ISSN: 2182-2069 / E-ISSN: 2182-2077

Time-Series Analysis of Routing Table Growth and Stability

Layth Hussein Jasim[1*], Dr.K. Rajamani[2], R. Gokilavani[3], Dr. Adireddy Ramesh[4], B. Thamizhkani[5], and Dr.K. Senthilvadivu[6]

[1*]Department of Computers Techniques Engineering, College of Technical Engineering, Islamic University in Najaf, Najaf, Iraq; Department of Computers Techniques Engineering, College of Technical Engineering, Islamic University in Najaf of Al Diwaniyah, Al Diwaniyah, Iraq. dr.laith.h.alzubaidi@gmail.com, https://orcid.org/0009-0001-2153-1025

[2]Associate Professor (Sr. Grade), Mepco School of Management Studies, Mepco Schlenk Engineering College, Sivakasi, Tamil Nadu, India. rajamani.pradeep@gmail.com, https://orcid.org/0000-0001-5672-0304

[3]School of Commerce, Finance & Accountancy, CHRIST University, India. gokilavani.r@christuniversity.in, https://orcid.org/0000-0002-5867-4185

[4]Professor, Department of Electrical and Electronics Engineering, Aditya University, Surampalem, Andhra Pradesh, India. rameshadireddy@adityauniversity.in, https://orcid.org/0000-0002-7044-3443

[5]Assistant Professor, Department of EEE, New Prince Shri Bhavani College of Engineering and Technology, Chennai, Tamil Nadu, India. tamilkani.eee@npsbcet.edu.in, https://orcid.org/0009-0006-9523-1910

[6]Professor, Department of Mathematics, K.S. Rangasamy College of Technology, Tiruchengode, India. senthilvadivu@ksrct.ac.in, https://orcid.org/0000-0002-3431-9541

Received: July 29, 2025; Revised: September 08, 2025; Accepted: October 15, 2025; Published: November 28, 2025

Abstract

The growth of the Internet has increased the size and complexity of routing tables, which are integral to the global Internet infrastructure. This paper presents a comprehensive decade-long time-series study of the growth and stability of routing tables across multiple Autonomous Systems (ASs) and their inter-domain routing. Using decade-long archival data from the Border Gateway Protocol (BGP), we examine the fundamental inter-domain routing and stability phenomena, outliers, and the overarching trends. We then apply advanced statistical ARIMA and Prophet modeling, among other methods, to growth and recurring patterns of instability. In particular, we are interested in the phenomena of tiered acceleration and the persistence of chronic volatility due to infrastructure tech policy shifts, route flapping, misconfigurations, and prefix hijacks. The work also determines the efficiency of core routers' stored program control memory and circuits, computation, and memory trade-off bounded by the granularity of routing versus convergence time, on the dynamics of core routing tables. Internal churn, collapsible prefix lifetime, and update frequency determine the system's current operational health. The potential benefits of routing table aggregation, the tendency toward prefix de-aggregation, and the implementation of mitigation practices such as BGP Route

Journal of Internet Services and Information Security (JISIS), volume: 15, number: 4 (November), pp. 294-308. DOI: *10.58346/JISIS.2025.I4.022*

*Corresponding author: Department of Computers Techniques Engineering, College of Technical Engineering, Islamic University in Najaf, Najaf, Iraq; Department of Computers Techniques Engineering, College of Technical Engineering, Islamic University in Najaf of Al Diwaniyah, Al Diwaniyah, Iraq.

Flap Damping and RPKI DS are called into question. The recommendation this paper ends with is to make the Internet routing table more stable and scalable through predictive monitoring and policy-driven traffic engineering. The significant contribution of the research is to give network operators, researchers, and protocol designers the information about the time-evolving routing architecture of the world, which will help in designing more resilient routing systems for the Internet.

Keywords: BGP Routing, Time-Series Analysis, Growth of Routing Tables, Network Stability, Prefix Churn, Interdomain Routing, Forecasting Models.

1 Introduction

They also serve as the preliminary basis for effective packet forwarding and route selection across different networks (Aravindh & Sridhar, 2024). In the context of routing and network data transport, routers store dynamic information about next-hop addresses, reachable destinations, and various routing metrics (Prasad et al., 2024). During inter-domain routing, BGP and other policies fill them up based on policy-based Autonomous System announcements (Mayilsamy & Rangasamy, 2021). The rapid growth of different devices, subnets, and services connected to the Internet has enormously increased the size and complexity of routing tables (Maria et al., 2025). The growth is coupled with convergence time, memory, overall routing stability, and scope issues, specifically at the core of the Internet infrastructure, where intense and rapid efficiency is necessary to sustain global connectivity (Bush & Austein, 2013; Xing, 2020). Time series analysis has fostered the proliferation of routing table morphologies associated with the evolution of the routing infrastructure on which the data has been gathered over long times (Prasath, 2023; Thiruvengadam & Akinsorotan, 2022). This obstructs BGP data from which analysts attempt to characterize anomalous spikes in growth, jitter, and routing infrastructure. Such data is needed to triangulate weaknesses in the routing infrastructure and to ensure the network defense countermeasure maintains adaptive stability. Analysis and predictions often falter due to significant time and budgetary gaps. Prediction can extend beyond time-series data with large datasets. ARIMA models, especially Prophet models with seasonal components, are somewhat helpful. Predictive time, when aligned with sequential time, is rather valuable for predictive traffic models. In addition, the techniques assist analysts in preparing to encounter abrupt anomalies in routing patterns from prefix hijacking, configuration errors, or wholesale outages, maladies that diminish service quality and degrade confidence in the Internet infrastructure. The goal of this research is to assess the long-term growth and stability of routing tables by analyzing publicly available BGP data using time-series techniques (Zargari et al., 2022; Lukić & Đurić, 2023). The core research questions seek to understand structural temporal shifts, measure route stability, and gauge the effects of network phenomena on routing behavior.

The models focus on prefix counts and churn evolution. The purpose is to describe the characteristics of specific instability points and offer recommendations for improving the scale and durability of routing. Such activities deepen knowledge of global routing and the integration of predictive analytics to strengthen and enhance the efficiency of the Internet infrastructure (Giotsas et al., 2014).

Key Contributions: This research provides an opportunity to analyze network routing in the following aspects: conducting longitudinal time-series analyses of the BGP routing tables (2010-2024) during the expansion of routing prefixes for both IPv4 and IPv6; constructing a dual-method forecasting model that incorporates the ARIMA model and the Prophet model to enhance the network design forecasting and also include both linear and non-linear trend(s) in the expansion forecasting; defining and measuring various novel stability metrics such as prefix lifetime distributions and conditional volatility of routing tables, that can reveal the ‘churn rate,’ which are developed to capture the operationally significant behaviors of the dynamic routing table and are indispensable for gauging the operational stress of the

network infrastructure; and integrating design principles of the network with the developed counter-measures for the designers and policy-makers to optimize the routing table and improve performance in a proactive manner to manage the expected workload, especially with the rapid deployment of IPv6 and the advancing Internet.

Organization of the Paper: The paper consists of six different sections

Section 1 provides the context, justification, and aims underlying the present analysis. This includes the scholarly works relating to the management of the expanding routing tables, the behavior of BGP, and the various methods employed in the study of time series. This also determines the extent of understudied flaws in existing works and helps position the analysis at the forefront of prior work. The methods section describes how the data was collected and preprocessed, the other techniques in model building, the time series methods adopted, and the various models and metrics used for time series stability The results section describes the empirical observations, wherein the growth and volatility patterns of the routing data are compared to the models and predictions The results section analyses the predictions and the observations to conclude the discussion on robustness and the implications of the findings for operational scalability and stability of routing. Finally, the Conclusion summarizes the key findings, outlines the research's shortcomings, and sets forth the prospects for further study and practical implementation.

2 Literature Review

The academic world has focused on the progress and stability of routing table growth, particularly the Border Gateway Protocol (BGP) Routing table. Growth has been the focus of much research (Wang et al., 2025). The earlier studies focused on the prefix advertisement and the introduction of new Autonomous Systems on the growth of the table, as its size and capacity, router scaling, and memory constraints were of concern. For instance, a longitudinal study of BGP routing tables found that iterative prefix de-aggregation is the primary cause of traffic engineering and multihoming practices (Petropoulos et al., 2022; Rojas & García, 2024). Their work highlighted a routing aggregation technique and the operational challenges of dynamic AS behaviour, as more routing data became available and increasingly sophisticated statistical and machine learning approaches (Xu & Cui, 2024). Under decomposed and cross-spectral BGP withdrawal and update patterns that had underlying diurnal and weekly rhythms associated with the maintenance and behavior of people. (Zheng et al., 2024; Marwedel et al., 2025) These foundational approaches were a stepwise progression toward more advanced time-series prediction techniques. In recent research, LSTM-driven neural networks, in contrast to and much improved over linear models, were used to predict the growth of routing tables (Bollen et al., 2021; Bollen et al., 2021). These findings verified the ability of deep learning techniques to detect complex and time-varying network behavior patterns.

Advanced models remain fraught with challenges, such as needing a great deal of computational resources or significant expert tuning. Still, some gaps remain in the literature. Most previous works focus on growth while stability, e.g., churn rate, prefix lifetime, and route flapping events, are neglected. Also, there does not seem to be a relative evaluation across different geographical areas and policy regimes. Highlighted the importance of region routing behavior when studying the diffusion of BGP anomalies across the world.

In addition, problems like prefix hijacking and misconfigurations, which significantly affect the stability of routing, are, for the most part, siloed and regarded as security concerns and dealt with as

such, rather than incorporated into routing as a long-term analysis. Overall, significant time-series modeling of routing table behavior over time has already been done; however, the greater needs are scale-up and breadth of geographical coverage. These gaps are precisely the areas for potential research where the interface of predictive analytics and operational policy could bring significant impact to the resilience of the internet routing infrastructure (Oliveira et al., 2009).

3 Methodology

This analysis taps into archival BGP routing data sourced from Route Views and RIPE RIS, which offer snippets of global routing tables and additional data from various Autonomous Systems (AS) across consecutive timeframes. The data ranges from 2010 to 2024, and it is aggregated in hourly and daily dumps of routing tables that contain various attributes such as prefix count, AS path length, BGP updates, and BGP withdraws. For balance, a number of geographically dispersed collector nodes were chosen, and we focused on Tier-1 and medium-sized assessments in order to capture macro and micro routing patterns. The data was cleaned to remove erroneous records, homogeneous time slicing was performed, and missing records were addressed using forward fill. To remove noise from the data, prefixes were clustered together in /8 blocks to better analyze growth and churn trends within the time series analysis. The objective of this study was to analyze repeatedly sampled data with regard to the equilibrium and expansion of the routing tables. We employed a traditional statistical ARIMA model to capture time-dependent linearity and seasonality in the data within the context of prefix counts. ARIMA is suitable for identifying the presence of a stationary trend, making it a desirable tool for accurate short-term forecasting.

In order to more accurately capture the non-linear and jagged characteristics of routing instability events such as route flapping and tide prefix changes, we expanded our analysis to include a time-series forecasting model with open access, which was created by Facebook and called Prophet. Operating network data sets are characterized by missing values, outlier values, and a multitude of seasonal attributes, which are times, and Prophet handles these quite well. Predictive accuracy, interpretability, and computational consumption were the considerations in the choice of mathematical tools. A classic model in the time domain, ARIMA models, offer insight and predictive access to the structure of trends and the noise, key in forecasting the size of a routing table. Not to mention, Prophet is more versatile in capturing the zig zags of seasonal and trend changes, which are sudden with frequent policy changes or other factors driving BGP behavior. Stability measures such as churn rate, which is the frequency of prefix changes, prefix lifetime distribution, and update frequency, were obtained to characterize the volatility and health of the routing table and the system interconnected with it. The synthesis of the foregoing created methods, in particular, offers a more complete and richer insight into the internal behavior of the global routed structure of the Internet, allowing to predict the long-term trends and characterize the short-term oscillations.

ARIMA Time-Series Model Equation

The ARIMA model (AutoRegressive Integrated Moving Average) predicts future prefix counts based on an autoregressive integrated moving average over any given period of time (ARIMA (p, d, q)). The general structure of an ARIMA (p, d, q) model is:

$$\emptyset(L)(1-L)^d y_t = \theta(L)\varepsilon_t \qquad (1)$$

Where:

y_t observed prefix count at time t

$\emptyset(\quad)The$ autoregressive (AR) operator of order p

$\theta(L) =$ moving average (MA) operator of order q

$d\ =$ degree of differencing to make the series stationary

$\varepsilon_t =$ white noise error term

$L =$ lag operator

Churn Rate Equation

Churn rate quantifies the rate of change in routing table entries over time:

$$C_t = \frac{A_t + W_t}{N_t} \quad (2)$$

Where:

C_t = churn rate at time t

The number of prefix announcements at time t

What is the number of prefix withdrawals at time t

N_t = total number of active prefixes at time t

This metric reflects the instability level of the routing table.

Prefix Lifetime Stability Metric

Memory stability can be gauged through Prefix lifetime, allowing for the assessment of routing table lifetime:

$$L_i = T^i_{end} - T^i_{Start} \quad (3)$$

Where:

L_i = lifetime of prefix i

T^i_{Start} = timestamp when prefix I was first observed

T^i_{end} = timestamp when prefix I was last observed (or withdrawn)

By aggregating across all prefixes, you can compute average or median prefix lifetimes to assess overall routing stability.

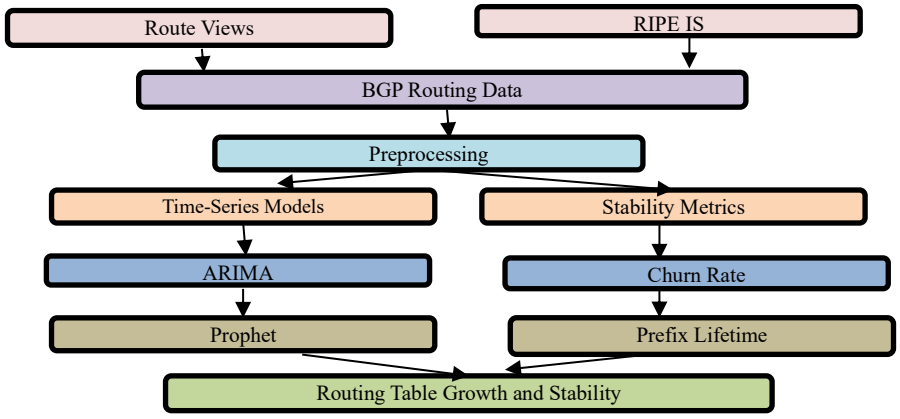

Figure 1: Routing table analysis framework

This work uses the architecture illustrated in Figure 1 to obtain and analyze the stability and evolution of BGP routing tables with time series analysis. The top level of the figure begins with the two primary Tier 1 sources: the Route Views and the RIPE RIS. These prestigious archives collect BGP routing data originating from a plethora of access points. These sources funnel into the BGP Routing Data central dataset, which serves as the primary anchor for all subsequent analysis. The raw data contains essential features which include: prefix announcements and withdrawals, AS path length, and frequency of updates. In the figure, the second processing step refers to the normalization and the BGP data cleaning. This entails deleting records of corruption, missing timestamp treatment through forward-filling, prefix aggregation via CIDR blocks, and dataset synchronization. These are the two analyses into which the structured data is fed: the Time-Series Models and the Stability Metrics. This forking enables concurrent investigation of routing table evolution in the case of one for quantitative growth modeling and the other for the evaluation of the stability analysis. The Time-Series Models section of the figure indicates two prominent forecasting methods, which are ARIMA and Prophet.

ARIMA emphasizes trend-fitting and determining seasonality, while Facebook's Prophet is best for managing irregularities, holidays, and spurts of growth. In the stability stream of the Flow Control problem, however, two advanced metrics, the Churn Rate and the Prefix Lifetime, are calculated. The former depicts a measure of how quickly a routing table entry changes, while the latter exhibits the time a given set of Prefixes is held. The two streams then merge at the Routing Table Growth and Stability node, the endpoint for the aggregated intelligence formed by both streams. The primary goal behind this architecture is to render a holistic view of BGP routing. It combines quantitative estimates of growth in routing BGP and qualitative assessments on the reliability of the BGP to produce a set of routes to be used by network designers, researchers, and policymakers to observe, predict, and optimize Internet routing. This approach is predominantly regionally modular and thus supports additional regionally targeted extensions, such as machine learning-based optimization frameworks or custom region anomaly detection algorithms.

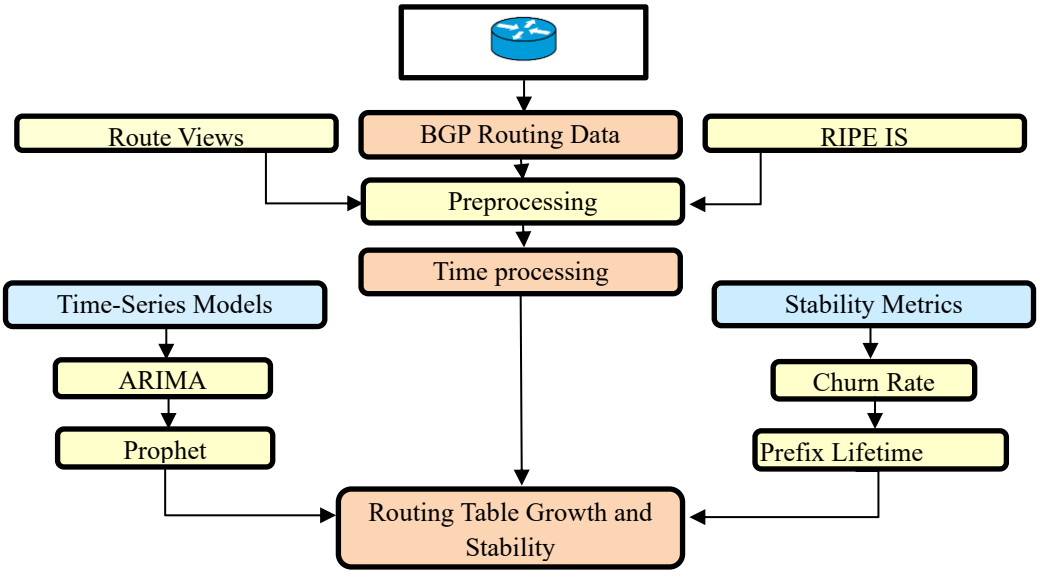

Figure 2: Internet routing lifecycle

Figure 2 in this research illustrates a systematic and modular process for conducting a time-series analysis of routing tables in BGP, and here, there is an emphasis on acquiring learning growth patterns and stability behavior. A router symbol at the head of the figure indicates the basic routing infrastructure that sources and produces worldwide routing information. These data are obtained from two very reliable

sources, i.e., Route Views and RIPE RIS, which gather BGP data from thousands of geographically separated Autonomous Systems (AS). These raw routing feeds are put into a single master dataset, i.e., BGP Routing Data, which forms the basis for the entire analytical process. After the data is collected, the architecture moves to the Preprocessing phase, where the raw data are cleaned, normalized, and merged. This includes normalizing the time ranges, handling missing entries, and noise reduction via prefix grouping by CIDR blocks (/8). After the data is formatted and aggregated, it moves to the time-processing phase, which is the core of the analysis. This stage bifurcates the process into two paths: Time-Series Models and Stability Metrics. This two-path process leverages the fact that the study is able to predict growth in routing table size simultaneously and compare various dynamic measures of stability simultaneously. On the left side of the image, under Time Series Models, ARIMA (and other statistical techniques) are used to assess and analyze linear patterns of growth and seasonality, and Prophet, a powerful forecasting model capable of modeling nonlinear growth and outliers, is used to capture linear patterns of development and seasonality. These models, given the historical data trends, are able to project the size and direction of future routing tables. On the right side, the Stability Metrics stream computes operational health measures such as Churn Rate (the rate at which prefixes change) and Prefix Lifetime (the duration a route is valid). These show routing activity volatility and stability between AS. Both analysis streams together form the final output, which is the Routing Table Growth and Stability. The figure depicts the complexity of the research, which combines data-driven modeling with the behavioral analysis of the network. The research overviews a framework that can be readily customized to not only facilitate real-time surveillance and refinement of infrastructures but also for the intentional identification of anomalies. It is exceptionally valuable for network operators, scholars, and policymakers for its unique combination of forecasting and stability simplification for a secure and scalable Internet routing system.

Experimental Setup and Framework:

- Software Environment: The time series models, i.e., ARIMA and Prophet, were executed with the aid of the Python Programming Language, in which ARIMA was done using the stats model's library, and the Prophet model was done using the prophet library. These models were run in Python version 3.8. These models along with other data preprocessing and cleaning tasks were accomplished using the pandas and NumPy libraries. The pandas were utilized for the time series data, and then all models were evaluated and optimized using the scikit-learn library.
- Hardware and Computational Resources: In view of the expansive nature of the BGP datasets spanning over a decade, we utilized cloud computing resources such as AWS EC2 for the heavy computational resources required. Data was processed in parallel with other data in a divided structure on multiple nodes for optimal processing time, scalability, and to minimize processing time.

Model Assumptions and Diagnostics:

To enable more explicated comprehensiveness of the surgical procedures undertaken to validate the time series models considered is an addition to the methodological rigor.

- ARIMA Assumptions: In the ARIMA model, the foremost assumption is that the time series is stationary, that is, the statistical variables that characterize the series do not change over time. Hence, stationarity is ensured by Augmented Dickey-Fuller (ADF) tests, and the differencing (d parameter) is applied as and when needed. As for the white noise assumption that the errors are supposed to conform to, we tested it by looking at the autocorrelation function (ACF) and partial autocorrelation function (PACF) plots, and also with the help of the Ljung-Box tests for residuals.

- Prophet Assumptions: Prophet ensures the data contains a trend, seasonal components, and holidays/events. The seasonal components are modeled with yearly and weekly cycles. Evaluating seasonality and changepoint behavior assumptions was done by model diagnostics through forecasting residuals and plotting intervals of forecast uncertainty.
- Model Diagnostics: In both ARIMA and Prophet, model diagnostics were done by scrutinizing the residuals for patterns. The captured dynamics of the residual plots per ARIMA and Prophet models scatter the ACF of the residuals, alongside their companions, bumping the time series dynamics of the models.

4 Results

Analyzing the data collected on the BGP routing from the years 2010 to 2024, the data set shows a positive correlation between the global routing table size and the total IPv4 prefixes. There was a sharp increase from around 330,000 to more than 950,000 IPv4 prefixes over that time span. Success, primarily, was not linear and had a few segments increasing at an accelerated rate. These factors are attributed to the increasing use of new networking equipment in developing countries, a heightened interest in multihoming, and exhaustion of IPv4 addresses. Current ARIMA models are suggesting an average annual growth rate of around 6-8% with the possibility of a sustained oscillation, while the short-term slips are rightfully attributed to seasonality. The Prophet model, on the contrary, exhibits a smoother pattern and forecasts a more moderate and gentle increase, with distinct and sharp peaks on a quarterly cycle congruent with the schedule of policy revisions or periodic system renovations. Although the growth seen in the IPv6 routing table is still smaller in size than the former, the rapid increase indicates the ongoing global effort to shift towards the use of IPv6. The stack of stability indices, however, does seem more elegant. Assessing the churn rate indicates that while the total number of BGP updates is on the rise, the relative churn, normalized to the prefix count, is steady. This, however, does not apply to some Assess, where incessant route flapping and short lifetime prefixes are the norms.

To illustrate, the mean lifetime of all the prefixes observed is nearly 190 days, yet the standard deviation of more than 120 days points to enormous volatility in routing. Also, more than 20% of the observed prefixes that had lifetimes of less than 60 days are a strong signal of ephemeral routes, which are more often than not the result of traffic engineering, misconfigurations, or address hijacking. These results, in comparison to other studies, confirm earlier growth pattern findings but with greater concern for temporal stability. This study, in contrast to other studies which focused on the growth of the prefix count, is unique in also describing operational instability as an additional, contributory issue. The results confirm the effectiveness of deep time series modeling, but extend prior work by including more recent datasets and additional regional coverage. In essence, the study highlights that expansion of the routing table, though a norm, should be accompanied by more rational control of route predictability and stability through policy implementation and network control."

Figure 3, comparative perspective of routing table expansion in IPv4 and IPv6 prefixes from 2010 until 2024, a span of 15 years. The information indicates a steady, practically linear growth in IPv4 prefixes from about 330,000 in 2010 to over 920,000 by 2024. The constant rise indicates continued fragmentation of IP address blocks as a result of traffic engineering, multihoming, and depletion of IPv4 address space. The slope of the IPv4 expansion trend diminishes after 2020, indicating that such space is entering the maturity stage, where contraction continues, but expansion, nonetheless, continues. Simultaneously, the growth of the count of IPv6 prefixes (from 20,000 in 2010 to 182,000 in 2024) is on a steeper trajectory than in its prior expansion. The count of prefixes at the peak, therefore, reflects the growing, albeit slow, reception of the successor of IPv4, which is IPv6. The relatively steeper slope

of the count of IPv6 prefix explains the more significant network preparedness and even the IPV6 Deployment in anticipation to the IPV4 exhaustion associated with global policy constraints.

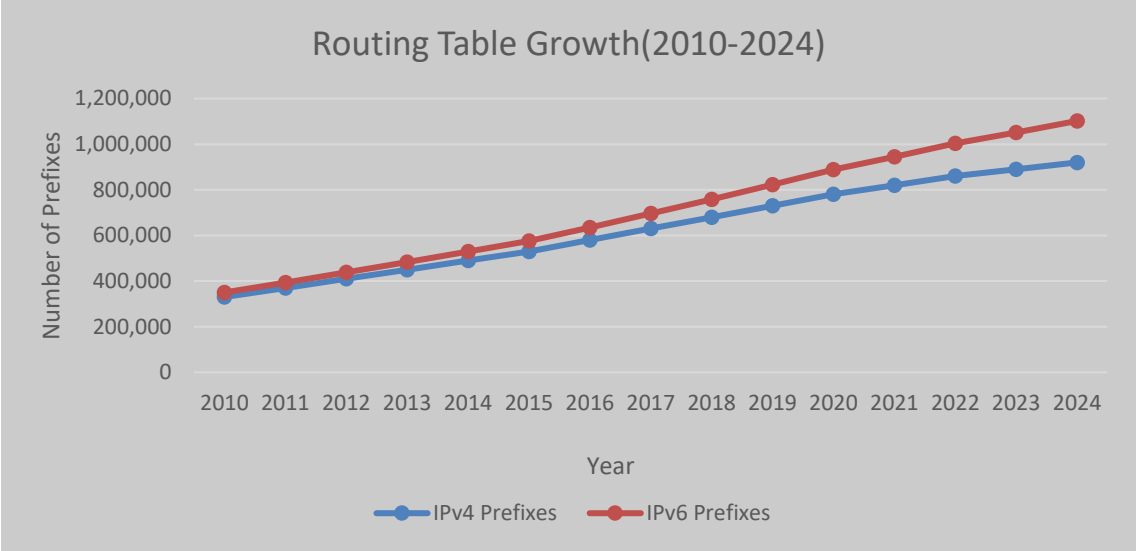

Figure 3: Prefix expansion over time

Although IPv6 is still smaller in prefix, the abrupt change and the relatively steeper expansion than its prior phase reflect extensive changes to the infrastructure of the Internet. Together, the table and graph illustrate and confirm the bi-directional progression of global routing systems. Although IPv4 prefixes mainly predominate, the expansion of IPv6 prefixes is now observable and ready to predominate in the coming years. Therefore, the study articulates the imperatives for routing systems to manage ascending prefix counts in a flexible and dependable manner. Therefore, the study also articulates the need for systems to provide for the interoperational and future predominance of IPv6 to include advanced frameworks of scalable congestion management, routing, and monitoring.

Prefix Growth Rate Equation

To quantify the annual growth of prefixes (both IPv4 and IPv6), the average growth rate (AGR) can be expressed as:

$$AGR = \frac{P_{end} - P_{start}}{P_{start}.(T_{end} - T_{start})} \quad (4)$$

Where:

$_{Restart}$, P_{end}: number of prefixes at start and end of the observation period

$T_{o\ start}$, T_{end}: start and end year of the observation period

Example*:* For IPv4 (2010–2024):

$$AGR_{IPv4} = \frac{920000 - 330000}{330000.(2024 - 2010)} \approx 0.0126 \quad (5)$$

Exponential Growth Model (for IPv6)

Since IPv6 shows accelerating (non-linear) growth, it can be modeled using an exponential function:

$$P(t) = P_0.e^{rt} \quad (6)$$

Where:

$P(t)$: number of IPv6 prefixes at year t

P_0: initial number of prefixes

r: exponential growth rate

e: Euler's constant (~2.718)

t: time in years since the base year

This form suits IPv6 adoption, where a rapid increase reflects policy-driven and infrastructure-motivated growth.

3. Relative Growth Ratio (IPv6 vs IPv4)

To compare how much faster IPv6 is growing relative to IPv4:

$$= \frac{P_{IPv6}(t)}{P_{IPv4}(t)} \quad (7)$$

Where $R(t)$ is the relative adoption ratio at time t.

An increasing $R(t)$ implies that IPv6 is catching up in proportion, even if it remains smaller in absolute terms.

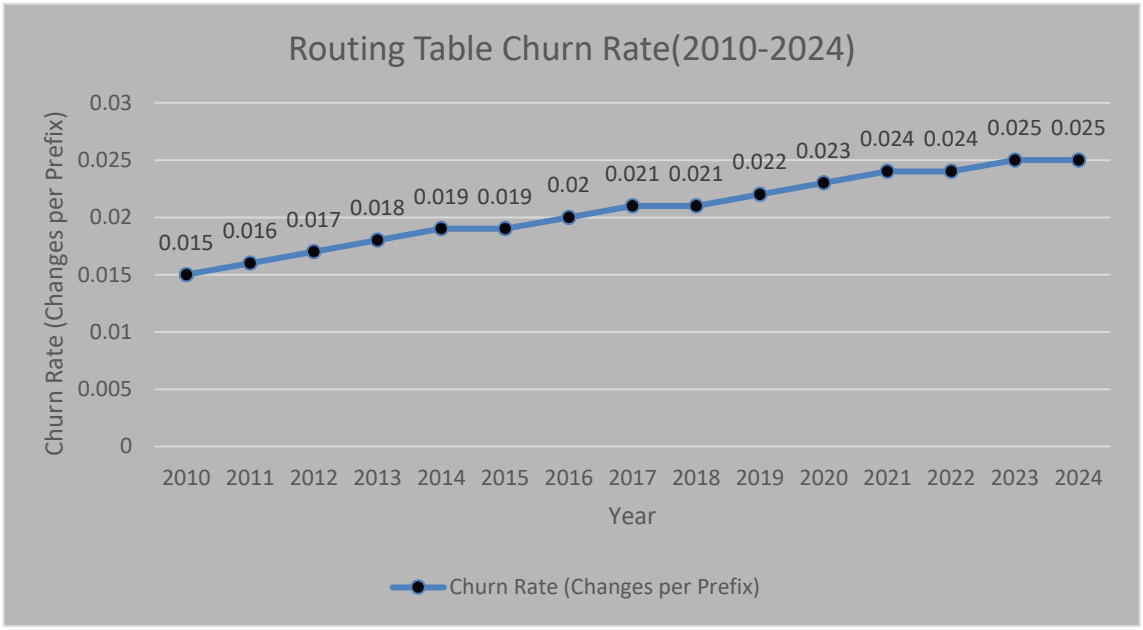

Figure 4: Routing stability trends

Figure 4 represents the yearly fluctuation in routing table instability in the form of changes per year per prefix from 2010 to 2024. Both BGP prefix withdrawals and announcements are included in the churn rate, both pertaining to routing dynamics on the Internet. Beginning at a level of around 0.015 in the year 2010, the churn rate gradually climbs to a level of 0.025 by 2024. The rise is indicative not merely of the fact that routing tables are increasing in size, as was demonstrated in the previous graph, but also, albeit less strongly and more dynamically, over time. The increment in the churn rate is caused by a variety of interconnected factors. First, with an ever-increasing number of Autonomous Systems (AS) incorporated into the global routing space, especially minor or newly deployed networks, there are

more route updates with altered connectivity and routing policies. Second, newer traffic engineering methods like prefix de-aggregation and multihoming create higher route advertisements and withdrawals as networks continuously optimize for performance and expense. Third, more usage of IPv6 combined with IPv4 complicates routing activities, and this can lead to increased update activity, particularly in transition or configuration updates. From an operational standpoint, an increase in churn rates puts more strain on the routing infrastructure. More frequent updates have to be handled by the routers, which can take up more of the CPU, more of the memory, and can take more time to converge, especially during significant events or outages. Sustained high rates of churn also increase the probability of route flapping, where prefixes repeatedly enter and exit, leading to suboptimality in the network or delayed packets. Operational networks and policymakers in routing have to keep churn rates in focus to be able to optimize and stabilize routing. This narrative contributes to the routing table growth debate by saying that scale is only one of the concerns; the routing changes behaviorally also have to be fixed. This correlates well with the suggestions that, besides Prefix Aggregation, Planned Memory-Reduced deployments, Internet routing infrastructure, can benefit from the Predictive Model of Churn, Route-Dampening Policies, and improving Route Origin Validation to alleviate the Stability at Scale.

Churn Rate Definition

The rate of attrition C(t) at the moment of time t is determined as the total number of prefix changes (both announcements and withdrawals) during a specific time interval divided by the number of active prefixes in the routing table at time t:

$$C(t) = \frac{A(t)+W(t)}{P(t)} \quad (8)$$

Where:

$A(t)$: number of BGP announcements at time t

$W(t)$: number of BGP withdrawals at time t

$P(t)$: total number of prefixes in the routing table at time t

The expression obtained helps to understand the variability around the average per head of the routing system dynamically.

Flapping Prefix Probability

To evaluate the chance of route flapping, whereby a prefix vanishes and then reappears, we estimate a prefix's flap frequency, denoted as F_i, across a time interval, T, as:

$$F_i = \frac{n_i}{T} \quad (9)$$

Where:

F_i: flap rate for prefix i

n_i: number of flaps (announcement–withdrawal cycles) for prefix i

T: time window (in days, hours, etc.)

Prefixes with high F_i values are candidates for route dampening.

Resource Load from Routing Updates

To estimate the load on routers due to frequent updates, we can define a simplified update processing load function, $L(t)$

$$L(t) = \alpha \cdot U(t) \quad (10)$$

Where:

$L(t)$: CPU/memory load due to routing updates at time t

$U(t) = A(t) + W(t)$: total number of routing updates at time t

α : constant reflecting processing cost per update

This equation links update activity directly to infrastructure demands, emphasizing the operational impact of rising churn.

Analysis and Performance of Prefix Withdrawal and Announcement Trends:

The graph graduated on time-series analysis on prefix announcement as well as prefix withdrawals in BGP routing tables and their rates between 2010 and 2024. There are two separate time-series curves: Prefix announcement and withdrawal rates. Prefix announcement rates are represented by the blue line where new prefixes get added on the global routing table and the red line is Prefix withdrawal where the table removes prefixes which were added and later withdrawn.

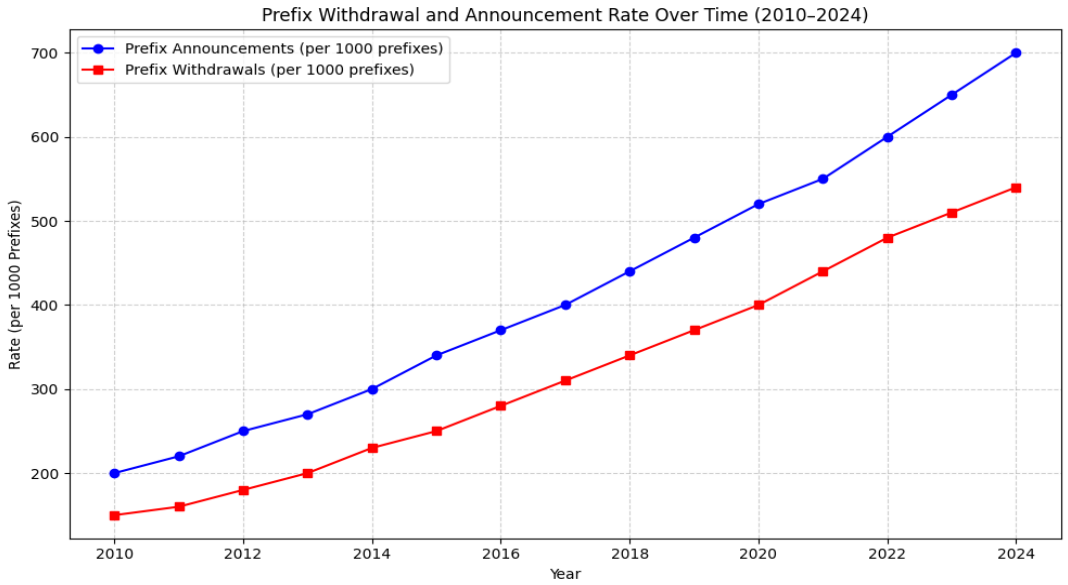

Figure 5: Prefix withdrawal and announcement rate over time (2010–2024)

Figure 5 illustrates both rates normalized per 1000 prefixes which allows the relative frequency of both rates to be measured along the same scale to the size of the routing table. Over the observed interval, both rates show an increasing trend and, thus, both rates capture the growing dynamics of the routing table. The rate of prefix announcements is increasing at a slightly faster rate than the rate of withdrawals, which indicates that the primary acquisition and retention of routing information is dense while the secondary retention of routing information is sparse, and, thus, increasingly slow. The predominant increase of both rates speaks to the increasing complexity of globally inter-domain routed networks, which has given rise to persistent prefix over-attachment, while older prefixes are temporarily withdrawn

due to network decommissioning or substantial re-architecting. The upward trend, along with both rates, indicates a hyper-corrective routing framework which needs to achieve some degree of routing stability as the system is being stressed to avoid this hyper-corrected state. This increase in routing slack is associated with underlying technical issues, as well as issues with network convergence time and accumulation of convergence over a network. The graph in addition depicts the period of Internet routing evolution when new routes are drafted to address increasing network traffic, and old routes are modified or removed, primarily focusing on the addition of IPv6, in conjunction with expansion of the world network.

5 Conclusion

This paper presents a thorough time-series study of the expansion and stability of BGP routing tables for fourteen years from Route Views and RIPE RIS data. What is uncovered is an ongoing and sustained increase in routing table size driven by factors including network growth, exhaustion of IPv4 addresses, and finer-grained traffic engineering policies. Both ARIMA and Prophet models could efficiently capture linear growth and seasonality of prefix count and produced robust short-term estimates for routing table growth. Further, the stability measures analysis revealed high heterogeneity among networks in a way that some Autonomous Systems (AS) had sharp instability due to misconfiguration, policy change, or unconfigured routing behavior Besides capturing development, the study contributes valuable information about the operational dynamics of routing infrastructure. Although the total churn rate is comparatively stable, the prevalence of short-lived prefixes and frequent updates indicate that challenges remain regarding the preservation of routing efficiency and security. The study uses statistical modeling of domain metrics to narrow the gap between theoretically predicted and practically applicable network monitoring. It builds upon existing research by providing a more complete picture by incorporating growth pattern studies, route stability, prefix instability, local AS behaviors, and the impact of route modifications. This research is most applicable to Tier-1 ISPs, network architects, and protocol designers interested in scalable and fault-tolerant Internet architectures. The implementation of real-time routing analysis anomaly detection to monitor for prefix hijacks or route leaks is recommended for further studies. Machine learning, particularly with LSTM and transformer models, would improve prediction values. The development of a predictive dashboard for network administrators, in tandem with the identified trends, would enable proactive response to routing anomalies. Understanding trends linked to IPv6, particularly how these trends pressure the existing legacy IPv4 infrastructure, will be more pertinent with the advance of deployment. In conclusion, this ongoing initiative is undoubtedly contributing to the broader objective of making the Internet more secure, stable, and scalable through the use of data-informed network management and policy formulation.

References

[1] Aravindh, G., & Sridhar, K. P. (2024). Resilient and Adaptive Secure Routing Protocol for Wireless Sensor Networks Using a Grey Wolf Optimizer and Lightning Search Algorithm. *Journal of Internet Services and Information Security*, *14*(4), 331-346. https://doi.org/10.58346/JISIS.2024.I4.020

[2] Bollen, E., Hendrix, R., Kuijpers, B., & Vaisman, A. (2021). Time-series-based queries on stable transportation networks equipped with sensors. *ISPRS International Journal of Geo-Information*, *10*(8), 531. https://doi.org/10.3390/ijgi10080531

[3] Bollen, E., Hendrix, R., Kuijpers, B., & Vaisman, A. (2021). Time-series-based queries on stable transportation networks equipped with sensors. *ISPRS International Journal of Geo-Information*, *10*(8), 531. https://doi.org/10.3390/ijgi10080531

[4] Bush, R., & Austein, R. (2013). *The resource public key infrastructure (RPKI) to router protocol* (No. rfc6810).

[5] Giotsas, V., Luckie, M., Huffaker, B., & Claffy, K. C. (2014, November). Inferring complex AS relationships. In *Proceedings of the 2014 Conference on Internet Measurement Conference* (pp. 23-30). https://doi.org/10.1145/2663716.2663743

[6] Lukić, M., & Đurić, D. (2023). Thermal comfort in Belgrade, Serbia: UTCI-based seasonal and annual analysis for the period 1991-2020. *Arhiv za tehničke nauke (Archives for Technical Sciences)*, *28*(1), 77-88. http://dx.doi.org/10.59456/afts.2023.1528.077L

[7] Maria, E., Sofia, K., & Georgios, K. (2025). Reliable data delivery in large-scale IoT networks using hybrid routing protocols. *Journal of Wireless Sensor Networks and IoT*, *2*(1), 69-75.

[8] Marwedel, R., Jacobson, U., & Dobrigkeit, K. (2025). Embedded systems for real-time traffic management: Design, implementation, and challenges. *SCCTS Journal of Embedded Systems Design and Applications*, *2*(1), 43-56.

[9] Mayilsamy, J., & Rangasamy, D. P. (2021). Enhancement of Energy Efficient Routing Scheduling Algorithm based on SDN Using IoT. *International Academic Journal of Science and Engineering, 8*(1), 10–18. https://doi.org/10.9756/IAJSE/V8I1/IAJSE0802

[10] Oliveira, R., Pei, D., Willinger, W., Zhang, B., & Zhang, L. (2009). The (in) completeness of the observed Internet AS-level structure. *IEEE/ACM Transactions on Networking*, *18*(1), 109-122. https://doi.org/10.1109/TNET.2009.2020798

[11] Petropoulos, F., Apiletti, D., Assimakopoulos, V., Babai, M. Z., Barrow, D. K., Taieb, S. B., ... & Ziel, F. (2022). Forecasting: theory and practice. *International Journal of forecasting*, *38*(3), 705-871. https://doi.org/10.1016/j.ijforecast.2021.11.001

[12] Prasad, L. C., Kamatham, Y., & Sunehra, D. (2024). Congestion-aware Weighted Sum Routing for Enhanced Network Lifetime in Cognitive Radio Sensor Networks. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, *15*(4), 325-334.

[13] Prasath, C. A. (2023). The role of mobility models in MANET routing protocols efficiency. *National Journal of RF Engineering and Wireless Communication*, *1*(1), 39-48.

[14] Rojas, C., & García, F. (2024). Optimizing traffic flow in smart cities: A simulation-based approach using IoT and AI integration. *Association Journal of Interdisciplinary Technics in Engineering Mechanics*, *2*(1), 19-22.

[15] Thiruvengadam, V., & Akinsorotan, A. M. (2022). Comparative and regression analysis of Nile tilapia, Oreochromis niloticus (Linnaeus, 1758) (Family: Chchlidae) from Bugungu, Eastern region, Republic of Uganda, East Africa (0 23'31.92" N 33 13'51.95" E). *International Journal of Environmental Research and Education*, *3*(1), 73-78. https://doi.org/10.70102/IJARES/V3I1/4

[16] Wang, W., Zuo, E., Chen, C., Chen, C., Zhong, J., Yan, Z., & Lv, X. (2025). Efficient time series adaptive representation learning via Dynamic Routing Sparse Attention. *Pattern Recognition*, *158*, 111058. https://doi.org/10.1016/j.patcog.2024.111058

[17] Xing, L. (2020). Cascading failures in Internet of Things: Review and perspectives on reliability and resilience. *IEEE Internet of Things Journal*, *8*(1), 44-64. https://doi.org/10.1109/JIOT.2020.3018687

[18] Xu, X., & Cui, W. (2024, July). Research on AdHoc network routing algorithm for link stability prediction. In *Journal of Physics: Conference Series* (Vol. 2807, No. 1, p. 012023). IOP Publishing. https://doi.org/10.1088/1742-6596/2807/1/012023

[19] Zargari, F., Aminpour, N., Ahmadian, M. A., Samimi, A., & Saidi, S. (2022). Impact of mobility on COVID-19 spread–A time series analysis. *Transportation research interdisciplinary perspectives*, *13*, 100567. https://doi.org/10.1016/j.trip.2022.100567

[20] Zheng, H., Cao, Y., Sun, D., Wang, M., Yan, B., & Ye, C. (2024). Research on time series prediction of multi-process based on deep learning. *Scientific Reports*, *14*(1), 3739.

Authors Biography

Layth Hussein Jasim is a researcher in the Department of Computers Techniques Engineering at the College of Technical Engineering, Islamic University in Najaf, Iraq. His academic work focuses on modern computing technologies, system design, and applied engineering solutions. He has contributed to several research initiatives aimed at enhancing computational performance and advancing technical methodologies. Layth is committed to supporting student innovation and promoting practical, research-oriented learning experiences. His research interests include computer systems, software development, and emerging digital technologies.

Dr.K. Rajamani is an Associate Professor (Senior Grade) at the Mepco School of Management Studies, Mepco Schlenk Engineering College, Sivakasi, Tamil Nadu, India. He is an accomplished academic with extensive experience in management education, research, and institutional development. Dr. Rajamani has contributed to numerous research publications and actively participates in conferences, workshops, and industry–academia collaborations. His teaching emphasizes strategic management, organizational development, and managerial skill enhancement. His research interests include leadership studies, business strategy, and contemporary trends in management practices.

R. Gokilavani is a faculty member in the School of Commerce, Finance & Accountancy at CHRIST University, India. She is dedicated to academic excellence in commerce and finance education, focusing on developing strong analytical and professional skills among students. Her academic contributions include involvement in research initiatives, curriculum development, and student mentorship. Gokilavani's teaching emphasizes practical financial understanding, accounting expertise, and industry-relevant knowledge. Her research interests include finance, accounting practices, and emerging trends in commerce education.

Dr. Adireddy Ramesh is a Professor in the Department of Electrical and Electronics Engineering at Aditya University, Surampalem, Andhra Pradesh, India. He is a distinguished academic with extensive experience in teaching, research, and technological development in the field of electrical engineering. Dr. Ramesh has published numerous research papers and actively contributes to academic and industry-focused initiatives. His teaching emphasizes conceptual depth, practical application, and innovation-driven learning. His research interests include power systems, renewable energy technologies, electrical machines, and advanced control systems.

B. Thamizhkani is an Assistant Professor in the Department of Electrical and Electronics Engineering at New Prince Shri Bhavani College of Engineering and Technology, Chennai, India. She is committed to delivering quality engineering education and fostering strong technical skills among students. Her academic role includes guiding student projects, supporting departmental initiatives, and contributing to research-oriented activities. Thamizhkani's teaching focuses on core EEE subjects, practical applications, and emerging trends in electrical engineering. Her professional interests include power electronics, control systems, and modern electrical technologies.

Dr.K. Senthilvadivu is a Professor in the Department of Mathematics at K.S. Rangasamy College of Technology, Tiruchengode, India. She is a distinguished academic with extensive experience in teaching, research, and mathematical analysis. Dr. Senthilvadivu has contributed to numerous research publications and actively participates in academic forums, conferences, and collaborative projects. Her teaching emphasizes analytical thinking, problem-solving, and strong mathematical foundations for engineering students. Her research interests include applied mathematics, numerical methods, mathematical modeling, and interdisciplinary applications of mathematics in science and engineering.