

On the Dynamics of Locators in LISP

Damien Saucez¹, Benoit Donnet²

¹ INRIA, Sofia Antipolis, France

² Université de Liège, Liège, Belgium

Abstract. In the Internet, IP addresses play the dual role of identifying the hosts and locating them on the topology. This design choice limits the way a network can control its traffic and causes scalability issues. To overcome this limitation, the Locator/Identifier Separation Protocol (LISP) has been introduced. In LISP, the addresses used to identify end hosts (i.e., identifiers) are independent of the addresses used to locate them (i.e., locators). LISP maps identifiers into a list of locators and provides a mean to transport the packets with the appropriate locator. A key feature of this separation is that several locators can be associated to a given identifier, leading to more control for an end-site on the path selection to reach a given destination.

In this paper, we show that the choice of the locator can have an impact on the performance and the reliability of the communication in a LISP environment. To this aim, we build a mapping between identifiers and locators as if LISP were deployed today. In addition, we extensively collect delay data between locators and demonstrate that the locator selection for a given identifier prefix impacts the performance of the LISP path in 25% of the cases. Finally, we measure the locators availability over time and demonstrate that it remains quite stable.

1 Introduction

During the last decade, the Internet has strongly evolved. Its natural growth combined with factors such as multihoming and interdomain traffic engineering has lead to an increase of the BGP routing tables [1, 2] and the BGP churn [3].

To cope with these problems, several solutions have been proposed. Most of them assume two different types of addresses: *identifiers* and *locators*. An identifier is used on an host to identify a connection endpoint while a locator refers to a node attachment point in the Internet topology. The proposals are divided in two categories: those attaching locators directly to hosts (HIP, SHIM6, . . . [4]) and those attaching locators to routers (see for instance LISP [5, 6]). It is worth to notice that, in the current Internet, an host address is at the same time its identifier and its locator.

A key advantage of the addresses separation is to offer the possibility of associating several locators to a given identifier. This implies the availability of multiple paths between two identifiers and, as shown by several studies, those paths often offer different characteristics in terms of delay, jitter, bandwidth, etc [7–10].




Fig. 1: LISP general behavior

In this paper, we quantify the effects of controlling paths between locators in the Locator/Identifier Separation Protocol (LISP). A key feature of LISP is the *mapping* between locators and identifiers. As LISP is not yet globally deployed³, we cannot rely on an existing solution for this mapping. Therefore, in this paper, we propose our own mapping as if LISP were deployed today. We generate our mapping based on BGP prefixes and Archipelago dataset [11]. Our mapping considers more than 14,000 locators associated to more than 1,800 identifiers.

We next evaluate the performance of each locator by extensively measuring the delay to each locator from a point located in Belgium. Based on this dataset, we demonstrate that the locator selection, for a given identifier, has an impact on the communication performance in 25% of the cases. We further evaluate the selected locators availability over time and show that it remains quite stable.

The remainder of this paper is organized as follows: Sec. 2 introduces how LISP works; Sec. 3 explains how we build a mapping between locators and identifiers as if LISP were deployed today and how we collected our dataset; Sec. 4 discusses the impact of the locator selection on the delay; Sec. 5 discusses reachability issues; finally, Sec. 6 concludes this paper by summarizing its main contributions.

2 LISP Background

The *Locator/Identifier Separation Protocol* (LISP) [5, 6] has been introduced to separate the identifier and the locator roles of IP addresses. LISP considers two independent address spaces: the *Routing Locator Space* (RLOC), that is globally routable, and the *Endpoint Identifier Space* (EID), that is only locally routable. With LISP, routers in the core of Internet are handling RLOCs like

³ A LISP testbed is available. See <http://www.lisp4.net/>

today and routes are maintained so that packets can be forwarded between any router. On the contrary, endpoints (typically hosts) receive EIDs. As an EID is only locally routable, it means that routers in the core of the Internet do not maintain routes towards EIDs.

LISP is used when packets are exchanged between two endpoints located in different sites. First, the source endpoint generates a regular IP packet to the destination endpoint and forwards it. When the packet arrives at the edge of the source endpoint site, it is processed by a LISP router called *Ingress Tunnel Router* (ITR). As the destination is an EID, it cannot be forwarded on the Internet. The ITR must first query the *mapping system* [6, 12, 13], a global database associating EIDs to RLOCs.⁴ Note that several RLOCs might be associated to a given EID. Once the ITR has received the mapping for the destination EID, it selects one RLOC and encapsulates the original message in a LISP packet. The source address of the LISP packet is the ITR RLOC and the destination address is the selected RLOC. As the source and destination addresses of this LISP packet are RLOCs, it can be forwarded over the Internet. The router with the destination RLOC is at the destination EID site and is called *Egress Tunnel Router* (ETR). ETRs are responsible of decapsulating LISP packets so that the original packet can be forwarded to its final EID destination. This behavior is illustrated in Fig. 1.

The mapping system is a key element of LISP. A mapping associates an EID prefix to a list of $\langle RLOC, priority, weight \rangle$ tuples. The priority and weight associated to each RLOC help the ITR in preferring an RLOC for reaching a given EID. RLOCs with the lowest priority are always preferred. When RLOCs share the same lowest priority, the weight is used for load balancing among themselves. One of the key feature of LISP mapping systems is that the mappings can be specific to each ITR. For instance, it is possible to return a different mapping to a site attacking you than the mapping returned to valid sources of traffic. Such incoming traffic engineering is not possible today with BGP.

The flexibility and control offered by LISP come at the cost of increasing the reachability detection complexity. Indeed, identifiers are not directly routable and have to be mapped to locators, themselves being assigned to the border routers. Therefore, it is no longer possible for the routing layer to move traffic to another border router in case of a router failure. The ITRs are thus in charge of checking the reachability of the selected ETR. To test the reachability, the ITRs can periodically probe the ETRs or monitor the traffic. Unfortunately, both solution provide poor detection times.

Reachability tests can be made in an active or passive way [5]. Active testing consists of sending probes to the ETR to ensure that it is reachable. This technique is implemented with the *RLOC Probing Algorithm* but does not scale for short time failure detection. On the contrary, passive reachability tests can be used when a short time scale detection is required. This is done by tagging a LISP packet with a nonce and the targeted RLOC has to tag the packets back to the source RLOC with the same nonce. The locator is considered as unreachable

⁴ To avoid requesting for mapping all the time, ITRs enable a mapping cache.




Fig. 2: Comparison between the number of locators per mapping as discovered with BGP AS paths and traceroutes analysis

if the nonce is not echoed within a given time frame. This technique is called the *Echo Nonce Algorithm* [5].

3 Methodology

Our work, in this paper, is based on LISP. However, except in the lisp4.net pilot network⁵, LISP is not yet deployed and, there is no clear clue on how to build the mapping between identifiers and locators. In Sec. 3.1, we explain how we construct the mapping as if LISP were deployed today. Next, in Sec. 3.2, we explain our delay measurement campaign.

3.1 Building a LISP Mapping

As LISP is not yet deployed, we have to build, on our own, the mapping between EIDs and RLOCs. In this section, we envision two ways for creating the mapping: the first one is based on BGP information and active measurements, while the second one only takes into account BGP.

For the first technique, we rely on three assumptions to determine how the mappings could look like: *(i)* as contiguous addresses tend to be used similarly [14], EID prefixes follow the current BGP prefixes decomposition; *(ii)* EIDs are used only at the stubs, not in the Internet core; *(iii)* ITRs and ETRs are deployed at the edge between the stubs and the providers, while locator addresses are allocated in a *Provider Aggregatable (PA)* mode.

We base our evaluation on several datasets. First, the BGP prefixes and their stub nature are extracted from the Route Views Project dataset [15] (labeled as *Route Views* in the remainder of this paper), we use BGP dump from the Oregon-IX collected on August, 12th 2010. Second, the mapping between BGP

⁵ See <http://www.lisp4.net>

prefixes and locators is done with the Archipelago dataset [11] (labeled as *Ark* in the following) collected on July, 24th 2010.

Following the first two assumptions (see above), we filter the prefixes from Route Views to extract the stub prefixes using the AS ranking provided by CAIDA [16]. We further take only into account the most specific prefixes. Less specific prefixes, typically used for resiliency, are thus filtered. By doing so, we have an upper bound on the number of EID prefixes to be supported by the mapping system. Indeed, filtering the less specific prefixes does not filter the prefixes deaggregated for load balancing reasons. After filtering out the Route Views dataset, 41.21% of the prefixes (i.e., 138,123 prefixes) can be considered as EID prefixes.

To determine the active locators for each EID prefix, we rely on the Ark dataset and the third assumption. In this paper, we consider the edge from a IP routing point of view: the router at the edge of a stub is the first router whose address it not inside the EID prefix. For each destination address in the Ark dataset, we determine its most specific BGP prefix, i.e., its EID prefix, and backtrack the traceroute until a hop address that does not belong to this BGP prefix. The hop address is then considered as one RLOC of the EID prefix. Computing this on all the completed traceroutes from the Ark dataset gives the list of RLOCs associated to a given EID prefix. As EID prefixes determined with Route Views are not always traceroutable, we were only able to extract 15,337 EIDs.

The second mapping construction technique is based only on BGP. We simply count the number of different neighbor ASes for each stub prefix. This gives us an approximation of the stub prefixes multi-homing degree. Neighboring ASes can be seen as the prefix locator from an AS-level point of view.

Fig. 2 provides the distribution, as a cumulative mass, of the number of locators (horizontal axis, in log-scale) associated to each EID for both mapping construction techniques described above in this section. The plain line labeled “BGP” has been obtained with the second technique, while the dashed line labeled “traceroute” exhibits results for the first one.

We first observe that, whatever the mapping construction technique considered, we face a large proportion of EID having a single locator: 61% for BGP and 75% for traceroute. This difference can be explained by the fact that some routers do not reply to traceroute probes while some others are not on the path to the EID during the traceroutes as they can be backup routers. However, the maximum number of locators associated to an EID is more important when traceroutes are considered: 13 for BGP and 194 for traceroute with an outlier network offering satellite broadcasting measured with 1,465 locators. This can be explain by the fact that some sites have many points of presence but these points of presence are all connected to the same providers.

In this paper, we are interested to see the impact of the locator selection on the performance and availability. We thus focus our study on prefixes that appear to have at least two locators in the traceroute evaluation. It is therefore important to observe that 68% of multihomed EIDs have either 2 or 3 locators.

This small number of locators limits the choices of locators for the EIDs. In the remainder of this paper, we will only consider mapping obtained with the first mapping construction technique (i.e., traceroute).

3.2 Delay Data Collection

In this paper, the RLOC selection is evaluated based on the *round-trip time* (RTT). We built our delay dataset by pinging each locator identified as explained above every five minutes from a vantage point located in Louvain-la-Neuve University campus in Belgium. The campus is single connected to the Internet and the provider is the national research network. It means that without LISP, the campus cannot do traffic engineering (neither incoming nor outgoing). If LISP was deployed, the campus could do performance based outgoing traffic engineering by selecting the preferred destination locators. The measurement campaign started on September, 3rd 2010 and lasted until September, 24th 2010.⁶ Measurements have been done in cycles, i.e., all locators are measured during a cycle and a new cycle started every five minutes leading to a total of 6,048 measurement cycles. As the measurements towards the whole set of locators lasted, at worst, five minutes, cycles never overlapped. We consider a locator as not reachable if its RTT is higher than one second. The decision of a five minutes sampling period is a tradeoff between the scalability of our measurements and operational reality where a five minutes granularity is often observed for billing and accounting purposes [17].

4 Impact of RLOC Selection

In this section, we investigate how the RLOC selection can impact the performance of the communications between two distant EIDs.

Fig. 3 shows the average distance between the locators and the best locator (i.e., the locator with the lowest delay), for each of the 1419 EID-RLOC mapping. This average distance has been computed as follows: for each EID prefix, at time t , we take the set of associated RLOCs, measure the distance (i.e., the delay) to each locator and, finally, consider the locator with the lowest delay as being the best locator at time t for the current mapping. The distance is then computed between the best locator and the remainder locators of the given EID as $delay_r - delay_b$, where $delay_r$ is the delay of the considered RLOC and $delay_b$ the delay to the best locator. We next compute the average distance and the 95% confidence interval for all distances over the whole measurement campaign, for all EIDs. The EID prefixes are ordered by average distance in Fig. 3.

We observe two trends in Fig. 3. First, for 90% of the EIDs, the average distance to the best locator is lower than 20ms and 77% are not distant by more than 10ms. This is an important result because, on average, for most of the EID

⁶ Our dataset is freely available at <http://inl.info.ucl.ac.be/content/locator-reachability-dataset>




Fig. 3: Average distance between each locator of a mapping and the locator presenting the lowest delay in the mapping




Fig. 4: Stability of the round-trip delay measured for each locator during 21 days represented as the cumulative distribution of the signal-to-noise ratio (i.e., average delay divided by standard deviation, in dB) of the round-trip delay

prefixes, the traffic can be balanced between all the locators without the risk of dramatically degrading traffic performance. Second, for 7% of the prefixes, the average distance is higher than 50ms and can reach more than 100ms. A so high average distance implies that performance can be degraded considerably if the locator has been badly selected.

Fig. 3 only gives an overview of the locator dynamics. The confidence intervals are close around the mean, suggesting that the delay is rather stable. This stability is confirmed by Fig. 4. It shows the cumulative distribution of the signal-to-noise (SNR) ratio (in dB) for the delay measured for each locator during the measurement campaign. The SNR is computed as $snr = \frac{\mu}{\sigma}$, where μ is the average delay and σ the standard deviation. Looking at Fig. 4, we see that a very few locators (i.e., 1.5%) present a negative SNR, meaning that they are completely unstable. On the contrary, 80% have an SNR above 20dB and are thus considered as stable. Said differently, those 80% of locators have a 10:1 ratio between the average and the standard deviation. Finally, 12% of the locators are very stable as they show an SNR over 32.04dB (i.e., a 40:1 ratio).

To refine our observations of the locators dynamics, we compute the *preferred sets* for each EID prefix at every measurement cycle, i.e., the set of locators for which the distance is at most τ percent higher than the locator with the lowest delay. We vary τ between 0%, 5%, 10%, and 100%. A value $\tau = 0\%$ would mean that only the best locator is considered while a value $\tau = 100\%$ would mean that locators with a delay up to twice the lowest one are considered as providing equivalent performance.

Fig. 5 shows the preferred set dynamics. The horizontal axis gives the number of times the preferred set has changed during the measurement campaign and

the vertical axis gives the cumulative distribution of this number of changes. A first interesting observation is that, even for a high tolerance of 100%, the preferred set changes over time. This implies that, even if delays are stable and locators relatively close to each others regarding delay, it is likely to find locators for an EID that are more than twice farer (in term of delay) than the best locator. However, for 75% of the EIDs, the number of changes is between 162 and 2,714, depending on the tolerance. It thus means that when only considering the best locator, the preferred set does not change 55% of the time. For a reasonable tolerance of 10%, the preferred set was unchanged 81% of the time. Therefore, when a network decides to potentially use several locators simultaneously, a tradeoff has to be achieved between the delay difference acceptable for the preferred locators and the frequency of the changes in the preferred set. This tradeoff can be observed with the tolerance.

Fig. 6 refines the understanding of the preferred set dynamics. Fig. 6 provides indeed the cumulative distribution of the *dissimilarity*. The dissimilarity, or the *Jaccard distance*, is a metric of the similarity of two sets. The dissimilarity of sets A and B is given by

$$dissimilarity(A, B) = 1 - \frac{|A \cap B|}{|A \cup B|}. \quad (1)$$

A dissimilarity of zero means that the two set are similar, i.e., both sets are made of the same locators. On the contrary, a value of one indicates that no element in the set is present in the other set, i.e., both sets do not share any locators. Fig. 6 shows the distribution of the dissimilarity between a preferred set at time t and the preferred set at time $t - 1$, for all the EID prefixes and all the cycles, for a tolerance of 10%.

We see that, in 85% of the cases, the preferred set does not change between a cycle and the next cycle. In addition, the preferred set was completely different from the previous preferred set in less than 1% of the observations. In other words, the preferred set tend to remain identical from one cycle to another and even if there is a change, this change is only partial. It is worth noticing a small peak for the 0.5 dissimilarity. This peak can be explained by the fact that the majority of the EIDs in our dataset has two locators.

The median value for the time a preferred set is kept is 22 hours and the 95th percentile is 5 days and a half. A low preferred set change rate means that the load of the mapping system because of preferred set remains limited.

5 Availability Issues

Dealing with router failures in LISP is fundamentally different than dealing with router failures in the current Internet. Currently, the path followed by a packet to reach an end-host only depends on the destination IP address. On the contrary, in LISP, the source network determines the ETR through which the packet must transit by selecting a destination locator. In other words, the source network (via the ITR) decides by which border router a packet enters




Fig. 5: Cumulative distribution of the number of times the set of preferred locators is changed per mapping for a tolerance of 10%




Fig. 6: Percentage of difference (i.e., dissimilarity) between the preferred set of a mapping at a given time with the preferred set of the same mapping at the previous cycle, expressed as the cumulative distribution of the dissimilarity for all the measured 10% tolerance preferred sets

the destination network. Without LISP, when a border router breaks down, the prefix of the network remains reachable as the network prefix is not bound to any specific router. In addition, as every router announces the same prefix, failure recovery is performed locally. Recovery can however not be performed locally with LISP. Indeed, the EIDs are hidden from the network and packets are sent to a particular router. As a consequence, failure recovery is performed by the ITR in LISP. First, the ITR must detect that it cannot reach anymore the desired locator. Once the failure has been detected, the ITR can bypass the failure by selecting another locator (when available). Failures are detected by the mean of active probing. For scalability issue, active probing can only be performed at a low frequency (i.e., a few times per minute). In this section, we evaluate the duration and frequency of locator failures to determine if end-to-end failure detection is acceptable in general.

Fig. 7 gives an insight of failures impacts. For each EID prefixes, we count the number of locator failures as well as the number of preferred set changes. Fig. 7 shows the cumulative distribution of the ratio of the number of locator failures to the number of preferred set changes. A value lower than 100% indicates that the preferred set changed more often than the number of observed failures. On the contrary, a value higher than 100% indicates that the preferred set changed less frequently than the failure occurrence. The plain line is obtained by only considering the failures of a locator that was in the preferred set at the time of its failure. On the contrary, the dashed line curve is obtained by only considering the failure of locators that were not in the preferred set at the time of the failure.




Fig. 7: Ratio between the number of failures and the number of preferred set changes per mapping for a 10% tolerance




Fig. 8: Cumulative distribution of the duration of the failures for every measured failure for both identifiers (e.g., EID) and locators (e.g., RLOC)

Looking at Fig. 7, it is interesting to see that, in nearly 80% of the cases, no more than 10% of the changes in the preferred set are due to the failure of a locator in the preferred set. This thus means that, most of the time, the preferred set changes because of performance reasons and not because of a locator failure. On the contrary, the dashed line curve shows that failures that do not impact the preferred set are more frequent than failures impacting the locator set. In 10% of the cases, it is even observed that the number of failures that does not impact the preferred set was higher than the number of preferred set changes, potentially by several orders of magnitude. In other words, a potentially important number of failure can have no impact on the forwarding of packets.

Fig. 8 gives the distribution of failures duration (in minutes) observed for the EIDs (curve labeled *EID*) and RLOCs (curve labeled *RLOC*). As we are measuring the availability once every five minutes (see Sec. 3.2), the minimum unavailability of a given EID (or RLOC) is five minutes. In addition, it is worth to notice that Fig. 8 considers in row unavailability.

Fig. 8 shows that in 68% of the cases, the EID has never been considered unavailable by more than one measurement cycle. This unavailability can come from network failure of less than five minutes but also because of the loss of the probe. It also means that 32% of the EIDs have been unavailable for ten minutes or more. Beside, 5% of the EIDs have been unavailable for at least one hour. If we compare this EID maximum consecutive unavailabilities with those from RLOCs, we can see that multiconnected EID sites reduce the unavailability duration.

In this section, we have seen that the failure of a locator is seldom and rarely lasts more that 10 minutes. This result is promising as it suggests that doing end-to-end failure detection with low frequency probes will be effective. However,

as we measured the network at low frequency, we might have miss short living failures. Understanding these short lived failures is left for future work.

6 Conclusion

The Internet is facing scalability issues. For that reason, LISP has been proposed. LISP improves the Internet scalability by separating the identifier and the locator role of IP addresses. In LISP, the core Internet is addressed with routing locators bound to the topology. On the contrary, end-hosts receive stable identifiers, independent of the topology. While the locators are globally routable, the identifiers are not. LISP then provides an encapsulation mechanism to transmit packets between identifiers. The association between the identifiers and their locators is provided by a mapping. A key feature of this separation is that several locators can be attached to a given identifier. This means that an end-site has more control on the path selection to reach a given destination by controlling the locators. However, paths are not of equal quality and the selected locator can have an impact on the traffic performance.

In this paper, we performed a three weeks delay measurement campaign to determine the impact of the locator choice on the traffic. Our first result is that for 75% percent of the stub networks, the traffic is not impacted by the locator selection as the measured delay does not vary by more than 10%. Hence, the traffic can be load balanced between the locators without severely impacting the communication performance. However, we also observed that for 25% of the network, a bad locator choice can severely impact the delay. The delay cannot thus be ignored when a locator must be chosen. Our measurement campaign also shows that locators presenting a short delay tend to keep this short delay for long periods. We are currently working on techniques to allow a network to efficiently determine if it can perform load balancing among the locators or if some locators must be avoided.

The ability for a network to control the paths it is using comes at the cost of a robustness loss. Indeed, LISP relies on encapsulation. Hence if a locator is not available, the network cannot recover the failure until the network that encapsulates the packets detects that failure. In LISP, failures are discovered with low frequency active probing. With our dataset, we showed that failures are seldom meaning that probing at low frequency can be effective. However, active probing will never be able to protect against short lived failures. For that reason, we are currently working on a local failure recovery techniques that do not require active probing but that can be applied directly to LISP.

References

1. Huston, G.: BGP routing table analysis reports (2004) See <http://bgp.potaroo.net>.
2. Meng, X., Xu, Z., Zhang, B., Huston, G., Lu, S., Zhang, L.: IPv4 address allocation and the BGP routing table evolution. *ACM SIGCOMM Computer Communication Review* **35**(1) (2005) 71–80
3. Meyer, D., Zhang, L., Fall, K.: Report from the IAB workshop on routing and addressing. RFC 4984, Internet Engineering Task Force (September 2007)
4. Savola, P.: A survey of IPv6 site multihoming proposals. In: *Proc. IEEE Internet Conference of Telecommunications (ConTel 2005)*. (June 2005)
5. Farinacci, D., Fuller, V., Meyer, D., Lewis, D.: Locator/ID separation protocol (LISP). Internet Draft (Work in Progress) draft-ietf-lisp-22, Internet Engineering Task Force (February 2012)
6. Meyer, D.: The locator identifier separation protocol (LISP). *Internet Protocol Journal* **11**(1) (March 2008) 23–36
7. Akella, A., A., S., Sitaraman, R.: A measurement-based analysis of multihoming. In: *Proc. ACM SIGCOMM*. (August 2003)
8. Quoitin, B., Iannone, L., de Launois, C., Bonaventure, O.: Evaluating the benefits of the locator/identifier separation. In: *Proc. ACM SIGCOMM MobiArch Workshop*. (August 2007)
9. de Launois, C., Quoitin, B., Bonaventure, O.: Leveraging networking performance with IPv6 multihoming and multiple provider-dependent aggregatable prefixes. *Computer Networks* **50**(8) (June 2006) 1145–1157
10. Zhou, X., Jacobsson, M., Uijterwaal, H., Van Mieghem, P.: IPv6 delay and loss performance evolution. *International Journal of Communication Systems* **21**(6) (June 2007)
11. claffy, k., Hyun, Y., Keys, K., Fomenkov, M.: Internet mapping: from art to science. In: *Proc. IEEE Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH)*. (March 2009)
12. Jakab, L., Cabellos-Aparicio, A., Coras, F., Saucez, D., Bonaventure, O.: Lisp-tree: a dns hierarchy to support the lisp mapping system. *IEEE J.Sel. A. Commun.* **28** (October 2010) 1332–1343
13. Fuller, V., Farinacci, D., Meyer, D., Lewis, D.: LISP alternative topology (LISP+ALT). Internet Draft (Work in Progress) draft-ietf-lisp-alt-10, Internet Engineering Task Force (December 2011)
14. Cai, X., Heidemann, J.: Understanding block-level address usage in the visible Internet. In: *Proc. ACM SIGCOMM*. (August 2010)
15. University of Oregon: Route views, University of Oregon Route Views project See <http://www.routeviews.org/>.
16. CAIDA: AS-Rank (2010) See <http://as-rank.caida.org/>.
17. Dimitropoulos, X., Hurley, P., Kind, A., Stoecklin, M.P.: On the 95-percentile billing method. In: *Proc. Passive and Active Measurement Conference (PAM)*. (April 2009)