APPENDIX

A. Collectors

The current version of I-seismograph utilizes the publicly available BGP data archive providers, including all the collectors from RouteViews and RIPE. Table I lists all the data collectors we use and their corresponding URLs.

Collector name	URL
RouteViews CHICAGO	ftp.routeviews.org/route-views.chicago/
RouteViews EQIX	ftp.routeviews.org/route-views.eqix/
RouteViews ISC	ftp.routeviews.org/route-views.isc/
RouteViews JINX	ftp.routeviews.org/route-views.jinx/
RouteViews KIXP	ftp.routeviews.org/route-views.kixp/
RouteViews LINX	ftp.routeviews.org/route-views.linx/
RouteViews NWAX	ftp.routeviews.org/route-views.nwax/
RouteViews PERTH	ftp.routeviews.org/route-views.perth/
RouteViews SAOPAULO	ftp.routeviews.org/route-views.saopaulo/
RouteViews SFMIX	ftp.routeviews.org/route-views.sfmix/
RouteViews SG	ftp.routeviews.org/route-views.sg/
RouteViews SOXRS	ftp.routeviews.org/route-views.soxrs/
RouteViews SYDNEY	ftp.routeviews.org/route-views.sydney/
RouteViews TELXATL	ftp.routeviews.org/route-views.telxatl/
RouteViews WIDE	ftp.routeviews.org/route-views.wide/
RouteViews 2	ftp.routeviews.org/
RouteViews 3	ftp.routeviews.org/route-views3/
RouteViews 4	ftp.routeviews.org/route-views4/
RouteViews 6	ftp.routeviews.org/route-views6/
RIPE RRC 00	data.ris.ripe.net/rrc00/
RIPE RRC 01	data.ris.ripe.net/rrc01/
RIPE RRC 02	data.ris.ripe.net/rrc02/
RIPE RRC 03	data.ris.ripe.net/rrc03/
RIPE RRC 04	data.ris.ripe.net/rrc04/
RIPE RRC 05	data.ris.ripe.net/rrc05/
RIPE RRC 06	data.ris.ripe.net/rrc06/
RIPE RRC 07	data.ris.ripe.net/rrc07/
RIPE RRC 08	data.ris.ripe.net/rrc08/
RIPE RRC 09	data.ris.ripe.net/rrc09/
RIPE RRC 10	data.ris.ripe.net/rrc10/
RIPE RRC 11	data.ris.ripe.net/rrc11/
RIPE RRC 12	data.ris.ripe.net/rrc12/
RIPE RRC 13	data.ris.ripe.net/rrc13/
RIPE RRC 14	data.ris.ripe.net/rrc14/
RIPE RRC 15	data.ris.ripe.net/rrc15/
RIPE RRC 16	data.ris.ripe.net/rrc16/
RIPE RRC 18	data.ris.ripe.net/rrc18/
RIPE RRC 19	data.ris.ripe.net/rrc19/
RIPE RRC 20	data.ris.ripe.net/rrc20/
RIPE RRC 21	data.ris.ripe.net/rrc21/

TABLE I: Complete list of all BGP data collectors used in I-seismograph

Note that collector RRC02, RRC08, and RRC09 have stopped updating their data archives and only provide historical data. We use such collectors for the analysis of historical events only.

B. BGP Databin Length Choice

I-seismograph's basic data processing unit is BGP *databin*, which is a summary of BGP activities over a constant time period. Deciding the length of this period is a tradeoff: It cannot be too short; otherwise, the BGP activities within every databin will always be too sparse, and it will be difficult to distinguish normal and abnormal level of activities. It cannot be too long either; otherwise, I-seismograph will suffer from a slow response time since it will take I-seismograph at least the length of one databin to process and report the impact on BGP.

We measured, for different lengths of BGP databin, how many BGP announcements, withdrawals, and updates usually occur. Fig. 1 shows the results. Clearly, one minute would be a reasonable choice.

C. Short-term Clustering Stopping Criterion

When conducting the short-term clustering for a reference period, we need to make sure the final s-normal cluster will contain at least 50% data from the original input, and the databins in the same cluster are much more similar than those from different clusters. There are indeed many different stop criteria for clustering, but in order to meet the requirements above, we found that it works best by checking if the intra-cluster distance is no more than 20% inter-cluster distance, as shown in Fig. 2.

D. Long-term Clustering Stopping Criterion

In the long-term clustering (Section III-E2), if the difference between the intra-distance of a cluster and that of its parent cluster is no more than a threshold, we decide the cluster is not "tighter" than its parent cluster, and we will not further divide the cluster. We experimented with 16 different long-term clustering processes and tested how all the processes may be affected by different threshold values ranging from 0.1% to 99%. As shown in Fig. 3, we found that every process generates basically the same number of clusters when the threshold is 2% or less, but the number decreases when it is more than 2%. Therefore, we choose 1% as a safe threshold value to ensure the long-term clustering obtains the largest number of clusters.

E. Impact Calculation Formula Basis

In calculating the deviation distance along each attribute, say A_i, we found that along each attribute the databins in the normal clusters (i.e., normal databins) are mostly within $[\mu_i - \sigma_i, \mu_i + \sigma_i]$ (as shown in Table II), where μ_i and σ_i are respectively the mean and standard deviation of all the databins from the normal cluster along attribute A_i. So we use the databin's absolute distance from $[\mu_i - \sigma_i, \mu_i + \sigma_i]$ as its deviation along attribute A_i.

	mean±std_dev
WW	100.0
WADup	100.0
Withdraw	92.978
WADiff	90.247
AW	90.557
Announce	92.879
AADiff	96.439
Update	93.034
AADup2	91.950
AADup1	91.950

TABLE II: Percentage of attribute values that fall into mean ± standard deviation.

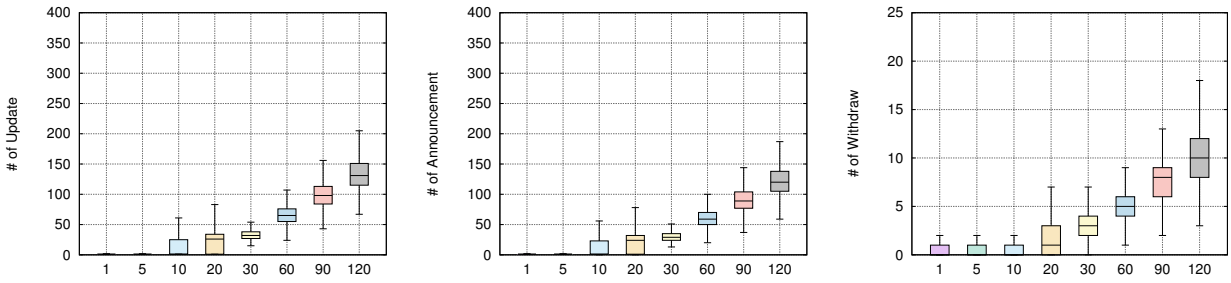

Fig. 1: Boxcharts of the # of three BGP dynamics attributes per databin with different lengths of databins, ranging from 1 second to 120 seconds.

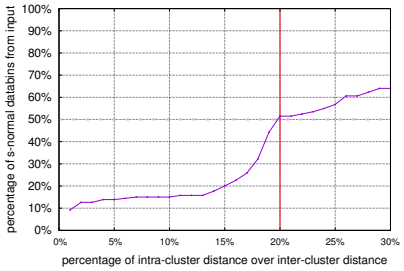

Fig. 2: Short-term clustering effect with different intra- vs. inter-cluster distance percentage.

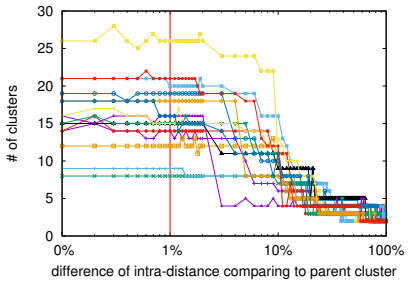

Fig. 3: Long-term clustering effect with different child- vs. parent-cluster intra-distance percentage.

F. Comparison of Short-term Normal Clusters from Different Days and Periods

We conducted an experiment to study how short-normal clusters from different days differ. We selected a four-week reference period for each year from 2008 to 2016, conducted short-term clustering over all these periods, and compared the s-normal clusters from the same reference period as well as short-term clusters from different reference periods. As shown in Fig. 4, clearly, the s-normal clusters are slightly different over different days from the same reference period, but can be more significantly different over different days from different reference periods.

G. Convergency Validation using Only Reference Periods Data

We also used the data from reference periods to test the convergency of I-seismograph. From Fig. 5, we can clearly see that the convergency test results using data only from reference periods are very close to the original results that use both the reference and monitoring periods, as shown in Section IV.B.

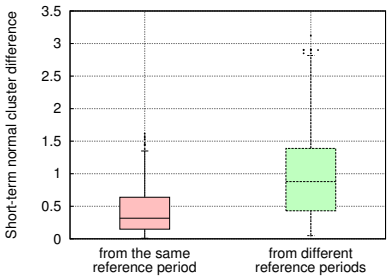

Fig. 4: The differences of short-term normal clusters for different days from either the same reference period or different reference periods.

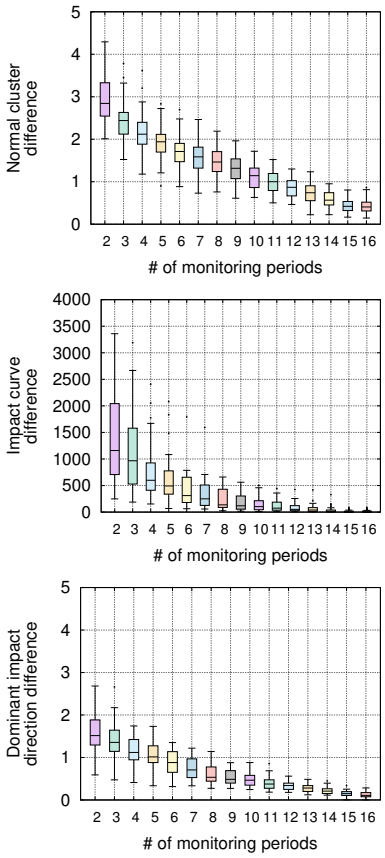

Fig. 5: The convergency of I-seismograph with data only from reference periods.