

On the Impact of Layer-2 on Node Degree Distribution

Pascal Mérindol, Benoit Donnet, Olivier Bonaventure
Université catholique de Louvain – ICTEAM
Louvain-la-Neuve – Belgium
firstname.name@uclouvain.be

Jean-Jacques Pansiot
LSIIT, Université de Strasbourg
Strasbourg – France
pansiot@unistra.fr

ABSTRACT

The Internet topology data collected through traceroute exploration has been extensively studied in the past. In particular, a remarkable property of the Internet, the power-law shape of node degree distribution, drew the attention of the research community. Several studies have since questioned this property.

In this paper, based on a large dataset collected using `mrinfo`, we show that the node degree distribution is strongly impacted by the presence of layer-2 (L2) networks, such as switches. L2 devices interconnect a large number of routers, themselves being also involved in multiple L2 interconnections. Such a situation induces nodes with very high degree when analyzing the layer-3 (L3) graph with traceroute probing. Considering the physical design of a network, our analysis provides a lower bound on the bias generated by using only an L3 view. We also provide a model that can be a first step towards L2 aware topology generation.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Network topology

General Terms

Measurement

Keywords

network topology, mrinfo, degree distribution, layer-2

1. INTRODUCTION

Several measurement systems aim at eliciting the Internet topology. CAIDA's *Archipelago* [1] uses 41 monitors tracing towards roughly 8.25 millions destinations. *iPlane* [2] constructs an annotated map of the Internet and evaluates end-to-end performances (latency, bandwidth, capacity, etc). *DIMES* [3] is publicly released as a daemon running on

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC'10, November 1–3, 2010, Melbourne, Australia.

Copyright 2010 ACM 978-1-4503-0057-5/10/11 ...\$10.00.




Figure 1: From star to clique: the effect of L2 device using a logical L3 vision

endhosts that traces a set of targets obtained from a central server. The recently introduced *Scamper* is able to conduct Internet measurement to large numbers of IPv4 and IPv6 addresses, in parallel, to fill a specified packets-per-second rate [4]. All of these efforts are based on one of the most common probing tool: *traceroute* [5].

The collected traceroute data is used for various purposes [6]. Among others, from a scientific point of view, the results underlie efforts to model the network [7]. In particular, the node degree distribution fascinates the research community, specially since the Faloutsos et al. seminal paper highlighting the power law shape of this distribution [8].

However, this power law distribution has been extensively questioned. Lakhina et al. [9] showed that, in simulations of a network in which the degree distribution does not at all follow a power law, traceroutes conducted from a small number of monitors can induce a subgraph in which the node degree distribution does follow a power law. Clauset and Moore [10] have since demonstrated analytically that such a phenomenon is to be expected for the specific case of the Erdős-Rényi random graphs [11].

Another potential reason influencing the node degree distribution are layer-2 (L2) clouds (such as Ethernet switches) and how they are used to interconnect layer-3 (L3) routers. A simple scenario involving one L2 device is illustrated on Fig. 1. Fig. 1(a) shows the real topology where four routers are interconnected through an L2 device. Each router on Fig. 1(a) has an L2 degree of 1. On the contrary, Fig. 1(b) provides a view of the same topology as it can be seen through a traceroute exploration: since the L2 device is not discovered (even not suggested), the topology appears as a full mesh and each router has an L3 degree of three. The impact of L2 devices on the node degree distribution has never been deeply investigated, mainly due to the difficulty of iden-

tifying those systems in a traceroute exploration. However, Pansiot and Grad [12], Clauset and Moore [10], and Willinger et al. [13] as well have already pointed the potential impact of L2 networks on the degree analysis.

In this paper, we propose several contributions in order to demonstrate and quantify the impact of L2 devices on node degree distribution. We base our work on a large dataset collected using `mrinfo` [14]. `mrinfo` is a tool that silently discovers all multicast enabled links belonging to an IPv4 router. It has been recently used for Internet topology discovery [15, 16]. One of the key advantages of `mrinfo` is that it allows one to infer the presence of an L2 network between routers. We demonstrate that our inference of broadcast L2 networks is coherent in 90% of the cases and seems to reveal Ethernet L2 devices such as switches. Considering the L3 view of the topology provided by our dataset, on average, 70% of the connections between routers can be classified as *point-to-multipoint* connections (i.e., three or more L3 routers connected through an L2 network). In contrast, a *point-to-point* connection refers to a direct link involving two routers.

Our first contribution is an analysis of the impact of L2 devices on the node degree distribution considering the whole network graph discovered by our `mrinfo` campaign. In particular, we show that the tail of the node degree distribution is heavier when using an L2 agnostic view than when distinguishing L2 and L3 layers. Our results also question the scale free property and the shape of such a distribution considering the physical topology of the Internet.

Second, we refine our view of the problem by focusing on several large Tier-1 and Transit autonomous systems (ASes). Our results show that, although each AS is topologically engineered in a different way, in all cases, the impact of L2 networks is really significant. We notice a great shift in the node degree distribution but specific to each AS. We analyze several key distributions to understand this property in each AS. Among others, we observe three phenomena:

1. the degree of L2 devices is generally greater than the one of routers,
2. routers are frequently connected to several L2 devices,
3. large access routers are involved in even more L2 connections.

To highlight the correlation implied by the third point¹, we provide a first analytical model that captures the heavy tail distribution of the L3 view. We also believe that this model can be a first step towards L2 aware topology generation.

Generally speaking, our results indicate that L2 networks must be carefully taken into account when analyzing physical topology properties such as path redundancy. The huge difference between the physical graph of the Internet and its L3 view can introduce a large bias in a path diversity analysis. Our L2 based topology extraction allows one to analyze a lower bound on such a bias.

The remainder of this paper is organized as follows: Sec. 2 describes our dataset, the methodology we follow for preprocessing the data, and our L2 inference mechanism; Sec. 3

¹There exists a strong correlation between the router point-to-point degree and the number of its L2 neighbors.

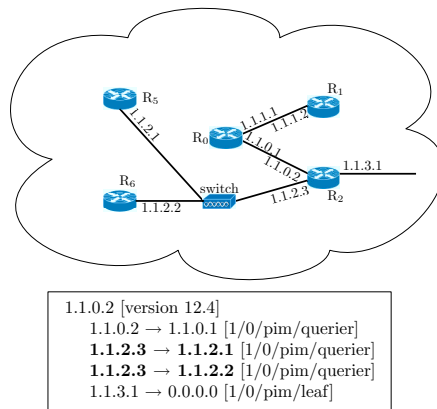


Figure 2: `mrinfo` example

provides a global view of the L2 impact on the node degree distribution; Sec. 4 refines this view by focusing on large Tier-1 and Transit ISPs; Sec. 5 provides a first analytical model capturing the dependence between L2 and L3 degrees; Sec. 6 positions our work regarding the state of the art; finally, Sec. 7 concludes this paper by summarizing its main results.

2. DATASET

In this section, we discuss the dataset used throughout this paper and how data was collected. We first present `mrinfo` (Sec. 2.1), a tool that is able to silently discover all the multicast interfaces of a router, and how we used it for gathering topology data. Next, we discuss how our dataset was preprocessed (Sec. 2.2). Finally, we explain how we can infer L2 networks from `mrinfo` data and how confident we are in our inference technique (Sec. 2.3).

2.1 `mrinfo`

`mrinfo` messages use the Internet Group Management Protocol (IGMP [17]). IGMP was initially designed to allow hosts to report their active multicast groups to a multicast router on their LAN. Most IGMP messages are sent with a Time-to-Live of 1. However, the Distance Vector Multicast Routing Protocol, DVMRP, has defined two special types of IGMP messages that can be used to monitor routers [18]. Although current IPv4 multicast routers do not use DVMRP anymore, they still support these special IGMP messages. Upon reception of an IGMP `ASK_NEIGHBORS` message, an IPv4 multicast router replies by sending an IGMP `NEIGHBORS_REPLY` message that lists all its multicast enabled local interfaces with some information about their state. Cisco and Juniper routers also report in the IGMP `NEIGHBORS_REPLY` message the version of their operating system. Fig. 2 shows an example of the usage of `mrinfo` to query the router `R2`, `1.1.0.2` being the responding interface of `R2`. `mrinfo` reports that this router is directly connected to `R0` (through interface `1.1.0.1`). We can also notice that `R2` is connected to routers `R5` and `R6` through an L2 network (labeled “switch” in Fig. 2) because interface `1.1.2.3` appears twice in the `mrinfo` reply (see bold text in Fig. 2). Finally, `mrinfo` reports that interface `1.1.3.1` has no multicast neighbor because the right IP address is equal to `0.0.0.0` (or is directly connected to a LAN,

as indicated by the “leaf” keyword). All this information is obtained by sending a single IGMP message. In practice, `mrinfo` provides information similar to the output of a `show` command on the router’s command line interface.

Our approach in probing the network with `mrinfo` is recursive and we call such a probing scheme `mrinfo-rec`. Initially, `mrinfo-rec` is fed with a single IP address corresponding to the first router attached to the `mrinfo-rec` vantage point. `mrinfo-rec` probes this router and recursively applies its probing mechanism on all the collected IP addresses. These recursive queries stop at unresponsive routers or when all discovered routers have been queried. The same process is run every day. It is worth to notice that an address not replying to an `mrinfo` probe during a given day will not be queried the days after except if it appears again in a list of captured addresses.

To illustrate this behavior, let us apply it on the topology depicted in Fig. 2. `mrinfo-rec` receives, as input, an IP address belonging to router R_0 . From R_0 , `mrinfo-rec` collects a set of neighbor IP addresses, i.e., $\{1.1.1.2, 1.1.0.2\}$. For all IP addresses in this set that were not previously probed, `mrinfo-rec` sends an IGMP `ASK_NEIGHBORS` message and, if the probed router replied, it again runs through the set of neighbor IP addresses collected.

This recursive probing scheme comes with the strong advantage of being very easy to setup as it initially requires only a single IP address as input. However, it has the drawback of limiting the `mrinfo-rec` probing spectrum. If the address set specified as input for a given day corresponds to non-responding routers (because, simply, the ISP filters IGMP messages), `mrinfo-rec` will not be able to discover any topological information. To overcome this limitation, the set of responding router addresses of a given day is given as input of `mrinfo-rec` the next day.

2.2 Data Processing

Since May 1st, 2004, we have been collecting topology data from a host located at the University of Strasbourg, France. In this paper, we consider the data collected until the end of December 2008. The entire dataset is publicly available [19] and has been collected using `mrinfo-rec`.

During the whole probing period, on average, `mrinfo-rec` daily discovered roughly 10,000 different routers while scanning 100,000 interfaces. For the remainder of this paper, we remove interfaces with non-publicly routable IP addresses or marked as down or disabled. The addresses that we consider as non-publicly routable are the special-use IPv4 addresses described in [20]. Specifically, we eliminate the private IP address blocks 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. We also remove the loopback address block 127.0.0.0/8 and the 0.0.0.0 address. On average, 25% of the interfaces collected by `mrinfo-rec` are categorized as non-publicly routable. We also remove all multicast tunnel interfaces detected with the keyword “tunnel” in the `mrinfo` raw output.

From our four year daily dataset, we arbitrarily select, each month, the largest `mrinfo-rec` output file, leading thus to 56 global topologies. Those 56 graphs form the dataset analyzed in this paper.

In the analysis provided in this paper, the inference of *L2 networks* is critical. In our context, by *L2 network*, we mean a technology allowing a router to transmit IP packets to several other IP routers through the same interface, i.e., a multi-access network. One often distinguishes between *Non*

Broadcast Multiple Access (NBMA) networks (ATM, Frame Relay, X25), and *broadcast networks* such as most LAN networks (Ethernet, token ring, FDDI).

These two kinds of networks behave differently as far as IP multicast is concerned. In particular, when using *Protocol Independent Multicast* (PIM) as a routing protocol in a broadcast network, each router is a PIM neighbor of each other, and only one of them is elected as the *querier* [21]. These properties can be easily checked with the `mrinfo` raw output (see the symmetry and querier rules in Sec. 2.3). In NBMA networks, IP packets are usually transmitted via circuits that behave as a collection of point-to-point or point-to-multipoint connections (which does not imply the rules defined in Sec. 2.3). Note that if these circuits locally use different subinterfaces and IP addresses, they cannot be discovered by `mrinfo`.

Furthermore, we can also study the specific case of MPLS networks [22]. Notice first that an MPLS network is not an actual *L2 network* since MPLS routers are usually IP routers, and MPLS signalling (LDP [23], RSVP [24], BGP [25]) uses IP. On the one hand, most techniques to transmit multicast packets over MPLS (for traffic engineering or multicast *L3 VPN*) use some sort of tunnels (Traffic engineering tunnels, or GRE encapsulation for PIM in *L3 VPN*), and since we discard tunnel interfaces from `mrinfo` output in our analysis, such MPLS networks should not introduce artifacts. On the other hand, in the case of basic MPLS, multicast packets are not sent through Label Switch Paths (LSP), contrary to unicast packets. Therefore, although `traceroute` may infer a false link between the two endpoints of an LSP, `mrinfo` will correctly provide the direct IP neighbor. However, the case of *L2 multicast VPN* (VPLS [26, 27]) is different since IP multicast routers (if any) of each customer site will see each other through a virtual *L2 switch* instead of the provider MPLS network. The `mrinfo` analysis will correctly discover this *L2 network*, but will not detect that it is an *L2 VPN* through the provider network. However this can happen only in customer ASes and will be seen only if `mrinfo` queries are not filtered between the VPN and Internet which is not very likely. In the following, we only study the impact of broadcast *L2 networks*: in practice, *ethernet networks* such as switches.

Among all interfaces discovered in our selected dataset (i.e., 56 global topologies), more than a million has a fully qualified domain name that allows us to determine the physical interface type (i.e., `ge0`, `tunnel0`, `vlan500`, `pos0`, ... , are typical naming conventions used by network administrators). About 42% of these names allow us to infer the interface type. Focusing on multi-access networks, we notice that about 99% of discovered interfaces seem to be Ethernet links of various data rates (from 10M to 10G) or VLAN interfaces (probably on top of Ethernet interfaces too).

2.3 L2 Inference

In Sec. 2.1, we have seen how an *L2 network* is suggested from the raw `mrinfo` output. However, this inference comes with two limitations. First, we are only able to detect *L2 networks* involving at least three routers. Second, the inferred presence of an *L2 network* may hide a “cascade” of *L2 networks*. However, the nature of the router degree distribution (i.e., the actual number of devices directly connected to it) is not impacted by those limitations. Only the degree of *L2 networks* might suffer from this. In this paper, we

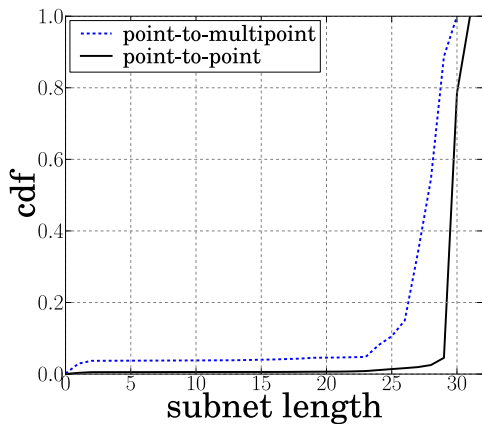


Figure 3: Subnet length distribution for point-to-point and point-to-multipoint connections

focus on the degree distribution of routers considering the difference between the L2 and L3 visions.

Here, we verify whether our inference technique is accurate and what kind of technology it seems to reveal. Indeed, there are several kinds of L2 networks such as ATM clouds or simple Ethernet switches. In our analysis, we particularly focus on broadcast L2 networks such as Ethernet.

There are several ways to validate our L2 inference technique. First, one could use small testbeds and access networks whose administrators give us the physical map. We performed such tests in our local AS in Strasbourg and found that our inference is correct in revealing L2 networks.

Second, the validation can be done “theoretically” based on three rules:

- *symmetry rule.* We verify that all the routers attached to a given potential L2 network have the same view. For instance, on Fig. 2, we are able to detect that router R_2 is connected to R_5 and R_6 through an L2 network. When probing R_5 and R_6 with `mrinfo`, the L2 connection to R_2 must also be visible in their `mrinfo` output as well as the connection between them. Such a symmetry rule suggests the detection of broadcast L2 network.
- *querier rule.* In a normal case, only one router per L2 network must be tagged as the IGMP “querier” (i.e., it won the querier election on the subnet [21]: it has the greatest IP address on the subnet). In the case depicted in Fig. 2, as interface 1.1.2.3 is tagged as “querier”, interfaces 1.1.2.1 of R_5 and 1.1.2.2 of R_6 should not be tagged as such.
- *subnet mask rule.* We verify the validity of the minimum subnet mask covering all IPs involved in the L2 network. Fig. 3 provides the cumulative distribution of the subnet lengths for point-to-point (i.e., a direct connection involving two routers) and point-to-multipoint (i.e., three or more L3 routers connected through an L2 network) connections. We notice that roughly 90% of the prefix lengths are very long: usually longer than /24 for point-to-multipoint and using a /30 or /31 prefix for point-to-point connections.

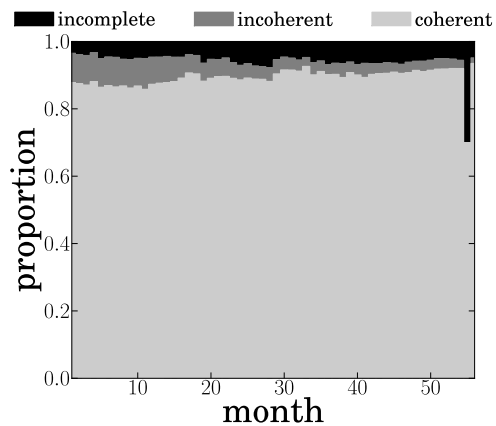


Figure 4: L2 inference rule verification with a /16 threshold for the subnet mask rule

Fig. 3 shows that more than 99% of the point-to-point connections (plain line) use a subnet mask longer than /24 and more than 95% use a subnet mask longer than /28. In a few cases the minimal subnet mask may yield a prefix shorter than /24 or even /16 (less than 5% for point-to-multipoint connections). This is probably an anomaly, possibly explained by the use of tunneling or a NBMA L2 WAN technology (ATM, frame relay, ...). We consider a /16 threshold for the minimal subnet mask rule applied to an L2 network.

Note that the symmetry and the querier rule need at least two responding L3 nodes connected to the same L2 network to be verified. Based on those three rules, we classify an L2 inference in one of three categories:

- *coherent.* In this case, the three rules are verified and we are sure that our L2 inference technique is accurate and suggests a traditional L2 broadcast network. This is the best case.
- *incoherent.* In this case, at least one of the three rules is not verified. This category gathers L2 networks that we are not able to accurately reveal (such as NBMA L2 WAN technology).
- *incomplete.* In this case, the L2 component is only seen by a single L3 device (the symmetry and querier rules cannot be verified).

Fig. 4 shows how our L2 inference technique verifies the three prescribed rules over the 56 months. Note that L2 inferred networks falling in the *coherent* category are typically “broadcast L2 network” such as Ethernet switches (thanks to the symmetry and querier rules).

Over the whole period and on average, we have the following results: 90.16% of the L2 inferred nodes are coherent and suggest broadcast capable L2 networks such as Ethernet switches. 5.84% of the inferred nodes are incomplete, meaning that only 4% are incoherent. Thus, in the vast majority of the cases, our L2 inference algorithm seems to be verified and coherent. In the remainder of our analysis, we ignore L2 components falling into incoherent or incomplete categories. Note that those problematic cases are not uniformly

Notation	Meaning
S	set of L2 nodes
R	set of L3 nodes
$N = R \cup S$	set of nodes (including L2 inferred ones)
$succ(a)$	set of neighbors of node a
$deg(a) = succ(a) $	degree of node a in the L2 graph
$r(x)$	number of L3 nodes of degree x
$s(x)$	number of L2 nodes of degree x
$t(x) = r(x) + s(x)$	number of nodes of degree x
$p(x)$	number of L3 nodes having a degree of x and at least one neighbor in S
$b(x)$	number of L3 nodes having x neighbors in S
$r'(x)$	number of routers having a logical L3 degree of x in the L3 graph

Table 1: Notations

distributed across the domains and time: for example, we notice that a large proportion of those cases comes from AS680 and AS5400 and seem to progressively and slightly disappear. We suppose that those cases correspond to “L2 circuits technologies” such as ATM.

Finally, note that the drop observed on Fig. 4 during the week 55 is due to a stronger filtering policy applied by the French National Research Network leading to a strong decrease in the amount of routers discovered by `mrinfo`.

3. HOW L2 NETWORKS IMPACT THE L3 NODE DEGREE DISTRIBUTION

In the remainder of the paper, we focus on the relative difference between an L3 and an L2 vision of the network. We do not claim to provide the actual degree distribution: our goal is rather to emphasize the difference in the apparent connectivity when adopting a vision of the network at a given layer.

Hence, using the same set of IP addresses, we construct two graphs, one corresponding to a `traceroute` like view and the other to the L2-aware view inferred thanks to `mrinfo`. Obviously, our analysis is based on a partial view of the Internet graph. Furthermore, we restrict the two constructed topologies to their connected and responding components, i.e., we remove all links connected to a router that does not respond to `mrinfo` so that we guarantee a set of active and existing links and routers. We assume that this partial vision does not impact the relative comparison although it may introduce some non coherent results in an absolute analysis, e.g., a large proportion of leaf routers. In practice, these routers might be located at the boundary of the multicast part of the probed network. Indeed, if their neighbors not already discovered are not multicast enabled, `mrinfo` will only report one active link per multicast border router: the one allowing `mrinfo-rec` to reach it.

3.1 Notations

In an L2-aware graph, there are two types of connections: L3 point-to-point connections that are correctly interpreted by `traceroute`, and L2 point-to-multipoint connections that

are seen as a full mesh of point-to-point connections using `traceroute`. Our goal is to measure the importance of this confusion using the node degree distribution as an indicator to understand the difference between the L3 and L2 graphs. Note that one may also use other indicators such as the clustering coefficient, the assortativity, or a path diversity indicator.

Table 1 provides the basic definitions and notations used in the remainder of this paper. In our analysis, the notion of *degree* is related to the L2 graph. In contrast, we use the term *layer-3 degree* to refer to the number of edges discovered by an L3 tool such as `traceroute`. A node n has a logical L3 degree of x if and only if

$$deg_r(n) + deg_s(n) = x. \quad (1)$$

Notations $deg_r(a)$ and $deg_s(a)$ respectively refer to the L3 point-to-point degree of node a , and to the L3 point-to-multipoint degree of a such that:

$$deg_s(a) = \sum_{\forall v \in succ(a) \mid v \in S} (deg(v) - 1). \quad (2)$$

Furthermore, we consider the IP multigraph issue: there might exist several links or IP address pairs connecting the same two nodes. In our analysis, the degree of a node is the number of the neighboring IP addresses, e.g., the number of active and replying IP addresses seen on the right part of the raw `mrinfo` output. Typically, an L2 node refers to a broadcast network such as an Ethernet switch, while an L3 node corresponds to an IP router.

Eqn. 2 is the key point of our analysis: L2 point-to-multipoint connections between a node a and its neighbors are seen as a full mesh using an L2 agnostic view, whereas using `mrinfo`, we can infer a star centered on L2 nodes instead of cliques.

In order to reveal the underlying L2 topology, when parsing our dataset, we generate an L2 aware graph containing inferred L2 nodes and links. The set of nodes in this L2 inferred graph is denoted N . Except considering the use of technologies such as VLAN, we can assume that this vision describes the L2 topology. Let us assume an L2 network connected to m routers ($2m$ directed links in the L2 graph), then it induces $m \times (m - 1)$ directed links in the L3 graph. With the use of VLAN, let us now consider k VLANs deployed on a L2 node connecting m routers (at most $2km$ directed links in the L2 graph instead of $2m$ in the physical topology), then, in the worst case, it can induce $km \times (m - 1)$ directed links in the L3 graph. Note that we provide a lower bound on the bias resulting from a L3 graph when considering the physical topology: $2m < 2km \leq km \times (m - 1)$, $\forall k > 1$, $m \geq 3$. Moreover, our analysis is conservative because `mrinfo` may underestimate the number of active IP interfaces: the probing view is restricted to multicast enabled links and `mrinfo` replies may be fragmented and then only partially interpreted. See the technical report [28] for details.

3.2 Results

In this section, we focus on the node degree distribution considering two points of view: the logical L3 graph and the underlying L2 graph. In particular, we would like to know how L2 devices might impact the degree distribution. The purpose of our study is to: (1) highlight the huge difference

between graphs describing L2 and L3 connections, and (2) explain why there exist nodes with very high degree in a L3 graph.

Using our L2 inference technique, we can easily compute $r(x)$ (i.e., the number of L3 nodes of degree x - see Table 1) and $s(x)$ (i.e., the number of L2 nodes of degree x - see Table 1). In contrast, with an L2 agnostic view such as with traceroute, the global degree distribution is purely logical: we can only compute $r'(x)$ (i.e., the number of routers having a logical L3 degree of x - see Table 1). Hence, using our `mrinfo` dataset, we can easily emphasize the difference between $t(x)$ (i.e., the sum of $r(x)$ and $s(x)$ - see Table 1) and $r'(x)$ because we are able to compute all key distributions defined in Table 1.

Fig. 5 shows the difference between those two distributions.² The horizontal axis gives the degree, while the vertical axis shows the proportion distribution. The shaded area represents the $t(x)$ distribution while the line refers to the $r'(x)$ distribution. Note that the maximum degree for $r'(x)$ is 115 but it is too tight (0.01%) to appear on the plot.

The degree distribution is strongly impacted using only an L3 perspective: the L2 agnostic distribution has a much more heavy tail than the one considering both L2 and L3 nodes. For example, using the L2 view ($t(x)$), we observe that there do not exist routers having a degree greater than 45 whereas the L3 views ($r'(x)$) suggests routers having degree larger than 100. This means that there are more nodes having large degree when L2 interconnections are translated to full mesh connections. Consequently, it also means that the shape parameter used to model the power law distribution of the node degree and the degree of connectivity in general are strongly biased if one wants to consider a physical view of a given network (which is generally the case when measuring path diversity characteristics for example).

This result is of the highest interest since many studies use L3 topologies such as those provided by *Rocketfuel* [29] to perform path redundancy analysis and to evaluate the quality of the re-routing protocol offered by link protection schemes [30, 31]. However, *Rocketfuel* topologies do not model the specific *Shared Risk Link Groups* (SRLG) [32, 33] of L2 networks. This leads to nodes with overestimated degree and thus more path diversity and redundancy than in the actual underlying graph.

A first attempt to model the high degree shifting of an L3 view is to consider that a router connected to an L2 network “will increase” its logical L3 degree by the number of neighbors connected to that L2 network. With a naive model considering that, in the vast majority of cases, L2 networks interconnect leaf routers (e.g., routers of degree 1, so that distributions $p(x)$ and $b(x)$ -see Table 1- are meaningless), we may obtain:

$$r'(x) \approx \begin{cases} r(x) + s(x+1) \times x & \text{if } x > 1, \\ r(1) - \sum_{s \in S} \text{deg}(s) & \text{otherwise.} \end{cases} \quad (3)$$

However, we can easily notice that this model is not sufficient to explain the heavy tail of the L3 distribution. Indeed, we can observe that obtaining such an amount of nodes hav-

²For the sake of clarity, we select only one of our 56 global topologies, the one being the most connected (i.e., with the maximum number of edges). The general behavior is roughly identical across our whole set of global topologies.




Figure 5: L2 ($t(x)$) vs. L3 ($r'(x)$) Node Degree Distribution Comparison

ing such a high degree requires point-to-multipoint connections through multiple L2 nodes. We will detail this issue in the next section considering a closer look at each key distribution (e.g., the L2 nodes degree distribution and the number of L2 neighbors) and according to the specifics of several routing domains.

4. DOMAIN SPECIFICITY

In Sec. 3, we have seen that the impact of L2 networks on the node degree distribution is considerable. This analysis was done on our entire dataset, without considering the potential particularities of domains. In this section, we refine our analysis, focusing on several ASes of interest: we want to know whether there exist some common patterns between domains.

We first discuss our methodology in extracting and selecting particular domains from the raw `mrinfo` dataset (Sec. 4.1) and, next, analyze the impact of L2 networks on the node degree distribution of selected ASes (Sec. 4.2).

4.1 Methodology

From our four year daily dataset, we arbitrarily select, each month, the largest `mrinfo-rec` output file, leading thus to 56 global topologies. For each of these 56 selected `mrinfo-rec` files, we extract the topologies of four particular ASes. This extraction is made possible with our *router-to-AS* mapping algorithm [16]. This algorithm assigns a unique AS number to each router and identifies AS border routers (AS-BRs). Our algorithm is based on a set of rules using common IP address allocation guidelines and probabilistic methods. To reinforce our mapping, it verifies, at each step, the consistency of the assignment returned by each rule. A subset of those rules are also used in [34].

Using this router-to-AS mapping, we have extracted a set of intra-domain topologies that can be classified as follows: *Tier-1* (such as Sprint), *Transit* (such as IUnet), and *Stub* networks (such as UNINETT). Those one-per-month topologies are freely available in various formats.³

Due to space limitations, we only consider four ASes taken from the set of extracted topologies, two being Tier-1 (AS1239

³See <http://inl.info.ucl.ac.be/content/mrinfo>.

Level	ASN	Name	$ R $	$ S $	α	$\beta - L3$	$\beta - L2$
Global	-	-	8,670 ± 448	1,100 ± 63	0.38 ± 0.01	0.69 ± 0.01	0.34 ± 0.01
Tier-1	1239	Sprint	649 ± 7	63 ± 0	0.58 ± 0.01	0.70 ± 0.01	0.34 ± 0.01
	3356	Level-3	254 ± 6	123 ± 7	0.67 ± 0.03	0.87 ± 0.01	0.69 ± 0.01
Transit	1267	IUnet	517 ± 16	225 ± 10	0.77 ± 0.04	0.85 ± 0.03	0.74 ± 0.03
	3269	Telecom Italia	822 ± 28	73 ± 2	0.69 ± 0.02	0.95 ± 0.01	0.67 ± 0.01

Table 2: Selected ASes global overview

and AS3356) and two being Transit (AS1267 and AS3269). We select those ASes because they are the largest topologies in our dataset. Furthermore, they are very well connected and seem more impacted by the presence of L2 networks compared to the global vision provided by the whole dataset (see Table 2). We performed the same analysis for other domains, results can be found in [28]. Note that we remove inter-domain links from those topologies reducing so the degree of connectivity of the routers in those ASes.

Table 2 provides an overview of the amount of L3 and L2 nodes in our dataset. Values have been averaged over the 56 considered `mrinfo` files. The confidence intervals are also given. In addition, Table 2 provides three ratios: (α) the proportion of routers connected to at least one L2 node in the L2 graph, and ($\beta - L3$ and $\beta - L2$) the proportion of connections using an L2 node. Note that $\beta - L3$ refers to the quantity of connections using an L2 point-to-multipoint network in the L3 graph whereas $\beta - L2$ refers to the same quantity in the L2 graph.

For the remainder of the analysis and for the sake of clarity, we select only one probing day per AS according to the size of the largest connected component inside each AS. [28, Sec. V] details this process: intuitively, we simply extract the *best topology* for each AS.

4.2 Experimental results

First, as shown in Table 2, the proportion of L2 nodes (i.e., $\frac{S}{N}$) is generally quite significant. On the whole dataset, approximately 11.2% of the nodes are inferred as L2 nodes. It is worth to notice that this proportion is not uniformly distributed across ASes. For instance, some large ASes, such as AS1267 or AS3356, verify: $|S| \approx 0.3 \times |N|$. This high proportion can obviously lead to an important difference when considering the connectivity degree in L2 and L3 graphs. However, another aspect impacts this difference: if the degree of L2 nodes is high, it means that a large number of routers are involved in L2 connections. For example, we can notice that, although less than 10% of the nodes in AS3269 are L2 networks, about 97% of the connections between routers modeled by a L3 graph actually involve L2 nodes. This phenomenon is due to the fact that L2 nodes interconnect a large number of routers. It is worth to notice that the number of logical links is much more important in the L3 graph than the number of “physical” links in the L2 one when most of L2 nodes involve more than three routers.

In the fashion of Fig. 5, and for the four considered ASes, Fig. 6 compares the node degree distribution when L2 nodes are taken into account ($t(x)$ - see Table 1) and when L2 nodes are ignored by an L3 only vision ($r'(x)$).

As mentioned in Sec. 3, we notice a great shift in the distribution of large degrees. However, we observe that this difference significantly varies from one AS to another. For

example, Level 3 is much more impacted than Sprint. In all cases, we can notice that the model given in Eqn. 3 cannot describe such a degree shifting. For instance, even for Sprint, the maximum router degree in the topology is 21 (see Fig. 7(a)). And if we assume that this router is connected to the largest L2 node of degree 21 (see Fig. 7(b)), we may obtain a L3 logical degree of $21 + 21 - 1 = 41$ whereas there exists a router having 47 neighbors in the L3 topology. Note that the Sprint topology given in the Rocketfuel dataset [29] presents approximatively the same degree distribution as the L3 one ($r'(x)$) given in Fig. 6(a): there also exists routers having a degree greater than 40. However, note that it is likely that this bias is also due to alias resolution limitations.

The bias resulting from an L2 agnostic view obviously depends on the $s(x)$ distribution, but also on the way L2 and L3 nodes are interconnected. Indeed, we can notice that the ratio $\frac{S}{N}$ is not the only factor explaining this bias. Fig. 6 shows completely different $t(x)$ and $r'(x)$ distributions even for topologies with less than 10% of L2 nodes, such as Sprint and Telecom Italia.

In practice, two other distributions could emphasize the influence of L2 nodes: $p(x)$ (i.e., the number of L3 nodes n having a degree of $x \mid deg_s(n) > 0$ - see Table 1), and $b(x)$ (i.e., number of L3 nodes having x L2 neighbors - see Table 1). Fig. 7(a), Fig. 7(b), Fig. 7(c), and Fig. 7(d) illustrate some key patterns useful to understand the difference between $t(x)$ and $r'(x)$.

First of all, we can notice among our set of selected ASes illustrated on Fig. 7 that network administrators do not seem to use a common engineering way to construct and deploy their topology. It highly depends on economic and technological factors not captured by classical models. Compared to the global analysis depicted in Fig. 5, we can highlight the difficulty of providing a model capturing the particularities of each domain.

Second, focusing on common properties, we can observe that r and p are slightly different distributions: keeping in mind that most routers counted in the r distribution also fall in p , we can emphasize the fact that peak values in r become still more important in p such that routers connected to L2 networks may strongly influence those peaks values. For example, in AS1239, it means that globally 40% of its routers have a degree of 1 while focusing on routers connected to L2 nodes, this proportion grows to almost 70%. In practice, on AS1239, we can conclude that the majority of the routers connected to L2 nodes are leaf nodes. However, this observation is highly topology dependent. The only common characteristic is the increase of the same peak values for each AS. Those results explain why the difference between the t and r' distributions shown in Fig. 6 is so high for most representative degree values.

Fig. 7(b) and Fig. 7(d) plot the degree distribution of L2


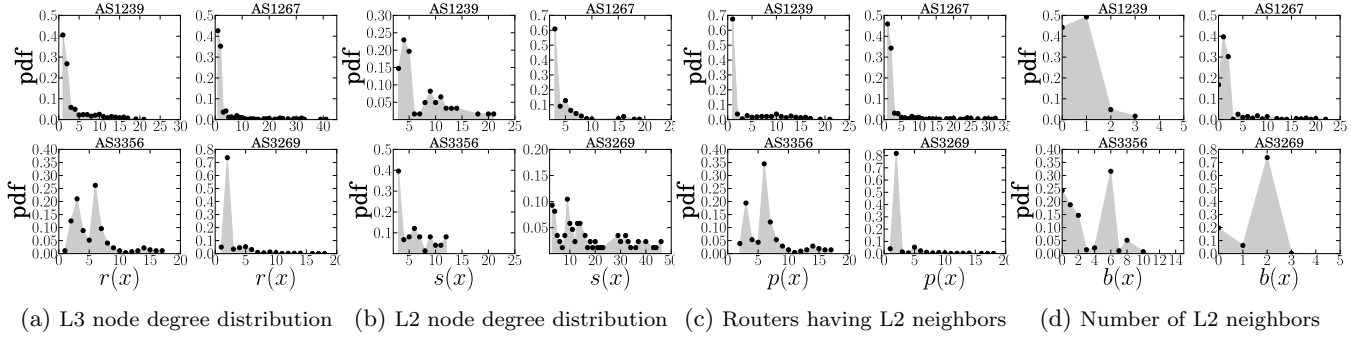


Figure 6: L2 ($t(x)$) vs. L3 ($r'(x)$) Degree distribution comparison for our set of selected ASes



(a) L3 node degree distribution (b) L2 node degree distribution (c) Routers having L2 neighbors (d) Number of L2 neighbors

Figure 7: Key distributions for our selected ASes

nodes, $s(x)$, and the distribution of the number of L2 neighbors of routers, $b(x)$. First, we can notice on Fig. 7(d) that the power law modeling the L2 node degree distribution is again highly dependent on the topology and not equivalent to the one modeling $r(x)$ for each given topology. The maximum degree of $s(x)$ plays a critical role in the absolute difference between $t(x)$ and $r'(x)$ but Fig. 7(c) highlights the higher importance of $b(x)$. First, we can see on Fig. 7(d) that most routers are connected to L2 nodes. Only 10% for AS3356 and 40% for AS1239 of routers are only connected through point-to-point links. Moreover, we observe that, for topologies such as AS1267 and AS3269, most routers are not connected to a single L2 node but to several ones: two for AS3269 and six for AS3356. Generally speaking, the number of L2 neighbors is an important factor explaining the heavy tail distribution of $r'(x)$. In practice, it means that a significant subset of the routers are connected to several high degree L2 nodes. Surprisingly, this experimental analysis shows that, for topologies such as AS3356 and AS3269, more than half of the routers are connected to more than two L2 nodes. We also notice that there exists routers connected to more than 20 L2 nodes in AS1267: in that case, it probably means that those routers are involved in VLANs crossing one or several switches.

Fig. 8 illustrates two practical cases leading to a large difference between $r'(x)$ and $t(x)$ distributions. Those scenarios are common in our dataset and may explain why the impact of L2 networks is so significant:

- (a) MAN access routers are generally connected to multiple metropolitan locations to distribute the access. The




Figure 8: Examples of L2 configuration

use of L2 networks is very common in order to facilitate and reduce the cost of small deployments.

- (b) The use of redundant access connections for load balancing purposes can increase of this effect.

On Fig. 8(a), the Internet access router has a logical L3 degree of nine while it has actually only three physical links. On Fig. 8(b), the two access routers are logically connected to each other router in order to ensure load balancing and resiliency. Note that a further analysis is necessary to understand the difference in terms of L2 usage between core and access part of a network. Indeed, the pure backbone part is probably differently impacted than the border of a network. To mitigate this effect, we need to characterize the core and access parts in order to divide networks in multiple areas. The use of traceroute probes may help us to perform

such a mapping. However, this is a huge challenge that is outside the scope of this paper.

5. WHY SUCH A GREAT SHIFT IN THE NODE DEGREE DISTRIBUTION?

The goal of this section is to model the relationships between the key L2 and L3 distributions explaining such an heavy tail distribution when considering a L3 view. We focus on the correlation between L2 and L3 devices that generates routers having a very high degree. In Sec. 4.2, we have seen that the degree distribution of L2 nodes ($s(x)$) and the distribution of the number of L2 neighbors ($b(x)$) are the first impacting factors.

5.1 L2 and L3 Correlations

In order to provide an accurate model, correlations between keys distributions should be considered. For routers involved in L2 interconnections, we denote p their maximum degree and b their maximum number of L2 neighbors. Using $r(x)$, $p(x)$ and normalizing $b(x)$ and $s(x)$ ⁶, we are able to propose a probabilistic model. For that purpose, we also need to quantify the nature of the dependence between the $p(x)$, $b(x)$ and $s(x)$ distributions.

Fig. 9 shows the normalized $p(x) \times b(x)$ dependence distribution (in practice, we have also plotted the distribution of routers which are not connected to L2 nodes). It allows us to understand the relationship between the $p(x)$ and $b(x)$ distributions. Obviously, $b(x)$ depends on $p(x)$ because a router cannot have more L2 neighbors than its total degree, as indicated by empty portions in upper triangles of Fig. 9. We further notice that large degree routers generally have a great number of their neighborhood being L2 nodes.

We thus observe a strong correlation between those two distributions meaning that $b(x)$ is not uniformly distributed among routers and their degree: a router having a large number of neighbors has a higher likelihood to be connected to a large number of L2 nodes. This correlation seems quite linear on most of the ASes that we analyzed.

We do not find any other linear dependency between other relationships: Figs. 10 and 11 do not suggest any kind of obvious dependence in relationships⁷. However, we cannot argue that there does not exist any kind of non linear dependence. Figs. 10 and 11 may also help to understand some specific behaviors seen in Fig. 6. Indeed, one can argue that $s(x)$ may be the main factor explaining the shift in the degree distribution. Potentially, if there exist some particular relations between $p(x)$ and $s(x)$, it may explain several shifting details in the degree distribution. For example, looking at the peak value at $x = 14$ in Fig. 6(b), we can explain this phenomenon with Fig. 11(b): we observe that a great

⁵Here, we also provide the degree distribution for routers which are not connected to L2 networks, i.e. $r(x)$ in general.

⁶Note that for $s(x)$, it implies to take into account the fact that large degree L2 nodes impact a greater number of routers than L2 nodes with small degree. For $b(x)$, we decide to model the dependence between this distribution and the one of $p(x)$.

⁷For these two figures we list the degree of L2 networks connected to routers and then analyze the potential dependence. Note that those distributions have been normalized according to the proportion of routers having a given x-axis value. Furthermore, one may also analyze the three dimensional relationship between p , s , and b distributions.

number of routers having a degree of 1 ($p(1)$ values) are connected to L2 nodes of degree 15 resulting in a large value of routers having a logical L3 degree of $15 - 1 = 14$. However, this kind of behavior is highly topology dependent and do not seem to follow any particular rule. Taking this kind of exceptions into account leads to lose the model generality. Here, we propose a high level model trying to capture global properties.

In this section, we investigate whether the $p(x) \times b(x)$ dependence seems sufficient to explain the large tail of the L2-agnostic degree distribution. To study this question, we decide to only model the linear correlation between $p(x)$ and $b(x)$.

5.2 Model

In this section, we only model the relationship between the number of L2 neighbors and the total number of neighbors of a router. Our goal is to highlight the fact that this relationship conditions the L2-agnostic degree distribution. Our model works as follows: considering a router having i neighbors with j being L2 ones ($j \leq i$), its L3 logical degree may become $x = i - j + \eta$ where η denotes the sum of the degrees (minus one) of its L2 neighbors. Because we do not model any relationship between s and the other distributions, we consider that the degree of its L2 neighbors only depends on $s(x)$. Generalizing this idea, the set of routers having an L3 logical degree of x is composed by all routers having a lower degree that may reach this degree thanks to their L2 neighbors. Our model is probabilistic in the sense that we consider the probability that a router having a degree of i connected to a set of j L2 nodes will increase its L3 logical degree by a certain amount, $x - i$, to appear directly connected to x neighbors.

In the following, the notation $bp(i, j)$ refers to the probability to have a router having a degree of i and j L2 neighbors ($0 < j \leq i$). In other words, $bp(i, j)$ corresponds to the probability distribution function given in Fig. 9 without considering cases where $j = 0$. Hence, our probabilistic model verifies:

$$r'(x) \approx \begin{cases} r(x) + \chi(x) - p(x) & \text{if } x > 1 \\ r(x) - p(x) & \text{otherwise.} \end{cases} \quad (4)$$

with $\chi(x)$ referring to the probability that L3 nodes with a lower degree than x reach a degree of x because of their connections to L2 nodes, such that:

$$\chi(x) = \sum_{i=1}^{\min(x-1, p)} \left[p(i) \times \sum_{j=1}^{\min(i, b)} bp(i, j) \times c_j^i(x) \right] \quad (5)$$

where $c_j^i(x)$ denotes the sum of all possible product combinations between a set of j elements $\{s(a_1), \dots, s(a_k), \dots, s(a_j)\}$, with $2 < a_k \leq x - i - j + 3 \forall k \in [1, j]$, and such that:

$$\begin{aligned} \sum_{k=1}^j a_k &= x - i + 2j \\ &= \omega. \end{aligned} \quad (6)$$

We denote δ the set containing a_k 's j -tuple, $\delta = [3, x - i - j + 3]^j$. Thus, it comes:

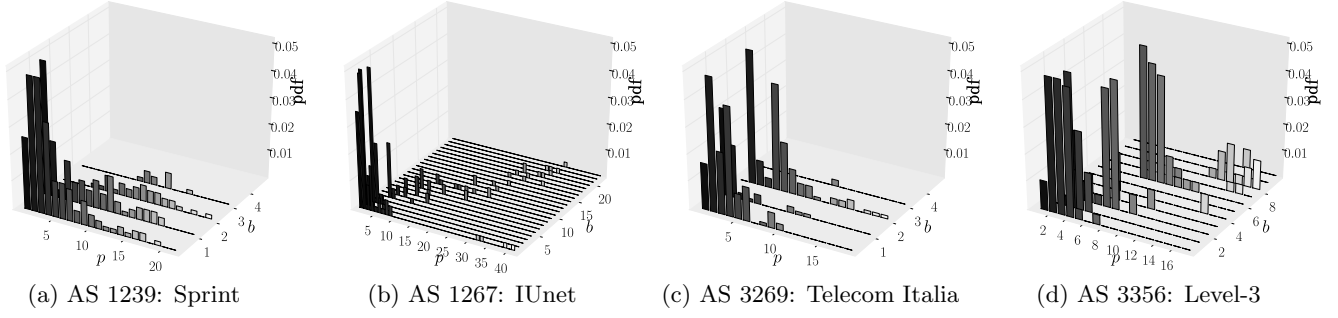


Figure 9: Dependence between p (degree of routers having L2 neighbors⁵) and b (number of L2 neighbors) distributions

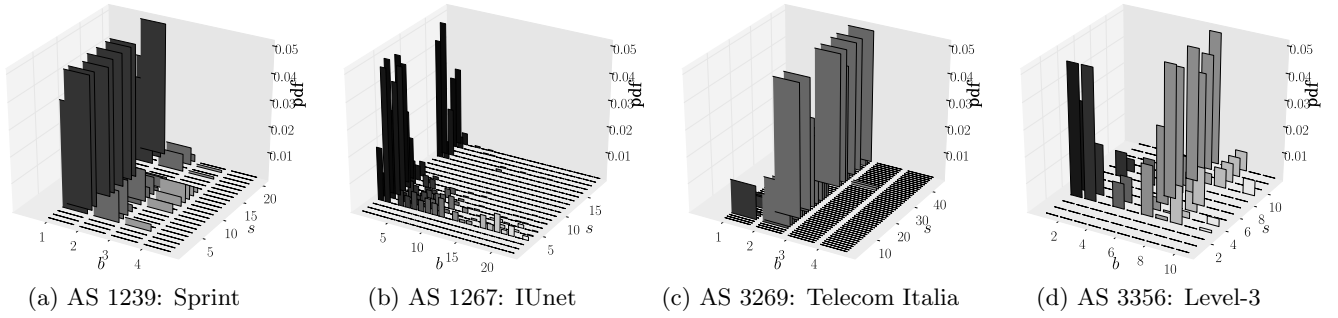


Figure 10: Dependence between b (number of L2 neighbors) and s (degree of L2 nodes) distributions

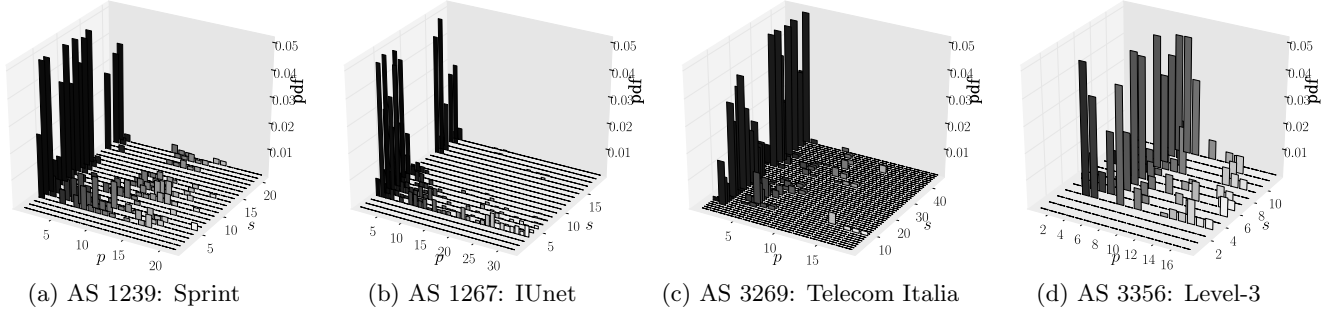


Figure 11: Dependence between p (degree of routers having L2 neighbors) and s (degree of L2 nodes) distributions

$$c_j^i(x) = \sum_{\forall j\text{-tuple}\{a_1, \dots, a_j\} \in \delta \mid \sum_{k=1}^j a_k = \omega} \prod_{k=1}^j s(a_k) \quad (7)$$

Note that in practice we do not have to explore all possible combinations in the space δ to perform the computation of $c_j^i(x)$. Indeed, we do not care about the degree apparition order because $c_j^i(x)$ is a sum of products, so the j terms are permutable in each sub-product. Let us denote Δ the set

of j -tuple $\{a_1, \dots, a_j\} \in \delta$ verifying Eqn. 6. In practice, we then just need to compute:

$$c_j^i(x) = \sum_{\substack{\forall \{a_1, \dots, a_l, \dots, a_j\} \in \Delta \\ \text{such that } a_l \leq a_{l+1} \forall 0 < l \leq j}} \prod_{k=1}^j s(a_k) \times \frac{k!}{\prod_{z=1}^w \sigma_z!} \quad (8)$$

where $\sigma_z, 1 \leq \sigma_z \leq j$, denotes the number of occurrence of each integer in the set Δ such that $\sum_{z=1}^w \sigma_z = j$. Thus, using this computation optimization, it is possible to compute $c_j^i(x)$ in a reasonable time, on the order of a few minutes




Figure 12: Degree distribution for real data and generated by our model using an L3 view ($r'(x)$)

even for $(i, j) \approx (40, 20)$ such as with AS1267. Results are given in Fig. 12: the lines labeled “ $r'(x)$ –model” correspond to the application of our model on our set of four ASes while the lines labeled “ $r'(x)$ –data” result from the experimental analysis. Depending on the AS, our model is more or less accurate to simulate the L3 view: it performs very well on Sprint but it is less accurate on Level-3. Indeed, the probabilistic nature of our model cannot capture sparse discrete values at the tail of the degree distribution. This values may be explained with some specific behavior of the $p \times s$ dependence. However, in all cases, our model captures the significant probability to generate a great shift in the node degree distribution. Considering a high level view of service providers networks using power laws shape to model their properties, we argue that our model is sufficient to provide a first step in the L2 aware topology generation. We claim that a probabilistic generation model taking into account parameters such that the similarity between r and p , the fact that s has a different shape than r and the linear correlation between p and b is able to provide the first step of the generation of L2 aware graphs. The second step consists in analyzing the connectivity nature of relationships between nodes: does it often exist a similar pattern (e.g., the degree) or some sort of correlation between directly connected nodes? This question implies the study of other graph properties such as the clustering coefficient in L2 aware networks.

5.3 Summary

Using the model given in Eqn. 4, we implicitly assume the independence between $s(x)$ and the two other distributions: $p(x)$ and $b(x)$. However, as shown in Fig. 12, our model captures the heavy tailed distribution $r'(x)$ on each AS. Although it does not integrate all stochastic details (Fig. 10 and 11 may help to solve peak values at the tail or differences at the head), it captures the significant probability to generate nodes with high degree. This result emphasizes the importance of the correlation between L2 and L3 degree properties and tends to show that the main critical relationship between key distributions is the one linking $p(x)$ and $b(x)$.

Thus, the answer to the question “Why such a great shift in the node degree distribution?” is: “Because a large proportion of the links of large degree routers are usually connected to L2 nodes”. In practice, on Fig. 9(a), we can interpret this answer like this: the largest router seen in Sprint has 21 neighbors with two being L2 networks. Then, considering

the L3 view, this router may reach, in theory (i.e., following the general idea of Eqn. 4), a degree of $21 - 2 + 2 \times (21 - 1) = 59$ (21 also being the maximum degree of L2 nodes). This phenomenon is exacerbated for topologies such as Level-3 (AS3356) because most large degree routers have many L2 neighbors (≥ 6). Note that, in the worst theoretical case, we might observe routers having a L3 logical degree of roughly 100 in this L3 topology.

Although there does not exist a common model describing all topologies, we have shown that the heavy tail of the node degree distribution using a pure L3 view is due to multiple factors (in particular $s(x)$ and $b(x)$ distributions), and the correlation between $p(x)$ and $b(x)$ helps to understand the shift in the degree distribution generated by an L3 view.

In order to generate L2-aware topologies, it is necessary to carefully consider several distributions, e.g., $r(x)$, $p(x)$, $s(x)$, $b(x)$, and their relationships to produce realistic models. Our model is a first basis for L2-aware graphs generation. According to the set of keys distributions and their similarities, it may allow for generating L2 topologies focusing on particular relationships between $p(x)$, $b(x)$, and $s(x)$ distributions. The next step being to model other properties, such as the assortativity, in order to capture the connectivity nature between nodes.

Finally, using L3 traces such as those obtained by trace-route campaigns and minimal L2 information, one may analytically refine our model to produce L2 aware topologies verifying several L2-L3 connectivity constraints. For example, it is possible to simplify the model and so reduce the number of distributions taken into account by considering some relationships between key them (e.g. the one between $r(x)$ and $p(x)$) and modeling them with specific power laws.

6. RELATED WORK

Modeling Internet topology has been a large subject of research since Faloutsos et al. seminal paper [8] in which they discuss the existence of power law relationships between several properties of the Internet graph.

If this work has been extended by others [35, 36, 37, 38, 39], it has also been strongly questioned. In particular, Lakhina et al. [9] show that if a graph has a node degree distribution that is very different from a power law, sampling from a small set of vantage points will yield a picture of a graph with a node degree distribution that follows a power law. One of the reasons evoked by Lakhina et al. is that the sampled graph has fewer edges than the genuine

graph. Clauset and Moore [10] explain analytically this bias. They demonstrate that, for sparse random graphs of large average degree, the apparent degree distribution displays a power law.

These biases have been further studied by Petermann and De Los Rios [40], and Dall’Asta et al. [41]. Guillaume and Latapy [42] have extended these studies to include the trade-off between the number of monitors and the number of destinations.

Li et al. have shown that the node degree distribution is not sufficient to describe network topology [43, 44]. They point out that there exist many different graphs having the same distribution of node degree. Some of them might be considered opposites from the point of view of network engineering. Instead, Li et al. propose models that incorporate technological constraints on router and link bandwidth and connectivity, together with abstract models of user demand and network performance.

Prior to our study, Pansiot and Grad already mentioned the influence of L2 networks on node degree as seen by traceroute [12]. Further, Clauset and Moore proposed another explanation for the emergence of power law degree distributions: the complex interaction of routers with the link-level topology (identical to our “L2 networks”). However, on the contrary to our work, they do not quantify the influence of L2 clouds. They further do not provide any model. Willinger et al. [13] also mention opaque L2 clouds as a source of potential bias when using traceroute based probing. In this paper, we go further by quantifying, on a large dataset, the impact of broadcast L2 networks, making them less opaque.

7. CONCLUSION

In this paper, we analyzed how layer-2 (L2) networks impact the node degree distribution resulting from a pure layer-3 (L3) vision (such as with data collected by traceroute). Our analysis is based on a large dataset collected using `mr-info`, a tool that silently discovers all multicast enabled IPv4 interfaces of a given router and that allows one to infer L2 networks interconnecting routers.

Our analysis showed that the power law shape usually describing the node degree distribution of IP networks is strongly impacted when considering the presence of L2 devices. In particular, the tail of the degree distribution is much less heavy when L2 networks are taken into account. Our results show that, unfortunately, there does not exist obvious common patterns between domains. However, for the ISP topologies studied in this paper, L2 networks play an important role in the network connectivity: it is very likely that large degree routers are connected to multiple L2 nodes having themselves a large degree. Considering the physical topology, this implies that the actual degree of connectivity (e.g., the number of redundant physical paths) is generally lower than the one discovered by a pure L3 campaign using traceroute.

In addition, we proposed a first probabilistic model capturing the interaction between L2 and L3 topological properties. Among others, we highlighted that the degree of L2 networks is generally higher than the degree of L3 nodes and large degree routers are involved in many L2 connections.

In a near future, we plan to investigate the impact of L2 devices in IP networks according to their access and core backbone parts. More generally, our dataset also questions the power law nature of the node degree distribution. Using

another probing methodology than traceroute, we might observe that the scale free property of the Internet may result from artifacts of a traceroute campaign.

Acknowledgements

This work has been partially supported by the European funded 216372 Trilogly project and the European funded 223936 ECODE project. Mr. Donnet’s work is supported by the FNRS (Fonds National de la Recherche Scientifique, rue d’Egmont 5 – 1000 Bruxelles, Belgium).

8. REFERENCES

- [1] k. claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov, “Internet mapping: from art to science,” in *Proc. IEEE Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH)*, March 2009.
- [2] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, “iPlane: An information plane for distributed services,” in *Proc. USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, November 2006.
- [3] Y. Shavitt and E. Shir, “DIMES: Let the internet measure itself,” *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, October 2005.
- [4] M. Luckie, “Scamper: a scalable and extensible packet probe for active measurement of the Internet,” in *Proc. USENIX/ACM Internet Measurement Conference (IMC)*, November 2010.
- [5] V. Jacobson et al., “traceroute,” UNIX,” man page, 1989, see source code: <ftp://ftp.ee.lbl.gov/traceroute.tar.gz>.
- [6] B. Donnet and T. Friedman, “Internet topology discovery: a survey,” *IEEE Communications Surveys and Tutorials*, vol. 9, no. 4, pp. 2–15, December 2007.
- [7] R. Pastor-Satorras and A. Vespignani, *Evolution and Structure of the Internet: a Statistical Physics Approach*. Cambridge University Press, February 2004.
- [8] M. Faloutsos, P. Faloutsos, and C. Faloutsos, “On power-law relationships of the internet topology,” in *Proc. ACM SIGCOMM*, September 1999.
- [9] A. Lakhina, J. Byers, M. Crovella, and P. Xie, “Sampling biases in IP topology measurements,” in *Proc. IEEE INFOCOM*, April 2003.
- [10] A. Clauset and C. Moore, “Traceroute sampling makes random graphs appear to have power law degree distributions,” arXiv, cond-mat 0312674, February 2004.
- [11] P. Erdős and A. Rényi, “On the evolution of random graphs,” *Publ. Math. Inst. Hung. Acad. Sci.*, vol. 5, pp. 17–61, 1960.
- [12] J.-J. Pansiot and D. Grad, “On routes and multicast trees in the Internet,” *ACM SIGCOMM Computer Communication Review*, vol. 28, no. 1, pp. 41–50, January 1998.
- [13] W. Willinger, D. Alderson, and J. Doyle, “Mathematics and the Internet: a source of enormous confusion and great potential,” *Notices of the*

- American Mathematical Society*, vol. 56, no. 5, pp. 586–599, May 2009.
- [14] V. Jacobson, “`mrinfo`,” 1995, see http://cvsweb.netbsd.org/bsdweb.cgi/src/usr.sbin/mrinfo/?only_with_tag=MAIN.
- [15] P. Mérindol, V. Van den Schriek, B. Donnet, O. Bonaventure, and J.-J. Pansiot, “Quantifying ASes multiconnectivity using multicast information,” in *Proc. ACM USENIX Internet Measurement Conference (IMC)*, November 2009.
- [16] J.-J. Pansiot, P. Mérindol, B. Donnet, and O. Bonaventure, “Extracting intra-domain topology from `mrinfo` probing,” in *Proc. Passive and Active Measurement Conference (PAM)*, April 2010.
- [17] S. Deering, “Host extensions for IP multicasting,” Internet Engineering Task Force, RFC 1112, August 1989.
- [18] T. Pusateri, “Distance vector multicast routing protocol version 3 (DVMRP),” Internet Engineering Task Force, Internet Draft (Work in Progress) draft-ietf-idmr-dvmrp-v3-11, October 2003.
- [19] J.-J. Pansiot, “`mrinfo` dataset,” see <http://svnet.u-strasbg.fr/mrinfo/>.
- [20] IANA, “Special-use IPv4 addresses,” Internet Engineering Task Force, RFC 3330, September 2002.
- [21] W. Fenner, “Internet group management protocol (IGMP), version 2,” Internet Engineering Task Force, RFC 2236, November 1997.
- [22] E. Rosen, A. Viswanathan, and R. Callon, “Multiprotocol label switching architecture,” Internet Engineering Task Force, RFC 3031, January 2001.
- [23] L. Andersson, P. Doolan, N. Feldman, A. Fredette, and B. Thomas, “LDP specification,” Internet Engineering Task Force, RFC 3036, January 2001.
- [24] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow, “RSVP-TE: Extensions to RSVP for LSP tunnels,” Internet Engineering Task Force, RFC 3209, December 2001.
- [25] Y. Rekhter and E. Rosen, “Carrying label information in BGP-4,” Internet Engineering Task Force, RFC 3107, May 2001.
- [26] K. Kompella and Y. Rekhter, “Virtual private LAN service (VPLS) using BGP for auto-discovery and signaling,” Internet Engineering Task Force, RFC 4761, January 2007.
- [27] M. Lasserre and V. Kompella, “Virtual private LAN service (VPLS) using label distribution protocol (LDP) signaling,” Internet Engineering Task Force, RFC 4762, January 2007.
- [28] J.-J. Pansiot, P. Mérindol, B. Donnet, and O. Bonaventure, “Internet topology discovery through `mrinfo` probing,” Université catholique de Louvain (UCL), Tech. Rep., October 2009, see <http://inl.info.ucl.ac.be/content/mrinfo>.
- [29] N. Spring, R. Mahajan, and D. Wetherall, “Measuring ISP topologies with Rocketfuel,” in *Proc. ACM SIGCOMM*, August 2002.
- [30] Y. Ohara, S. Imahori, and R. Van Meter, “MARA: Maximum alternative routing algorithm,” in *Proc. IEEE INFOCOM*, April 2009.
- [31] K.-W. Kwong, L. Gao, R. Guérin, and Z.-L. Zhang, “On the feasibility and efficacy of protection routing in IP networks,” in *Proc. IEEE INFOCOM*, March 2010.
- [32] B. Rajagopalan, J. Luciani, and D. Awduche, “IP over optical networks: a framework,” Internet Engineering Task Force, RFC 3717, March 2004.
- [33] Network Protection, “Shared risk link group (SRLG),” February 2010, <http://www.network-protection.net/shared-risk-link-group-srlg/>.
- [34] B. Huffaker, A. Dhamdhere, M. Fomenkov, and k. cuffy, “Toward topology dualism: Improving the accuracy of AS annotations for routers,” in *Proc. Passive and Active Measurement Conference (PAM)*, April 2010.
- [35] D. Magoni and J. J. Pansiot, “Analysis of the autonomous system network topology,” *ACM SIGCOMM Computer Communication Review*, vol. 31, no. 3, pp. 26 – 37, July 2001.
- [36] A. Broido and k. claffy, “Internet topology: Connectivity of IP graphs,” in *Proc. SPIE International Symposium on Convergence of IT and Communication (ITCom)*, August 2001.
- [37] A. L. Barabási and R. Albert, “Emergence of scaling in random networks,” *Science*, vol. 286, pp. 509–512, October 1999.
- [38] A. Medina, I. Matta, and J. Byers, “On the origin of power laws in internet topologies,” *ACM SIGCOMM Computer Communications Review*, vol. 30, no. 2, pp. 18–28, April 2000.
- [39] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, “The origin of power laws in internet topologies revisited,” in *Proc. IEEE INFOCOM*, June 2002.
- [40] T. Petermann and P. De Los Rios, “Exploration of scale-free networks,” *The European Physical Journal B*, vol. 38, p. 201, 2004.
- [41] L. Dall’Asta, I. Alvarez-Hamelin, A. Barrat, A. Vázquez, and A. Vespignani, “A statistical approach to the traceroute-like exploration of networks: Theory and simulations,” in *Proc. Combinatorial and Algorithmic Aspects of Networking (CAAN) Workshop*, August 2004.
- [42] J. L. Guillaume and M. Latapy, “Relevance of massively distributed explorations of the internet topology: Simulation results,” in *Proc. IEEE INFOCOM*, March 2005.
- [43] L. Li, D. Alderson, W. Willinger, and J. Doyle, “A first-principles approach to understanding the Internet’s router-level topology,” in *Proc. ACM SIGCOMM*, August 2004.
- [44] D. Alderson, L. Li, W. Willinger, and J. Doyle, “Understanding Internet topology: Principles, models, and validation,” *IEEE/ACM Transactions on Networking*, vol. 13, no. 6, pp. 1205–1218, December 2005.