Monitoring network topology dynamism of large-scale traceroute-based measurements

Thomas Bourgeau
LIP6 Computer Science Laboratory, UPMC Sorbonne Universités and CNRS, Paris, France.
<thomas.bourgeau@upmc.fr>

***Abstract*—Network topology discovery with distributed traceroute-based measurement systems is important to monitor, measure, diagnose and capture IP-level network topology dynamism. Depending on the discovered topology size and the captured topology dynamism accuracy, a compromise has to be done regarding the measurement time granularity and the scale of these measurement systems. In this paper, we present our large-scale measurement dataset, and analysis of the network topology dynamism captured in a real measurement scenario. We also quantify the missed dynamism information with coarser measurement time granularity inferred by our proposed algorithm. These results confirm that probing less frequently, as it is the case of most of the existing measurement systems today, can dramatically affect the dynamism information captured.**

***Index Terms*—Active measurements, Network topology dynamism, Monitoring**

I. Introduction

Large view of the Internet topology is essential to analyze its key properties [1], [2], understand its structure [3], [4] observe its evolution [5], [6] and model its dynamism [7], [8]. Different approaches building Internet topology are proposed in the literature [9], [10], [11], [12], in our work, we focus on capturing IP-level network topology dynamism [1], [13], [14] as discovered by large-scale distributed traceroute-based measurement systems [15], [16], [17], [18], [19] offering end to end observation and analysis of the underlying topology. Most of these measurement systems expose different characteristics in terms of scalability (number of measurement sources used to discover the path toward a set of destinations) and measurement time granularity (the period τ needed by these systems to accomplish a complete discovery cycle). Thus depending on the size of the monitored network topology and the accuracy of the topology dynamism captured, a compromise has to be done regarding the measurement time granularity and the number of deployed measurement sources of these measurements systems. Furthermore, scaling up these measurement systems with broad distribution of carefully chosen measurement agents and destination points is mandatory to yield good estimates of network graph properties [20], [11], [21]. However, while the scaled up measurement systems offer to discover additional topological information during a discovery cycle, they might provide a poor representation of the topology dynamism that occurs as the duration of the discovery cycles could be longer than the dynamism duration to be captured. For example, a state of the art measurement system, ARCHIPELAGO [15], needs about two days to cover millions of destinations (all the Internet network prefixes) from 45 vantage points with traceroute like tool. We argue that two days measurement will not reach a high accuracy in catching the whole network topology dynamism. In addition, traceroute-based measurement systems might suffer from measurement inaccuracy when probing toward load balancers [22] leading to false topology dynamism interpretation if new measurement methods [23], [24] are not applied.

In this paper, we propose to analyze the network topology dynamism information captured in a real measurement scenario and quantify the impact of coarser time granularity on the dynamism information missed.

The remainder of this paper is organized as follows. First, we define our processed measurement dataset and explain our algorithm to infer less frequent measurements in the network topology inference section. Then, we present our metrics and analyze the impact of measurement time granularity on dynamism information missed. Finally, we discuss our observations and provide perspectives for our work on how to better capture Internet topology dynamism.

II. Network topology graph inference

In this section we explain our procedure and experimental setup to build network topology graphs. Then we introduce our notations and inference algorithm related to our approach.

A. Fine grained measurement dataset

Capturing Internet topology dynamism at the IP-level with large-scale distributed traceroute-based measurement systems requires accurate and fast discovery of the network. Most of the existing systems are not accurate enough since they might discover false dynamism information and measurement artifacts (false links, false loops, etc.) inferred by regular traceroute tool when probing toward load balancers [22]. In addition, it is difficult to capture most of the dynamism information as the measurement time granularity of these system is too long. Consequently, we built our own fine grained measurement dataset respecting the accuracy and time granularity requirements as mentioned previously. More precisely, we avoid false dynamism information and correct the inaccuracy of the legacy traceroute tool in our measurement dataset, by processing high frequency measurement using Paris traceroute tool [24] during 2 months over 580 nodes of PlanetLab overlay [25] probing 800 PlanetLab destinations with a time granularity of $\tau = 1$ hour which produced our fine grained measurement dataset that we use to build incremental graph topology of the network over the time. Note that we have only kept network discovery results from 230 measurement agents that have been continuously probing during the entire experiment period corresponding to two months continuous discovery cycles or rounds [1].

B. Network topology graph representation.

Having explained how we generate our measurement dataset, here we define how to aggregate these measurements to create directed graphs of the discovered IP-level topology $G_k^\tau = (V_k^\tau, E_k^\tau)$ composed of graph elements that are either node sets V_k^τ (distinct nodes formed by the IP addresses) and link sets E_k^τ (distinct successive IP address pairs) as discovered by the measurement system for each measurement round k with a time granularity of τ. Note that $G^\tau = \{\bigcup_{j=1}^{K} G_j^\tau\}$ represent all unique graph elements that appeared at least once during the entire experiment time (K= 1488 rounds). Moreover, we introduce a state value $\delta_k^\tau(n)$ representing the status of each graph elements $n \in G^\tau$ seen at a given measurement round k to be either a graph element presence when $\delta_k^\tau(n) = 1$ or a graph element absence when $\delta_k^\tau(n) = 0$.

$$\forall n \in G^\tau,\, \delta_k^\tau(n) = \begin{cases} 1, & \text{if } n \in G_k^\tau \\ 0, & \text{if } n \notin G_k^\tau \end{cases}$$

Finally, we propose to summarize the entire state values by a state matrix $M_\tau(G^\tau)$ where each element $m_\tau(n,k) = \delta_k^\tau(n)$ represent the status of graph elements $n \in G^\tau$ for each measurement rounds of the experiment.

C. Inference algorithm

Having explained how we generate the state matrix from our fine grained measurement with a time granularity of τ (1 hour) for each measurement round k, we propose to infer what would have been observed with longer discovery cycles. It is intuitive that having very frequent discovery cycles might inform us about longer discovery periods, but not at the inverse. So, for all graph elements $\forall n \in G^\tau$, to infer the potential state matrix that would have been measured with agents probing $\tau_y = y \times \tau$ less frequently (where $y > 1$), we propose to generate synthetic discovery rounds from our fine grained dynamism representation of the state matrix $M_\tau(n)$ to infer the virtual state matrix $M_{\tau_y}(n)$ for agents probing $\tau_y = y \times \tau$ times slower. We build this virtual state matrix as follow; for any virtual round $k_y = \|\frac{k}{y}\|$ we select randomly a state value $m_\tau(n,k)$ between rounds of the fine grained measurement such as $\exists\, m_{\tau_y}(n,k_y) \,|\, m_{\tau_y}(n,k_y) \in \{m_\tau(n,k), \ldots, m_\tau(n,k+y)\}$. The construction of the inferred state matrix and virtual rounds is depict in Fig. 1. For simplicity, we use the measurement frequency notation in the rest of this document. Thus, for an inferred probing measurement round that is $y \times \tau$ slower than our fine grained measurement, we say that it's frequency is F_y.

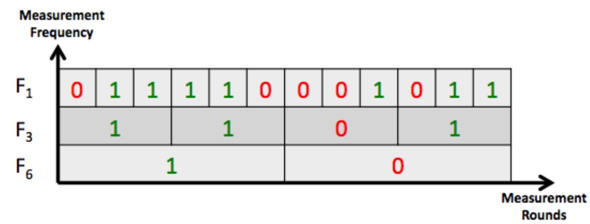

Fig. 1. Building inferred virtual rounds

III. Dynamism analysis

In this section we present our formalism and analysis of network topology dynamism information captured by introducing the dynamism features and metrics.

A. Dynamism features

We feature out the network topology dynamism captured during a given observation window with dynamism events identified by the appearance or the disappearance of graph elements between consecutive rounds and with static states defined as the observation of the duration of continuous presence or continuous absence of these graph elements over a number of rounds.

1) Dynamism event: To analyze the dynamism events that occur between successive rounds, we compare the state value of graph elements $n \in G^\tau$ at measurement round k to its value at a measurement round $k+1$, indicating which graph element has appeared or disappeared. We observe a dynamism event when $m_\tau(n,k) \neq m_\tau(n,k+1)$ and we define $\psi_\tau^n(k,k+1)$ to be the dynamism events information gathered for a given graph elements $n \in G^\tau$ between two consecutive measurement rounds k and $k+1$ in the experiment.

$$\psi_\tau^n(k,k+1) = \begin{cases} 1, & \text{if } m_\tau(n,k+1) - m_\tau(n,k) > 0 \\ 0, & \text{if } m_\tau(n,k+1) - m_\tau(n,k) = 0 \\ -1, & \text{if } m_\tau(n,k+1) - m_\tau(n,k) < 0 \end{cases}$$

If $\psi_\tau^n(k,k+1) = 1$ it is a graph element appearance otherwise, if $\psi_\tau^n(k,k+1) = -1$ it is a disappearance. Note that in case of $\psi_\tau^n(k,k+1) = 0$ there is no dynamism event but a consecutive presence or absence; static state formulated hereafter.

[1]We have carried our measurement study between the 25th of May to the 25th of July 2010 starting with 580 agents. Due to some instability on PlanetLab nodes only 230 agents were running continuously during the experiment period.

2) Static state: We define a static state to be either a continuous presence following an appearance event and ending with a disappearance event or a continuous absence that follows a disappearance event and ends with an appearance event. The dynamism state duration of size N given bellow characterizes the length of a continuous presence or a continuous absence for a graph element $n \in G^\tau$. Note that this static state that lasts N rounds, starts at round $k+1$, is preceded by a dynamism event at round k, and is succeeded by a dynamism event at $k+N+1$. Thus $\alpha^n_\tau(k, N) = 1$ denote a consecutive presence of length N otherwise if $\alpha^n_\tau(k, N) = -1$ it is a consecutive absence of length N

$$\alpha^n_\tau(k, N) = \frac{1}{2}[\psi^n_\tau(k, k+1) - \sum_{j=k+1}^{k+N} \psi^n_\tau(j, j+1)] = \{1, -1\}$$

B. Dynamism metrics and observation

In the following, we introduce metrics to analyze the dynamism features describe earlier. We focus particularly on the observation of the number of dynamism events and on the duration of the static state either from our fine grained measurement or from the inferred measurement. Furthermore, we quantify the missed network topology dynamism when inferring slower measurement rounds.

*1) **M1 - Number of dynamism events**:* This first metric aims at counting the number of dynamism events that occurs at each rounds of the entire measurement experiment. We propose to count the number of dynamism events (appearance and disappearance) at each measurement round. This information is also interesting to evaluate the proportion of dynamism events compared to the entire graph elements measured and its evolution in time.

As explained earlier, an appearance or a disappearance of a graph element can be captured only between two consecutive rounds. If we consider to study the entire dynamism events to be the sum of appearance and disappearance of the entire graph elements between successive rounds, we define $\Phi_\tau(k, k+1) = \sum_{j=1}^{n} |\psi^j_\tau(k, k+1)|$ to be the full dynamism events number between two consecutive rounds.

We present in Fig. 2(a), 2(d) the results of our metric for the entire number of dynamism events observed for our fine grained and inferred measurements. For each rounds of the fine grained measurement, we have obtained in average a graph of 14,322 ips and 40,850 links. Furthermore, the average number of total dynamism events observed for the graph elements at each round is equal to 200 ips and 1500 links. This results reflects that the dynamism events represent a small proportion of the entire graph as it involve only an average of 1.4% of all ips and 3.6% of all links.

If we consider the results of M1 applied to all inferred virtual rounds, we observe that probing less frequently increases the number of dynamism events captured. For example, we found that the average number of dynamism events increased to 5.5 % for ips and to 12.2 % for links when inferring measurements every 48 hours. As expected, these results confirm that probing less frequently accumulate a portion of the dynamism events that occur during the measurement period which might falsify the dynamism results captured. Nevertheless, we might miss dynamism events with coarser measurement period if they occur an odd-numbered of time; for example if an ip disappears and reappears during the measurement period we won't consider it as a dynamism event. Moreover, if a dynamism event occurs many times during a coarser measurement period, we might underestimate the entire number of dynamism events captured. Finally, the observed results show a periodical behavior for dynamism events that remains with the same periodicity with an average period of 160 rounds (7days) for all fine grained and inferred network topology.

*2) **M2 - Occurrence of dynamism state**:* This second metric is about measuring the length duration of consecutive presence for all graph elements. Thus, we define $d_p(N, \tau) = \sum_{j=1}^{n}\sum_{k=1}^{K}(\alpha^n_\tau(k, k+N) = 1)$ to be the number of presence duration of length N observed for all graph elements $\forall n \in G^\tau$ at all rounds $k \in [1, 1488]$. This should tells us if there is really much detail to observe at finer time granularity and perhaps indicate where the richest dynamism information are. Based on this metric, we propose to analyze in Fig. 2(b), 2(e) the proportion of continuous presence observed. For this purpose, we calculate the CDF function $P(d_p(k, \tau)) \leq k$ of the continuous presence probability for all graph elements during the entire experiment both for fine grained and inferred virtual measurements. We found that in our fine grained measurement typically 10% of the ips and 40% of the links are present less than half a day (12 rounds). Then, as we probe less frequently, we observe that the length of presence for ips and links get longer; as for F_{48}, 10% of the ips have a presence length less than 20 days (500 rounds) and 40% of the links have a presence length less than 37 days (900 rounds). Thus, this confirms that it is important to probe at a finer time granularity as the information on presence length can be dramatically distorted. Moreover, probing every two days as in ARCHIPELAGO might miss short presence ips and links thus been less exhaustive for reporting dynamism features.

*3) **M3 - Probing frequency effects on dynamism feature**:* The last metric that we propose to analyze is dedicated to quantify the dynamism information missed by probing slower than our fine grained measurement. For that we compare the number of dynamism events captured at a specific virtual measurement round to the sum of dynamism events captured on consecutive fine grained measurement rounds which correspond to the virtual measurement round. The mathematical formulation of this

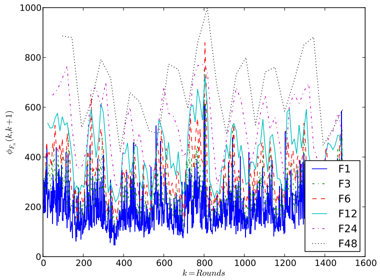

(a) Total dynamism for all ips

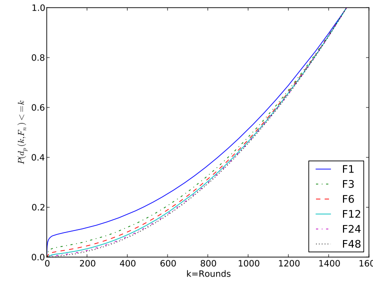

(b) CDF presence length for ips

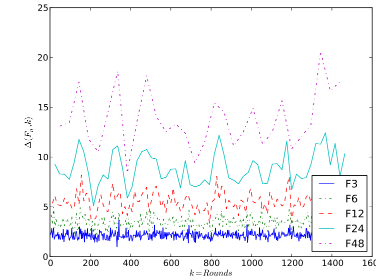

(c) Proportion of missed dynamism events for ips

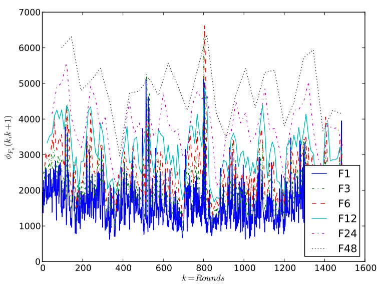

(d) Total dynamism for all Links

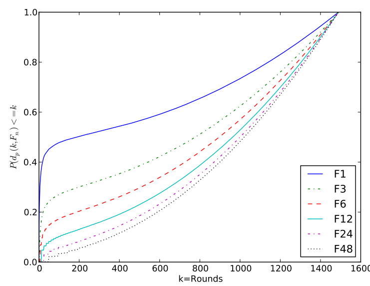

(e) CDF presence length for links

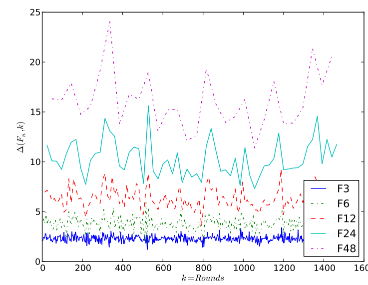

(f) Proportion of missed dynamism events for links

Fig. 2. Results of the total topology dynamism events, presence static states and proportion of missed dynamism events observed for all the graph elements at 6 different virtual frequencies ($F_1, F_3, F_6, F_{12}, F_{24}, F_{48}$)

metric is describe bellow.

$$\Delta(\tau_y = y\tau, k) = \frac{\sum_{j=k}^{k+y} \Phi_\tau(j, j+1)}{\Phi_{\tau_y}(k, k+1)}$$

We observe in Fig. 2(c), 2(f) that half of dynamism events observed at F_1 for links and ips are missed when probing 3 times slower. Furthermore, as we probe less frequently, we miss more and more graph elements events; for instance probing every 2 days (as does ARCHIPELAGO) may reveal 15 times fewer dynamism events compared to probing every hour. Thus, these results point out the need for probing at a finer time granularity to avoid missing dynamism events and to better capture Internet topology dynamism.

IV. Conclusion and perspectives

In this paper, we present our analysis and results of the network topology dynamism features and metrics. We provided the important requirements to design an optimized large scale measurement system, that is capturing the most of Internet topology dynamism either short or long term while lowering measurement load. We have seen from our fine grained measurements, that we run at a very frequent time basis over a large-scale measurement system in Planetlab network, that dynamism events are a marginal portion of the entire topology graph elements as it represent 1.4% of all ips and 3.6 % of all links. However, this information is crucial for understanding network topology behavior (40% of the links and 10% of the ips have a presence length less than half a day) and for helping network operators to better identify network topology dynamism. Furthermore, lowering the probing frequency can dramatically affect the observed dynamism as probing at ARCHIPELAGO's frequency is missing 15 times the dynamism events that really occurs. These results provide incentives to improve measurement systems with more frequent probing to capture the most topology dynamism details. Based on that, we follow new directions to increase probing frequencies without increasing probing load mainly by exploiting measurement redundancy and dynamism aware probing as it is specified in our future work where we design new efficient algorithms to better capture Internet topology dynamism with very frequent measurement frequency, low measurement load and better accuracy.

Acknowledgments

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n°224263-OneLab2. It was conducted in collaboration with Jordan Augé, under the supervision of Timur Friedman.

References

[1] M. Latapy and C. Magnien, "Complex network measurements: Estimating the relevance of observed properties," in *Proc. of IEEE INFOCOM*, 2008.

[2] R. Pastor-Satorras, A. Vázquez, and A. Vespignani, "Dynamical and correlation properties of the Internet," *Physical Review Letters*, vol. 87, no. 25, p. 258701, 2001.
[3] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *ACM SIGCOMM*, 1999, pp. 251–262.
[4] M. E. J. Newman, "The structure and function of complex networks," *SIAM REVIEW*, vol. 45, pp. 167–256, 2003.
[5] G. Siganos, M. Faloutsos, and C. Faloutsos, "The Evolution of the Internet: Topology and Routing."
[6] S. Dorogovtsev and J. Mendes, "Evolution of networks," *Advances in Physics*, vol. 51, no. 4, pp. 1079–1187, 2002.
[7] D. J. Watts and S. H. Strogatz, "Collective dynamics of small-world networks," *Nature*, vol. 393, no. 6684, pp. 440–442, 1998.
[8] A.-L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.
[9] B. Donnet and T. Friedman, "Internet topology discovery: a survey," in *IEEE Communications Survey and Tutorials*, 2007.
[10] "University of oregon route views project," http://www.routeviews.org.
[11] N. Spring, R. Mahajan, and D. Wetherall, "Measuring isp topologies with rocketfuel," in *Proc. ACM SIGCOMM*, 2002, pp. 133–145.
[12] K. Keys, ""iffinder", a tool for mapping interfaces to routers," *See http://www. caida. org/tools/measurement/iffinder*.
[13] A. Hamzaoui, M. Latapy, and C. Magnien, "Detecting events in the dynamics of ego-centered measurements of the internet topology," in *Proc. IEEE WiOpt.* IEEE, 2010, pp. 505–512.
[14] Í. Cunha, R. Teixeira, D. Veitch, and C. Diot, "Predicting and Tracking Internet Path Changes," in *Proc. ACM SIGCOMM*, Toronto, Canada, Aug. 2011.
[15] K. Claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov, "Internet mapping: from art to science," in *Proc. of CATCH 2009*, March.
[16] Y. Shavitt and E. Shir, "Dimes: let the internet measure itself," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 5, pp. 71–74, 2005.
[17] 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. of OSDI 2006.*
[18] K. Chen, D. Choffnes, R. Potharaju, Y. Chen, F. Bustamante, D. Pei, and Y. Zhao, "Where the sidewalk ends: Extending the Internet AS graph using traceroutes from P2P users," in *Proc. of ACM CoNEXT 2009.*
[19] M. Alves, L. Corsello, D. Karrenberg, C. Ogut, M. Santcroos, R. Sojka, H. Uijterwaal, and R. Wilhelm, "Providing active measurement as a regular service for isp's," *Proc. PAM 2002.*
[20] Y. Shavitt and U. Weinsberg, "Quantifying the importance of vantage points distribution in internet topology measurements," in *Proc. IEEE INFOCOM*, 2009.
[21] A. Lakhina, J. W. Byers, M. Crovella, and P. Xie, "Sampling biases in ip topology measurements," in *Proc. of IEEE INFOCOM 2003.*
[22] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding traceroute anomalies with paris traceroute," in *Proc. of IMC 2006.*
[23] B. Augustin, T. Friedman, and R. Teixeira, "Measuring load-balanced paths in the internet," in *Proc. of the 7th ACM SIGCOMM conference on Internet measurement*, ser. IMC '07. New York, NY, USA: ACM, 2007, pp. 149–160.
[24] "Paris traceroute website," http://www.paris-traceroute.net.
[25] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, "PlanetLab: an overlay testbed for broad-coverage services," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 3, pp. 3–12, 2003.