Country-wide Long-term Event Detection and Classification Mechanisms Using Spatiotemporal BGP Prefix Data

S. P. Meenakshi, *Member, IAENG*, M. J. S. Raman and V. Kamakoti

***Abstract*—In a network of autonomous systems (ASes), prefix reachability can be affected by events such as link and node failures, router misconfiguration, route flaps and prefix hijack attacks. Furthermore, with the increase in global organizations deploying their data centres and content services in other countries that have good networking infrastructures, studying prefix reachability is of significance. Since detecting these events at the end user level is challenging, by monitoring the spatially distributed routing table features, such as AS path distributions and spatial prefix reachability distributions, the events can be detected at the control plane level. In this work, a measure and a method to detect long term events at a country level AS topology are proposed. To understand the occurrence of control plane level events, temporal pattern analysis over the distributed peer prefix announcements of a country-level AS topology was conducted. Five Asian countries are considered as a representative set to study the occurrence of events. To capture and measure the spatially occurring events in a single temporal pattern, we proposed a counting-based measure using prefixes announced by *x* % of *n* spatially distributed peers. Our method to detect events employs mean change point detection technique with normal and CUSUM test statistics over the proposed measure. Other statistical techniques such as regression estimation and K-means clustering are used in our method to quantify the impact and duration of long-term control plane events. The detected events are validated using average path pattern correlation and Fisher scores of different path length features. We also validated the events using the SEA-Me-We4 cable cut event manifestations. The comparison results with other event detection techniques demonstrate the efficacy of the mean change point technique with normal assumption used in our method.**

***Index Terms*—Long Term events, BGP Prefix change, X percentage of N Peers, Mean change point detection, K-means Clustering**

I. Introduction

With global organizations expanding the deployment of data centres and content services outside their countries, determining prefix reachability that ensures 24x7 availability to these servers from global destinations is important. The impact of prefix reachability due to Internet outages and other events affects bandwidth and latency. This situation is already causing concern for business organizations [1]–[3]. For instance, in June 2012, the Sea-Me-We4 submarine cable interconnecting India, Europe and the Middle East was cut between Malaysia and Thailand [3]. During this period, traffic was diverted using an alternate route, which increased the latency of data packets.

Prefix non-reachability occurs due to events such as link and node failures, router misconfiguration, route flaps and prefix hijack attacks on the Internet autonomous system (AS) level topology. Such events have an impact on the control plane, whereby AS paths that are used by routing protocols, such as Border Gateway Protocol (BGP), can be affected. When the data plane traffic is unable to reach the intended destination, the outages are immediately perceived and reported by the user. However, in the control plane, any event that increases the latency leading to degraded performance is generally not perceived by the user. To capture such an event, a geographically distributed control plane monitoring infrastructure is required. By monitoring the control plane features, such as AS path length distributions, spatial prefix reachability distributions and covering to overlap route ratio, we can detect control plane disruptions. In our work, we consider a simple and computationally efficient feature, namely, spatiotemporal prefix reachability, which encompasses the potential spatiotemporal events.

Using the control plane data, we perform a detailed analysis on the manifestation of events over different spatiotemporal aggregation levels. We select the best aggregation as our measure for event detection. Identifying such events and their impacts on the originating prefixes of a country is crucial for successful business in the content, data and call-centre industries. The impact could exist for either the short term or the long term. An event that has a performance impact for less than a day is classified as a short-term event, whereas if the performance impact of an event persists for more than a day, then it is classified as a long-term event. Such a classification can be achieved by using the prefix count time series. However, prefix count time series data labelled with country-level short- and long-term events are not readily available. Hence, supervised classification methods that depend on labelled training data to construct the class models cannot be used in this case. Therefore, we have to depend on unsupervised learning methods that leverage the mined knowledge about the events from available control plane data.

In particular, the mining technique called clustering is used to extract the knowledge in our work. The K-means clustering algorithms is employed to segment the first difference of the prefix time series data to learn about the events based on their impacts. We use the first difference of the prefix time series data to remove the trend information. After removing the trend information from the time series, the K-means algorithm is applied on the first differenced prefix count to extract the knowledge on events. This knowledge is used in

Manuscript received February 11, 2017; revised May 24, 2017.

S. P. Meenaksh is with the Department of CSE, IIT Madras, Chennai, Tamilnadu, India and with the School of Computing Sciences, VIT University, Vellore, Tamilnadu, India. (email:spmeenakshi@vit.ac.in)

M. J. S. Raman and V. Kamakoti are with the Department of CSE, IIT Madras, Chennai, Tamilnadu, India.

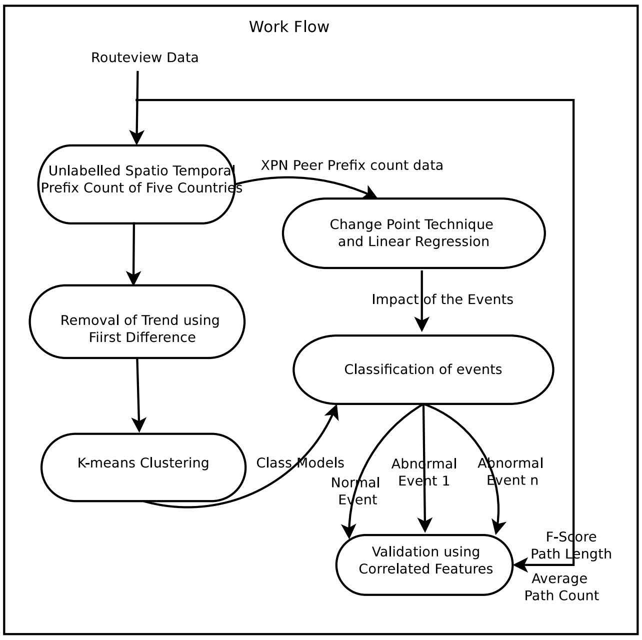
Work Flow
Routeview Data
Unlabelled Spatio Temporal Prefix Count of Five Countries
XPN Peer Prefix count data
Change Point Technique and Linear Regression
Removal of Trend using Fiirst Difference
Impact of the Events
Classification of events
K-means Clustering
Class Models
Normal Event
Abnormal Event 1
Abnormal Event n
Validation using Correlated Features
F-Score Path Length
Average Path Count

Fig. 1: Event Detection and Classification Process Workflow

the classification process, where the events are detected using the mean change point detection technique and regression estimation that are part of the proposed method. The workflow of our event detection and classification process is shown in Fig. 1.

The control plane data source considered in this work is publicly available route-view data [4], which is a collection of route information bases (RIBs) retrieved from different geographically distributed peers. The prefix reachabilities for the Indian AS topology and for four other Asian countries, namely, China, Japan, South Korea and Taiwan, were extracted from the retrieved data, and various aggregation levels were studied to select the suitable one for event detection.

As a case study, we rigorously analysed the temporal prefix reachability patterns for the Indian AS as follows: the control plane data were retrieved from all spatially located peers of the route-view data, as shown in step 1 of Fig. 1. A spatiotemporal prefix measure from X % of N (XPN) peers that captures the Internet-wide control plane activities using a single pattern was proposed. The proposed measure was applied with cumulative sum (CUSUM) and the normal statistics-based mean change point detection technique to detect long-term events. Because prefix announcement deviant behaviour is the basis for detection, the mean change in the pattern is considered to be an indication of a long-term behavioural change. The actual trend of the proposed measure estimated through regression and the mined cluster knowledge on impact was used to classify the event to the appropriate impact classes. We also validated the method applied on the proposed feature using the average path pattern, Fisher scores on path lengths and media event reports. Similar analyses were conducted for other countries using the identified potential event possessing spatiotemporal prefix patterns.

The remainder of this paper is organized as follows: Section II presents the BGP anomalous behaviour studies reported in the literature for different large-scale Internet events. Section III discusses route-view data and the country-level analysis and mining of events from country-level AS topologies. Section IV discusses the feature selection criteria. A detailed empirical study on country-level prefix data is presented in Section V. Section VI presents the proposed method. In Section VII, we discuss the validation and performance of our method. We present the conclusions and directions for future work in Section VIII.

II. Related Work

As discussed previously, control plane data have been used to detect events such as prefix non-reachability. In the literature, the deviant behaviour that occurs with the attributes of BGP has been studied to detect abnormal events. For example, the BGP updates and routing tables can be monitored to obtain relevant attributes for the analysis or detection of an event. Since obtaining live updates can be difficult, one can also use the data archived at RIPE and route-view monitors to verify that such attributes can be used to detect events based on the problem under consideration. In these cases, the events that have occurred are already known; hence, a correlation between the deviant behaviour and events can be established. We will now review such works available in the literature.

The studies on BGP instabilities and dynamics [5], [6] considered 10 distinct BGP attributes, such as AADiff, WADiff and WADup, to summarize the BGP dynamics. In these methods, the current and decade-old BGP dynamics were compared using the ten distinct attributes. It was found that the healthier forwarding dynamics are dominant and that the associated updates are 84 % of the total dynamics. At the same time, the updates related to pathological behaviour are 15 %. Using these results, they concluded that the BGP infrastructure is healthier today than it was a decade ago. One of the attributes used in the study is the prefixes that are announced or withdrawn by BGP. Upon the occurrence of any unusual event, BGP activities exhibit major changes in the number of announced or withdrawn prefixes from the ASes. This results in statistical changes on forwarding and pathological updates. The work on the BGP routing stability of popular destinations [7] showed that most of the update events are concentrated on few prefixes that do not receive much traffic. The majority of the popular prefixes tend to have shorter events in addition to having fewer events. Moreover, 0.1 % of all prefixes had an average event duration exceeding 40 seconds due to unreachability of the prefixes.

In another case, worm attack events [8] in two edge ASes with less than 0.25 % prefix announcements of the routing table entry were observed with a surge in the total update messages by 6 % at the monitoring vantage points (VPs). The origin AS change attack event was detected using an information visualization technique [9] that integrates mechanical analysis and human intelligence to improve the detection efficiency. Visualization-based BGP anomaly detections in real time and using mining techniques were presented in [10]–[12]. The misconfiguration event studied in [13] revealed that the origin misconfiguration event includes failure to summarize address space, prefix hijacks and propagation of private prefixes. Policy changes such as traffic engineering, multihoming and provider switch-induced new routes last longer than new routes induced by misconfiguration. The failure (link/node) signature is also the opposite of the signature of the misconfiguration. The forensic framework proposed in [14] applied data mining techniques on the labelled updates to extract rules for detecting normal and abnormal events, such as worm attacks and blackouts.

The three-step dynamic threshold-based method proposed in [15] to cluster BGP updates into routing events showed how the threshold values learned dynamically in the refined clustering step enhanced the clustering of events. The time-based dynamic and static thresholdings were leveraged in [7], [15], [16] for locating the routing events from the BGP update messages. The time threshold value was used to cluster consecutive BGP updates of the same prefix into one routing event. The number of routing events, prefixes involved in the routing events and duration of the routing events were reported in these studies. Prefix hijack event detection was studied in [17]–[20]. The detected prefix hijack and topology disorder anomalies were further analysed in [21] to establish their reliability by correlating them with other publicly available information, such as Internet exchange points (IXPS), which is the database and looking glass of an AS. The improved outcome is then provided as input to their classification algorithms. An online mechanism to detect BGP instability events pertaining to the router level was proposed in [22]. The features extracted from the update messages were considered to be a time series. The adaptive sequential segmentation algorithm was employed on the time series data to detect instability events.

In the process of understanding the BGP router requirements in the future scalable environment, [23], [24] investigated the trends in BGP update growth (churn). In particular, the authors in [23] considered the BGP update temporal data as a time series and measured its various components. The studies [16], [25]–[28] analysed streams of BGP update messages from vantage points throughout the Internet, with the objective of inferring the root cause and source of the routing changes. Route flapping, peer session loss, policy changes and link failures were the events inferred as major triggers for AS route changes. The changes could be detected from the withdrawn and announced prefixes of the update messages. The suspect set of ASes that induce route changes were narrowed down in the work [25] by correlating the update messages from different VPs. The root cause event that might have triggered the changes was also deduced from the suspected set of AS updates.

A. Contributions of Our Work

1) Although our work adopts the definitions for events given in [23], our focus is on detecting and quantifying long-term events that have an impact on the originating prefixes of ASes at a country level. A spatiotemporal originating prefix count measurement is extracted from bi-hourly BGP routing table data for further processing. The primary benefits of considering routing table data are reduced computational complexity due to the lower volume of data compared to the BGP update message volume and the elimination of processing duplicates associated with the update messages.
2) Rather than correlating the VP measurements, we consider aggregated spatiotemporal prefix count statistics, particularly 95 % of the stable prefixes from XPN VPs (stable prefixes as referred to in [25]) over which change point detection statistical techniques are applied to identify and quantify long-term events. We argue that to detect long-term events, the prefix count at day-level granularity is sufficient rather than processing highly frequent update messages at second-level granularity (less than the route convergence time).
3) The detected events are validated through inference using the changes in AS path features. The changes in AS path features are based on the long-term event specifications presented in [13], [26]. In [26], the events were inferred from the path patterns clustered using prefixes and peers. In our work, the inferences are based on the frequent path patterns clustered based on the originating AS. The originating AS-based frequent path clustering helps to identify the affected country-level ASes with minimal computational requirements.
4) In the work performed in [29], a mechanism to measure the impact of various Internet earthquakes was proposed. Multiple BGP attributes and temporal aspects were considered to cluster the data as normal and abnormal. With respect to a normal baseline value, the impact of various known Internet earthquakes was measured for the known event duration. Our work differs from this previous study by employing automatic event detection of the a priori unknown events and its duration and impact quantification using spatiotemporal data. The originating prefix changes for a country level are measured using aggregated spatiotemporal prefix counts to quantify the effects of long-term events.
5) It has been established that finding a plausible reason for every event from unlabelled spatiotemporal prefix data is difficult and error prone [23], [30]. Therefore, we propose mechanisms to detect the durations and quantify the impact of the events that can play the role of an indicator to understand the country-level reachability. For this purpose, the feature selection F-score measure used in [22], [31] is used in our work to identify AS path features for inferring the events.
6) In-depth knowledge of country-level reachability is essential for organizations that are deploying their service-based branches in other countries. The majority of the events reported in the literature are short-term events related to route convergence or persistent route oscillations [7], [13], [32]. To the best of our knowledge, this work is the first attempt of long-term event detection and inference at a country level using spatiotemporal feature data. Our work can be considered as an extension of [29].

III. Route-view Data and Country-level Analysis

Our intuition for detecting the events is based on spatiotemporal variations. Spatiotemporal variations could either be due to a volume measure, such as prefix count change, or could be reflected in a short-term or long-term impact on the AS path. Therefore, we extract the features relevant to prefix changes and the AS path changes and measure the features. This measurement provides an indication of an event that has occurred at the control plane. Therefore, we use geographically distributed route-view monitor data.

TABLE I: Peer Spatial Location Distribution

Asia	Europe	North America	Australia	Africa
2	10	23	1	2

Route-view monitors collect and archive route information retrieved from a maximum of 37 VPs placed across 34 agreed upon ASes of different geographic locations. VPs [33] are peer routers that are assigned by the ASes to show the routes to the collector for which they are advertising. In this work, we use the term peers synonymously with VPs. The BGP feeds collected between every two-hour interval are stored as time-stamped snapshots in the route-view monitors [4]. The attributes of interest to our work are prefixes, next-hop and AS path. In the archive, the prefixes are entered in ascending order of their numerical values. The next-hop attribute indicates the IP address of the peers assigned by the associated AS from which the monitoring point collects the data. If there are two different next-hop addresses to an AS, then it is interpreted as the BGP feeds are collected from two AS border routers (ASBR) of the corresponding AS.

The AS path entries for each prefix are sequences of ASes starting with the peer AS, followed by intermediate ASes in the path and ending with the AS that originates the prefix. The route entries for a prefix by different peers are arranged consecutively in the snapshot. This order helps to count the number of peers from which a prefix can be reachable to an originating AS during the snapshot period. Furthermore, the peers can be grouped geographically based on continent. The distributions of peer locations in five continents extracted from the data are presented in Table I. The continents Asia, Australia and Africa have less than ten percent of the total peers, whereas South America has no peer representation. In the Internet, there are hundreds of ASBRs interconnected with each other. However, the route-view data are presented from 37 routers. Nevertheless, to demonstrate our technique, we use the existing peer representations from the route-view data. Since the number of routers is a small percentage of the hundreds of ASBRs and because we can still detect events using this method, we can safely state that we obtain a lower bound on the actual occurrence of events.

Grouping the peers geographically can lead to certain constraints. Such captured spatiotemporal events could be proportional to the number of peers. Hence, our proposed method expects the peers to be evenly distributed among the geographic locations. In practice, this is not the case. As previously mentioned, more peer representations from different geographical locations might provide a better view on the events occurring in those regions.

Five representative Asian countries, namely, India, China, Japan, South Korea and Taiwan, are considered for our country-level analysis. The rationale for selecting these countries is as follows: first, two of the countries possess larger economies, and the next three countries are technologically advanced nations. To extract the originating prefixes of the respective country-level AS topologies, the list of AS numbers assigned to the associated countries is necessary. We have used the methodology specified in [34] to obtain the assigned AS numbers from the APNIC [35] Regional Internet Registry (RIR). The APNIC RIR contains the static record of AS numbers and allocated IPv4 and IPv6 prefixes for countries in the Asia-Pacific region. The extracted country-level AS numbers [36] are used to filter the prefixes that have these AS numbers at the end of the AS-PATH attribute. The unique prefixes counted at the five-country level and at the global level are presented in Table II. As shown in this table, 13 % of the global prefixes are associated with these five Asian countries, and the peer count varies across countries. The variation in the peer count may be attributed to the policies exercised by the peers in announcing the prefixes of different countries.

Daily snapshots that are time stamped with 0000 hours from 01/01/2012 to 31/12/2012 were considered in this work. The unique prefixes announced at the beginning of the year from various peers to the ASes of the five countries as well as the global level are presented in Table II. For the global AS level, the unique prefixes of the complete snapshot are considered. A maximum of four peers of the total of 37 peers announce less than 50 % of the total unique prefixes. Excluding the prefixes announced by these peers, the average number of prefixes announced from the remainder of the peers and the standard deviation in the prefix announcement are computed. The peers announce an average of 3,80,797 unique prefixes with a standard deviation of 4,150 prefixes at the global level. With respect to the Indian AS country-level topology, the average unique prefixes announced by the peers are 16978 with a standard deviation of 565. For the other countries, beginning with China, the average unique prefixes announced from the peers are 9531 with a standard deviation of 306; for Japan, the average unique prefixes announced from the peers are 5742 with a standard deviation of 81; for South Korea, the average unique prefixes announced from the peers are 14015 with a standard deviation of 316; and for Taiwan, the average unique prefixes announced from the peers are 3062 with a standard deviation of 138. From the standard deviations of the prefix announcements of peers at the country level, we can infer that prefix announcement peer policies vary across countries. This result leads to a further inference that during events, the impact on the prefix announcements would also vary across countries. Hence, a country-level analysis of prefix announcements from various peers is necessary to quantify the prefix reachability impact of the events. This inference is further supported by the temporal unique prefix announcements and total AS paths (aggregated from all peers) of different countries.

Statistical characterizations of the temporal prefixes and AS path features of the countries were performed, and the results are presented in Table III. Each country has different statistical measurements for both of the features. Among the five countries, India is relatively high in value in the considered statistical aspects except for correlation. Similarly, Taiwan is less in value in all the statistical aspects. AS path deviations are more or less equal to the prefix deviations for all the countries except for Taiwan, which has fewer path deviations compared to prefix deviations. To further understand the temporal relationship between AS paths and prefixes, the correlation statistics were computed. The correlation between AS paths and prefix counts is higher and similar to India, Japan and South Korea and higher than China and Taiwan. The positive correlation indicates that prefix and AS path features together either increase or decrease temporally.

TABLE II: Announced Prefix Counts: Global and Five Asian Countries

Date	Type	Peer Count	Peers	Unique Prefix Counts
01-01-2012	Global	37	286 293 701 812 852 852 1221 1239 1299 1668 2152 2497 2905 2914 3130 3130 3257 3303 3356 3549 3549 3561 5056 5413 6539 6762 6939 7018 7660 8001 8492 11537 11686 13030 22388 31500 39756	380992 384469 379432 373237 381070 381067 381843 379792 376669 379936 382751 381453 **2474** 380729 381143 381135 380480 **150422** 378602 380738 378630 379754 381341 380096 379826 381433 382521 379882 384557 364665 385139 **13435** 388438 382100 **13575** 389785 382585
	India	37	286 293 701 812 852 852 1221 1239 1299 1668 2152 2497 2905 2914 3130 3130 3257 3303 3356 3549 3549 3561 5056 5413 6539 6762 6939 7018 7660 8001 8492 11537 11686 13030 22388 31500 39756	17135 17160 17136 17139 17132 17132 17150 17133 14915 17135 17141 17153 **1** 17134 17150 17150 17148 **4677** 17132 15565 17134 17133 17103 17115 17116 17134 17134 17134 17178 15256 17158 **289** 17159 17149 **289** 17472 17148
	China	36	286 293 701 812 852 852 1221 1239 1299 1668 2152 2497 2914 3130 3130 3257 3303 3356 3549 3549 3561 5056 5413 6539 6762 6939 7018 7660 8001 8492 11537 11686 13030 22388 31500 39756	9533 9896 8997 9512 9567 9567 9909 8999 9521 8998 9872 9894 9516 9529 9529 9465 **5944** 8982 9533 9533 8999 9530 9524 9508 9535 9555 9530 10249 9095 9771 **583** 9899 9532 **581** 9898 9551
	Japan	36	286 293 701 812 852 852 1221 1239 1299 1668 2152 2497 2914 3130 3130 3257 3303 3356 3549 3549 3561 5056 5413 6539 6762 6939 7018 7660 8001 8492 11537 11686 13030 22388 31500 39756	5752 5764 5752 5742 5752 5752 5756 5752 5751 5752 5767 5772 5746 5752 5752 5752 **4489** 5752 5757 5757 5747 5752 5752 5753 5752 5758 5752 5813 5292 5753 **668** 5767 5752 **668** 5767 5752
	South Korea	36	286 293 701 812 852 852 1221 1239 1299 1668 2152 2497 2914 3130 3130 3257 3303 3356 3549 3549 3561 5056 5413 6539 6762 6939 7018 7660 8001 8492 11537 11686 13030 22388 31500 39756	13988 14005 13986 **8742** 13999 13999 14007 13986 13986 13986 14006 14000 14272 14270 14270 13990 **4389** 13994 13994 13995 13985 13994 13989 13997 13992 13998 13985 14017 12696 14003 **457** 14006 13997 **457** 15092 14001
	Taiwan	36	286 293 701 812 852 852 1221 1239 1299 1668 2152 2497 2914 3130 3130 3257 3303 3356 3549 3549 3561 5056 5413 6539 6762 6939 7018 7660 8001 8492 11537 11686 13030 22388 31500 39756	3024 3279 3000 2993 3040 3040 3078 2997 2993 2985 3266 3093 2916 3001 3001 3009 **1931** 3000 3004 3004 2985 3000 3025 3032 3076 3041 3000 3371 2724 3359 **316** 3266 3025 **316** 3372 3040

Using this characteristic, the normal and abnormal events can be identified in a straightforward manner by employing threshold values. Other than the positive correlation, in all other cases, the relationship between prefixes and AS paths is unclear. To identify the occurrence of events in these cases, higher granularity details are required for both prefixes and AS paths. Identifying higher granularity details, computing the duration of the occurrence of events and determining whether an abnormality is caused due to single or multiple events are some of the challenges associated with this problem. We address the first two challenges in this work. We propose a spatiotemporal measure on prefix count announced by XPN peers to identify the events. The change point detection technique is used to measure the duration of events, and linear regression is employed to quantify the impact of the events. We also use AS path changes and various path length F-scores [31], [37] to validate the occurrence of events.

The temporal relationship between spatially aggregated unique prefixes and AS paths for the five countries is shown in Figs. 2, 3, 4, 5 and 6. The positive correlation between AS paths and prefixes computed and given in Table III is visually observable in these figures. Moreover, a linear trend in prefix count is identifiable using the fitted regression line for all the countries. The temporal prefix growth rate for each country is computed and shown in the Table III. For the year 2012, India has the highest prefix growth rate among the countries. In addition to the linear trend, seasonal variations are found in the prefix graph. In the case of the AS path feature, for the time duration between 1 and 200 days, long-term path drops are manifested in the graphs for all the countries. Particularly, the path drops between 120 to 180 days have similar patterns for all the countries. Three types of relationships are observed between AS paths and prefixes from the graphs. In the first type of relationship, both features increase or decrease together. In the second type, when there is no seasonal variation in the prefix count, there is still an increase in the path count. Finally, in the third type, the path count decreases while there is no seasonal variation in the prefix count. As mentioned previously, the second and third cases require further exploration using higher granularity details to understand the events.

With respect to the Indian AS topology, we observed significant changes in the unique prefix counts between the beginning and end of the year. There is a 10 % increase in the unique prefix count at the end of the year, which is an indication of a positive growth trend. In the case of unique prefixes announced by $>$ 90 % of peers, it is 76 % of the total unique prefixes at the beginning of the year and 60 % at the end of the year. There is a 16 % decrease in prefixes announced by $>$ 90 % of peers over this one year period. However, the prefixes announced by the 81-90 % peer range accounts for 22 % of the total unique prefixes at the beginning of the year and increases to 39.5 % of the total unique prefixes at the end of the year. This peer range registers an increase of 17.5 % over the span of one year.

When the prefix counts of both the $>$ 90 % and 81-90 % peer ranges are considered, it is 98 % of the total unique

TABLE III: Country-Wise Temporal Statistics on AS Paths and Prefixes

Country	AS Paths		Prefixes			Cor(AS Paths, Prefixes)
	Mean	Std Deviation	Mean	Std Deviation	Growth Rate	
India	561088.5	17961.18	18222.72	622.58	5.54	0.76
China	314179.4	12146.69	12727.94	487.84	2.7	0.34
Japan	192621.8	6481.77	6096.83	228.47	1.73	0.78
S.Korea	451173.8	13068.36	15308.4	449.69	1.73	0.78
Taiwan	102432.5	2493.30	3596.10	115.23	0.78	0.43

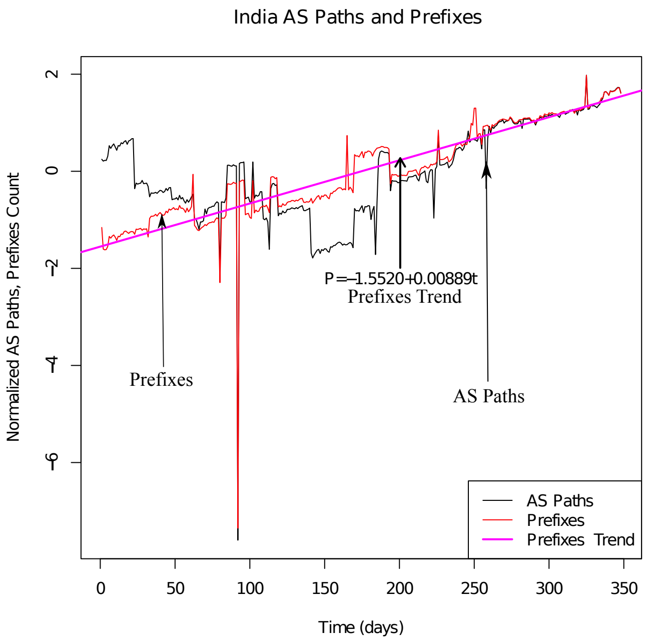

Fig. 2: Temporal Dynamics of India AS Paths and Prefixes

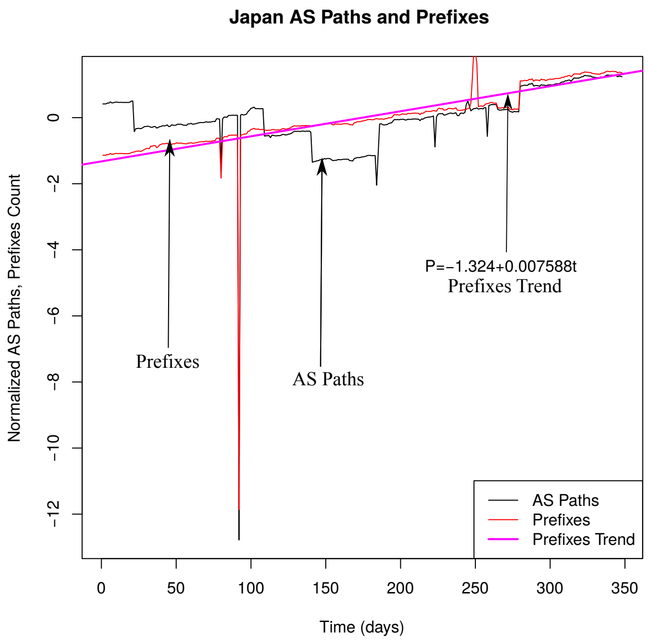

Fig. 4: Temporal Dynamics of Japan AS Paths and Prefixes

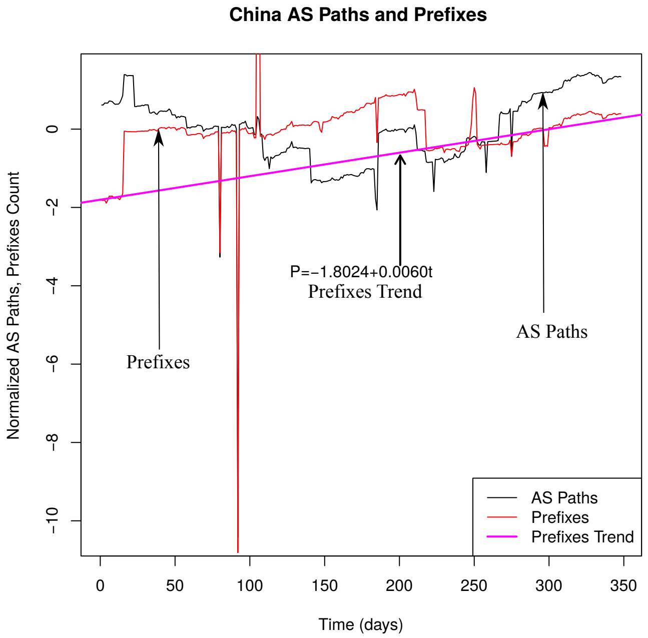

Fig. 3: Temporal Dynamics of China AS Paths and Prefixes

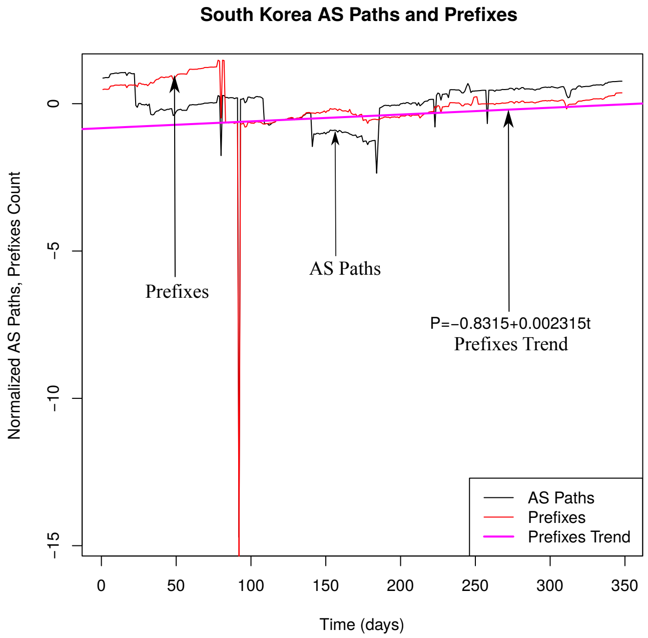

Fig. 5: Temporal Dynamics of South Korea AS Paths and Prefixes

prefixes at the beginning of the year and 99.5 % of the total unique prefixes at the end of the year. The combined prefix counts exhibit less difference than when the two ranges are considered separately. From the observed combined prefix statistics, we can infer that the increase in combined prefixes results from a significant increase in the 81-90 % peer range and that a few percent of prefixes shift from the > 90 % peer range. This shift in prefix count should have been caused as a result of events that occurred in the AS topology and the peer dynamics.

To understand the contribution of peer dynamics in the

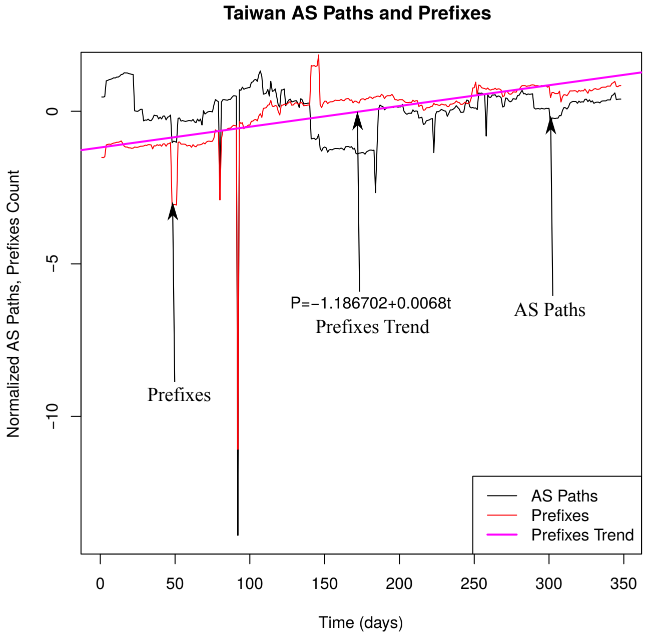

Fig. 6: Temporal Dynamics of Taiwan AS Paths and Prefixes

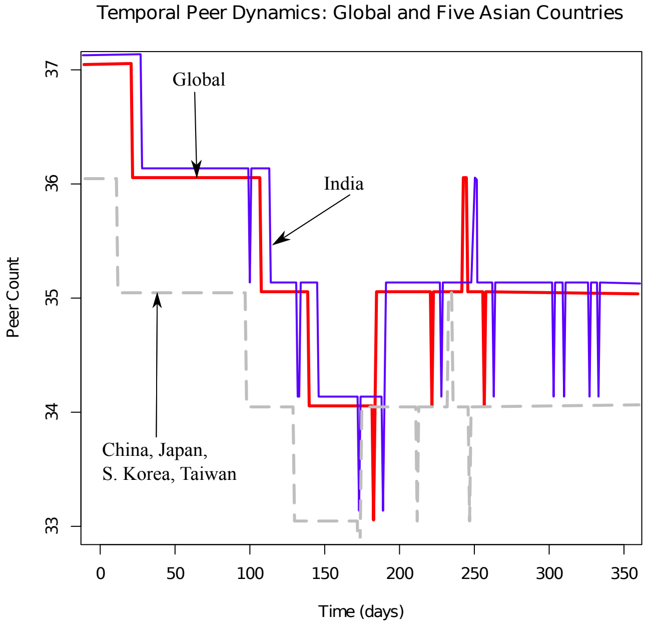

Fig. 7: Temporal Peer Dynamics: Global and Five Asian Countries

prefix announcement temporal variations, we analysed the number of peers involved in the daily snapshots. It was found that the number of peers vary temporally in the snapshots. This result might be due to peering failures induced by congestion, route flaps or policy changes. A maximum of 37 peers were observed in a span of one year.

We also observed from the route-view snapshots that the peers were dropped or added with varying time intervals. For instance, the peer AS located in Romania, AS 39756, ceased peering from the 27^{th} of January, 2012, to the 28^{th} of August, 2012, and resumed peering on the 29^{th} of August, 2012. The peer AS 3741 located in South Africa started peering only from 22/07/2012. The peer's temporal dynamics with respect to the global AS topology and to the AS topologies of the five Asian countries is shown in Fig. 7.

As shown in Fig. 7, there is an occasional peer count drop of one to the Indian AS topology but not at the global level. For the other countries, the peer drops are similar to the global peer drops but with less than one in the total peer count. After analysing the data, the peer involved in this drop for the Indian AS topology was identified as AS 2905 located in South Africa. This peer announces only one prefix, which belongs to Google corporate network AS 45566, and it occasionally drops the prefix announcement. This AS does not announce any prefix for the other four countries. Other than these add/drop peer dynamics, few peers announce a very small percentage of prefixes. Overall, we can observe that either one peer is occasionally added or dropped. The peer dynamics might have an impact on the increase or decrease in the total AS path. However, unique prefix counts are computed by considering the prefix announcements of all the undropped peers (union of prefix sets of each undropped peer); hence, the impact of peer drops reflected in the unique prefix count is very low.

In our work, we first establish the occurrence of events through the prefix count patterns of different peers and propose mechanisms for event detection, quantification and validation. To justify our proposed method for country-wide event detection and quantification, we present a detailed analysis and results on the Indian AS topology data and present the comparison results of other countries. We explore the unique prefix counts extracted from 37 peers individually for India and for the global level to establish the occurrence of different types of events in the Indian AS topology. Both the counts were compared after performing normalization. In the normalized scale, 25 of 37 global-level prefix temporal patterns exhibit a linear trend with only one significant drop. This drop occurred on the 17^{th} of April, 2012. The remainder of the peer prefix temporal patterns exhibited different increases and drops. In the case of India, the extracted prefix count pattern varies significantly in 17 of 37 peers by exhibiting different increases and drops. At the same time, the remaining 20 peer prefix announcements have similar temporal linear trend patterns with three visibly significant variations. Among them, the first is a long-term drop from the 15^{th} of March to the 5^{th} of April, 2012. The second is a stochastic drop that appears on the 17^{th} of April, 2012. The third is another long-term increase that appears between the 7^{th} of July and the 29^{th} of July, 2012. The second and third events coincide with the time of a cable cut that occurred in the SEA-Me-We4 submarine cable. We refer to these patterns (20 of 37) as dominant patterns, and we refer to the peers associated with it as the majority set (20 peers out of 37). This set is placed in Group I, and the remainder of the peers with different prefix temporal patterns are placed in Group II. The Group I and Group II peers are presented using their AS numbers along with their geographical locations in Table IV.

The representative temporal prefix count patterns for the global, Group I and Group II categories for India are presented in Fig. 8. The prefix temporal patterns of peers belonging to Group I exhibit high similarity in time and volume; hence, they are assumed to be caused by the same type of events. However, the prefix temporal patterns of Group II vary in time and volume and are therefore assumed to be caused by different types of events.

TABLE IV: Peer AS Groups and Locations

Type	Peer's ASes and Locations
Group I	1221 (Australia), 2497 (Japan), 293, 701, 1239,1668, 2152, 2914, 3130, 3356, 5056, 6939,7018 (USA), 852, 6539 (Canada), 286 (Netherlands), 3257 (Germany), 5413 (UK), 6762 (Italy)
Group II	7660 (Japan), 3549, 8001, 11537, 11686, 22388 (USA), 812 (Canada), 1299 (Sweden), 8492, 31500 (Russia), 3303, 13030 (Switzerland), 39756 (Romania), 2905, 3741 (South Africa)

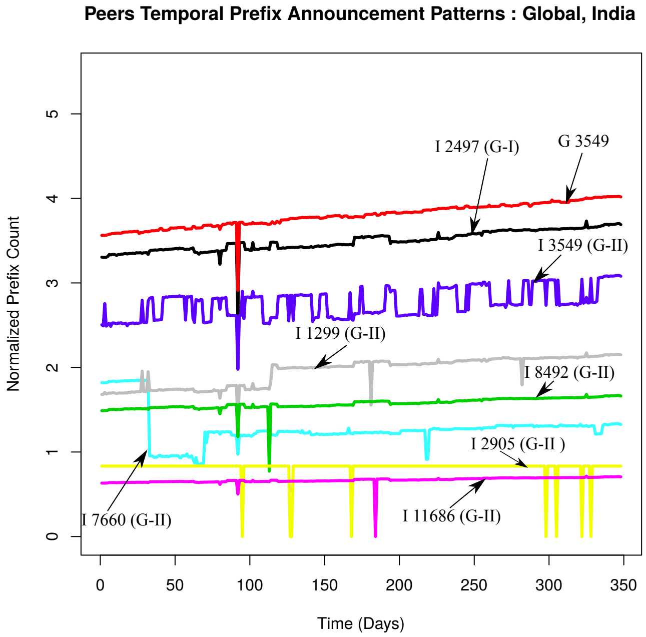

Fig. 8: Prefix Announcement Temporal Dynamics From Different Peers

The varying temporal patterns of individual peers lead to the intuition that the stochastic and long-term prefix drop or increase occurs network wide with different magnitudes (prefix counts) due to various events. The events are manifested in one or multiple peer temporal patterns. The global prefix temporal pattern is shown using the aggregated unique prefixes announced to all the countries. From the global and individual peer patterns, we can observe that events are clearly manifested in individual peer patterns rather than in the aggregated global-level prefix pattern. Rather than exploring all the individual peer prefix temporal patterns to identify the abnormal events, we sought to find an aggregate prefix pattern that manifests the maximum number of events. From a country-wide perspective, we are interested in finding answers to the following questions to understand the stability and robustness of the topology.

- How do the prefix announcements vary temporally and across peers?
- What is the percentage of unique prefixes that are reachable to a country-level AS topology from the maximum number of peer locations?
- In what percentage of peers is the maximum temporal variation manifested? (to aggregate the prefixes from the percentage of peer announcements)
- How are long-term events detected and quantified from the aggregated prefix temporal variations?
- How can the detected events be validated?
- How does the duration between successive long-term events vary across countries?

The statistical characterization of aggregated AS paths, prefixes of different countries and the global as well as country-level individual peer prefix temporal pattern analysis lead to the intuition for proposing a spatiotemporal prefix counting measure (specific percentage of aggregation) that could be used to find answers to some of the aforementioned research questions.

IV. Measure Selection Criterion and Its Characteristics

In this section, we discuss the rationale behind choosing the spatiotemporal aggregated measure [38], particularly the prefix count announced by XPN spatially distributed peers. Network events such as infrastructure failures and policy changes are consequences of some activities during a certain period of time and at a particular location. The events are observable when a significant transition occurs in the states of the activities. Due to the distributed nature of the events, the transitions (evidence) are to be observed in space and time to detect and track them. Hence, a spatiotemporal measure was selected to obtain better detection accuracy. Another reason for choosing this measure is based on a cause-effect relationship. If there is an event at one location, then it could affect the activities of other locations. For example, when the SEA-ME-WE4 cable cut occurred in June 2012 near Singapore, voice and Internet services were affected in India through the eastern route. However, the services were restored with an alternate route, which increased latency until the eastern route was re-established. This event could be observed in space and time as a drop in prefix announcements of one location that might cause an increase in prefix announcements of other locations for a certain period of time. At each location, the event could be identified using its specific temporal pattern. Our proposed spatiotemporal measure is a specific combination of the temporal patterns aggregated from different locations into a single pattern.

The combined spatiotemporal pattern has both spatial and temporal characteristics. The prefix count time series with components such as long-term trend and random irregular variations elicit the temporal characteristics. The long-term trend arises due to the evolution in the number of new content and data centre service networks. The irregular variations of prefix counts could have occurred due to node/link failures, popular events such as World Cup Football and policy changes. Our primary interest is to identify the irregular variations that are long term in nature caused by various events at a country level. Furthermore, we seek to determine the times at which the event is occurring and classify it to an impact class.

A. Prefix Increase Scenarios

At a country level, new ASes are added to the topology or already existing ASes themselves increase the number of prefixes. If new ASes are included, then the prefix counts will increase or remain the same due to prefix aggregation. It is the same case when already existing ASes increase the number of prefixes. When the prefix count increases beyond a normal threshold value, then it is an abnormal event. Due to misconfiguration, prefix hijack attacks, or new popular events, the prefix count for an AS increases beyond the normal threshold value. In a study on popular BGP prefix events [7], it was observed that the majority of prefix events have an average duration of 20 seconds. This observation suggests that most of the events resulted from routing convergence rather than long-term oscillations. They hypothesized that most of the inbound and outbound prefix events are due to route changes. From this study, we can infer that during prefix events, either a route change occurs or routes are completely withdrawn for the existing prefixes. In the work [13], it was observed that 45 $\%$ of prefixes have a route time length duration of one day. These route time durations are mostly caused due to different types of misconfiguration related to the de-aggregation of prefixes, related origin and foreign origin prefixes. Route time length durations of more than one day and less than 7 days are observed for 30 $\%$ of the prefixes. These types of routes may be due to worm attacks. The remaining 30 $\%$ of the prefixes have route time lengths of more than 7 days and may be induced by policy changes and peering failures. The peering events presented in [32] showed that a single peering in the route itself carries traffic for a greater number of prefixes of many originating ASes. Hence, a peering failure may result in a greater number of path changes. Thus, determining whether a simultaneous path change has occurred to different prefixes belonging to many origin ASes passing through that peer can be used as a heuristic to detect peering failures. We used this heuristic to validate the long-term events detected by our proposed detection method.

B. Prefix Decrease Scenarios

ASes drop prefixes that are not further used by their customers. In practice, ASes drop prefixes temporally with a small value. When an originating AS fails or stops functioning, then the prefix count for that AS drops completely in the routing table data. During prefix hijacking, the prefix counts of the hijacked ASes decrease. Country-level resource failures such as power blackouts or policy-level censorships [39] lead to a greater number of prefix drops. During the time of policy misconfiguration in the subset of ASes placed in a geographical location, prefixes are dropped and restored (oscillations) in short time intervals. In AS peering failures or policy changes in a geographical location, prefixes of ASes that pass through the peering AS are dropped partially or completely and have different long-term manifestations in the AS paths. Partial prefix drops confined to a region due to policy changes appear as a reduced number of paths, variations in the path length or changes in the peers of the path without affecting the length of the path. In the case of peering failures, alternate paths or backup paths are used. Using AS path characterization, long-term policy changes and peering failures can be inferred. In the case of policy changes, it can be of two types. The first type is at the originating AS level. An intermediate AS completely excludes all the prefixes of a specific originating AS from forwarding; hence, the paths passing through the intermediate AS are withdrawn for the particular originating AS but are still available for other originating ASes. The second case is at the country level. All the prefixes of a specific country are excluded from forwarding by an intermediate AS. All the paths that pass through the intermediate AS to a country will be withdrawn. However, the paths that include the intermediate AS to other countries exist in the routing table. In the case of peering failure in an intermediate AS, it cannot forward any of the prefixes. None of the paths in the routing table will have that intermediate AS in the paths.

Using the knowledge gained from the literature on long-term events, we formalize our proposed spatiotemporal measurement. The peers from which the routes have been collected are located in different spatial locations with varying dynamics. Noticeable deviations in the number of reachable prefixes from XPN spatially distributed peers reveal the spatial characteristics of the measure. The value for "X" is determined based on the number of maximum prefix announcement peers at the country level. The set of unique prefixes $P_1, P_2, \ldots, P_x$ announced by x out of n peers that constitute X $\%$ are combined spatially as $P_1 U P_2 U \ldots U P_x$ at time t. This spatially aggregated temporal measure might have linear or non-linear trend characteristics. The trend characteristics, the heuristics behind choosing the "X" percentage peers and the value for "x" are identified through the country-level empirical study discussed in Section V.

To understand the behaviours of normal and abnormal events, K-means clustering is performed on the spatiotemporal prefix count measure in our work. Univariate clustering on the first difference of the measure provides different cluster ranges. The ranges could be used to determine normal and abnormal events of different impact values. After the detection of long-term events, the impact values are quantified with the help of an estimated regression line. The regression lines are computed based on the trend characteristics of the measure. Then, the quantified impact values are classified to one of the classes identified by the K-means clustering algorithm. Each detected event is further analysed using the nature of the AS path feature changes to infer the possible causes of the events. We have also used publicly available information on long-term events [3] to validate the detected event. The inference of the cause may be further validated using email surveys from affected providers, as specified in [13], as well as other publicly available information, such as the looking glass of an AS [21], which is beyond the scope of our work.

C. AS Path Feature Selection

Events that occur in the backbone are distributive in nature and manifested in the BGP monitoring points located at different geographical locations. In the case of long-term events, the traces are manifested both temporally and geographically. Failures, policy changes and new popular occurrences might cause long-term changes in both prefix volume and AS paths. Volume-based features are considered

TABLE V: Prefix Volume and AS Path-Based Features

Feature	Definition	Category
1	Maximum unique prefix count	volume
2	Single upstream change prefix count	volume
3	Multiple upstream change prefix count	volume
4	Self de-aggregated prefix count	volume
5	Related origin prefix count	volume
6	Foreign origin prefix count	volume
7	XPN peer prefix count	volume
8	Average AS path length	AS-Path
9	Maximum AS path length	AS-Path
10	Average unique AS path length	AS-Path
11	Average edit distance	AS-Path
12	Maximum edit distance	AS-Path
13	Minimum AS path length	AS-Path
14	Minimum edit distance	AS-Path
15	Peer Counts of 2 . . . MaxPathLength edit distances	AS-Path

to be more efficient for detecting the BGP anomalies and are used in constructing machine learning-based classifier models in [31], [37], [40]. Important features that can be extracted from the monitoring point RIB data are presented in Table V.

In the enumerated features, we consider the volume-based spatiotemporal feature: XPN peer prefix counts for detecting the long-term events. For each originating AS at a country level, during events, the minimum, maximum and average AS path lengths and peer counts associated with 2 . . . MaxPathLength edit distances are modified depending on the type of event. Hence, relevant AS path-based features are selected using Fisher scores to validate and infer the cause of the events. The occurrences of events are verified using the selected AS path features and publicly available real-time events.

D. Fisher Scores

Fisher scores measure the correlation and relevancy existing among the features. F_i is score of a set of i=1 . . . 6 AS path features, and each feature is represented as a column vector. The feature measurements during change point day (CPD) and one day before (BCPD) are used to compute the scores of the feature relevance using equations 1 and 2. Here, measurements on CPD are considered to belong to the anomaly class, and those on BCPD are considered to belong to the regular class. In the case of prefix count increase or decrease events, the mean and variance of the CPD and BCPD are used in equations 1 and 2 to compute the scores.

Fisher Score Equation for Prefix Count Increase Events

$$F_{Score} = \frac{m_{CPD} - m_{BCPD}}{S^2_{CPD} + S^2_{BCPD}} \tag{1}$$

Fisher Score Equation for Prefix Count Decrease Events

$$F_{Score} = \frac{m_{BCPD} - m_{CPD}}{S^2_{BCPD} + S^2_{CPD}} \tag{2}$$

The significance of using the Fisher score is that it maximizes inter-class separation $(m_{CPD} - m_{BCPD})$ and minimizes intra-class variances $(S^2_{CPD} + S^2_{BCPD})$. The features are ordered in decreasing order of the score, and the feature with the highest score is selected for validation because it has the highest discrimination capability between two classes among the other features.

V. Empirical Study

In this section, we study the prefix counts $P_{1,...,x}$ of x peers computed as $P_1UP_2U \ldots UP_x$ by varying x from the extracted Indian AS topology data. Based on the analysis, we estimate the appropriate x value, its corresponding XPN spatially distributed peers and the trend of $P_{1,...,x}$. The identified prefix counts $P_{1,...,x}$ announced from the XPN range of peers are further considered for event detection. The impact quantification of the event is performed based on the trend. The routing table data contain the reachability information of unique prefixes from each of the peers to the Indian AS topology. We formally denote the peers that are spatially distributed at time t as Y_1, Y_2 $\ldots Y_n$, where n is the total number of peers announcing the prefixes. Let P be a vector of prefixes P_1, P_2 $\ldots P_n$, announced by n peers at time t. The maximum unique prefix value P_{max} at time t is computed as P_1 U P_2 ...U P_n. The P_{max} values vary temporarily and are $P_{max(t+1)}$ $\leq$ or $\geq$ $P_{max(t)}$ due to the trend and events. We are interested in estimating a particular peer percentage range (referred to as bins in the histogram analysis of outlier detection) by taking into account two specific criteria. The first criterion requires the peer percentage range that announces the majority of the prefixes P_{maj}, and the second criterion imposes the percentage range that manifests more spatial events.

To measure the P_{maj} for an XPN, we created different percentage range (peers) bins. Each bin represents the percentage range of peers. For each peer percentage associated with a particular range, the numbers of unique prefixes announced by the peers are counted and summed to obtain P_{maj}. The equation for computing P_{maj} is given in 3.

$$P_{maj} = \sum_{X_i \in XPN} P_{uniq\_cnt}(X_i) \tag{3}$$

The $P_{uniq\_cnt}(X_i)$ is computed by counting the number of unique prefixes announced by X_i percentage of peers. Consider the cases that could occur with this type of computation. When X_i percentage of peers increase the number of prefixes that are announced at time t, then $P_{uniq\_cnt}(X_i)$ of the bin that belongs to the XPN range would also increase. Alternatively, $P_{uniq\_cnt}(X_i)$ drops could occur in the XPN peer range when X_i percentage of peers drop prefixes. The temporal variation in P_{maj} of various peer bin ranges is an indication of events. The changes will be manifested in the *temporal patterns* of each peer bin range.

Since the maximum unique prefix count P_{max} is subject to variations due to the dynamics in the number of peers n, we impose the first criterion to determine the XPN that announces at least 95 % of P_{max}. Furthermore, to identify the XPN that manifests more spatial events, the second criterion is imposed. Such an XPN peer range is estimated using an iterative procedure. Initially, the peer bins are assigned the following percentage ranges.

1) Above 90 %
2) 81 to 90 %

TABLE VI: Prefix Reachability From Range of Spatial Locations

Date	Total Unique Prefixes	> 90 % Peer Prefixes	81-90 % Peer Prefixes	51-80 % Peer Prefixes
01-01-2012	17502	13286	3813	33
02-01-2012	17225	13256	3847	16
03-01-2012	17215	13350	3744	30
29-12-2012	19294	11456	7701	34
30-12-2012	19292	11428	7730	34
31-12-2012	19222	11499	7592	40

TABLE VII: Prefix Reachability From Range of Spatial Locations

Date	Unique Prefixes	28-50 %	14-27%	6-13 %	< 5 %
01-01-2012	17502	13	47	23	287
02-01-2012	17225	14	46	28	18
03-01-2012	17215	10	46	22	13
29-12-2012	19294	16	38	35	14
30-12-2012	19292	14	37	35	14
31-12-2012	19222	8	36	33	14

3) 51 to 80 %
4) 28 to 50 %
5) 14 to 27 %
6) 6 to 13 %
7) Below 5 %

The P_{maj} prefixes reachable from each initial peer bin range for the AS topology of India are presented in Tables VI and VII. The initial bin ranges are then refined using the empirical data under consideration by enforcing the two criteria.

As shown in Table VI, for the sum of P_{maj} for the peer XPN ranges, greater than 90 % and 81-90 % constitute approximately 99 % of the unique prefixes. This can be interpreted as 77 % of the unique prefixes are announced by greater than 90 % of the peers from different locations. In addition, another 22 % of the unique prefixes are announced by 81-90 % of the peers. Cumulatively, 99 % of the prefixes are announced by more than 80 % of the peers. In the observation period, maximum temporal dynamics have occurred in these two ranges. When there is a decrease or increase observed in the 90 % range, in most cases, an equal value of simultaneous decrease or increase is observed in the 81-90 % range. From this result, we can infer that less than 10 % of peers are subjected to temporal dynamics in most cases.

Considering the temporal peer dynamics observed in the two ranges, we have computed the prefixes announced by each percentage of peers above 80 % at time t. The announced P_{maj} for various XPN peers are given in Fig. 9.

It is observed from the calculations that > 88 % of peers announce more than 95 % of the prefixes. Moreover, for the same temporal pattern, we could observe two long-term (more than a day) drops and five short-term drops (one day) that have an impact of more than 10 %. Two long-term increase events with a 1-2 % increase in P_{maj} are also observed. From the comparisons of temporal patterns presented in Fig. 9 with individual peer temporal patterns presented in Fig. 8, it has been visually observed that the prefix (increase/decrease) events that occurred in individual peer patterns are captured more in the above 88 % peer range than in the remaining XPN peer ranges. Moreover, when statistical properties are considered, the temporal pattern of this range is positively correlated with a correlation coefficient of 0.5 with the total unique prefix pattern. From this result, it can be interpreted that 50 % of the time, no positive correlation exists between the patterns due to events. The XPN peer ranges and their correlation to the total unique prefix pattern are given in Fig. 10.

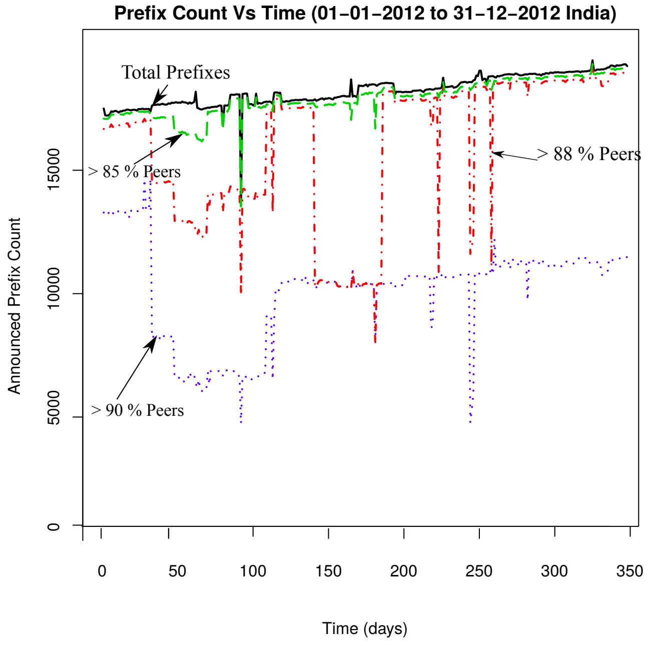

Fig. 9: Prefix vs Time

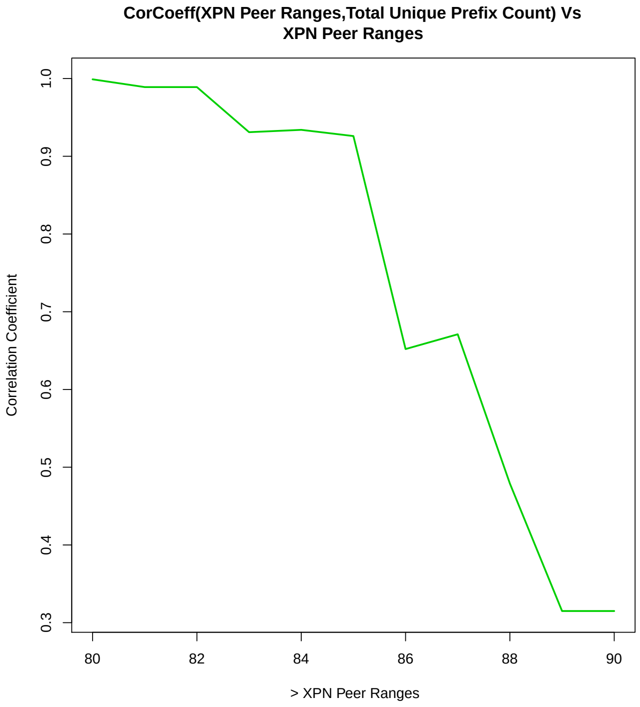

Fig. 10: Correlation Coefficient With Total Unique Prefixes for Various XPN Ranges

Based on these two results, the 88 % peer percentage range is chosen as the optimal XPN, and the P_{maj} pattern for this range is used for the identification of control plane events to the country India. Following a similar analysis, the XPN peers with greater than 95 % of prefixes are identified

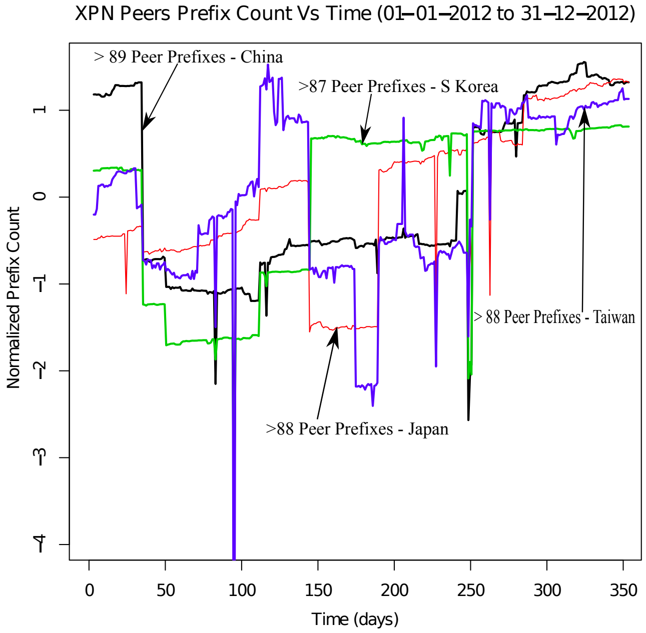

Fig. 11: XPN Peer Prefixes vs Time for Four Countries

for the other four countries, and the normalized patterns are shown in Fig. 11. Each country pattern manifests different long-term events and a linear trend. The identified patterns also have correlation coefficients of 0.0068, 0.811, 0.038, and 0.49 with total unique prefix counts for the countries China, Japan, South Korea and Taiwan, respectively. The optimal peer percentage that captures all the events requires further investigation and is beyond the scope of this paper.

VI. Method

In this section, we describe the techniques used in our proposed method for detection and classification of long-term events over the identified XPN peer range prefix. In general, the time series exhibits the characteristics of trend and long- and short-term deviations from the trend and seasonal patterns. We have considered the following characteristics in this work: long-term deviations and trend. The long-term deviations in the time series indicate a parameter, that is, the mean change in the generation process of the time series. The trend indicates the non-stationarity (non-mean reverting) property of the time series. For instance, we could observe drops and increases for short and long durations with an upward trend in the P_{maj} measure of all the peer ranges from Fig. 9. The properties present are due to factors such as new services, node or link failures and attacks. The prefix count evolves temporally with a long-term trend due to new services. However, it will have long and stochastic variations during the case of node or link failures and attacks. Since P_{maj} is affected by many external factors, it is reasonable to expect long-term prefix count deviations along with the trend. Hence, the mean variant characteristics of the measure during long-term events are utilized for detecting events by employing the mean change point detection technique. The long-term trend in the measure is identified as linear. Hence, a linear regression model is fitted to estimate the parameters of the trend line. The trend line is further utilized to quantify the impact values of the change point segments, that is, long-term events.

A. Change Point Detection

Change point detection is a segmentation mechanism on time series data for detecting mean invariant segments. It is the name given to the problem of estimating the point at which the statistical properties of a sequence of observations change. Here, segments are assumed to have the same mean value. The segments are identified using the Segmentation Neighbourhood Search (SegNeigh) [41] algorithm, which employs dynamic programming to optimize over a cost function. The algorithm optimizes for a given number of k segments based on optimal solutions for the k-1 segments. To accomplish this task, a search is performed over all previous change point locations within 1 to t-1 time units and selects the one that provides the optimal minimum cost segmentation up to t time units. In our work, the likelihood ratio test statistic (normal assumption for data) and the CUSUM statistic are computed on the time series data for hypothesis testing in the SegNeigh algorithm to detect the mean change points. When a distribution-free assumption is considered for the data, the CUSUM statistic can be chosen. We are interested in finding the k-1 change points that represent the k optimal segments with different mean values from the given n data points.

The detection of the mean change point can be posed as a hypothesis testing problem of which single or multiple change points can be identified. H_0 is the null hypothesis that corresponds to no change point at position m. H_1 is an alternate hypothesis that corresponds to a change point at position m when the null hypothesis is rejected. The testing of the hypothesis is performed using the likelihood ratio test statistic if the data are assumed to have a normal distribution. When no distribution assumption is made for the data, the CUSUM test statistic is computed to apply the hypothesis test.

The SegNeigh search algorithm optimally identifies k-1 change points by minimizing equation 4, where C is the cost function for a segment and βf(m) is a penalty value. The test statistic is used in the cost function. To guard against overfitting of the data in the case of detecting multiple change points, a penalty value is used. The numerical value assumed for the penalty controls the type I error, i.e., false detection.

$$\sum_{i=1}^{m+1} [C(y_{(ti-1)+1:ti})] + \beta f(m) \tag{4}$$

Since our search algorithm is deterministic, it will explore $2^{(n-1)}$ possibilities to identify the optimal k-1 change points. When the number of change points k-1 is already known, the search space will be reduced by $\binom{n-1}{k-1}$. In our work, we did not make any assumptions on the number of change points and allowed the algorithm to analyse all $2^{(n-1)}$ possibilities. By reusing the information computed for k-2 change points stored using the dynamic programming technique, the optimal segments for k-1 change points are identified. The incorporation of the dynamic programming technique reduces the time complexity of the search algorithm from $\mathcal{O}(2^n)$ to $\mathcal{O}(Qn^2)$, where Q is the maximum number of change points to identify.

The change point detection mechanism that employs the SegNeigh algorithm is applied on the P_{maj} prefix data associated with the above 88 % peer range. The prefix data time

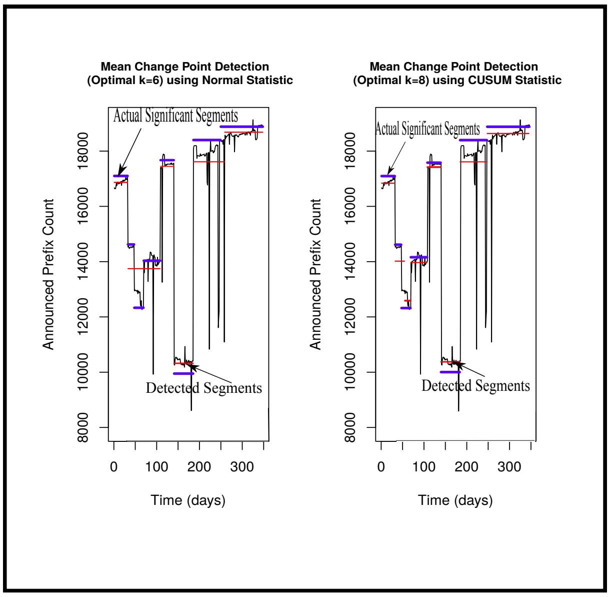

Fig. 12: Mean Change Point Detection Using the Segmentation Neighbourhood Algorithm

series is segmented based on the minimization of equation 4 with k=2 as starting value. Using visualization technique, k value is incremented to detect all the event segments. Further, the detected segment can be classified as significant event segment, redundant segment, suboptimal segment and containment segment. The significant event segment denotes either normal or abnormal event segment. The redundant segments are defined as single event(normal/abnormal) segment divided into multiple event segments that are classified into same impact class. Suboptimal segments are defined as segments that have less than 3 days duration. Also two or more significant mean change segments contained within another segment are defined as containment segment. Detection of suboptimal and containment segments lead to errors. The significant segments are identified with minimum error using normal statistic for k=6,and using CUSUM statistic for k=9 . Further increase in k value produces either redundant or suboptimal segments in normal statistic. Also one containment segment is present when normal statistics is used. In case of CUSUM statistic, two delayed detections are present but there is no further segmentation for increase in k value. The actual significant segments, the detected segments using normal statistic and using CUSUM statistic on the considered time series are shown in Fig. 12.

B. Clustering and Class Labelling

The long-term event is assumed to occur between any two consecutive change points l-1 and l. Hence, the identified change point segments j (cps_j) that lie between the change points l-1 and l are to be classified to one of the events. Since knowledge of the class of events that caused the changes is not readily available, we use statistical properties to characterize the event classes. The K-means algorithm is employed on the first difference of the 88 $\%$ XPN prefix count data. The number of cluster values can be passed as a parameter to the algorithm. The clusters are formed by iteratively minimizing the intra-cluster distance and maximizing the inter-cluster distance values. When no further improvement on the inter- and intra-cluster distances is possible, the algorithm terminates and outputs the results. The resulting clusters were occasionally locally optimal because of the initial cluster centre assumption. To overcome this problem, we performed the K-means algorithm 20 times for a k value and selected the best clusters using two indices, namely, goodness of variance fit measure and tabular accuracy, which are computed using Jenk's tests. For a set of k values from 2 to 15, we ran the algorithm for the previously mentioned number of times and obtained the indices for each k value. By plotting the k value against the indices, the optimal k value is chosen.

The clusters identified for the optimal k value consist of normal and event-driven clusters. The inference on normal clusters is based on the assumption that the first difference data follow a normal distribution asymptotically. According to this assumption, the cluster that contains greater than 90$\%$ of the data is classified as a normal cluster (NC), and the remaining clusters are classified as event-driven clusters (EC). We label the normal and event-driven clusters based on their value range, i.e., NC(L_i, U_i) and EC($Ł_i, U_i$). Since actual labels for events are not available, we label the clusters based on the value range itself and finally arrive at impact-valued classes. We refer to each cluster as an impact-valued class since it is formed on the first difference of prefix count data, which reflects the impact that occurred on a day. The results of the K-means algorithm on the considered data are presented in Table VIII.

C. Linear Regression and Classification

The impact value of each change point segment j (cps_j) is computed with respect to a regression line. It was observed that the temporal P_{maj} patterns of various XPN peer ranges exhibit a linear trend with a positive regression coefficient in the range of five to six. A similar type of trend with a regression coefficient of 5.5 is also exhibited by the total unique prefix count pattern. This leads to the inference that the prefix count increases linearly with a regression coefficient between five and six. Consequently, a regression line for the 88 $\%$ peer range is fitted for the regression coefficient 5.5, and the intercept parameter has been estimated. The estimated regression lines for the total unique prefix count and 88 XPN peer range prefix patterns are given in Fig. 13.

To quantify the impact values of event segments, the corresponding regression line segment mean value is subtracted from the mean value of each cps_j. The difference is considered to be the impact value of each segment. When the difference is greater than a threshold range, the segment cps_j is defined as significant change segment and is classified to one of the event clusters EC($Ł_i, U_i$). The threshold range (-292.5 to 292) was obtained from the normal cluster impact range of the K-means clustering algorithm. Event classification and the impact values of the events shown in Fig. 12 are presented in Table IX.

In the six identified segments, three segments (2,4, and 5) have impact value differences that are greater than the

TABLE VIII: K-means Clustering Results

S. no	k-value	Goodness of variance, Tabular accuracy	Cluster Interval	Label (E-Event) (N- Normal)
1	7	0.97,0.75	(-292.5,292)	$NC_{(-292.5-292)}$
2			(292,2032)	$EC_{(292,2032)}$
3			(2032,5342.5)	$EC_{(2032,5342.5)}$
4			(5342.5,7682)	$EC_{(5342.5,7682)}$
5			(-2053,-292.5)	$EC_{(-2053,-292.5)}$
6			(-5439.5,-2053)	$EC_{(-5439.5,-2053)}$
7			(-7331,-5439.5)	$EC_{(-7331,-5439.5)}$

TABLE IX: Mean Change Point Time and Prefix Reachability Events

S. no	Time Period	Regression Segment mean	$>$ 88 % Mean	Impact Value	Long Term Event	Duration (Days)
1	01/01 - 05/02	16769	16867	98	No	36
2	06/02 - 04/05	17110	13701	- 3409	$EC_{(-5439.5,-2053)}$	89
3	05/05 - 06/06	17451	17425	- 26	No	33
4	07/06 - 21/07	17694	10333	-7361	$EC_{(-7331,-5439.5)}$	45
5	22/07 - 20/09	18029	17532	-497	$EC_{(-2053,-292.5)}$	61
6	21/09 - 31/12	18540	18535	-5	No	102

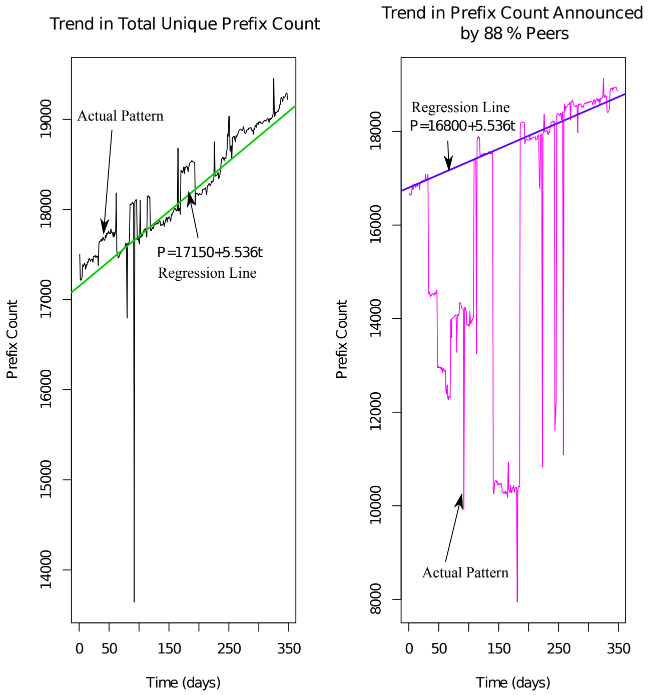

Fig. 13: Estimated Regression Lines With Parameters

threshold range and can thus be classified as long-term events. In particular, the negative sign indicates a drop in prefix announcements in the 88 % peer range. Furthermore, the time period corresponding to those segments can be considered as a long-term event duration. For a total duration of 195 days, long-term events with significant prefix drops that correspond to three different event classes occurred in the Indian topology.

VII. Discussion and Validation

The change point detection technique identifies k segments in the considered time series. The k-1 points are computed by exploring all the combinations of split points that minimize the cost function. When a significant mean change is not detected by the method, it is called as type II error or missed detection. To minimize type II errors, the parameter value k can be increased [42]. The suboptimal segments may lead to false detection or type I errors. To reduce the false detection, a penalty value can be specified as a parameter to the method. The Schwarz information criterion (SIC) is used as a penalty value in our method to minimize type I errors. The presence of short term events in the data introduce higher variance in the identified segments that violate the constant variance assumption of the mean change point detection technique. The standard deviations of k=6 segments shown in Fig. 12 using normal statistic are given in table X. From the computed standard deviations, huge variance difference is observed across the segments. Analysis on the data associated with the segments reveals that the short term events present in the segments are contributing to such variations. The short term events are removed and imposed with a variance threshold on the identified segment. The threshold value is derived from the normal cluster interval of kmeans result. After imposing the constraint, to identify an appropriate k value that minimizes Type I and Type 2 errors, we explored k values starting from 2 and analysed whether the technique captures all the significant changes using the statistics. The analysis result shows that the optimal k value for normal statistics is 8 and CUSUM statistics is 9.

To compare our method with other methods specified in the literature for detecting change segments, we consider the type I, type II errors and delayed/early detection as performance metrics. For this purpose, two techniques used in change detection methods from the literature are chosen. The first technique is the generalized log likelihood ratio (GLR) test to detect variance change points (VCPT) specified in [22]. The second technique is the exponential weighted moving average (EWMA) [43] used to detect small process mean shifts. We applied these techniques on our data after tuning the parameters such as standard deviation and smoothing value of the techniques. The comparison results are shown in Fig. 14. Our method with normal statistics detected all the 8 significant segments shown in Fig. 12 as actual significant segments without any error. In case of

TABLE X: Standard Deviation of the Detected Segments using Normal statistic

Segment	1	2	3	4	5	6
Std.Dev.	115.12	840.76	774.25	383.73	970.59	766.41

TABLE XI: Performance Comparison of the Techniques Over Different Countries' Data

S. no	Country	Our Method (CUSUM)	Our Method (Normal)	GLR - VCPT	EWMA
1	India	1 Type I, 1 early detection	No error	2 Type II	3 Type II
2	China	1 Delayed Detection, 2 Type II	1 delayed detection	2 Type II, 1 delayed detection	3 Type II
3	Japan	2 Early Detection, 1 Type II	1 Type II	2 Type I 1 Type II	2 Type II
4	S. Korea	2 Type I, 1 delayed detection	1 Type I	2 Type II	2 Type II
5	Taiwan	No error	No error	3 Type I 2 Type II	No error

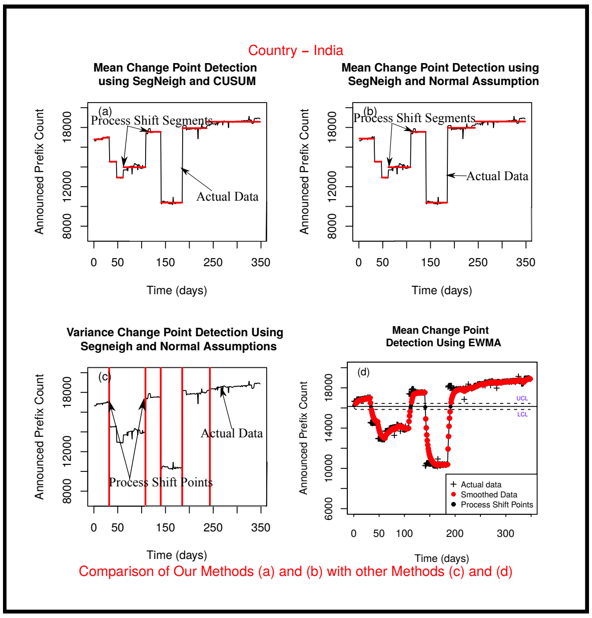

Fig. 14: Comparison of Different Methods in Change Detection

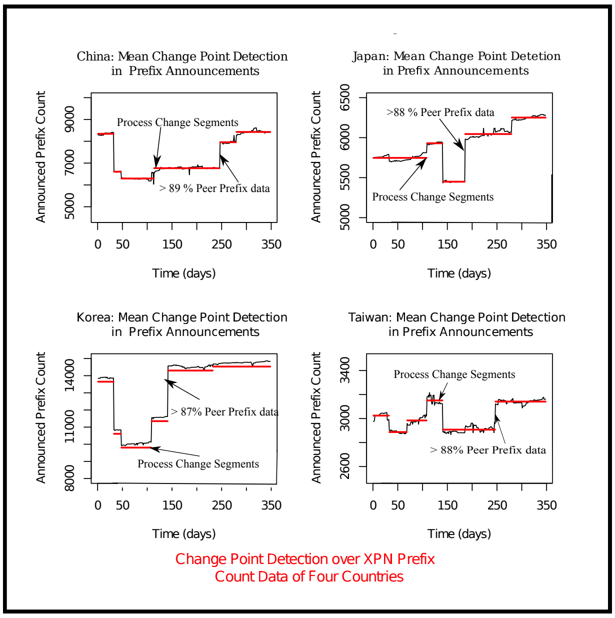

Fig. 15: Change Point Detection Over Different Countries' XPN Peer Prefix Data

CUSUM test statistics, 9 segments are detected out of which second segment is false detection and ninth segment is early detection. The GLR-based variance change detection method missed segments 3 and 4. Similarly, EWMA missed detecting event segments 3,4 and 8. The change point detection method with the imposed variance constraint has been applied to the XPN peer prefix count data of the other four countries and shown in Fig. 15. The data for each country reflected different k optimal segments. From the optimal segments, long-term events are classified. The countries China and South Korea are classified with 4 long-term events. The other two countries Japan and Taiwan are classified with three long-term events.

The detection performances of all the four techniques over the five countries' data are shown in Table XI. Our method with a normal distribution assumption of the data detects the significant change segments with a minimum number of type I, type II errors and delayed detection. The k value has been explored from 2 for each country to reduce type II errors. The penalty value using SIC and variance constraint are imposed to reduce the type I error. We observed that our change point detection method is sensitive to higher short-term variations. The presence of such an event immediately after a long-term event affects the detection process by shifting the change points that lead to delayed detection errors. To further improve mean change point detection performance, higher short-term variation points are replaced with suitable neighbouring data point values. In the Indian XPN peer range prefix patterns, after smoothing higher short-term variations the number of event segments have increased when compared to before smoothing. The second segment in before smoothing has been split to three significant segments. All the detected segments have the variance within the range between -300 to +300.

Using the change point detection technique with normal assumption, we identified five long-term events for Indian XPN peer range prefix pattern. The impact of each event is classified based on the K-means cluster ranges. Since each event is identified using the statistical properties extracted through mining, we attempt to validate the detection, duration and classification performance of our methods using media reports and another correlated feature, namely, total paths. Legitimate route change events due to link status change typically affect a large number of prefixes [17], [18], whereas route changes due to attacks such as prefix hijacking particularly target specific network prefixes. Although most of the link events could be rectified in a single day, events such as underwater cable cuts might take longer than a month to resolve. In our case, four of the detected events have more than a 15 % prefix reachability reduction and prevail for more than 30 days. Hence, the classified events could be related with a link status change induced by cable cuts or peer policy change. The time period associated with the fourth event (segment 6) correlates with the SEA-Me-We4 cable cut and restoration period [44]. The change point mean for this segment indicates that there is an impact of a 42 % drop in the P_{maj} announced by the $>$ 88 % range of peers. However, the prefixes dropped by this peer range are announced by a lower number of peer ranges, as is indicated by the unique prefix count per time unit measurement. The inference here is that during the period from 07/06/2012 to 21/07/2012, few peers are affected by the physical cable event with a mean drop of 7361 prefixes.

Another indicator used to validate the occurrence of long-term events is total AS paths. To normalize the prefix count, we averaged the total AS paths by dividing it with the average number of peers. The average numbers of paths and prefixes announced by $>$ 88 % of peers are given in Fig. 16. The average path pattern also exhibits a significant drop during the event detection period and establishes the occurrence of a long-term event during this period. The average path pattern itself provides necessary visual evidence for the duration of long-term events, but it is not sufficient to detect the long-term event as precisely as our proposed P_{maj} prefix measure for XPN peer ranges. The events occurring in geographically distributed peers are well captured by our measure in the identified XPN peer range. Moreover, events are classified based on the impact ranges. The five long-term events with three different impact ranges and durations provide the indication for the reachability of the Indian AS topology from different geographical regions. If we consider the first four events (cps_2, cps_4) as true positives due to their substantial impact range, for a duration of approximately four months, there is more than a 15% drop in prefixes announced from few geographically distributed peers.

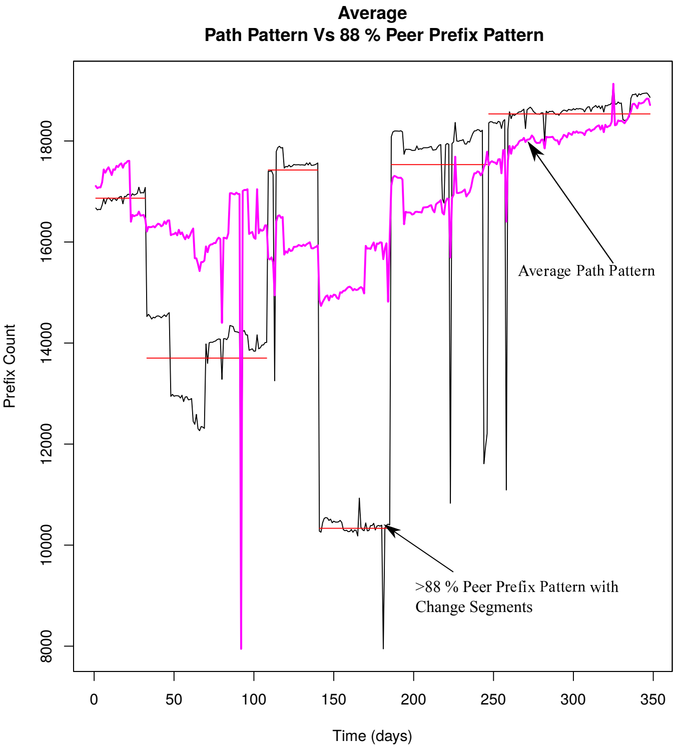

Fig. 16: Pattern Comparison Between Prefix Counts and Average Paths

A. Validation Using Features Based on Fisher Scores

At a country level, from the monitoring peers to each of the originating ASes, many distinct AS paths with different path lengths exist. Representative distinct path patterns from the peers AS286, AS293 and AS701 to the originating AS 4755 are given in Table XII. The paths are announced by the peers to the originating AS. We can observe that among the distinct paths, the shortest path is used by the peers to announce the maximum number of prefixes to the originating ASes. The number of distinct paths and the path length from each peer to the originating AS exhibit variations. During long-term peering events and failures, the path length and the intermediate ASes in the path change. Based on the heuristic discussed in IV-A, we consider AS path features such as maximum AS path length, minimum AS path length and average AS path length from all the peers to each originating AS of a country and employ Fisher scores to show the discriminating nature of the features to long-term events.

We compute the Fisher scores of the considered features for 3 days prior and 3 days after, including the event day specified in Table IX. The scores are shown in Table XIII. The prefix event identified on 06/2/2012 is also manifested in path length and is evident through the Fisher score for that event. The Fisher score based on the AS path length feature is not distinguishable for the remainder of the events. Hence, we can infer that during the time of those events, there might not be any significant path length changes.

From the features that we have used for validating the captured prefix events, we can infer that a single feature alone is not sufficient to validate all the identified events. As the root causes of the events are unknown, one has to rely on the manifestation of the events in different features to validate the occurrence of the events. Hence, we have to

TABLE XII: Distinct Path Patterns

Date	Announced Prefixes	Peers	Distinct Paths
03/02/2012	1177	286	286 6453 4755
	310		286 1239 6453 4755 4755
	14		286 6453 4755 4755
	6		286 2914 4755 4755 4755
	4		286 1239 6453 4755
	3		286 3561 9498 9730 7633 7633 18101 4755
	1180	293	293 6453 4755
	325		293 6453 4755 4755
	6		293 2914 4755 4755 4755
	560	3303	3303 6453 4755
	29		3303 15412 6453 4755
	6		3303 2914 4755 4755 4755
	5		3303 6453 4755 4755
	1		3303 1299 6453 4755 4755
	1		3303 1299 6453 4755

TABLE XIII: Fisher Scores for the Features

Dates	AS Path Length		
	Maximum	Minimum	Average
03-04 /02/2012	0.086	-0.001	-0.0148
04-05 /02/2012	-0.042	-0.044	-0.0157
05-06 /02/2012	**1.04**	**0.150**	**0.174**
06-07 /02/2012	0.020	0.049	0.055
07-08 /02/2012	-0.029	-0.070	-0.065
04-05 /06/2012	0.071	-0.0248	-0.032
05-06 /06/2012	-0.073	-0.117	-0.064
06-07 /06/2012	0.099	-0.006	0.019
07-08 /06/2012	-0.118	-0.092	-0.008
08-09 /06/2012	0.068	0.224	0.124
19-20 /07/2012	0.062	-0.064	-0.070
20-21 /07/2012	0.062	0.000	0.009
21-22 /07/2012	0.005	0.044	0.025
22-23 /07/2012	0.022	-0.017	-0.006
23-24 /07/2012	-0.254	-0.018	-0.027

explore the features shown in Table V to select the features that better discriminate the identified abnormal events from normal occurrences.

VIII. Conclusion

In this paper, we have proposed, validated and evaluated a method for identifying events based on control plane data. Specifically, we have analysed the temporal prefix pattern of five Asian countries collected from different geographically distributed peers to understand the reachability events. Further, the peer distribution analysis on different continents showed that the number of representative peers in Asia, Australia and Africa is less than 10 % of the total peers involved in route collection and to understand the events occurring in these regions, more number of representative peers are required. The temporal prefix count pattern analysis from individual peers reveals the occurrence of significant pattern changes with respect to location, duration and impact. To capture the events that occur spatially, in a single temporal pattern and quantify the impact and duration, we proposed a counting-based spatiotemporal measure on prefix reachability from XPN peer ranges. Using this measure in our proposed method, we detected long-term events through the mean change point detection technique. The impact of each event is quantified using the trend line estimated through linear regression. The sensitivity and specificity of the techniques were validated using the Sea-Me-We4 cable cut events reported in news blogs, the average number of paths and the Fisher scores of the path lengths. The events were classified to specific impact classes identified through K-means clustering and provide the indication for the prefix reachability to a country-level AS topology from different locations. When compared with event detection techniques of other methods, our mean change point technique with normal distribution assumption on the data outperforms with minimum number of type I, type II errors and delayed/early detection. In addition to error performance, the prefix impact values and the duration are quantified for long-term events using our method. The event quantification can be used as an indicator variable to measure the reliability of a country-level AS topology. From the validation perspective, more number of features are to be analysed for selecting the features used as evidence in verifying the prefix reachability events. In our future work, we plan to explore various feature selection methods to select the best set of features that discriminate the identified abnormal prefix events.

References

[1] GoogleNews, “Indian Outsourcing Sector Hit by Internet Disruption,” 2008. [Online]. Available: http://afp.google.com/article/ALeqM5gadarOY4frbM-XSyXS0_wMfxpHig

[2] TimesOfMalta, “GO Submarine Cable Fault Part of Wider Disruption between Italy and Egypt,” 2008. [Online]. Available: http://www.timesofmalta.com/articles/view/20081219/local/go-submarine-cable-fault-part-of-wider-number-between-italy-and-egypt.237909

[3] BBCNews, “Bangladesh Suffers Internet Disruption After Cut Cable,” 2012. [Online]. Available: http://www.bbc.co.uk/news/technology-18366007

[4] Routeviews, “University of Oregon Route Views Project,” 2005. [Online]. Available: http://archive.routeviews.org/oix-route-views/

[5] C. Labovitz, G. R. Malan, and F. Jahanian, “Internet Routing Instability,” *IEEE/ACM Transactions on Networking*, vol. 6, no. 5, pp. 515–528, Oct. 1998.

[6] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkranz, “BGP Routing Dynamics Revisited,” *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 2, pp. 5–16, Mar. 2007.

[7] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, “BGP Routing Stability of Popular Destinations,” in *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement*, ser. IMW ’02, 2002, pp. 197–202.

[8] M. Lad, X. Zhao, B. Zhang, D. Massey, and L. Zhang, “Analysis of BGP Update Surge During Slammer Worm Attack,” in *5th International Workshop on Distributed Computing - IWDC*, Kolkata, India, 2003 Dec 27-30, pp. 66–79.

[9] S.-T. Teoh, K.-L. Ma, S. F. Wu, D. Massey, X.-L. Zhao, D. Pei, L. Wang, L. Zhang, and R. Bush, “Visual-Based Anomaly Detection for BGP Origin AS Change (OASC) Events,” in *Self-Managing*

Distributed Systems: 14th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management, DSOM 2003, 2003, pp. 155–168.

[10] S. T. Teoh, K. Zhang, S. Tseng, K. Ma, and S. F. Wu, "Combining Visual and Automated Data Mining for Near-Real-Time Anomaly Detection and Analysis in BGP," in *Workshop on Visualization and Data Mining for Computer Security (VizSEC/DMSEC 2004)*, Washington, DC, USA, Oct 29 2004, pp. 35–44.

[11] T. Wong, V. Jacobson, and C. Alaettinoglu, "Internet Routing Anomaly Detection and Visualization," in *2005 International Conference on Dependable Systems and Networks (DSN 2005)*, Yokohama, Japan, 2005 Jun 28 - Jul 1, pp. 172–181.

[12] S. T. Teoh, S. Ranjan, A. Nucci, and C.-N. Chuah, "BGP Eye: A New Visualization Tool for Real-Time Detection and Analysis of BGP Anomalies," in *Proceedings of the 3rd International Workshop on Visualization for Computer Security*, ser. VizSEC '06, Alexandria, Virginia, USA, 2006, pp. 81–90.

[13] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP Misconfiguration," *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4, pp. 3–16, 2002.

[14] J. Li, D. Dou, Z. Wu, S. Kim, and V. Agarwal, "An Internet Routing Forensics Framework for Discovering Rules of Abnormal BGP Events," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, pp. 55–66, 2005.

[15] X. Wu, X. Yin, Z. Wang, and M. Tang, "A Three-Step Dynamic Threshold Method to Cluster BGP Updates Into Routing Events," in *2009 International Symposium on Autonomous Decentralized Systems*. Athens, Greece: IEEE, Mar. 2009, pp. 1–6.

[16] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs, "Locating Internet Routing Instabilities," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 4, pp. 205–218, Aug. 2004.

[17] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A Prefix Hijack Alert System," in *Proceedings of the 15th Conference on USENIX Security Symposium - Volume 15*, ser. USENIX-SS'06, Vancouver, B.C., Canada, 2006, pp. 153–166.

[18] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis, "A Light-Weight Distributed Scheme for Detecting IP Prefix Hijacks in Real-Time," in *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM '07, Kyoto, Japan, 2007, pp. 277–288.

[19] H. Ballani, P. Francis, and X. Zhang, "A Study of Prefix Hijacking and Interception in the Internet," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4, pp. 265–276, Aug. 2007.

[20] X. Hu and Z. M. Mao, "Accurate Real-Time Identification of IP Prefix Hijacking," in *Proceedings of the 2007 IEEE Symposium on Security and Privacy*, ser. SP '07, 2007, pp. 3–17.

[21] M. Wübbeling, M. Meier, and T. Elsner, "Inter-AS Routing Anomalies: Improved Detection and Classification," in *6th International Conference on Cyber Conflict, Tallinn, Estonia*, 2014, pp. 223–238.

[22] S. Deshpande, M. Thottan, T. K. Ho, and B. Sikdar, "An Online Mechanism for BGP Instability Detection and Analysis," *IEEE Transactions on Computers*, vol. 58, no. 11, pp. 1470–1484, Nov. 2009.

[23] A. Elmokashfi, A. Kvalbein, and C. Dovrolis, "BGP Churn Evolution: A Perspective From the Core," *IEEE/ACM Transactions on Networking*, vol. 20, pp. 571–584, 2012.

[24] G. Huston and G. Armitage, "Projecting Future IPv4 Router Requirements From Trends in Dynamic BGP Behaviour," in *Australian Telecommunication Networks and Applications Conference (ATNAC)*, 2006. [Online]. Available: http://caia.swin.edu.au/pubs/ATNAC06/Huston_1m.pdf

[25] M. Caesar, L. Subramanian, and R. Katz, *Towards Localizing Root Causes of BGP Dynamics*. Berkeley, CA, USA: Computer Science Division, University of California, Berkeley, 2003.

[26] D.-F. Chang, R. Govindan, and J. S. Heidemann, "The Temporal and Topological Characteristics of BGP Path Changes," in *11th IEEE International Conference on Network Protocols (ICNP)*. IEEE Computer Society, 2003, pp. 190–199.

[27] M. Lad, A. Nanavati, D. Massey, and L. Zhang, "An Algorithmic Approach to Identifying Link Failures," in *10th Pacific Rim International Symposium on Dependable Computing (PRDC)*, Mar 2004, pp. 25–34.

[28] T. Wong, V. Jacobson, and C. Alaettinoglu, "Making Sense of BGP," in *NANOG 30*, Miami, FL, 2004 Feb 8-10. [Online]. Available: http://www.nanog.org/meeting-archives/nanog30/presentations/wong.pdf

[29] J. Li and S. Brooks, "I-seismograph: Observing and Measuring Internet Earthquakes," in *INFOCOM*, 2011, pp. 2624–2632.

[30] R. Teixeira and J. Rexford, "A Measurement Framework for Pin-Pointing Routing Changes," in *Proceedings of the ACM SIGCOMM Workshop on Network Troubleshooting*. ACM, Sep 2004, pp. 313–318.

[31] N. M. Al-Rousan, S. Haeri, and L. Trajkovic, "Feature Selection for Classification of BGP Anomalies Using Bayesian Models," in *International Conference on Machine Learning and Cybernetics, ICMLC 2012*, Xian, Shaanxi, China, 2012 Jul 15-17, pp. 140–147.

[32] D.-F. Chang, R. Govindan, and J. Heidemann, "On the Origins of BGP Path Changes," in *IEEE International Conference on Network Protocols*, 2003. [Online]. Available: http://www.isi.edu/div7/oneil/publication_files/origins_of_bgp.pdf

[33] M. Luckie, "Spurious Routes in Public Bgp Data," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 14–21, Jul. 2014.

[34] X. Xu, Z. M. Mao, and J. A. Halderman, "Internet Censorship in China: Where Does the Filtering Occur?" in *International Conference on Passive and Active Network Measurement (PAM)*, 2011, pp. 133–142.

[35] APNIC, "Asia Pacific Network Information Center," 1992. [Online]. Available: ftp://ftp.apnic.net/pub/stats/apnic/delegated-apnic-extended-latest

[36] S. P. Meenakshi and S. V. Raghavan, "Forecasting and Event Detection in Internet Resource Dynamics Using Time Series Models," *Engineering Letters*, vol. 23, no. 4, pp. 245–257, 2015.

[37] Y. Li, H. Xing, Q. Hua, X. Wang, P. Batta, S. Haeri, and L. Trajkovic, "Classification of BGP Anomalies Using Decision Trees and Fuzzy Rough Sets," in *2014 IEEE International Conference on Systems, Man, and Cybernetics, SMC 2014*. San Diego, CA, USA: IEEE, 2014 Oct 5-8, pp. 1312–1317.

[38] A. Rude and K. Beard, "High-Level Event Detection in Spatially Distributed Time Series," in *International Conference on Geographic Information Science*, ser. Lecture Notes in Computer Science, N. Xiao, M.-P. Kwan, M. F. Goodchild, and S. Shekhar, Eds., vol. 7478. Springer Berlin Heidelberg, 2012, pp. 160–172.

[39] A. Dainotti, C. Squarcella, E. Aben, K. Claffy, M. Chiesa, M. Russo, and A. Pescap, "Analysis of Country-Wide Internet Outages Caused by Censorship," in *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*, Nov 2011, pp. 1–18.

[40] M. Cosovic, S. Obradovic, and T. Ljiljana, "Classifying Anomalous Events in BGP Datasets," in *Proceedings of the 29th Annual IEEE Canadian Conference on Electrical and Computer Engineering (CCECE 2016)*. Vancouver, Canada: IEEE, May 2016, pp. 697–700.

[41] I. E. Auger and C. E. Lawrence, "Algorithms for the optimal identification of segment neighborhoods," *Bulletin of Mathematical Biology*, vol. 51, no. 1, pp. 39–54, 1989.

[42] S. Vorobeychikov, Y. Burkatovskaya, and E. Sergeeva, "TAR(p)/ARCH(1) Process with an Arbitrary Threshold: Guaranteed Parameter Estimation and Change-Point Detection," *IAENG International Journal of Applied Mathematics*, vol. 46, no. 3, pp. 353–366, 2016.

[43] S. Deshpande, M. Thottan, T. Ho, and B. Sikdar, "A Statistical Approach to Anomaly Detection in Interdomain Routing," in *3rd International ICST Conference on Broadband Communications, Networks, and Systems*. IEEE, 2006, pp. 1–10.

[44] Renesys, "SMW4 Cut Shakes Up South Asia," 2012. [Online]. Available: http://www.renesys.com/