paper is not to study the small timescales, the decision was made to use a granularity of 15 minutes for the finest timescale.

Once we have a model of the routing of GÉANT, we compute for each Netflow entry the corresponding AS path the traffic takes to reach its destination, and attribute the traffic seen to each AS-level link along the path. We call an *edge* ² e of the AS graph G , a pair $ASX - ASY$ appearing as two consecutive and distinct ASs in the AS path computed by our model of GÉANT. We attribute to each *edge* e the amount of traffic it carries during each time interval. For more details about this data, we refer to [22].

We study a contiguous 26 days period between May 5 2005 and May 31 2005, corresponding to 2592 15-minutes time intervals.

3 Distances

3.1 Distance between two topologies

In this paper, we define the distance between two graphs G_0 and G_1 as follows:

² We use the terms *edge* and *link* interchangeably in this paper, but they always refer to an AS-level edge. An AS-level edge does not correspond to a physical link of the router-level graph, but may correspond to several physical links on the topology.

$$DG(G_0, G_1) = 1 - \frac{I(G_0, G_1)}{U(G_0, G_1)} \quad (1)$$

where $I(G_0, G_1)$ represents the number of AS-level edges in the intersection of G_0 and G_1 and where $U(G_0, G_1)$ represents the number of AS-level edges in the union of G_0 and G_1 . A graph distance of 0 means that the two graphs are identical. A distance of 1 means that the two graphs do not have a single AS-level edge in common.

3.2 Distance between two traffic topologies

As we are not only interested in the AS-level topology, but the traffic that crosses each AS-level edge, we define a distance between two graphs weighted by the traffic seen on AS-level edges:

$$DG_{traf}(G_0, G_1) = 1 - \frac{I_{traf}(G_0, G_1)}{U_{traf}(G_0, G_1)} \quad (2)$$

where

$$I_{traf}(G_0, G_1) = \sum_{e \in I(G_0, G_1)} \min(TR_e(G_0), TR_e(G_1)) \quad (3)$$

and

$$U_{traf}(G_0, G_1) = \sum_{e \in U(G_0, G_1)} \max(TR_e(G_0), TR_e(G_1)). \quad (4)$$

$TR_e(G)$ denotes the amount of traffic that edge e has on graph G . $I_{traf}(G_0, G_1)$ is equivalent to the intersection of the two graphs $I(G_0, G_1)$, but where we consider that the intersection is defined by the sum of the minimum amount of traffic common to all edges in the graph intersection $I(G_0, G_1)$. $U_{traf}(G_0, G_1)$ is defined similarly, as the sum of the maximum amount of traffic of all edges in the graph union $U(G_0, G_1)$.

4 Distance between individual time intervals and global traffic topology

Global traffic patterns in the Internet have typically been studied without checking whether the traffic properties do depend on the considered timescale [5, 3, 15]. Those studies have concluded that a few popular source-destinations (end-hosts or networks) do account for the majority of the traffic. [19] has shown that this picture of traffic over-simplifies reality. In practice, only a subset of the source-destination pairs is stable on timescales smaller than hours. We thus expect that the AS-level topology spanned by the traffic on small timescales will differ from the topology spanned over large timescales.

4.1 Graph similarity

To compare the topology spanned by traffic over short and large timescales, we build the graphs spanned by traffic for each 15 minutes time interval over the 26 studied days, denoted by $G_i, i = 1, \dots, 2592$. We also build the graph from the traffic over the 26 days of the study, denoted by G_{global} . We then compute for each G_i the graph distance (see equation 1) between G_i and G_{global} .

Figure 1 shows the cumulative distribution of the distance between the G_i and G_{global} for the 2592 time intervals. For all time intervals, the distance is larger than 0.57. Less than 43% of the AS-level edges known by G_{global} appear during any 15 minutes time interval. The distance can be as large as 0.72, hence sampling only 28% of the existing AS-level edges of G_{global} .




Fig. 1. Distribution of graph distance between the G_i 's and G_{global} .

The graphs of traffic during 15 minutes time intervals are thus very different from the global traffic over large timescales. The G_i 's and G_{global} cannot be considered as topologically similar.

4.2 Traffic similarity

G_{global} contains all AS-level edges for which traffic has been observed over the 26 studied days. Now, we want to compute the distance between the G_i 's and G_{global} , but in terms of the amount of traffic. Our traffic distance defined in equation 2 compared the traffic on each edge of the two compared graphs. As edges of G_{global} cumulate traffic over a far longer time period than the G_i 's, we divide the amount of traffic seen on each edge of G_{global} by 2592, i.e. we average

traffic over time for each edge. We denote G_{global} where the traffic of each edge has been averaged by G_{global}^{traf} . The graphs for each 15 minutes time interval where traffic is attributed on each edge are denoted by $G_i^{traf}, i = 1, \dots, 2592$. Then, we compute the traffic distance as in equation 2 between each G_i^{traf} and G_{global}^{traf} .

Figure 2 shows the cumulative distribution of the distance between the G_i^{traf} and G_{global}^{traf} for the 2592 time intervals. For most (99%) of the time intervals, the distance is larger than 0.82. This indicates that the global traffic distribution is very different from the short-term traffic distribution. As already hinted in [19], the topological traffic distribution observed over large timescales is not representative of the traffic distribution over shorter time intervals.




Fig. 2. Distribution of graph distance between the G_i^{traf} 's and G_{global}^{traf} .

5 Distance between consecutive time intervals

In Section 4, we showed that the traffic over 15 minutes time intervals and over the whole studied period differs very much, as seen through the graph distance. The implications of Section 4 reinforce the findings of [19]. These implications do not mean that modeling Internet traffic on the AS-topology is out of reach. Rather, the long-term traffic distribution does not represent well the short-term one, so that short-term traffic changes should be taken into account in a traffic model. To better understand the short-term dynamics of the traffic on the AS-level graph, in this section we study changes between consecutive time intervals.

5.1 Graph distance between time intervals

In Section 4.1, it was shown that traffic on the AS-level graph varies much, at least when distance was with respect to G_{global} . Instead of computing the distance between the G_i and G_{global} , we compute the distance between G_i and G_{i+1} , for $i = 1, \dots, 2591$. The cumulative distribution of those distances is shown on Figure 3.




Fig. 3. Distribution of graph distance between G_i and G_{i+1} .

Contrary to the figures in Section 4.1, the graph distance between consecutive time intervals is small, between 0.1 and 0.2. Consecutive AS-level graphs spanned by traffic over 15 minutes time intervals are thus close to each other. This means that the graph of traffic evolves relatively smoothly over time over such timescales.

5.2 Traffic distance between time intervals

If we compare the consecutive G_i^{traf} instead of the G_i , we obtain the distribution shown on Figure 4. On this figure, we obtain even smaller distances for the traffic between consecutive time intervals, typically between 0.06 and 0.1. Only very few consecutive time intervals have large distances, up to 0.5.

The distances between consecutive time intervals give a far more optimistic picture of traffic variability on the AS topology than found in Section 4.1. Modeling traffic dynamics should thus require relatively small changes over time. However, as shown in [19], the traffic on different parts of the AS topology has different dynamics. In Section 6, we will analyze this dynamics of the traffic on AS-level edges.




Fig. 4. Distribution of graph distance between G_i^{traf} and G_{i+1}^{traf} .

6 Traffic dynamics on AS-level edges

In this section, we seek to find out explanations for the rather large distances between the G_i 's and G_{global} , and the small distances between consecutive G_i 's. The dynamics of the traffic on different AS-level edges should explain those distances between the graphs spanned by the traffic. In Section 6.1 we study the relationship between the *lifetime* of AS-level edges and the amount of traffic they carry. In Section 6.2 we perform a multi-resolution analysis of the traffic dynamics on AS-level edges.

6.1 Amount of traffic vs. lifetime

First, we look at the relationship between the amount of traffic seen by an AS-level edge and for how many 15 minutes time intervals this edge has traffic at all. Previous work has shown that traffic observed by an AS has a tree-like structure rooted at the observing AS and whose leaves are the destination ASs [18, 19], and on average edges farther away from the root see less traffic. We thus expect that different edges observed different traffic dynamics.

We call the total number of 15 minutes time intervals that an AS-level edge is observed to carry traffic its *lifetime*. The x-axis of Figure 5 gives the lifetime. The y-axis gives the percentage of traffic, in logarithmic scale. The dots in Figure 5 give the percentage of traffic that edges having a given *lifetime* represent. We see that most of the dots correspond to large lifetimes. The solid curve in Figure 5 gives the cumulative traffic as a function of edge lifetime. On this curve, we see that edges that have a small lifetime do not represent a significant fraction of the

traffic. About 80% of the traffic is carried by those AS-level edges that appear almost all the time.

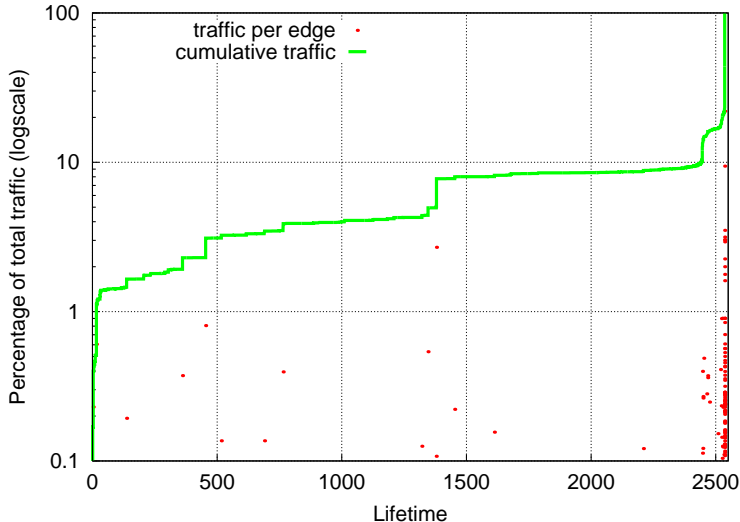


Fig. 5. AS-level edges' life time and the amount of traffic they carry.

6.2 Edge variance decomposition

From Section 6.1, we know that only edges having a large enough lifetime should be considered, as other edges do not represent a significant fraction of the total traffic. Now, we would like to better understand the traffic dynamics on those edges that capture most of the traffic on the AS topology. Because of known non-stationarity of Internet traffic [1, 20], we do not rely on spectral analysis but wavelets [2]. Wavelets belongs to multi-resolution analysis and allow to decompose the variance of the traffic on each edge into the respective contribution of each timescale.

Figure 6 provides the breakdown of the traffic variance within each edge across the different timescales, as computed through the wavelet coefficients. Timescales go from 30 min (scale 1) to about 5 days (scale 9), and are indicated with different colors. Independently for each edge, we stack the relative contribution of each timescale to the total variance of the traffic of this edge, by starting from the smallest timescale and successively adding the contribution of larger timescales.

The x-axis of Figure 6 gives the edges, ordered by decreasing amount of traffic. We observe that edges having most traffic (left of Figure 6) have on average more of their variance within the larger timescales (8 hours or more). For edges that do not have much traffic, the lowest three timescales (between 30




Fig. 6. Decomposition of traffic variance among timescales.

minutes and 2 hours) account for almost 30% of their variance. Edges that see a lot of traffic are thus less bursty on small timescales than edges that see less traffic. The burstiness of the traffic varies much across edges.

This behavior is consistent with previous studies in the networking literature that have debated on the traffic variability on different types of links. Studies of large backbone links have concluded that traffic burstiness tends to a non-stationarity Poisson process as link capacity increases [1]. Studies of smaller links and networks on the other hand have found that self-similar processes better describe traffic [6, 7, 11, 17, 9, 24]. Figure 6 shows that the process that best describes traffic burstiness on a given edge has much to do with the amount of traffic observed on this link.

From Figure 6, we do not have a feeling of what timescales are really important if we want to explain the dynamics of most of the traffic. For this, we turn to Figure 7, where we weight the variance at each timescale by the amount of traffic seen for the considered edge. As in Figure 6, edges are ordered by decreasing amount of traffic on the x-axis. We observe on Figure 7 that the lower 4 timescales do not contribute to a significant fraction of the total traffic-weighted variance. Scale 9 (“128 hours”) accounts for about 50% of the traffic-weighted variance. Scales 6 to 9 account for more than 90% of the traffic-weighted variance. This means that even though burstiness appears at small timescale below hours, most of the traffic dynamics happens for large timescales.

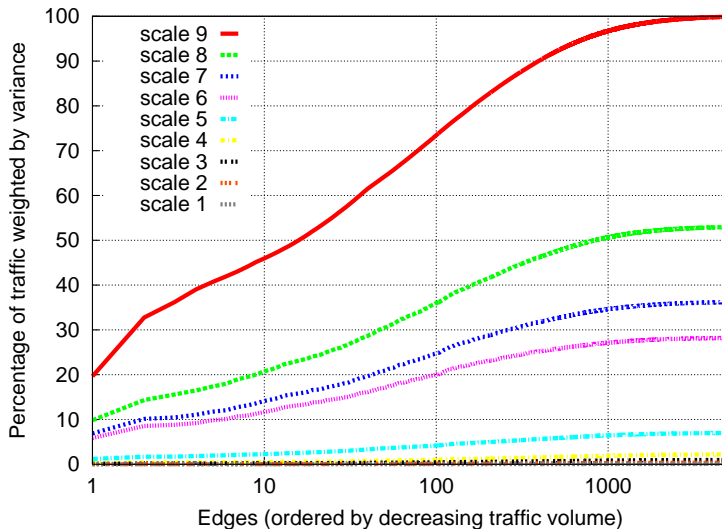


Fig. 7. Traffic-weighted distribution of the variance.

We are now in a position to explain the behavior of the graph and traffic distances observed in Sections 4 and 5. As most of the traffic dynamics is contained within large timescales, the distance between the traffic graph during a small time period (e.g. G_i) and the global graph (e.g. G_{global}) will be large. Unless two graphs are close in time, e.g. consecutive G_i 's, the distance between two traffic graphs will be significant due to edge dynamics. Models of Internet traffic on the AS topology need to consider relevant timescales, e.g. hours or more, unless they will have to deal with complex traffic burstiness that is not important to reproduce for traffic dynamics on the Internet topology.

7 Conclusion

In this paper, we combined routing and traffic, and studied the evolution over time of the traffic on the Internet topology. We relied on the traffic observed by a large transit provider for almost a month, to measure the changes of the topology spanned by the traffic.

We computed distances between the traffic graph over small and large timescales. We found that the traffic observed at large timescales differs from traffic observed at small timescales. However, variations between consecutive time periods are relatively limited, i.e. the topology spanned by the traffic from one time period to the next is small. Small timescales, i.e. less than a few hours, do not account for a significant fraction of the traffic dynamics. Most of the traffic dynamics on the Internet topology happens for timescales of several hours. The slowly changing traffic pattern is responsible for large distances observed between the traffic graphs on small timescales and the global traffic graph.

There are several implications of this paper on complex networks. First, models of the Internet traffic on the topology should concentrate on large timescales, and try to reproduce the long-term variations of the traffic pattern on the topology. Second, other complex networks undergo complex dynamics like the Internet, e.g. road traffic networks or biological networks. Studying the topological dynamics of those systems will help understand the global behavior of those systems, and in turn help to understand the functions implemented within them.

Acknowledgements

We thank DANTE for making the GEANT traffic and routing data available. This work was funded by the Federal Ministry of Education and Research of the Federal Republic of Germany (support code 01 BK 0805, G-Lab). The authors alone are responsible for the content of the paper.

References

1. Cao J., Cleveland W., Lin D., and Sun D.: On the nonstationarity of Internet traffic. In: International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS), pp. 102–112, ACM press, New York (2001)
2. Daubechies I.: Ten Lectures on Wavelets. Volume 61 in CMBS-NSF Series in Applied Mathematics. SIAM press, Philadelphia (1992)
3. Fang W. and Peterson L.: Inter-AS traffic patterns and their implications. In: Global Telecommunications Conference (GLOBECOM), pp. 1859–1868, IEEE press, New-York (1999)
4. He G. Y., Faloutsos M., and Krishnamurthy S.: Quantifying routing asymmetry in the Internet at the AS level. In: Global Telecommunications Conference (GLOBECOM), pp. 1474–479, IEEE press, New-York (2004)
5. Kleinrock, L., and Naylor, W.E.: On measured behavior of the ARPA network. Proc. of the 1974 National Computer Conference, Vol. 43, pp. 767–780, AFIPS Press, Arlington, Va. (1974)
6. Leland W. and Wilson D.: High time-resolution measurement and analysis of LAN traffic: Implications for LAN interconnection. In: Conference on Computer Communications (INFOCOM), pp. 1360–1366, IEEE press, New-York (1991)
7. Leland W., Taqqu M., Willinger W., and Wilson D.: On the self-similar nature of Ethernet traffic. IEEE/ACM Transactions on Networking 2(1), pp. 1–15 (1994)
8. Cisco NetFlow services and applications. <http://www.cisco.com/warp/public/732/netflow>
9. Park K. and Willinger W.: Self-Similar Network Traffic and Performance Evaluation. Wiley-Interscience (2000)
10. Park K. and Willinger W.: The Internet as a Large-Scale Complex System. Oxford University Press (2005)
11. Paxson V. and Floyd S.: Wide-Area Traffic: The Failure of Poisson Modeling. IEEE/ACM Transactions on Networking 3(3), 2267–244 (1995)
12. Quoitin B., Uhlig S., Pelsser C., Swinnen L., and Bonaventure O.: Interdomain traffic engineering with BGP. IEEE Communications Magazine 41(5), pp. 122–128 (2003)

13. Quoitin B. and Uhlig S.: Modeling the Routing of an Autonomous System with C-BGP. *IEEE Network Magazine* 19(6), 12-19 (2005)
14. Rekhter Y., Li T., and Hares S.: A Border Gateway Protocol 4 (BGP-4). Internet RFC 4271 (2006)
15. Rexford J., Wang J., Xiao Z., and Zhang Y.: BGP Routing Stability of Popular Destinations. In: *Internet Measurement Workshop*, pp. 197-202, ACM press, New-York (2002)
16. Tangmunarunkit H., Govindan R., Shenker S., and Estrin D.: The Impact of Routing Policy on Internet Paths. In: *Conference on Computer Communications (INFOCOM)*, pp. 736?-742, IEEE press, New-York (2001)
17. Thompson K., Miller G.J., and Wilder R.: Wide-area internet traffic patterns and characteristics. *IEEE Network Magazine* 11(6), 10?23 (1997)
18. Uhlig S. and Bonaventure O.: Implications of interdomain traffic characteristics on traffic engineering. *European Transactions on Telecommunications* 13(1), pp. 23-32, Wiley, New-York (2002)
19. Uhlig S., Magnin V., Bonaventure O., Rapier C., and Deri L.: Implications of the topological properties of Internet traffic on traffic engineering. In: *Symposium on Applied Computing*, pp. 339-346, ACM press, New-York (2004)
20. Uhlig S.: Non-stationarity and high-order scaling in TCP flow arrivals: a methodological analysis. *Comput. Commun. Rev.* 34(2), 9-24 (2004)
21. Uhlig S. and Quoitin B.: Tweak-it: BGP-based interdomain traffic engineering for transit ASes. In: *Next Generation Internet Networks*, pp. 75-82, IEEE press, New-York (2005)
22. Uhlig S., Quoitin B., Lepropre J., and Balon S.: Providing public intradomain traffic matrices to the research community. *Comput. Commun. Rev.* 36(1), 83-86 (2006)
23. Uhlig S.: *From the traffic properties to traffic engineering in the Internet*. VDM Verlag Dr. Müller (2008)
24. Willinger W., Paxson V., Riedi R., and Taqqu M.: Long-Range Dependence and Data Network Traffic. In: Doukhan P., Oppenheim G. and Taqqu M. (eds.) *Long range Dependence : Theory and Applications*, pp. 373-408, Birkhäuser (2003)
25. GNU Zebra Routing Suite. <http://www.zebra.org>