Application-aware Oscillation Detection in Internet Paths

Vinay Kolar (vinkolar@cisco.com) – Principal Engineer, JP Vasseur (jpv@cisco.com), PhD - Cisco Fellow Engineering Lead - February 2022

Abstract: ***Application-aware path-selection protocols are being employed widely in enterprise networks. Such path-selection protocols offer steering of application traffic based on how a path supports the application experience. An important aspect of designing such protocols is to avoid path oscillations. Oscillations occur when two paths may alternate between providing good and bad application experiences for an application. The paper first studies the different type of oscillations using network telemetry collected from large and diverse set of networks and using Service Level Agreement (SLA) templates for applications. Oscillations along three main networking KPI metrics (loss, latency, and jitter) is presented. We show that loss and jitter are more prone to oscillations. The paper identifies two categories of oscillations: Near Boundary Oscillations (NBO) and Wild Oscillations. It then proposes data-driven approach for detecting and avoiding NBOs. The paper also proposes visualization approach called Path-Cubes to recognize undesirable path oscillations. Such techniques can used by predictive routing engines to avoid oscillations. It can be used by network administrators to recognize paths are oscillating for which applications and due to which networking KPI metric.***

1. Introduction

Path-selection protocols are becoming more application aware. Several path-selection protocols, such as overlay routing protocols in SDWAN, offer facilities such as Application Aware Routing (AAR) [1], where the admin can specify Service Level Agreement (SLA) templates. Administrators can specify SLA for path Quality of Service (QoS) metrics, such as loss, latency, and jitter, that are good for an application. For example, SLA for voice applications can be specified to be loss <= 3%, two-way latency <= 300ms, and jitter <= 50 ms. If the path is deemed to not provide SLA to the application, the path selection process will switch the path to one that respects SLA, if one exists. Such a configuration enables administrators to configure different SLA thresholds for different applications. This greatly reduces the number of path changes that need to be applied by the path-selection protocol as an action is taken only when the paths fail to meet SLA thresholds.

This paper illustrates the problem of path-oscillations and path-choice under application aware path selection. Such oscillations are different than the oscillations traditionally studied in underlay path selection protocols such as the ones used in inter-domain routing [2]. The objective for application-aware path selection is to detect paths where application SLA is oscillating between SLA requirements being met (or not) and for the path selection process to select a more appropriate path.

Many paths are either violating or not-violating SLA for an application – for most times. The forwarding engine can differentiate such paths easily. A key requirement though is detect potential path oscillations due to SLA oscillating between boundaries. If the path selection protocol does not prevent such oscillations in an intelligent manner, there will be consistent issues to application experience of the users since such oscillations are not non-disruptive (triggers jitters, potential packet reordering, …).

This paper studies path oscillation using a data-oriented approach. First, the paper proposes mechanisms to detect path oscillation along different QoS metrics for a given application profile. The study categorizes the path oscillations into two categories: Wild Oscillations and Near Boundary Oscillations (NBO). NBOs are paths which isolate along the SLA boundaries; path selection can easily isolate such paths and takes a decision to convert them into two classes: desirable and non-desirable. However, NBO paths should be avoided. Second, the paper presents ways to visualize undesirable path oscillations using a three-dimensional cube and density plots to show how a path oscillates.

Cisco's Predictive Internet solution [3] makes use of a large and diverse path data available to detect undesirable oscillations, and to ensure providing better application experience to its users. This paper shows the diversity of oscillation across multiple customer networks. Most paths oscillate along loss or jitter for voice applications. The paper then shows examples of path behaviors for different type of oscillations, and propose techniques to identify and remediate such path-oscillations.

2. Detecting application quality oscillations

In application aware path selection, the application experience of a path is often mapped to the different telemetry measured for the path. It is a common practice to map application experience to path QoS metrics such as loss, latency and jitter that are measured through probes such as Bidirectional Forwarding Detection (BFD). These are referred to as *SLA thresholds* for an application. For example, the table below show SLA thresholds may be configured for different applications.

Table 1: SLA thresholds for different applications

Application	**Packet loss threshold**	**Latency threshold (two-way)**	**Jitter threshold**
Voice apps	3%	300 ms	50 ms
Sharepoint app	3%	300 ms	300 ms
Go-to-meeting	1%	300 ms	50 ms

It is not necessary to always map path QoS metrics as a proxy for application experience; experience can be mapped to other telemetry, if available. For example, it can be mapped to the feedback directly computed by the application, such as concealment time measured at the codec's of voice application or any other score emitted by the application.

We will now define oscillation and associated metrics to measure oscillations. A QoS metric qos (e.g., loss, latency or jitter) for path P is said to have *SLA violation* at time t for application A if its value is greater than a said threshold, i.e.,

$$slav_{(P,A,qos)}(t) = \begin{cases} 1, & if\ qos(t) > T_{(qos,A)}, \\ 0, & otherwise, \end{cases}$$

where $T_{qos,A}$ is the threshold for SLA violation for qos. A path P is said to have *SLA violation* for an application if any of the QoS metrics has SLA violation.

$$slav_{(P,A)}(t) = \begin{cases} 0, & if\ (qos(t) \leq T_{(qos,A)})\ \forall\, qos \in \{loss, latency, jitter\}, \\ 1, & otherwise. \end{cases}$$

A QoS metric qos for a path P and application A is said to have *QoS-oscillation* at time t if *qos* value jumps from an SLA violation state to non-violated state, or vice-versa, i.e.,

$$qosc_{(P,qos,A)}(t) = \begin{cases} 1, & if\ \left|slav_{(P,A,qos)}(t) - slav_{(P,A,qos)}(t-1)\right| > 0 \\ 0, & otherwise. \end{cases}$$

A path P is said to SLA-oscillate (or simply, *oscillate*) for an application A at time t if it goes from SLA violated state to non-violated state, or vice-versa, i.e.,

$$osc_{(P,A)}(t) = \begin{cases} 1, & if\ \left|slav_{(P,A)}(t) - slav_{(P,A)}(t-1)\right| > 0 \\ 0, & otherwise. \end{cases}$$

Note that QoS-oscillation does not always mean that the path oscillates for the application. Figure 1 shows two examples. Example 1 shows time-series of loss fraction (0 to 1, representing 0% to 100% loss respectively), latency (in ms) and jitter (in ms) for two paths. The red dashed line shows the threshold for different QoS metrics for voice application ($T_{(qos,voice)}$). Vertical gray lines show times when the path has oscillation, i.e., $osc_{(P,voice)}(t) = 1$. The left path oscillates in loss. The center path oscillates in all three QoS metrics (loss, latency and jitter); the path oscillation can be caused due to any of the QoS-oscillations. Example 2 shows a special case where the path is violating SLA due to high jitter for most of the time. Hence, even though the path has loss-oscillations and latency-oscillations, it has very few oscillations since the path is already in an SLA violated state.

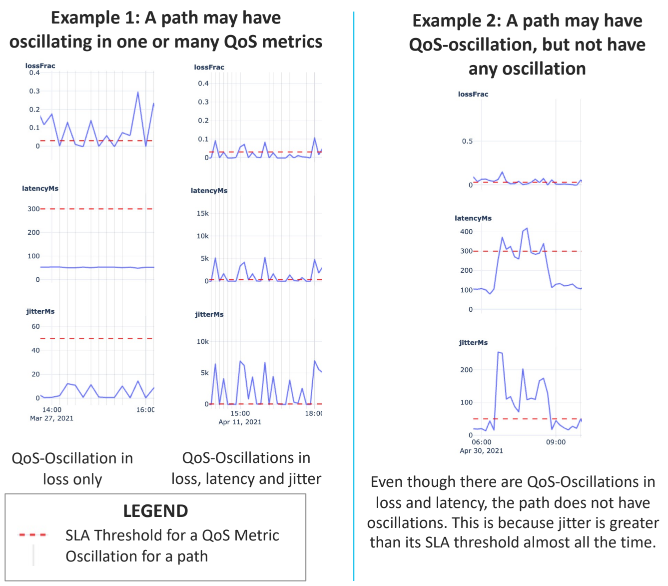

Figure 1: A path may oscillate along any QoS metric (loss, latency or jitter) as shown in Example 1. However, as shown on Example 2, oscillation along one QoS does not necessarily mean that the path is oscillating.

3. Categories of oscillation

This section quantifies two metrics to measure the frequency of oscillations and the intensity of oscillations paths may have.

Fraction of oscillation for a QoS metric qos is defined as the fraction of time-steps the path has observed oscillations for a given qos. If the path P is observed for n timesteps, and it has violated qos for $m_{(P,A,qos)}$, then $\mathrm{frac}_{(P,A,qos)} = m_{(P,A,qos)}/n$.

Now we define frequency of oscillation for a path across all QoS metric. *Fraction of oscillation* for a path P is defined as $\mathrm{frac}_{(P,A)} = m_{(P,A)}/n$, where $m_{(P,A)}$ is the number of oscillations for the path.

While $\mathrm{frac}_{(P,A)}$ illustrates how often a path experiences oscillations, it does not quantify how much the path oscillates. For example, a path may oscillate by around 20 ms in latency or 1000 ms in latency. These are two different types of behaviors that needs entirely different actions to be taken. A path-selection protocol or network policy change can avoid reacting to small oscillations such as 20 ms latency oscillations. However, a 1000 ms oscillation in latency needs to be addressed to avoid severe user experience degradation. Jump in oscillations is defined to quantify such intensity of oscillation.

A *jump in QoS-Oscillation* $\delta_{(P,A,qos)}(t)$ for a path P and application A at time t is defined as the absolute difference between QoS-metric qos at time (t-1) and t, if there is a QoS-Oscillation at t. At times when there is no QoS-oscillations, it is undefined, i.e., $\delta_{(P,A,qos)}(t) = |qos(t-1) - qos(t)|$, for all times t when there is a QoS-oscillation for a QoS metric qos. The jump measures how much did the QoS metric jump when it crossed the boundary. Based on the distribution of the jumps in QoS metrics, the oscillations for a path can be categorized into:

- Wild Oscillations: These are oscillations where the qos metrics usually jump by a large amount. Such oscillations cannot be ignored by the path-selection protocol or the network administrator since it has a large effect of user-experience.
- Near-boundary oscillations (NBO): These are oscillations where the QoS metric varies slightly around the SLA thresholds. For example, the loss in Example 2 of Figure 1 shows near-boundary oscillation.

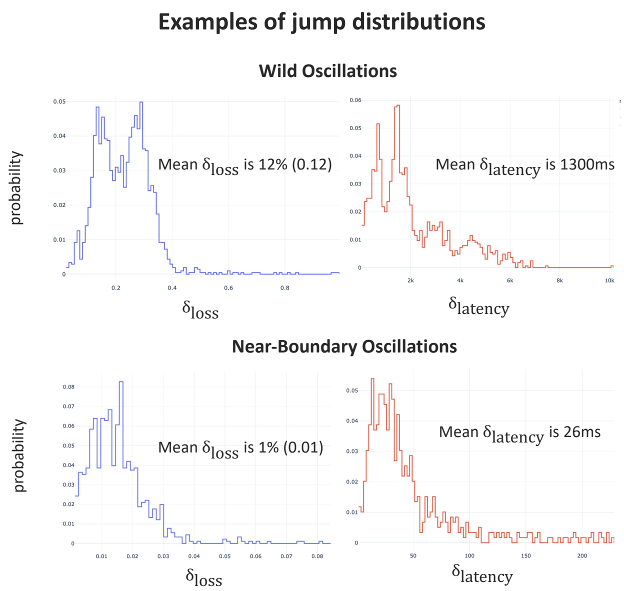

Figure 2: Examples of Wild and Near-boundary oscillations defined by studying jump distributions.

Figure 2 shows examples of Wild and Near-boundary oscillations for loss and latency for various paths. Each plot shows the histogram of loss- and latency-jumps with y-axis as fraction of samples observed. A path can be tagged as having Wild or NBO based on the jump distribution (or statistic such as mean jump) of the main QoS metric that is causing the oscillation.

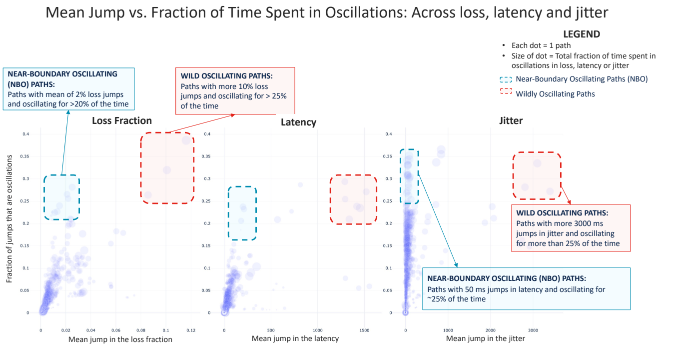

Figure 3: A scatter plot of fraction of jumps and mean jump for various QoS metrics on one sample network. The paths (circles) in top-right region indicate wild oscillations, where as the ones in top-left region of the graph indicate NBOs.

The oscillation dynamics for a given network and the number of paths that have wild, or near-boundary oscillations can be seen in one example network shown in Figure 8. Each graph shows the fraction of jumps for a given qos for voice application ($\mathrm{frac}_{(P,voice,qos)}$) on Y-axis and the mean jump ($\overline{\delta}_{(P,voice,qos)}$) on X-axis. Each dot signifies one path, and the size of the dot shows the fraction of SLA-oscillation $\mathrm{frac}_{(P,voice)}$. The paths (dots) in the top-right corner for every graph signifies paths with high mean jumps and high fraction of jumps; these are wild oscillation paths. The dots on the top-left corner of the graph are the ones with near-boundary oscillations for that QoS metric.

4. Visualizing oscillations via Path Cubes

Metrics discussed in the previous section provides the statistics to infer the frequency, intensity, and the type of oscillation the path is experiencing. While this is usually sufficient to quantify, it does not provide more details into how the path is behaving. This section discusses a novel 3-dimensional visualization of path oscillations by looking at all QoS metrics jointly.

The path dynamics can be shown by time-series as shown in Figure 1 or by looking at various metrics as shown in Figure 2. However, it is relatively harder to look at three time-series across time or distributions of jumps. At Cisco, we have created QoS dynamic visualization tool called Path-Cubes, which enables interactively viewing the path behavior and reasoning on what type of behavior may cause such dynamics. Such tools are extremely useful to network experts and data-scientists alike to troubleshoot diverse paths seen on the Internet.

Figure 4 shows three ways to visualize a path with wild oscillations in latency and jitter. The left graph (a) shows the typical time-series of the QoS metrics. Graph (b) shows Path-Cube trajectories. Here the the 3 QoS metrics (loss, latency, and jitter) is represented as a 3-dimensional plot. The trajectory of the three QoS metrics is joined by a line. The green-box shows the region where SLA is not violated; if a path has loss, latency and jitter within this green-box the application experience is supposed to be good. Figure 4 shows the green-box for for the SLA thresholds for voice (as described in Table 1). An oscillation in this visualization refers to the line crossing the boundary of the green-box. Note that the loss, latency and jitter are clipped at 9%, 900ms and 150 ms, respectively (three times the SLA thresholds for voice).

Figure 4 (b) shows that the path is provides great user experience sometimes (well within the green-box). However, it frequently moves from well within the green-box to extreme values of latency, jitter and/or loss. Such figures enable clearly seeing how much the path oscillates and how often it oscillates – especially when interactively seeing by rotating and zooming. The video interactivity is also provided in the tool where the evolution of path-cube across different time-steps is shown as animated frames.

Figure 4 (c) shows where the path spends most of the time by visualizing the density in 3d plots. The 3d-space is broken down into smaller cubes, and within each cube, the time spent by the path (in terms of fraction of time-steps) is measured; this is indicated by the color. Instead of the green-box, just a cube with black-lines is shown here. In the below example, it can be seen that the path spends most time in: (i) well within the green box; (ii) almost all jitter values; and (iii) very high jitter and latency values. By looking at the (b) and (c), it is clear that this path has wild oscillations.

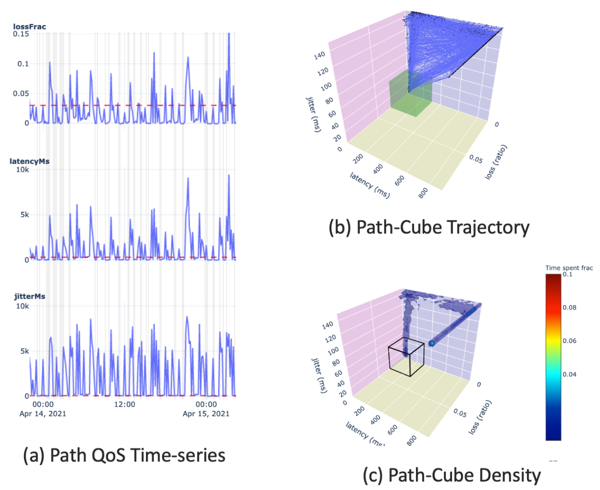

Figure 4: A path with Wild Oscillations in latency and jitter. It is represented as a time-series (a), or Path-cube trajectory and density (b and c).

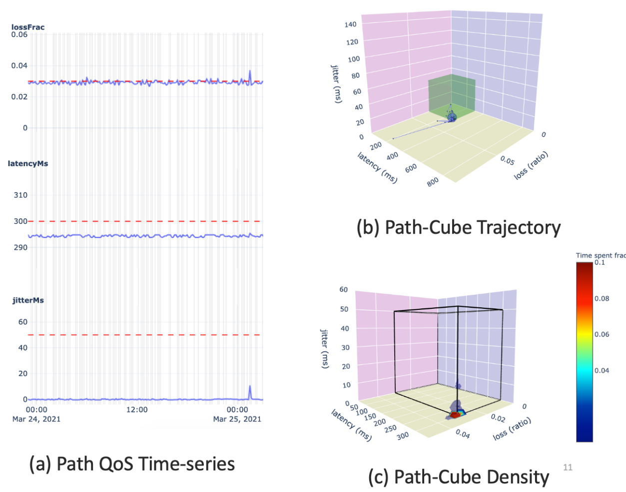

Figure 5: An example of Near-Boundary Oscillating Path

Figure 5 shows a path with NBO, where the path is highly probable to oscillate in loss and latency. This can be clearly seen by the Path-cubes density and trajectory plot, where it spends most of the time in the boundary of the SLA-violation box.

5. Oscillations in different networks

The previous section provided metrics and examples of path oscillation. This section discusses the dynamics of path oscillations seen across multiple networks.

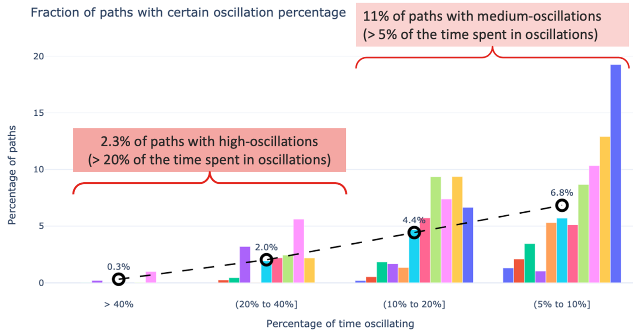

Figure 6: Percentage of paths that have various degrees of oscillations for voice application across different networks

The oscillations summary is presented here for thousands of paths, deployed across the world on different networks for 12 operational networks. Two months of data is studied. Figure 6 shows the overall percentage of paths (Y-axis) that have various amounts of oscillations (X-axis). The X-axis shows the four bins of oscillation ranges ranging from 5% to 40% of the time spent in oscillations (SLA-oscillations). Each colored bar is for one network. The dashed black line shows the average percentage of paths across all networks. It can be seen that around 2.3% of the paths spend more than 20% of the time in oscillations; such paths can be deemed as high-oscillation paths. Other 11% of the paths spend 5% to 20% of its time in oscillations, which is significant time where experience varies for the user between good and bad.

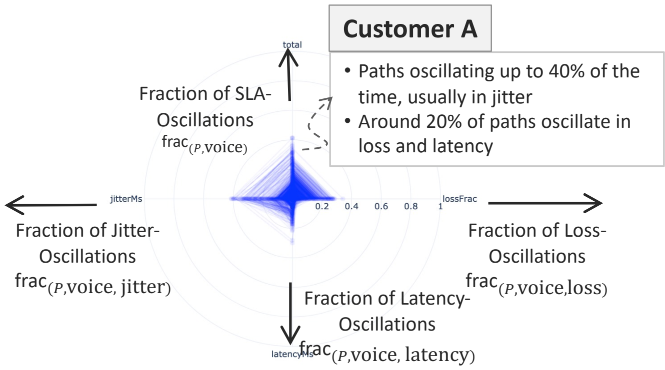

Figure 7: Oscillations depicted for one customer A for voice application.

Figure 7 show a polar chart to show the fraction of oscillations, and main QoS causing oscillation across a network of thousands of paths. Each axis shows fraction time spent in SLA-oscillation (vertical top axis), or QoS-oscillation (other three axes). Each path is represented as a semi-transparent line that which has some value of SLA- and QoS-oscillation. Note that only top-5000 paths with $frac_{(P,voice)} > 0.1$

Is shown in the figure.

From the above figure, it is clear that the customer has a few paths that are that have fraction of SLA-oscillations > 0.3, and most of these are caused due to jitter. There is also a lot of paths that have QoS-oscillations in loss, and some in latency.

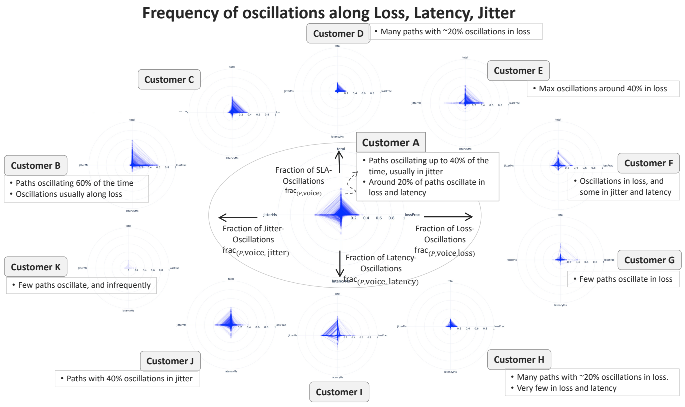

Figure 8: Oscillation diversity across networks. Each network may have oscillations due to different QoS metrics.

Figure 8 shows different networks with oscillations for voice-applications. The figure clearly shows that each customer may have a different signature for oscillation. While some networkks have high-oscillations due to loss (Customers B, C), others have relatively medium oscillating paths due to loss (Customers D and E). Some networks like Customer G have very few paths with high-oscillations. Other networks like Customer K may have almost no highly oscillating paths. Few other networks like Customer A, I and J can have oscillations due to multiple QoS metrics such as loss and jitter. Depending on the type of the network, the network administrator can take appropriate action to configure the network. Such graphs help the customer to not only see which paths are oscillating, but also enable to find high-level problems across multiple paths in the network.

6. Conclusion and Next Steps

This paper describes the phenomenon of path oscillations for application-aware path-selection. Path oscillations for application-aware path-selection is much different than the ones defined for intra-domain routing. Here, the oscillations are defined as number of times a path switches between good and bad application experience for a user. While most products and studies take a protocol-approach to identify such path-oscillations, this paper takes a data-oriented study; it quantifies the type of path by observing the oscillations for different applications by analyzing the telemetry over several weeks or months of data. Such an approach helps in two main use cases. First, predictive routing engines such as Cisco's Predictive Internet can utilize long term data to accurately classify path isolations, and then to remediate the ill-effects. Second, it can be used by network administrators to measure how many paths are oscillating for which applications and due to which QoS. These help in identifying the possible root-causes for bad user-experience, and configuring the network to rectify the problems.

In terms of specific contribution, the paper first identifies application-aware oscillations using telemetry and SLA thresholds. It then quantifies the frequency and intensity of oscillations. Based on these metrics, two categories of oscillations are identified: Wild and Near-Boundary oscillations. Each category has specific action to rectify. While path-selection protocols or network policy should switch the routing away from paths with Wild oscillations, Near-Boundary oscillations can be avoided by slightly adjusting the SLA thresholds without possibly affecting user-experience. The paper then discusses Path-Cubes, an interactive visualization approach to study path trajectory and density. Finally, the paper shows oscillations in different sample networks. Each network has specific fraction of paths with oscillations which are caused by different QoS metrics. This diversity is shown by studying 12 sample networks.

References

[1] Cisco, "Cisco SD-WAN: Application-Aware Routing Deployment Guide," Cisco, 21 July 2020. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/cisco-sdwan-application-aware-routing-deploy-guide.html. [Accessed 28 Aug 2021].

[2] K. Varadhan, R. Govindan and D. Estrin, "Persistent route oscillations in inter-domain routing," *Computer networks,* vol. 32, no. 1, pp. 1--16, 2000.

[3] J. P. Vasseur, "Towards a Predictive Internet A new world, with new challenges," Cisco Systems, 2001.

[4] P. Montero-Manso, G. Athanasopoulos, R. J. Hyndman and T. S. Talagala, "FFORMA: Feature-based forecast model averaging," *International Journal of Forecasting,* vol. 36, no. 1, pp. 86-92, 2020.

[5] J. Vasseur and V. Kolar, *Large-scale Internet Path modelling and applications,* 2021.