Consolidated Review of
Network Fingerprinting: TTL-Based Router Signatures

1.Strengths

This paper makes two contributions. The first is to point out that some routers use different initial TTL values for different types of packets, specifically ICMP packets caused by different events. It is pointed out that this can in principle give an extra dimension when trying to identify they type of router based on initial TTL measurements. The second contribution is to use this fingerprinting to extend the authors' prior work on identifying and classifying MPLS tunnels.

Fairly lightweight and can piggyback on traceroute. I believe that fingerprinting routers is useful, because routers of different classes have different behaviors, and so it might make sense to treat them in different ways when, say, probing. The authors present a simple active probing-based technique that appears to work reasonably well as a pretty coarse signature of devices.

The main contribution appears to be pointing out that the authors' previous estimate of the number of "invisible" MPLS tunnels was inaccurate. The signature mechanism appears to give a little bit of insight into the prevalence and nature of MPLS tunnels.

The paper was straightforward and easy to follow

2.Weaknesses

The paper was not very ambitious: it only used a 2-tuple signature and only briefly investigated a single use. Fine for a short paper, but could have done more.

The paper only checks for one type of consistency, whether multiple vantage points return a consistent fingerprint for a common target. It should also check that aliases return consistent fingerprints.

The arguments for why the TTL signature should give some insight into MPLS characteristics are not very strongly supported (but this is a 6 page paper). Tie-in with MPLS deployments is not clear. The second contribution, Section 3, is poorly written. It relies too heavily on knowledge of the authors' previous work, [10].

3.Comments

Overall, this is a nicely written paper, with an interesting use for I think the paper is timely. I enjoyed reading it. The paper makes an important contribution and could lead to many interesting follow-ons. I've talked to a number of researchers recently about cases when these types of signature could come in handy. Earlier papers like DisCarte and Justine Sherry's timestamp work had similar signatures (and should probably be cited in that regard), but this method seems lighter weight. I think the paper is useful, explores most of the basic important issues, and somewhat slight. It seems like a good fit for a short paper.

The main piece I believe you need to add is a section on Signatures Consistency that looks at consistency across the various aliases of a router. This is very straightforward to perform and will enrich the paper and strengthen the claim that the signatures often correspond to router vendors. Less important, but I would also like to see how consistent prefixes / ASes are. I would think that an individual network would generally deploy a single type of router in a given role. Is there reason to expect more than 2 different initial values if you increase n? I would think a router might have one for responses generated in software and one for responses generated in hardware. You posit that longer signatures can "provide better distributions among router's OS." Do you have any evidence that that will work? For which potential applications do you need a finer-grained understanding of what the device is? Do you have thoughts on how to fingerprint at a finer granularity? It would also be great to list more applications for which your current approach suffices.

Overall, I think this work shows some promise for giving an indirect way to help characterize MPLS tunnel deployments. I'm not sure how useful TTL signatures may be more generally, and I wish the focus of the paper had been more directly on the MPLS application as it would have enabled the authors to go a little more deeply into the various issues and testbed experiments they performed. In section 2.1, it wasn't clear from the description how you know that the IP addresses you choose for probing are routers. Is the Ark target list guaranteed to only contain routers? Section 3.2 is pretty dense --- there is a lot of detail and a number of points are not particularly clear. For example, the end of paragraph 1 states "... it seems that 255,64 routers are more attractive for MPLS operations..." I don't understand what's being said. The experiment with the testbed routers should really be explained better in that section, too, since the point is quite important to the inferences you're making. The statement in section 2 that "it is worth to notice that #hops < 30" needs a citation.

My comments are mainly intended to help improve the paper's final version or the authors' future work on this topic:

- I felt that the paper could have looked at the implications of including more entries in the tuples. What would be the additional fields stand for? Would additional fields provide more discriminating information? Or would the general classification remain the same. Some discussion of this coupled with a few results would have made this paper more interesting.
- Detecting the nature of MPLS deployment is an important use-case for this technique, yet I did not find the MPLS section very clear or convincing. It also seemed incremental wrt the authors' prior work.

The authors should consider these issues in preparing the final version of their paper: exactly how are the TTL signatures helping? Can the authors provide more authoritative measurement results? In general, the writing could be improved, as several sub-sections were difficult to follow.

Section 3 is quite unclear, especially the discussion of how the abundance of <255,255> fingerprints for opaque tunnels invalidates the previous conclusions about invisible tunnels. As another example of lack of clarity, it is not clear how the three attributes listed at the top of page 5 combine to allow tunnels to be classified. What is qTTL? More background from both [10] and RFC4950 should be included. For example, does 4950 specify that the router must put the label stack in the time-exceeded

message, or that it may? The acronym LER should be spelled out, especially since it looks like a typo for LSR.

4.Summary from PC Discussion

In the PC meeting discussion, the reviewers noted that they liked approach of using different initial TTL values as a router signature, though there was discussion regarding the fact that the contribution was limited and incremental over prior work. There was also some discussion about the difficulty in reading Section 3, and the fact that this aspect of the work was not especially convincing at this stage. Still, the reviewers felt that the router signature contributions were interesting and may be leveraged by other researchers, thus they recommended accepting the paper.

5.Authors' Response

We are grateful to our anonymous reviewers for their relevant feedback. We took most of their comments into account for the camera-ready version of our paper.

In particular, we improved the use-case section about MPLS to make it clearer and stand-alone (without necessarily reading ref [10]). Besides, we show how our fingerprinting method can be used beyond the MPLS scope: since it allows to determine router brand and OS distribution, it can be useful to determine if measurements are biased due to router type dependency. For example if measurements are done on some sample networks (e.g., ASes) that are not representative of the Internet router mix (e.g., sample containing only Cisco routers), then these measurement results cannot be extended to the Internet as a whole if the measured values are router-type dependent. So our method, or its extension, could be used both to determine if a sample is representative of the router mix, and if a given feature (e.g., BGP border router) is independent of the router type or not. Alias resolution is again a good illustration of our method usefulness: if some routers do not react as usual to some alias resolution technique, it is not necessary to apply this technique to them since results will be biased and probably wrong.

Regarding router scale consistency of the TTL signature, we plan as a future work to conduct a robust and large scale alias resolution campaign to verify this assumption. As a preliminary step, we want to extend our work considering other types of ICMP replies and other IP header fields that may be discriminant to achieve a better OS/brand partition. Finally, considering this refined distribution, our goal is to understand whether IP networks are heterogeneous in terms of hardware and software. Analyzing the OS deployment evolution in large IP networks can be a nice study to check common accepted statement such as "old devices are pushed at the edges while new ones are deployed in the core".