Network Fingerprint and *MPLS tunnel discovery* measures campaign results

Daniele Formichelli

March 2015

1 Introduction

In this document are presented the results of a measures campaign performed to validate the *Network Fingeprint* and *MPLS Tunnels Discovery* tools of *Portolan* [6] desktop client. For further details see [4].
The targets of the measures campaign are all the address that are in Portolan target list. That list is composed of 36183 targets belonging to 36183 different ASes distributed all over the world. The campaign has been run from a network located in San Miniato (Pisa, Italy) which belongs to the AS 3269 (ASN-IBSNAZ Telecom Italia s.p.a.). Mappings between IP addresses and their ASes has been done using data from Isolario [5].

2 Network Fingerprinting

2.1 Introduction

Network Fingerprinting is the collection of configuration attributes of remote devices during network communication. The combination of these parameters can be used to infer some characteristics of the remote machine such as its operating system. Moreover fingerprints can be used in the process of router alias resolution.
Portolan fingerprint classifies devices basing only on the initial TTL of their packets as described in [7]: the fingerprint of a router is composed by a triplet containing the initial TTL of different types of sent IP packets, in particular three types of ICMP packets are considered:

ICMP TTL Expired packets received during traceroutes when TTL reaches 0

ICMP Echo Reply packets received in response to ICMP echo request packets

ICMP Port Unreachable packets received in response to UDP probes

Fingerprint	Portolan	Vanaubel	Ratio
<255, 255>	52%	56%	1.08
<255, 64>	13%	11.5%	0.88
<255, -1>	8%	15%	1.87
<128, 128>	2%	3%	1.5
<64, 64>	22%	11.5%	0.52
Others	3%	3%	1

Table 1: Comparison of fingerprints obtained by *Portolan* with the ones reported in Vanuabel's paper

When an ICMP packet is received the original TTL is estimated as one of four possible values (32, 64, 128, 255) among which the smallest one that is greater than the received value is taken. Since more than 99% of the paths are less than 30 hops the estimation is very accurate. A TTL of -1 is associated to unresponsive hops (in this case 6 probes are sent to be sure that the router doesn't reply).
Network equipments can be classified basing on their TTL fingerprint [7] (it considers only first two fields of *Portolan* fingerprint):

<**255,255**> includes Cisco routers

<**255,64**> includes Juniper routers running JunOS

<**128,128**> includes Juniper routers running JunosE

<**64,64**> includes many desktop OS

2.2 Results

During the campaign the fingerprints of 87963 interfaces have been discovered.

Table 1 and Figure 1 show the distribution of fingerprints obtained with *Portolan* compared with the one reported in [7].

Results confirms the predominance of Cisco routers ($< 255, 255 >$) with about 50% of market share followed by Junos ones ($< 255, 64 >$ or $< 128, 128 >$) with about 15%. About 20% of the devices use the recommended value for TTL which is 64.
The obtained values are quite similar to the ones obtained by [7] except for $< 255, -1 >$ which is almost an half and $< 64, 64 >$ which is almost the double of the value reported in Vanuabel's paper. However Portolan measure includes a sample which is very small with respect to the one reported in [7] (88k vs 335k

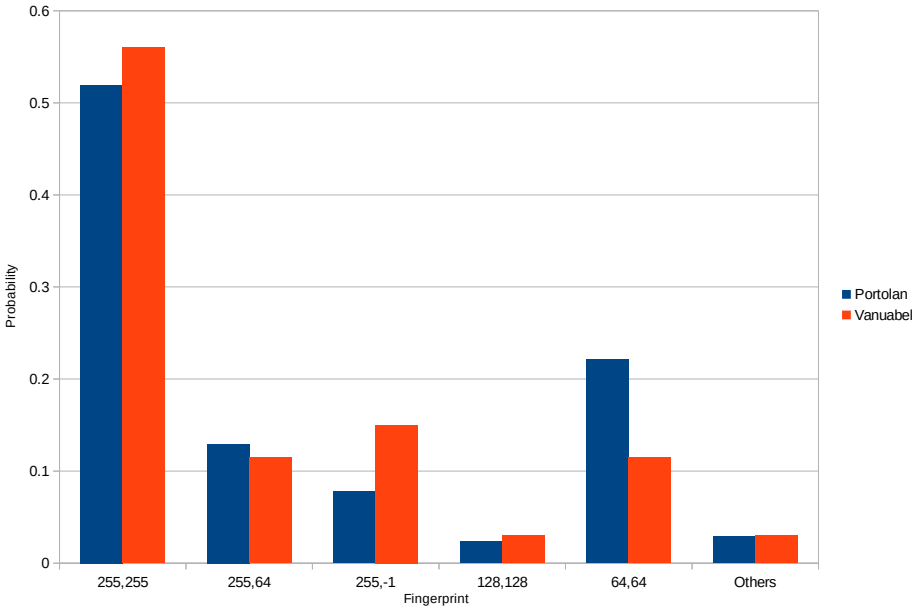

Figure 1: Comparison of fingerprints obtained by *Portolan* with the ones reported in Vanuabel's paper

discovered interfaces) and all the measures are taken by just one vantage point (vs 200 VP around the world in [7]). For these reasons for the purpose of the validation the results can be considered good enough.
Using *Portolan* the obtained fingerprints contain also a value for the source TTL of ICMP destination unreachable packets in response to UDP probes. Table 2 and Figure 2 show the distribution of that part of the fingerprint with respect to the first part which is relative to ICMP probes. The dominant response is -1 (i.e. no response) and the remaining percentage is almost entirely 255 for $< 255, 255 >$ and $< 255, 64 >$, 128 for $< 128, 128 >$ and 64 for $< 64, 64 >$. That confirms (as hypothesised in [7]) that probably adding this third field to the fingerprint doesn't add much information to the two fields fingerprint.

Figure 3 shows the distribution of the fingerprints for the ASes with at least 150 discovered interfaces, descending ordered by number of discovered interfaces. The distributions are quite heterogeneous but in most cases there is prevalence of $< 255, 255 >$, $< 255, 64 >$ or $< 255, -1 >$ which confirms the prevalence of Cisco and Juniper routers.

	255	128	64	32	-1
<255, 255>	47.26%	0.01%	0.11%	0.02%	52.59%
<255, 64>	71.43%	0%	0.04%	0%	28.53%
<255, -1>	3.51%	0%	0%	0%	96.49%
<128, 128>	0.05%	17.83%	0.39%	0%	81.73%
<64, 64>	2.58%	0.14%	56.79%	0%	40.49%
Others	18.59%	0.35%	8.59%	1.21%	71.25%
Average	23.90%	3.05%	10.99%	0.21%	61.85%

Table 2: Distribution of ICMP destination unreachable fingerprints in response to UDP probes with respect to the two fields fingerprint related to ICMP probes

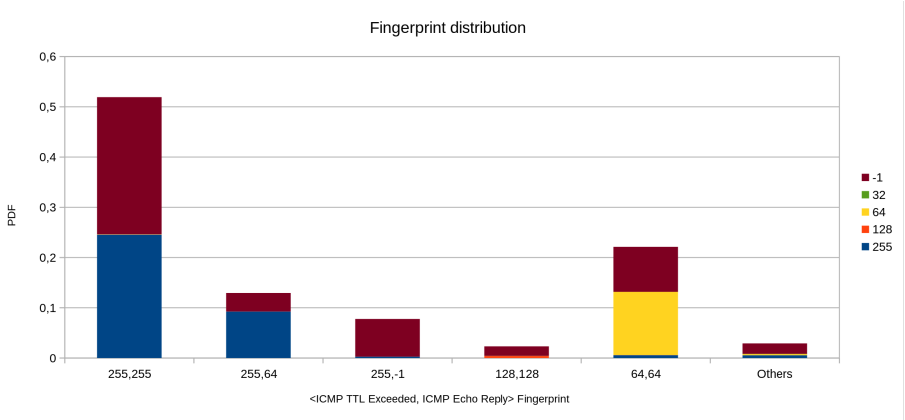

Figure 2: Distribution of ICMP destination unreachable fingerprints in response to UDP probes with respect to fingerprint related to ICMP probes

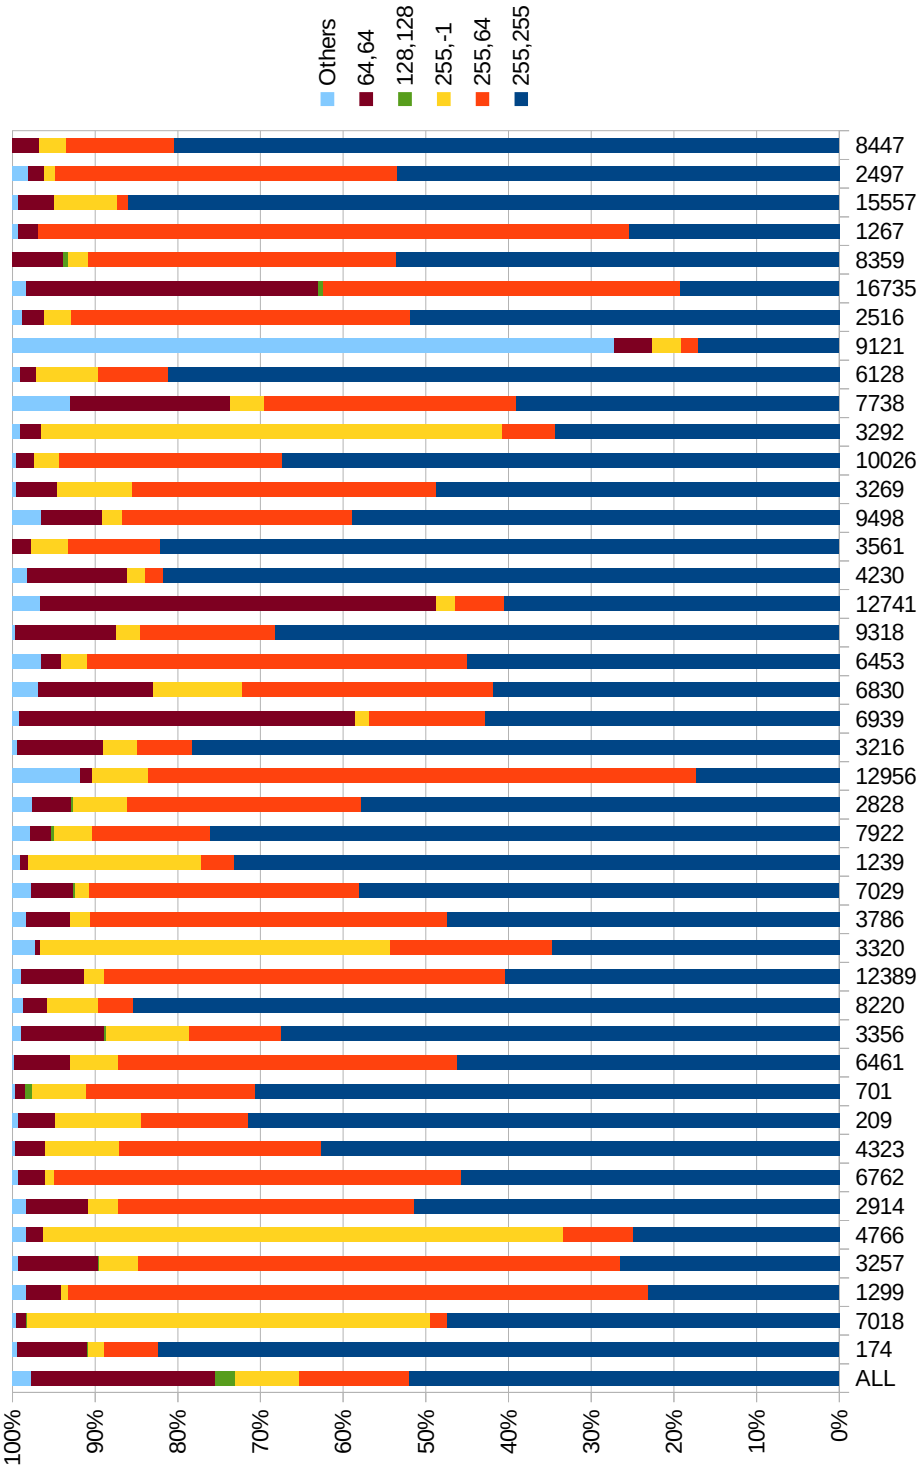

Figure 3: Distribution of fingerprint for AS with at least 150 discovered interfaces

	RFC4950 enabled	**RFC4950 disabled**
ttl-propagate enabled	Explicit	Implicit
ttl-propagate disabled	Opaque	Invisible

Table 3: MPLS tunnel classification

3 MPLS Tunnels Discovery

3.1 Introduction

Portolan MPLS Tunnels Discovery is an implementation of the techniques to discover and classify Multi Protocol Label Switching (MPLS) tunnels on the path to a destination which are described in [3].
In an MPLS tunnel (Figure 4) different components can be identified:

Label Edge Router an LER is a router at the edge of an MPLS tunnel i.e. a router that receives plain packets and transmits labeled packets or viceversa (R1 and R5 in the figure)

Label Switched Path an LSP is the path between two LER (R2-R3-R4) in the figure

Label Switching Router an LSR is a router that belongs to an LSP (R2, R3 and R4 in the figure)

When an IP packet reaches the LER of a LSP there are two different options that can be active or not for that path:

ICMP MPLS extension a router that implements RFC 4950 [1] includes in the ICMP error packet the MPLS stack of the received probe (as an ICMP extension as defined in RFC 4884 [2]). When an ICMP error message with this extension is received it is clear that the source belongs to an MPLS path

ttl-propagate if activated, the ingress LER copies the TTL field from the IP header to the LSE-TTL field of the MPLS header so each LSR of the LSP will be discovered by a traceroute. Viceversa, if ttl-propagate is not active only the the last LSR of the LSP will be discovered

From the combination of these two configurations four different types of tunnels arises (Table 3):

explicit tunnels all LSRs in the tunnel are discovered by traceroute and they attach the MPLS stacks to ICMP error messages

implicit tunnels all LSRs in the tunnel are discovered by traceroute but they don't attach the MPLS stacks to ICMP error messages

opaque tunnels only the last LSR will be discovered but it attaches the MPLS stack to the ICMP error message

invisible tunnels LSRs in the tunnels are not visible at all

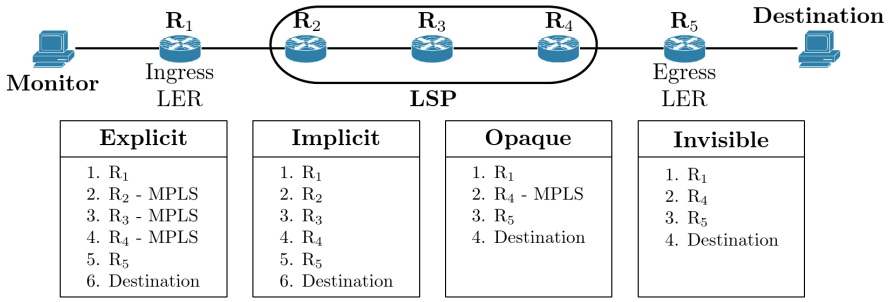

Figure 4: Taxonomy of MPLS tunnel configurations and corresponding traceroute behaviours

While performing a traceroute when an ICMP error message is received the following objects are collected:

- attached MPLS stack (if any)
- TTL of the received IP packet
- TTL of the IP packet quoted in the received ICMP error messages (from now on q-ttl)

An *explicit* tunnel is detected when there are two or more consecutive hops with an attached MPLS stack.
An *opaque* tunnel is detected when there is just one hop with an attached MPLS stack. Moreover the length of the tunnel can be estimated as 256 - LSE-TTL. Unfortunately in most cases the LSE-TTL field is set to 1 or 255 which cannot be the real expected value.
Invisible tunnels cannot be detected.
To infer *implicit* tunnels there are two methods:

- *q-ttl* is the easiest and consists in checking the value of the quoted TTL. If it is greater than 1 then the ICMP Time Exceeded packet can be caused only by the LSE-TTL field that reached zero. Moreover an increase in the values of *q-ttl* should be observed while traversing the tunnel (see Figure 5). Unfortunately some routers report a quoted TTL equals to 1 even in this case and for this reason not all *implicit* tunnels can be discovered in this way. This method can have false negatives (i.e. it doesn't discover some implicit tunnels) but all discovered tunnels are real (no false positives)

- *u-turn* is more complex and requires additional probing. It consists in comparing the ttl of the response packet to two types of probe: the first one generates an ICMP Time Exceeded message and it is generally routed to the destination passing through the last hop of the LSP. The second one generates an ICMP Echo Reply or Port Unreachable message and it is usually routed directly to the destination. For this reason the TTLs of the response packets will differ and probably that hop belongs to an *implicit* tunnel. Moreover, since both the direct path and the path through the last hop of the LSP usually pass through the ingress LER, the difference between the TTLs should follow a pattern like X, $X-2$, $X-4$, ..., 4, 2, 0, where X is two times the tunnel length (see Figure 5). Six probes per interface are sent to be sure that there are no load balancers in the return path (in this case this technique cannot be used). *u-turn* tunnels of length one are discarded because they are probably false positives [3].

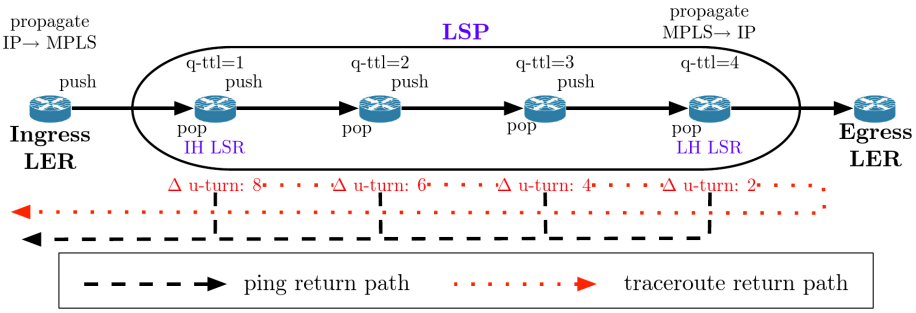

Figure 5: MPLS implicit tunnel detection using *u-turn*

3.2 Results

Table 4 shows the diffusion of MPLS tunnels among the discovered interfaces. Table 5 shows the diffusion of MPLS tunnels among the paths of the campaign. Table 6 shows the distribution of the types of tunnels among the ones discovered during the campaign.

Some of the discovered tunnels start in one AS and end into another AS. The pairs of ASes with at least one of these tunnels are:

- 286 (KPN Internet Backbone, NL) - 12469 (INFONET-NETHERLAND KPN B.V., NL)
- 3741 (ISInternetSolution, ZA) - 30988 (ISInternetSolutions, NG)
- 4270 (Red de Interconexion Universitaria, AR) - 10834 (Telefonica de Argentina, AR)

Tunnel Type	Interfaces	Percentage
Explicit	2228	2.53%
Implicit	1093	1.24%
Opaque	636	0.72%
Any	3957	4.5%

Table 4: Diffusion of MPLS tunnels per interface

Tunnel Type	Paths	Percentage
Explicit	6445	17.8%
Implicit	3391	9.4%
Opaque	4746	13.1%
Any	11381	31.5%

Table 5: Diffusion of MPLS tunnels per path

Tunnel Type	Unique Tunnels	Percentage	Average Length
Explicit	1507	45.31%	3.3
Implicit	758	22.79%	2.8
Opaque	1061	31.90%	2.8
Total	3326	100%	3.1

Table 6: Distribution of MPLS tunnel types and their average length

- 6389 (BellSouth.net Inc., US) - 7018 (AT&T Services Inc., US)
- 8968 (BT-Italia BT Italia S.p.A., IT) - 12797 (ASN-ATLANET BT Italia S.p.A., IT)
- 22742 (CT-ED-NET - State of Connecticut, US) - 25691 (CTSTATEU - Connecticut State University System, US)

As could be expected these pairs belong to ASes owned by the same company or related companies. All other tunnels are completely contained within one AS.

To evaluate the effectiveness of implicit tunnels inference methods the number of explicit tunnel discovered also by implicit techniques can be used, that result can be compared with the one reported in [3] to validate the implementation of that methods:

qttl 32% of explicit tunnel are discovered also using qttl inference method. The remaining tunnels are discovered probably due to routers that copy MPLS-TTL field into IP TTL field before sending the ICMP error package

uturn 3% of explicit tunnel are discovered also using uturn inference method. The remaining tunnels are not discovered probably due to routers that either don't use the LER to forward ICMP error packages or don't reply to pings

any 33% of explicit tunnel are discovered also using one of the implicit tunnel inference method (some tunnels are discovered by both methods). If we consider that this effectiveness is the same for the discovery of real implicit tunnels then the estimated number of implicit tunnels in this measurements campaign is more than 2000

Figure 6 shows the distribution of the tunnel types for the ASes with at least 20 discovered tunnels, descending ordered by number of discovered tunnels.

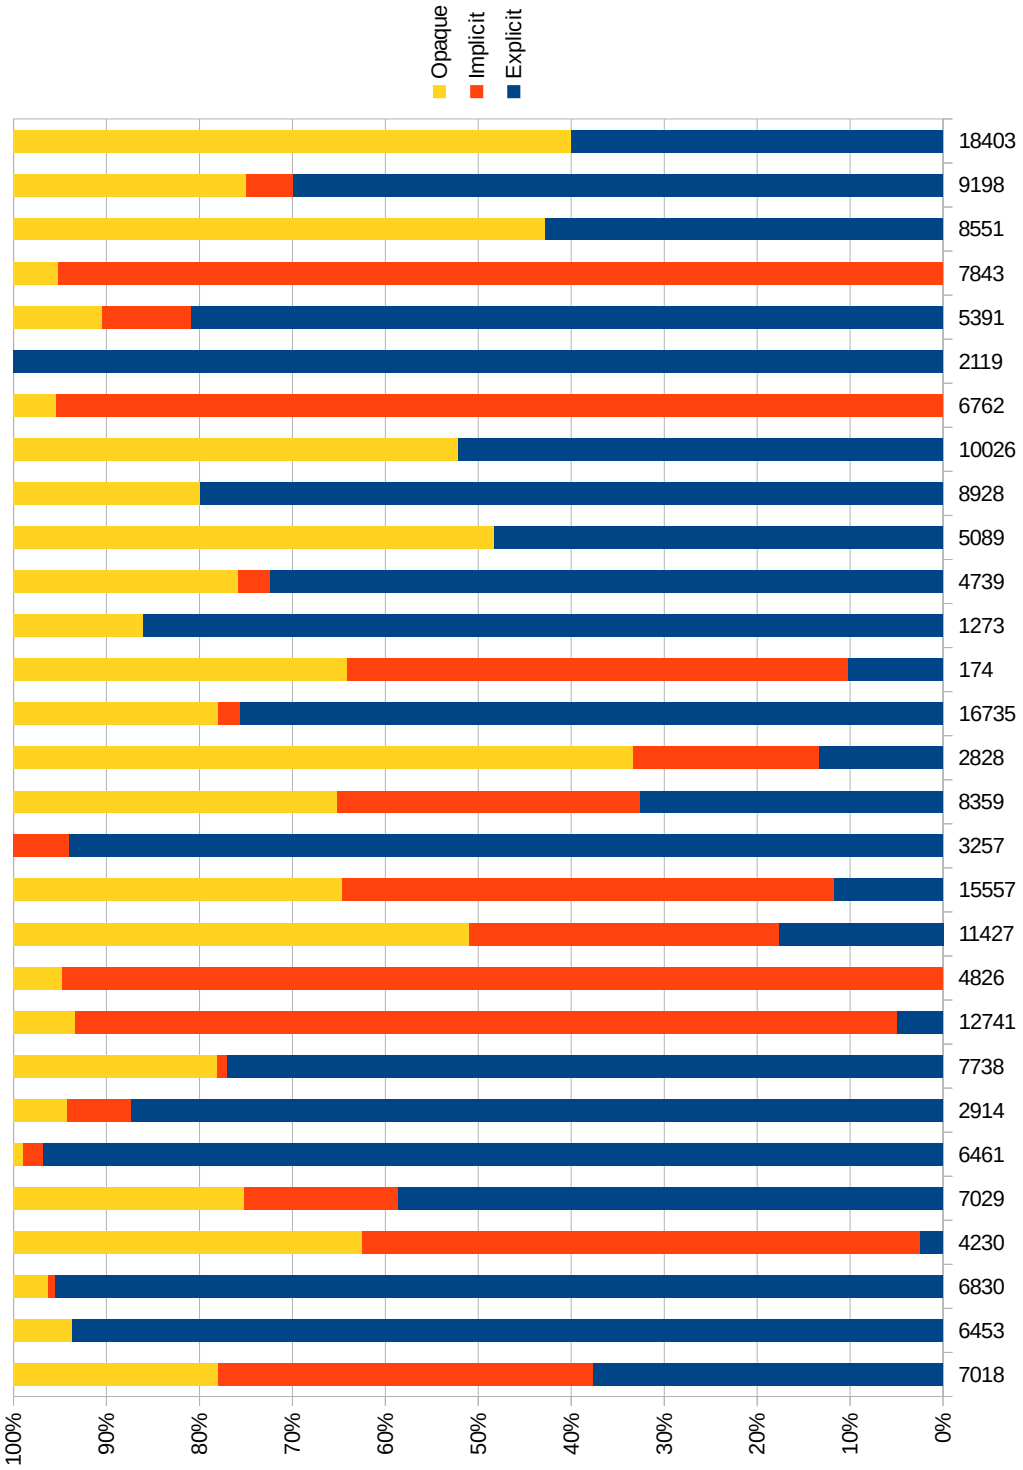

Figure 6: Distribution of MPLS tunnel types for AS with at least 20 discovered tunnels

References

[1] R Bonica, D Gan, D Tappan, and C Pignataro. Extended icmp to support multi-part messages. *draft-bonica-internet-icmp-16 (work in progress)*, 2007.

[2] R Bonica, D Gan, D Tappan, and C Pignataro. Icmp extensions for multiprotocol label switching. *Internet Engineering Task Force, RFC*, 4950, 2007.

[3] Benoit Donnet, Matthew Luckie, Pascal Mérindol, and Jean-Jacques Pansiot. Revealing mpls tunnels obscured from traceroute. *ACM SIGCOMM Computer Communication Review*, 42(2):87–93, 2012.

[4] Daniele Formichelli. Design and implementation of portolan for desktop operating systems linux, windows and os x. Master's thesis, Universit di Pisa, Italy, 2015.

[5] Informatics and Telematics Institute of the Italian National Research Council (IIT-CNR). Isolario.

[6] University of Pisa, Informatics, and Telematics Institute of the Italian National Research Council (IIT-CNR). Portolan.

[7] Yves Vanaubel, Jean-Jacques Pansiot, Pascal Mérindol, and Benoit Donnet. Network fingerprinting: Ttl-based router signatures. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 369–376. ACM, 2013.