On the extent of correlation in BGP updates in the Internet and what it tells us about locality of BGP routing events

Andrey Sapegin[a], Steve Uhlig[b]

[a] *Hasso Plattner Institute (HPI), University of Potsdam, Germany. Email: andrey.sapegin@hpi.uni-potsdam.de*
[b] *Queen Mary, University of London, UK. Email: steve@eecs.qmul.ac.uk*

Abstract

The Border Gateway Protocol (BGP) is the core routing protocol in the Internet. It maintains reachability information towards IP networks, called prefixes. The adoption of BGP has come at a price: a steady growth in the routing table size [1] as well as BGP updates [2].

In this work, we take a different look at BGP updates, by quantifying the amount of prefix correlation in the BGP updates received by different routers in the Internet. We design a method to classify sets of BGP updates, called spikes, into either correlated or non-correlated, by comparing streams of BGP updates from multiple vantage points.

Based on publicly available data, we show that a significant fraction of all BGP updates are correlated. Most of these correlated spikes contain updates for a few BGP prefixes only. When studying the topological scope of the correlated spikes, we find that they are relatively global given the limited AS hop distance between most ASs in the Internet, i.e., they propagate at least 2 or 3 AS hops away. Most BGP updates visible from publicly available vantage points are therefore related to small events that propagate across multiple AS hops in the Internet, while a limited fraction of the BGP updates appear in large bursts that stay mostly localised.

Our results shed light on a fundamental while often misunderstood aspect of BGP, namely the correlation between BGP updates and how it impacts our beliefs about the share of local and global BGP events in the Internet. Our work differs from the literature in that we try as much as possible to explicitly account in our methodology for the visibility of BGP vantage points, and its implication on the actual claims that can be made from the data.

Keywords: interdomain routing, BGP update correlation, BGP propagation

1. Introduction

Today's Internet grows steadily and this growth often creates operational problems. The most urgent and well-known example is running out of IP addresses [3]. Although IPv6 [4] solves this specific problem, it does not address the scalability of the routing system, especially of its main routing protocol, BGP.

BGP is a path-vector protocol through which neighbouring routers exchange reachability information. For BGP, a reachable network is represented by a network prefix, i.e., a block of IP addresses. Based on the set of alternative paths it learns from its neighbours, a BGP router selects its best path to reach prefix it knows, and informs its neighbours about these path changes through BGP update messages. Contrary to link-state protocols, BGP messages do not provide direct information about which part of the topology underwent a change that triggered the BGP message.

Since BGP is the glue that binds all networks of the Internet together, it suffers from Internet growth more than other routing protocols. Each BGP router has a table where it stores routes to all possible destination networks (prefixes) in the Internet: more than $400,000$ routes nowadays [5]. With this large number of networks that make the Internet, the number of BGP messages grows as well [5], and seems to mostly depend on the growth of the Internet itself [6]. Further, sometimes BGP routers send to each other up to 100 times more update messages than usual [7, 8]. However, the causes and the amount of locality in BGP updates is still mostly unknown. With the imminent deployment of IPv6 and the trend towards network virtualisation, the Internet routing system might be under even more pressure in the future.

Multiple previous works have investigated BGP dynamics. Lad et al. [9] proposes a graphical tool to capture BGP routing dynamics. Flavel et al. [10] present a mathematical model of the evolution of the number of BGP entries in the BGP routing table, at a given vantage point. Mao et al. [11] rely on BGP beacons to better understand the global BGP dynamics in the Internet by injecting controlled BGP updates in the network. Feldmann et al. [12] use passive monitoring for locating routing instabilities. Park et al. [13] analyse the inter-dependencies of BGP updates received from one peer inside an AS.

The seminal paper in the area from Feldmann et al. [12] has developed, back in 2004, a methodology for locating routing instabilities causing update bursts. Despite this work and many others [14, 15, 16, 17, 18, 19, 20, 21], we believe that the current state-of-the-art, including our own work, is far from have reached a sufficient understanding of the locality of BGP routing events.

In this work, we analyse BGP update streams from multiple public observation points. We study how much of the BGP prefix updates received in the Internet are dependent on each other. We also estimate the locality and spread of routing events and classify streams of BGP prefix updates caused by these events as either correlated or not, i.e., as either global or local. Coming up with such a classification is not as straightforward as might appear at first. Indeed, developing the necessary methodology reveals multiple challenges in analysing BGP streams in time and space, such as distinguishing between spikes and noise, matching the different streams of BGP updates efficiently, estimating the distance (in AS hops) as well as the maximum inter-arrival time between two correlated spikes. We believe that our methodology as well as our results will help further work to develop a much better methodology to understand the spatio-temporal dynamics of BGP, and the actual locality of BGP routing events.

Our contributions are the following:

- We analyse the correlation of prefixes in BGP update messages when seen from multiple vantage points.
- We find that a significant part of BGP updates in fact contains correlated information about routing events that are seen by many other routers - up to 94%. The number of such correlated updates strongly depends on the number of peering links among the other routers that also observe these events.
- Using correlation in BGP update messages together with AS topology information, we propose a methodology[1] capable of identifying the locality of routing events caused by bursts of BGP updates, by dividing them into either global and local events.

[1]The source code is available on demand from the authors.

- Our results show that a large majority of BGP updates observed from publicly available vantage points are related to routing events that propagate multiple AS hops and involve a limited number of BGP prefixes. Further, large bursts of BGP updates in terms of BGP prefixes tend to be more localised than small bursts.

The remainder of this paper is structured as follows. We start in Section 2 by describing the data used. We examine situation when routers receive identical updates ("correlated spikes") from their neighbours in Section 3. Section 4 presents our classification of BGP updates based on the notion of "correlated spikes". In Section 4, we provide a basic classification which is refined in Section 5. We discuss the wider implications of our work in Section 6 and summarise our work in Section 7. An appendix (Appendix B) provides details on the spike classification methods used in Sections 4 and 5.

2. Data

In this paper, we use BGP data from RouteViews [22] and RIPE NCC [23], as well as AS-level maps of the Internet from the Internet Topology Collection [24] and CAIDA [25].

The RouteViews Project [22] collects real-time information about the interdomain routing system. Another public source of BGP data is available from RIPE NCC [23] - Regional Internet Registry. Both RouteViews and RIPE route collectors rely on multi-hop BGP sessions to peer with different ASs at multiple locations. All received updates, as well as collectors' RIBs[2] are stored in MRT format. Each update has a timestamp in Unix time format (route collectors are NTP-synchronised).

In total, we have analysed 37 million globally distributed BGP update messages during a 24 hours period on June 1 2009 by 20 collectors: 7 at RouteViews [22] and 13 at RIPE NCC [23]. Table 1 summarises the collectors used, their location, the number of monitors[3] from which the collectors received BGP data, and the number[4] of BGP prefix updates we used from them.

[2]Routing Information Base, database containing routes to all known prefixes

[3]the number of "active" (constantly sending updates to the route collector) monitors could be less than shown in the table

[4]after conversion of dumps to ASCII machine-readable format with "route_btoa" utility from "mrt" package

Collector	Location	Number of monitors	Number of prefix update
route-views2.routeviews.org	University of Oregon, Eugene, USA	32	10128729
route-views4.routeviews.org	University of Oregon, Eugene, USA	6	1371562
route-views.eqix.routeviews.org	Equinix, Ashburn, USA	7	1857771
route-views.isc.routeviews.org	ISC (PAIX), Palo Alto, USA	12	2849866
route-views.linx.routeviews.org	LINX, London, UK	16	6868316
route-views.wide.routeviews.org	DIXIE (NSPIXP), Tokyo, Japan	5	464883
route-views.kixp.routeviews.org	KIXP, Nairobi, Kenya	1	24474
rrc00	RIPE NCC, Amsterdam, Netherlands	15	2018173
rrc01	LINX, London, UK	82	45717
rrc03	AMS-IX, Amsterdam, Netherlands	99	659075
rrc04	CIXP, Geneva, Switzerland	15	852655
rrc05	VIX, Vienna, Austria	53	1644456
rrc06	JPIX, Otemachi, Japan	7	182455
rrc07	NETNOD, Stockholm, Sweden	19	1079721
rrc10	MIX, Milan, Italy	19	1048222
rrc11	NYIIX, New York, USA	36	2212838
rrc12	DE-CIX, Frankfurt, Germany	53	2655014
rrc13	MSK-IX, Moscow, Russia	24	1432168
rrc15	PTTMetro-SP, Sao Paulo, Brazil	12	448168
rrc16	NOTA, Miami, USA	7	1864

Table 1: BGP collectors.

In addition, we provide a historical overview of correlated spikes' rates based on the data shown in the table 2.

date	route-views2		route-views4		route-views.eqix		route-views.isc		route-views.linx		route-views.wide	
	monitors	prefixes	monitors	prefixes	monitors	prefixes	monitors	prefixes	monitors	prefixes	monitors	prefixes
01.06.2004	34	10632554	n/a	n/a	1	99898	12	2406783	8	2024018	4	724403
01.06.2005	32	9702122	n/a	n/a	7	1374402	12	2667821	13	4395336	4	468155
01.06.2006	37	14555094	n/a	n/a	9	1782556	13	3053438	13	5859430	5	428558
01.06.2007	39	13417022	n/a	n/a	9	2019160	14	2271952	16	6806687	4	253855
01.06.2008	38	14668210	n/a	n/a	8	523652	12	1639538	14	2814849	4	169436
01.06.2009	32	10128729	6	1371562	7	1857771	12	2849866	16	6868316	5	464883
01.06.2010	18	8058356	3	422399	8	996500	13	2592709	20	10320928	5	550650
02.06.2011	12	4416961	2	728313	5	615974	3	437543	6	4215134	n/a	n/a
28.05.2012	11	5626518	2	512704	5	569179	2	878920	6	2838248	n/a	n/a

Table 2: Collectors selected for historical overview.

To calculate the distribution of number of origin ASs in Section 5.4, we also bring into play the RIBs for the each route collector. We used the files mentioned in table 3 to restore AS Path and, therefore, origin AS for prefix withdrawals. To achieve this goal, for each collector we selected the RIB dump at the closest date before 00:00, June 1 2009. We load[5] the RIB for every collector before loading of update messages. During reading the update

[5]after conversion to ASCII machine-readable format with "bgpdump" utility

messages, we update the information in every RIB to be actual at any time we find a prefix withdrawal.

Collector	RIB file	Number of routes
route-views2.routeviews.org	rib.20090601.0000	11160773
route-views4.routeviews.org	rib.20090601.0000	1538768
route-views.eqix.routeviews.org	rib.20090531.2248	1425491
route-views.isc.routeviews.org	rib.20090531.2225	2631297
route-views.linx.routeviews.org	rib.20090531.2259	6364523
route-views.wide.routeviews.org	rib.20090531.2334	1149802
route-views.kixp.routeviews.org	rib.20090531.2324	260
rrc00	bview.20090531.1559	3425210
rrc01	bview.20090531.1559	2844752
rrc03	bview.20090531.1559	5134915
rrc04	bview.20090531.1559	1775275
rrc05	bview.20090531.1559	2369055
rrc06	bview.20090531.1559	571827
rrc07	bview.20090531.1559	1438606
rrc10	bview.20090531.1559	1630815
rrc11	bview.20090531.1559	2239387
rrc12	bview.20090531.1559	2488528
rrc13	bview.20090531.1559	2565264
rrc15	bview.20090531.1559	618140
rrc16	bview.20090531.1559	586080

Table 3: RIB files used.

We augment the AS-level topology provided by the BGP data from the considered collectors with AS topologies obtained from the Internet Topology Collection [24], also from June 1 2009. The Internet Topology Collection [24] provides daily AS-level Internet topologies based on BGP data (routing tables and updates), route servers and looking glasses. Each node/link in their database is annotated with the time it was first observed and the time it was last observed [24]. Information about the inferred geographic location of each link in the ASPATH is also included.

To estimate the impact of visibility on route collectors (Section 4), we also utilise the "IPv4 Routed /24 AS Links Dataset" from CAIDA [25], which "provides regular snapshots of AS links derived from traceroute measurements" [25]. We used three map files from early June: t1 and t2 collected between May 31 and June 2 2009, and t3 collected between May 31 and June 5 2009. These files cover the period closest to the date of analysed BGP updates, June 1st 2009. Both direct (marked with "D") and indirect links (marked with "I") were extracted from the files.

Table 4 provides a summary of the number of ASs and AS links visible from the AS topology maps. Note that despite its modest size compared to

the Internet Topology Collection, the CAIDA AS topology maps are useful as they have been shown to better sample the dense interconnectivity of the core of the Internet [26].

Source	Number of AS-level links	Number of ASs
Internet Topology Collection	135,627	33,549
CAIDA t1	41,346	16,304
CAIDA t2	41,875	16,419
CAIDA t3	38,209	16,319

Table 4: AS topology maps.

3. Correlated spikes

BGP updates received by routers across the Internet are expected to contain some amount of correlation. In the early days of the commercial Internet, some correlation was shown to be related to mistakes in the immature BGP implementations of router vendors [27]. Nowadays, with the maturity of BGP implementations, such causes of BGP correlation are very unlikely. This does not mean that correlation in BGP updates does not happen in the Internet [13]. Indeed, events such as router or link failures can cause both intra- and inter-AS BGP route propagation. In such cases, routers may receive BGP information related to such an event from multiple neighbours.

The sheer size of today's Internet coupled with the complexity of the interdomain routing system makes it very difficult to infer the cause of BGP updates across the Internet [12]. In this paper, we make a step forward towards better understanding the local and global—at the topological level—nature of routing events through the viewpoint of correlation in BGP updates. We define the notion of a "correlated spike" of prefix updates (announcements or withdrawals), which captures cases where the BGP updates received by a router from its neighbours contain information about the same prefixes. In this section, we quantify the amount of visible correlation in BGP updates as seen by publicly available vantage points and show that the vast majority of all BGP prefix updates are correlated.

3.1. Illustrative example

Before quantifying the amount of correlation in BGP updates, we introduce the notion of a "correlated spikes". For this, we rely on an illustrative example where a router receives BGP updates for the same prefix set from different neighbours.

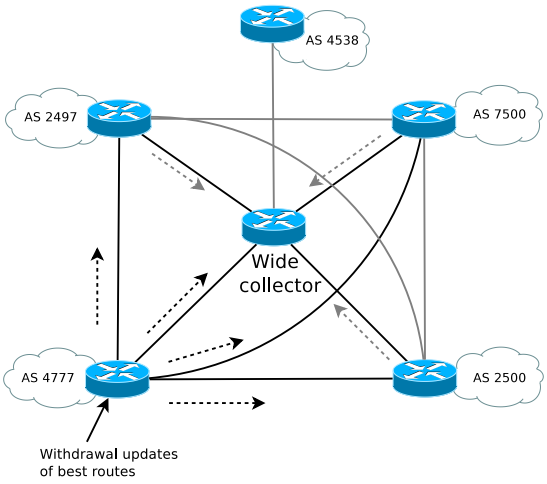

Figure 1: Propagation of prefix update set causes correlated spikes seen on "Wide" collector.

Assume that a BGP router receives a large number of different prefix withdrawals at the same time. For each prefix, it deletes the corresponding entry from its Adj-RIB-In[6], i.e., only for the peer from which the update was received. If the withdrawn route is currently the best route stored in the Loc-RIB[7], the router chooses an alternative best route and if none exists, marks the network as unreachable. As the best path for that prefix has changed, the router must send an update to all concerned peers. As shown on the Figure 1, when the BGP router in AS4777 receives a withdrawals for its best routes, it will forward withdrawals to its peers - routers from AS2500, AS2497 and AS7500, as well as the "Wide" RouteViews collector. If these withdrawals also affect best routes, routers from AS2500, AS2497, AS7500 will forward

[6]The Adj-RIB-In is a routing information base which containfs unprocessed routing information that has been received by the local BGP speaker from its peers [28].

[7]The Loc-RIB is a routing information base which contains the routes that have been selected as best by the local BGP speaker's Decision Process [28].

updates further to their peers, including "Wide". In this case we may see those "correlated spikes", that are visible for the "Wide" collector - the same amount of same prefix updates from each of its neighbours: AS4777, AS2497, AS2500, and AS7500.

In this paper, we are interested in understanding the prevalence of such situations, by identifying spikes of prefix updates and quantifying the redundancy between different update spikes. Correlated spikes are interesting not on their own, but because they reflect BGP propagation and shed light on its locality.

3.2. Definition of correlated spikes

To quantify correlation in the raw BGP data, i.e., a dump of updates received by a collector, we need to group updates into spikes coming from different monitors. Each spike represents a group of prefix updates received at the same Unix timestamp—accurate within a second—from one of monitors (one neighbouring AS). We have compared such spikes received from different ASs by one collector. Figure 2 shows all BGP updates received by router "wide.routeviews.org", AS 6447 on June 1 2009, grouped into spikes received from different ASs.

The x-axis shows the time (in seconds) starting from the beginning of the observation period. The y-axis shows the number of BGP updates received from one of the neighbouring ASs over time.

The spiky nature of BGP updates time series makes it hard to distinguish individual spikes from Figure 2. However, many update spikes on the plot are *correlated* and are strongly clustered in time. They contain similar numbers of BGP updates and are received almost at the same time by the Wide router from different neighbours. Such *correlated* spikes of BGP updates are actually pretty common when inspecting the time series. These spikes are caused by situations when ASs propagate update messages received from their neighbours and these neighbours are interconnected.

We call a set of spikes of BGP updates **correlated** if they (1) were observed over a given period of time, (2) were received from multiple routers (whether the routers are in the same AS or not is to be specified) and (3) contain information about the same prefixes. However, all prefix updates in such spikes do not need to be identical, i.e., a subset of the prefixes need to be the same in the two compared spikes. The reason for this decision hides in the meaning of correlated spike: a correlated can be due to one or more

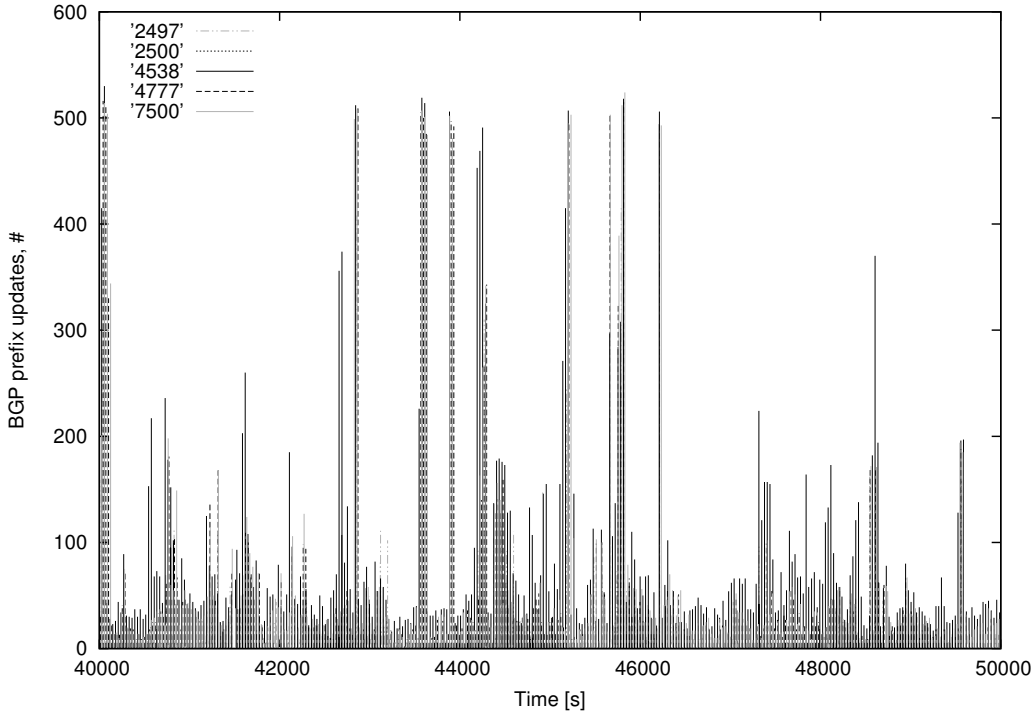

Figure 2: Time series of BGP updates.

routing events [8]. While checking spikes for correlation, if one spike contains prefixes from at least one other spike, we call them correlated. We require that a large enough fraction of prefixes are common to both spikes to consider them as correlated.

We also choose to neglect the type of BGP updates and require only common prefixes to be present, as a given routing event may lead to different streams of updates and withdrawals to be observed by different routers in the Internet [19].

3.3. Update spikes characteristics

Before comparing BGP update spikes, let us first understand their general characteristics. Are most spikes small or large? Are most BGP updates contained in small or large spikes? Figure 3(a) answers the first question,

[8] We are interested to capture the cases when one spike presents a superposition of different routing events, i.e. different sets of correlated spikes; see Section 4 for details.

while Figure 3(b) answers the second. Figure 3(a) provides the cumulative distribution of spikes as a function of their sizes (in number of prefixes). Figure 3(b) provides the cumulative number of BGP updates that belong to spikes smaller than a given size (in prefixes). From Figure 3(a), we observe that most spikes contain a relatively small number of BGP prefixes, e.g., 90% of them have less than 100 prefixes. On the other hand, Figure 3(b) shows that less than half of BGP updates actually belong to spikes containing up to 100 BGP prefixes. Most bursts of BGP updates are therefore relatively small, and a significant fraction of BGP updates belong to relatively large spikes, i.e., spikes that contain hundreds or thousands of BGP updates.

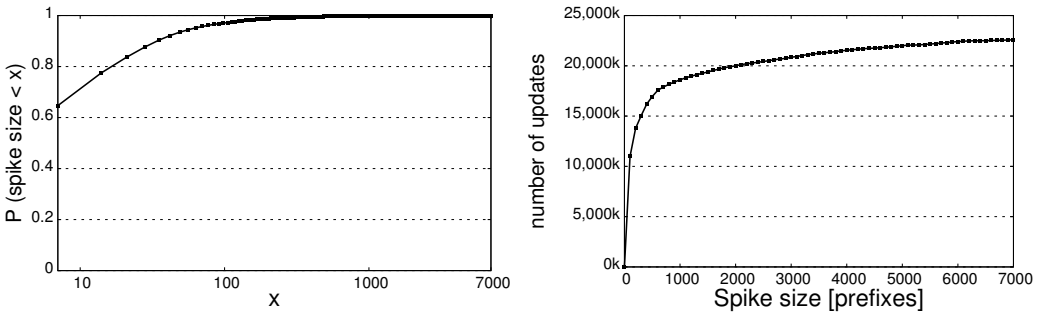

(a) Spike size distribution.

(b) Cumulative number of updates in BGP spikes.

Figure 3: Spike distribution

The implications of these results are twofold. First, one cannot ignore small BGP update spikes, as they make the majority of spikes and represent a significant fraction of all BGP updates. Second, one cannot ignore large update spikes, as they represent a large fraction of the BGP updates. All in all, this section implies that an analysis of BGP updates must include all bursts of updates to provide a representative view of the data.

3.4. Quantifying correlation in BGP updates

Now we compare spikes to find out which ones are correlated versions of the others. In this paper, we say that two spikes are correlated if they share a given fraction of prefixes and are separated in time by less than a given duration. To find correlated spikes in BGP data, we need to compare each spike with the other spikes received from other ASs within the considered time interval.

For example, when considering a time interval of 5 seconds between spikes, we limit the comparison of spikes to those received within 5 seconds before and after the time when the considered spike was received. When comparing two spikes, if the smallest of them contains more than a predefined percentage of the prefixes from the other, both spikes will be marked as correlated.

During the comparison process, we have used different values for the two parameters - *time interval* and *minimum percentage of common prefixes* in the two spikes. For the minimum percentage of common prefixes, values of 10%, 40%, 70% and 99% were used. For the time interval, we decided to use values closely related to MRAI[9] timer values, which has been shown to play a critical role in BGP convergence [29, 27, 30, 31]. Two update messages sent by a BGP speaker to a peer and related to a given prefix must be separated by a time at least equal to the value of the MRAI timer [28]. The suggested default value for the MRAI Timer on eBGP sessions is 30 seconds [28]. We decided to use values not larger than this default value for the MRAI Timer (5, 15 and 30 seconds) as the maximum time interval between two compared spikes, to avoid marking spikes containing update messages caused by route flapping as correlated.

Furthermore, we mark out correlated spikes received from neighbours directly interconnected with each other. For example, among the neighbours of observation point "wide.routeviews.org", AS4777, AS2500, AS2497 and AS7500 are connected, see Figure 1. In this case, the Wide observation point is highly likely to receive correlated spikes from AS4777, AS2500, AS2497 and AS7500.

We present on Figure 4 the fraction of BGP updates in correlated spikes, as seen by 6 vantage points around the world, for different time interval values[10] and minimum percentage of common prefixes in two spikes to be considered correlated. The three plots of Figure 4 are similar. Larger values of the time interval increases the fraction of BGP updates that will be included in the correlated spikes, as expected. Indeed, a larger time interval will allow more BGP updates to be included in the comparison. Only one observation point (LINX in London), does not show a noticeable impact of the value of the time interval, due to the very high fraction of correlated spikes.

[9]Minimum Route Advertisement Interval Timer [28].

[10]We also verified that the results shown hold for higher values of the time interval, e.g., 60s or 120s, in case MRAI would affect them.

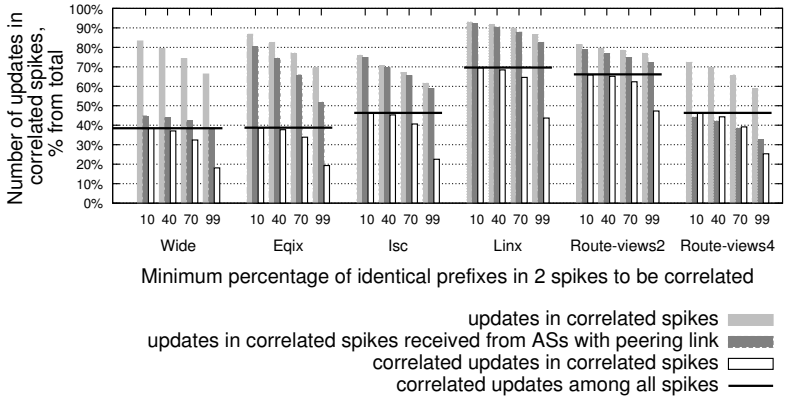

(a) 5s time interval.

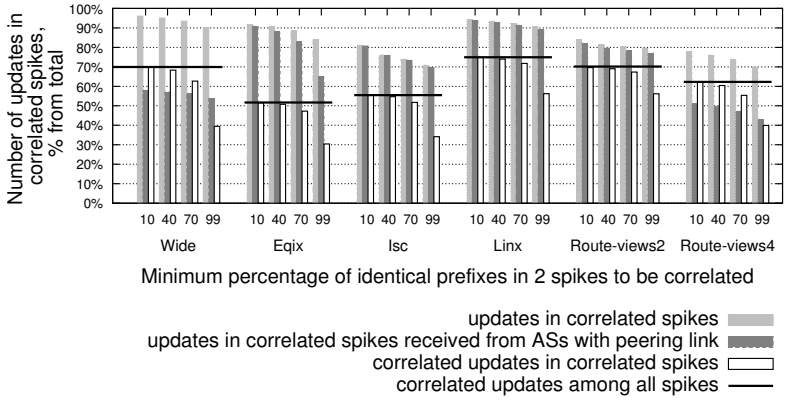

(b) 15s time interval.

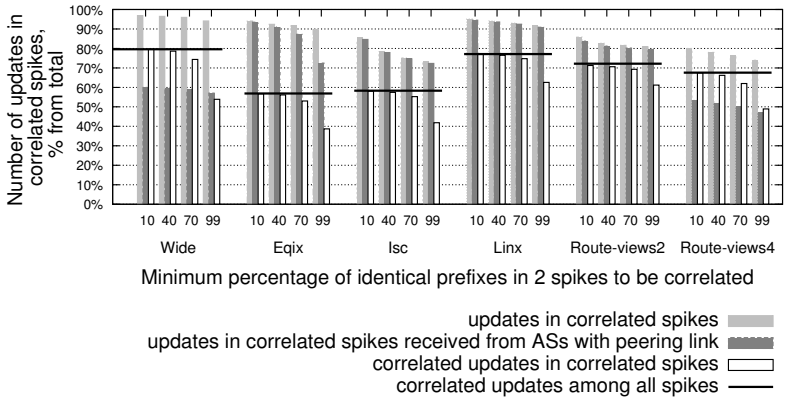

(c) 30s time interval.

Figure 4: Number of BGP updates in correlated spikes, for different time intervals and minimum percentage of common prefixes.

The minimum percentage of common prefixes in two spikes to be correlated is clearly relevant for all observation points: the higher this percentage, the smaller the fraction of BGP updates in the correlated spikes. When a larger fraction of common prefixes is required for two spikes to be considered correlated, less and less BGP updates spikes satisfy the requirements. The noisy nature of BGP propagation and updates [32] makes it difficult to clearly distinguish subsets of BGP updates that actually correspond to information about similar BGP events. To estimate the amount of updates that should not be considered as correlated in correlated spikes, we compare the fraction of prefixes in correlated spikes with the number of prefixes actually marked as correlated (white bars on Figure 4), as well as with the number of correlated prefixes among all spikes (the black horizontal lines). The difference between the light-grey and white bars on Figure 4 corresponds to updates that belong to spikes classified as correlated, but that actually do not belong to correlated prefixes in the spike. We observe that the number of correlated prefixes among all spikes are a higher bound on the number of correlated updates in the correlated spikes. Figure 4 also distinguishes between the percentage of prefix updates in correlated spikes (light-grey bars), and the same percentage when considering only BGP updates received from neighbours that are connected to each other (dark-grey bars). We notice that three vantage points, namely ISC, LINX and Route-Views2, are not sensitive to whether all BGP updates are considered, or just those from neighbours that are interconnected. This is because at these vantage points, most neighbours are interconnected. This hints at the importance of interconnectivity among the neighbours of the vantage points. We will come back to this soon.

The historical overview on Figure 5 shows that rates of prefix updates in correlated spikes remains relatively stable from 2004 to 2009, except for the Wide vantage point. However, the rate highly depends on the connectivity among the neighbours of the monitors. For example, the decrease of active monitors among route collectors' neighbours causes the percentage to plummet in 2010-2012, see table 2, though less for routeviews2.

The main take-away message from Figures 4 and 5 is that a major fraction of BGP updates are correlated, irrespective of the tuning of the parameters used to identify correlated spikes. Rates of correlated prefix updates are stable across years, despite dips caused by changes in the connectivity between neighbours of the vantage points. This consequence is important for BGP analysis, as it means that a significant amount of prefix correlation exists in publicly available BGP data, which might be exploited to better understand

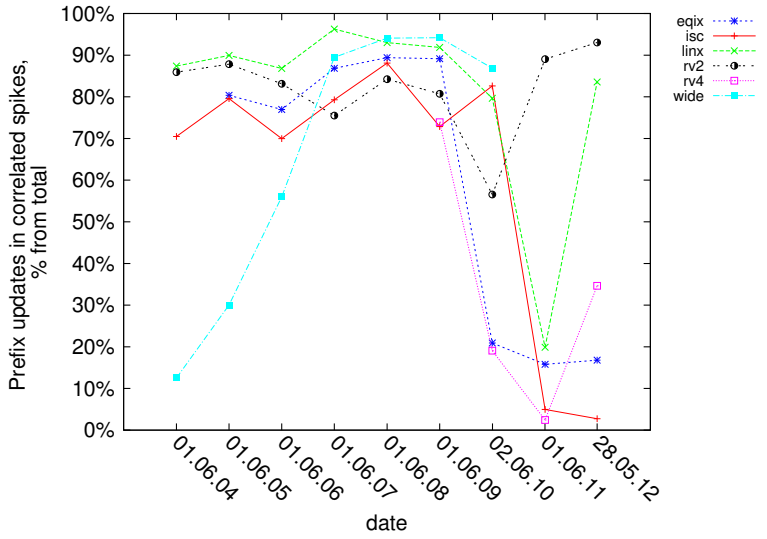

Figure 5: Historical changes in the rates of correlated update messages for selected route collectors, 30s time interval, 99% same prefix updates in two spikes to be concerned correlated.

the origin of BGP routing events [12] as well as the propagation of BGP update messages in the Internet [33].

We now come back to the relation between the amount of BGP updates in correlated spikes and the interconnectivity among neighbours from which a given observation point receives BGP updates. To illustrate the dependency between the correlated spikes and interconnection between neighbouring ASs, Table 5 provides—for the 6 previously considered observation points—the average number of BGP updates per spike, the average number of AS relationships between the neighbours of this route collector, the percentage of correlated spikes, as well as the percentage of BGP updates in correlated spikes. Observation points in Table 5 are ordered by increasing average number of AS relationships between their neighbours. We observe that as the average connectivity between neighbours increases, the fraction of correlated spikes also increases, irrespective of the time interval used. Note that when the percentage of BGP updates is considered (last 3 columns), not the percentage of spikes, the numbers are not increasing with the AS relationships with neighbours. This is to be expected, as spikes of BGP updates may contain arbitrarily small or large number of BGP updates (see Appendix A).

Observation point	Updates per spike	AS relationships between neighbours	AS relationships per neighbour	Correlated spikes, 5s interval	Correlated spikes, 15s interval	Correlated spikes, 30s interval	Updates in correlated spikes, 5s interval	Updates in correlated spikes, 15s interval	Updates in correlated spikes, 30s interval
Route-Views4	13.48	4	0.57	69%	85%	91%	59%	70%	74%
WIDE	18.53	6	1.2	55%	86%	93%	66%	90%	94%
EQIX	20.89	30	2.73	72%	88%	93%	70%	84%	90%
ISC	26.47	64	4.27	81%	92%	95%	61%	70%	73%
LINX	22.27	101	5.61	92%	95%	96%	88%	91%	92%
Route-Views2	17.48	270	7.11	95%	97%	98%	77%	79%	81%

Table 5: Relation between the AS connectivity of route collectors and correlated spikes (99% of common prefixes in two spikes to be correlated).

3.5. Summary

We found in this section that a significant fraction of BGP spikes as well as BGP updates are correlated, when seen from multiple vantage points. These correlated spikes make up to 94% of all BGP updates seen at a given observation point. Despite variations in the fraction of correlated spikes and updates across vantage points or parameters of the methodology, we conclude that our methodology is able to identify correlation, and is not too sensitive to the values of the parameters.

We identified that the amount of correlation is strongly related to the connectivity between neighbours of a given observation point.

4. Correlated spikes and locality of BGP propagation

Correlation in BGP updates is related to the local or global nature of BGP propagation for the corresponding routing events. A routing event that is very local, i.e., only a few routers will receive prefix updates caused by the event, will lead to less observed correlation compared to a global event, where prefix updates related to the event will be propagated through the majority of ASs in the Internet causing changes in routing tables. Therefore, in this section, we analyse correlated spikes to better understand how they can help us find out what fraction of BGP updates stem from local events, and what fraction from global events. Of course, we cannot expect to reverse-engineer routing events, as it has been shown that locating routing events is a very difficult problem [12]. Therefore, our goal in this paper is to use the notion of correlated spikes to help understanding the propagation of BGP updates in the Internet, and more specifically their locality.

4.1. BGP spikes: single or correlated?

In this Section, we want to distinguish between spikes caused by local routing events, from those caused by global events. It is possible to use correlated spikes for this purpose, as they reflect the BGP update propagation process. Indeed, during the propagation of routing changes caused by one event or set of events, correlated spikes containing updates for the same prefixes will be received by multiple BGP routers, including publicly available vantage points. It is reasonable to assume that global events will lead to correlated spikes that are observed from all vantage points. On the contrary, correlated spikes caused by local events will be observed by very few vantage points.

For each spike, we identify correlated spikes (see Section 3.2). As opposed to the previous section, we compare spikes using data from all collectors. We therefore change the considered time interval to 120 seconds before or after the spike to correspond to BGP convergence time [34]. To achieve certainty considering correlated spikes, we always require them to have 99% of identical prefixes.

If we fail to find correlated spikes for a given spike, this spike might have been caused by a local routing event. On the contrary, if we find correlated spikes, this spike might have been caused by a global routing event. However, the presence or absence of correlated spikes does not have to constitute a signature of local and global events. Indeed, a limited number of observation points might not have a good enough visibility to conclusively distinguish between local and global routing events. Therefore, our goal is first to classify spikes (see the classification method in the Appendix B.1) according to whether we find compelling evidence that this spike comes as a set of correlated spikes, or if the spike is single.

From our observations in Section 3.3, we found that the BGP spikes that need to be considered in our analysis span all possible sizes from small to large. In the remainder of this paper, we group spikes according to bins with a step size of 100. All spikes of size 0..99 belong to the first bin. We ignore all spikes containing 7500 or more BGP updates, as they are too few and account for only 2.8% of the total number of BGP updates. The last considered bin therefore contains spikes of sizes 7400..7499.

Now, we analyse all BGP updates and classify them according to whether they come as single or correlated spikes. Figure 6 shows—for each individual bin—the fraction of BGP updates that were classified as single or correlated spikes. We observe from Figure 6 that most spikes containing few BGP

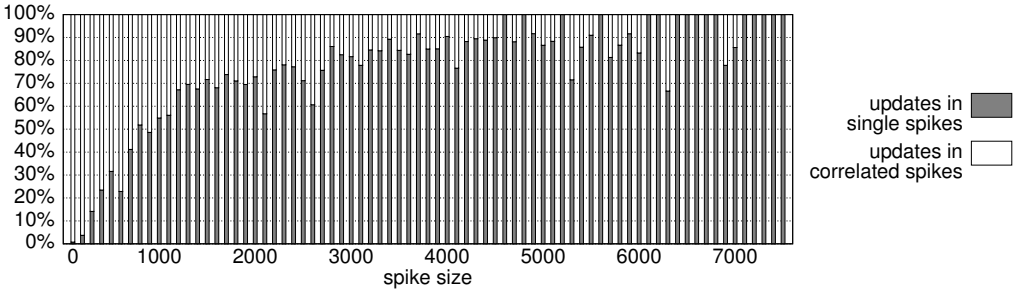

Figure 6: Fraction of BGP updates in single and correlated spikes (120s time interval, minimum 99% common prefixes in correlated spikes, 0.1 threshold), on a per-bin basis.

updates are correlated spikes. The larger spikes on the other hand tend to be single. When considering the total fraction of BGP updates, see Figure 7, we observe that most BGP updates are indeed correlated (80.58%). Small spikes make up a significant fraction of the total BGP updates, and are mostly correlated. Larger spikes, while containing a majority of single spikes, are not numerous enough to make single spikes represent a significant fraction of the total BGP updates, with only 19.42% of all BGP updates. Note that changing the value of the parameters, e.g., the minimum percentage of common prefixes in correlated spikes, does not question the conclusions of this section given that the majority of BGP updates belong to correlated spikes (see Appendix B).

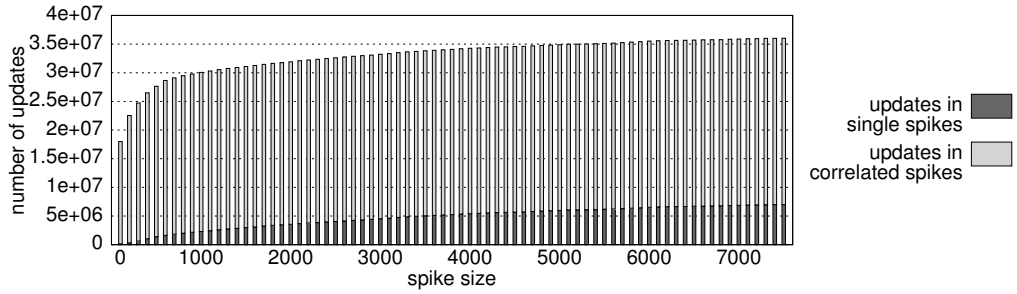

Figure 7: Cumulative number of BGP updates in single and correlated spikes (120s time interval, minimum 99% common prefixes in correlated spikes, 0.1 threshold), on a per-bin basis.

4.2. Impact of visibility in AS connectivity

So far, we classified spikes into correlated or single, assuming that our observation points have sufficient visibility of the connectivity between the

ASs from which our observation points receive BGP updates. The percentage of correlated spikes we detect depends on our visibility of the connectivity between ASs. We now estimate how our results relate to the visibility of the vantage points used.

Table 6 presents the visible connectivity between the 172 ASs[11] from which our 20 observation points receive BGP updates, based on the Internet Topology Collection (ITC) [24, 35], as well as the CAIDA AS topology [25], and their intersection. The second column of Table 6, labeled "Neighbours with connectivity", shows how many ASs, among the 172 from which the considered observation points have at least one neighbour, are seen from the different AS topology datasets. The last two columns of Table 6 provide the number of ASs among the 172 that have neighbours which also belong to these 172 - at least one neighbour (third column, "interconnected neighbours") or all neighbours (fourth column, "fully visible neighbours"), according to the different AS topology datasets. In other words, we say that an AS has "fully visible neighbours" only when both this AS and all its neighbours are connected to observation points. In this case we say that our visibility for this AS is full.

AS topology map	Neighbours with connectivity	Interconnected neighbours	Fully visible neighbours
ITC	172	172	1
CAIDA	155	142	13
ITC $\cap$ CAIDA	155	140	15

Table 6: Visibility of AS from which our observation points receive BGP updates.

As can be observed from Table 6, both the ITC and the CAIDA datasets have a similar view of the ASs from which the collectors obtain BGP data, at least in terms of which of them are connected to at least one other of these 171 ASs. The interest of the CAIDA dataset from our viewpoint is revealed in the fully visible neighbours. Indeed, from the CAIDA dataset, we get a larger number of ASs among the 172 whose neighbours are all in the 172 from which we receive BGP data. Of course this does not mean that the CAIDA data has a better coverage compared to the ITC one. It means that because

[11] number of "active" ASs that actually sent data to route collectors during the observation period

the ITC dataset has so much more coverage, it is much more difficult to have a good visibility of the BGP data of the neighbours of a given AS, which is important to have confidence in the amount of observed correlation. Note that the number of fully visible neighbours from the intersection of the ITC and the CAIDA datasets is higher than for the individual datasets. Indeed, these ASs have more limited connectivity when restricting ourselves to the intersection of the datasets, so there are more of them.

If we consider only the 15 ASs for which we have BGP data from all their visible neighbours, the percentage of BGP updates in correlated spikes comes to $82,58\%$ (for one-second spikes with size less than 7500 prefix updates). This observation supports our results from section 4.1.

4.3. Summary

On top of the quantification of correlation in BGP updates, this section has shown the complex relationship between AS-level connectivity and the extent to which correlation can be observed. We believe that our quantification is meaningful, as it is actually a lower bound on the actual correlation that would be observed if more observation points were considered and a more complete AS-level topology was available. We therefore argue that studies of BGP propagation, e.g., [36, 13], need to consider the impact of the current limitations in available data [37] on the relation between observed BGP updates and routing, as already noted by previous work [12, 38].

5. Global propagation of correlated spikes

In the previous section, we classified BGP spikes into correlated or single. Classifying a spike as correlated does not tell us directly how global the corresponding routing event is. It only says whether a set of BGP updates was received from multiple ASs. If an event is purely local to a given AS, we should not be able to see correlated spikes for it. At the other end of the spectrum, a global event will be seen by all observation points and should be received from all ASs. In this section, we refine the previous analysis by classifying spikes, both correlated and single, and we attempt to relate the refined classification to the locality in the BGP propagation.

5.1. Exploiting AS distance information

To provide better insight into the global nature of correlated spikes, we compute the AS hop distance between any two ASs from which the group of

correlated spikes were received. We compute the shortest AS path between two ASs based on AS topology maps (see Section 2). For each group of correlated spikes, we then compute the maximum over all the shortest AS paths between any two ASs in the group.

This maximum AS hop distance corresponding to a given group of correlated spikes will be used to infer the degree of locality of the corresponding routing events. For example, if a group of correlated spikes was received from routers separated by 4 AS hops, we can safely assume that the corresponding routing event (or the set of events) that caused the given spike is global and spread across a significant fraction of the Internet. On the contrary, when a group of correlated spikes was observed only from 2 routers belonging to 2 directly interconnected ASs, we can assume that the routing event or set of events that caused this spike is relatively local.

Given that we know that our classification of single spikes might be biased by limited visibility of the AS topology, we decided to refine our classification into the following classes:

- Single with 0-33% of neighbours visible
- Single with 34-66% of neighbours visible
- Single with 67-100% of neighbours visible
- 1 hop correlated
- 2 hops correlated
- ...
- n hops correlated

We divided the single spikes into 3 groups, depending on the visibility we have about connectivity between the ASs that sent these single spikes. Indeed, if a spike is classified as single, the likelihood that this spike is really single increases as the visibility of the connectivity between the neighbours of the corresponding AS increases. Further, as the number of hops between correlated spikes increases, the likelihood that the underlying set of routing events are global also increases. In this way, we can assess the extent of the locality of routing events in the Internet.

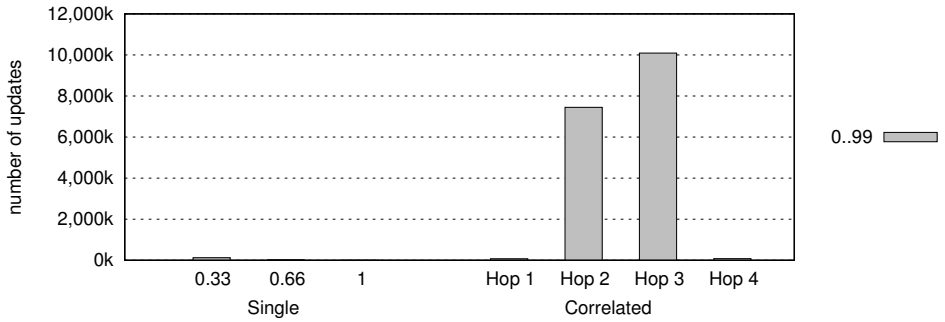

Figure 8: Classification of spikes with size from 0 to 99. Time interval - 120 seconds before and after given spike. 99% of same prefixes in 2 spikes to be correlated.

5.2. How much of BGP propagation is global?

We now turn to the results of our classification (see Appendix B.2 for a detailed description). From Section 3, we know that the majority of BGP updates observed from publicly available data are correlated. We also know that small spikes tend to contain a larger fraction of correlated updates compared to large spikes. Together, these observations suggest that the majority of BGP updates might actually stem from routing events that propagate globally, but that do not involve a large number of BGP prefixes. In this section, we aim to quantify the extent to which the underlying routing events propagate globally, by using the maximum AS hop distance among correlated spikes as a measure.

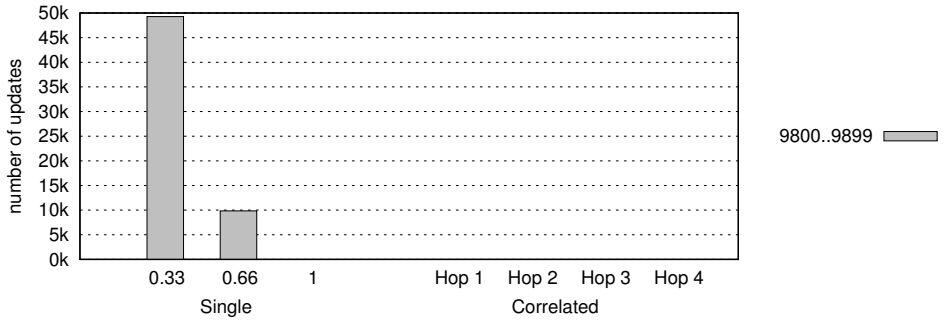

Figure 9: Classification of spikes with size from 9800 to 9899. Time interval - 120 seconds before and after given spike. 99% of same prefixes in 2 spikes to be correlated.

We first look at the small spikes, made of up to 99 BGP prefixes, as these spikes represent about half of all BGP updates we observe. Figure 8 displays

the breakdown of spikes of size up to 99 BGP prefixes into the different maximum AS hop distances. We observe that most spikes have propagated 2 or 3 AS hops. Given that most BGP paths in the Internet are 3 to 5 AS hops long [23, 22], the underlying routing events are definitely not local. On the other hand, they do not have to be Internet-wide.

Spikes of size between 9800 and 9899 BGP prefixes show a very different picture, see Figure 9. These spikes are mostly classified as single, and are learned from ASs for which our visibility is limited (less than a third of their neighbours are observed by the collectors). Note that the limited visibility is partly a consequence of the relatively good coverage of the ITC dataset. The question of how different the picture would be if we had a better coverage of the Internet from BGP data is open. We partially addressed this question in the previous section, when we restricted the considered ASs to those for which we have good visibility from BGP. However, the number of such ASs for which we have good visibility is too small to reach a conclusion.

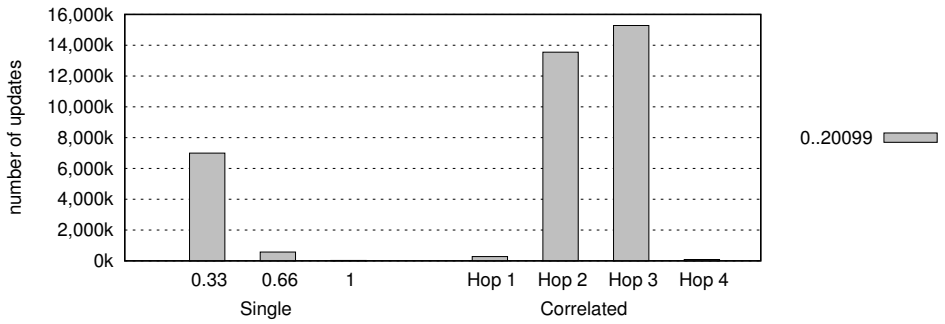

Figure 10: Classification of spikes with size from 0 to 20099. Time interval - 120 seconds before and after given spike. 99% of same prefixes in 2 spikes to be correlated.

The results for all spike sizes (Figure 10) confirm the two previous trends, with a majority of BGP updates in correlated spikes that relate to non-local routing events, and the rest being single spikes that are likely to be caused by local routing events.

We can conclude that unless publicly available data gives us a misleading picture of spikes correlation due to incomplete AS-level data and not enough observation points, most BGP updates are indeed correlated, belong to bursts of BGP data containing few prefixes, and caused by routing events that propagate a few AS hops away, i.e., not local.

5.3. Leveraging propagation time

Given the possible limitations of the AS-level data, we provide in this section a different type of evidence to corroborate the previous results. BGP propagation in the real Internet takes time, due to the BGP timers such as MRAI [27, 29, 30, 31, 39, 33]. Therefore, by studying the timing of BGP spikes, it is possible to check whether our classification is confirmed by differences in the time at which correlated spikes were received. Indeed, inter-arrival time of updates caused by local routing event (or set of events) should be relatively small when compared to the propagation time of global routing events.

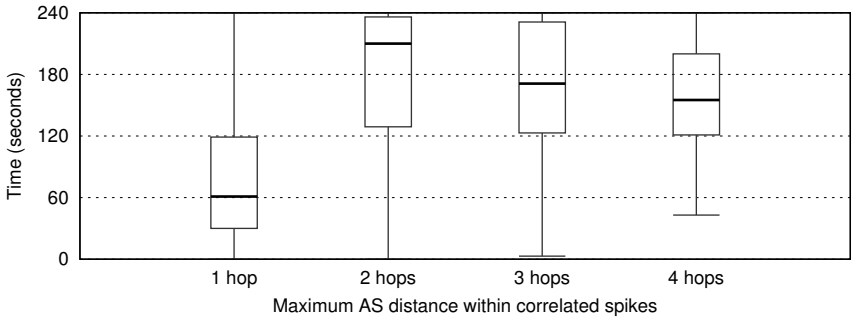

Figure 11: Distribution of maximum time interval between 2 spikes in a group of correlated spikes. Time interval - 120 seconds before and after a spike. 99% of same prefixes in 2 spikes to be correlated.

Figure 11 shows, using box plots, the distribution of inter-arrival time between the first and the last spikes in the group of correlated spikes, using a 120 seconds time interval. As for the distance calculation, we consider only the $\langle$*observation AS*, *observation router*$\rangle$ pairs with a sufficient number of updates in their spikes in our analysis (using conditions B.4 and B.5, see Appendix B). The box plots show the minimum and maximum, as well as the 25^{th}, 50^{th}, and 75^{th} percentiles of the distribution.

From Figure 11, we observe the strong difference in the median time between spikes for those separated by one AS hop, compared to the others. When correlated spikes are separated by strictly more than one AS hop, the BGP propagation time is well above 2 minutes for most spikes, as expected from previous studies about BGP convergence [27, 29, 30]. Note that due to the widely different configurations of MRAI timers in the Internet, as well

as the specific nature of some events, some spikes are able to propagate very fast.

5.4. Origin ASs

To shed some light on the locality of correlated spikes, we inspect the number of different origin ASs that appear in them. If a correlated spike corresponds to a single origin AS, we expect that the routing event is likely to be related to this very origin AS. Indeed, routing events that take place in the middle of the AS path would likely affect a set of prefixes whose best paths go through this particular part of the network.

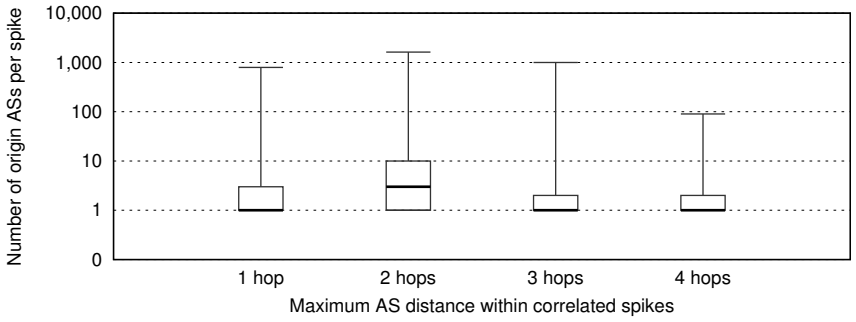

Figure 12: Distribution of number of origin ASs in spikes considered as correlated. Time interval - 120 seconds before and after a spike. 99% of same prefixes in 2 spikes to be correlated.

On Figure 12, we provide, for each AS hop distance, the number of origin ASs present in correlated spikes. For each AS hop distance, the box plots show the distribution of the number of origin ASs. Note that for some prefix withdrawals we were not able to identify the origin AS. We therefore have excluded such withdrawals from the analysis. We observe on Figure 12 that most correlated spikes correspond to a single origin AS. Very few correlated spikes involve a large number of origin ASs, as would happen in the case of a routing event that takes place in a transit network for example, likely involving a significant number of best path changes in the network. For an AS hop distance of 2, 25% of the spikes involve more than 10 origin ASs, and are unlikely to be related to the origin ASs. Therefore, we conclude that most correlated spikes are related to local events related to the origin AS, but that propagate across the Internet.

6. Discussion

The current mental model of the research community regarding the propagation of BGP updates is that most of BGP updates relate to routing events that propagate globally in the routing system. This belief originates in previous work on routing dynamics [40, 41, 6] and the use of address space [1], suggesting that a large number of BGP updates are unduly caused by specific types of networks, namely multi-homed stub ASs, which then propagate globally. Contrary evidence was brought in [2], showing that no specific type of network contributes more than its fair share of BGP updates, and that prefix de-aggregation or traffic engineering cannot be claimed as the reasons behind a major fraction of the BGP updates.

In this paper, we provide yet another view of the BGP routing system, observing that most BGP updates do belong to relatively small—in terms of number of prefixes—bursts of BGP updates that propagate a few AS hops away. Large bursts of BGP updates on the other hand stay relatively localised. Our observations partly support the current mental model that most of the observed BGP updates are indeed caused by globally propagating events. However, publicly available data does not allow us to be sure about the exact amount of correlation, as the data may miss local events. If such information was available from at least a few networks, we could increase our confidence in the relative share of local and global routing events. Given the methodological challenges of analysing BGP data, and the incompleteness of publicly available data [37], especially at the many Internet Exchange Points [42, 43], it remains open how much exactly of the reality of BGP is known and how much of it can be observed. We therefore cannot provide a definitive answer to the question of the amount of BGP updates that are triggered by local events that do not propagate globally, and its consequence on the health of the interdomain routing system.

While correlation in BGP updates is not forcibly very insightful in terms of its direct applications, we showed that by combining it with other aspects of BGP routing, e.g., convergence time, there is still a lot of open space in better understanding the global BGP routing system. For example, we believe that analysing correlated spikes from a sufficient number of vantage points, studying the inter-arrival times of BGP updates and relating it to the AS topology, will enable tracing update bursts during propagation of routing events and therefore identify the likely location of a routing event.

7. Summary

In this work, we took a fresh look at BGP updates, namely the amount of correlation in the BGP updates received by different routers in the Internet. Thanks to this new way of looking at BGP, we developed a method to classify sets of BGP updates, called spikes, into either correlated or non-correlated (single). Based on publicly available data, we showed that a significant fraction of all BGP updates belong to correlated spikes, which are seen at multiple locations across the Internet. Further, these correlated spikes are actually relatively small when measured in terms of the number of BGP prefixes involved.

We related BGP updates correlation to the topological scope of BGP convergence, which allowed us to determine areas in the Internet topology affected by common prefix updates (likely caused by common routing events) and classify update spikes depending on their topological extent. We showed that correlated spikes are relatively global, i.e., they propagate at least 2 or 3 hops away. Correlated spikes are therefore the consequence of routing events that are likely to propagate across the whole Internet. Large bursts of BGP updates on the other hand tend to contain a larger share of BGP updates created by events that are local, i.e., which do not propagate globally in the Internet.

References

[1] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang. IPv4 Address Allocation and the BGP Routing Table Evolution. *ACM Computer Communication Review Magazine*, 35(1):71–80, 2005.

[2] L. Cittadini, W. Muhlbauer, S. Uhlig, R. Bush, O. Maennel, and P. Francois. Evolution of Internet Address Space Deaggregation: Myths and Reality. *IEEE Journal on Selected Areas in Communications*, 28(8):1238–1249, Oct 2010.

[3] Neal Leavitt. IPv6: Any Closer to Adoption? *IEEE Computer Magazine*, 44(9):14–16, 2011.

[4] S. Deering and R. Hinden. *RFC 2460 Internet Protocol, Version 6 (IPv6) Specification.* Internet Engineering Task Force, December 1998.

[5] G. Huston. BGP Routing Table Analysis Reports. http://bgp.potaroo.net/.

[6] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkranz. BGP Routing Dynamics Revisited. *ACM Computer Communication Review Magazine*, 37(2):5–16, 2007.

[7] Mohit Lad, Xiaoliang Zhao, Beichuan Zhang, Daniel Massey, and Lixia Zhang. Analysis of BGP Update Surge during Slammer Worm Attack. In *IWDC*, pages 66–79, 2003.

[8] Earl Zmijewski. Reckless Driving on the Internet, February 2009.

[9] M. Lad, D.Massey, and L. Zhang. Link-rank: A graphical tool for capturing BGP routing dynamics. In *Proceedings of IEEE/IFIP Network Operations and Management Symposium*, pages 627–640, 2004.

[10] A. Flavel, M. Roughan, N. Bean, and O. Maennel. Modeling BGP Table Fluctuations. In *Proceedings of the International Teletraffic Congress*, pages 141–153, 2007.

[11] Z. Mao, R. Bush, T. Griffin, and M. Roughan. BGP Beacons. In *Proceedings of ACM Internet Measurement Workshop*, pages 1–14, 2003.

[12] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs. Locating Internet routing instabilities. In *Proceedings of ACM Conference of the Special Interest Group on Data Communication (SIGCOMM)*, pages 205–218, 2004.

[13] J. Park, D. Jen, M. Lad, S. Amante, D. McPherson, and L. Zhang. Investigating occurrence of duplicate updates in BGP announcements. In *Proceedings of the Passive and Active Measurement Conference*, pages 11–20, 2010.

[14] x. Wu, X. Yin, Z. Wang, and M. Tang. A three-step dynamic threshold method to cluster BGP updates into routing events. In *Proceedings of International Symposium on Autonomous Decentralized Systems*, 2009.

[15] Renata Teixeira and Jennifer Rexford. A measurement framework for pin-pointing routing changes. In *Proceedings of the ACM SIGCOMM workshop on Network troubleshooting: research, theory and operations*

practice meet malfunctioning reality, NetT '04, pages 313–318, New York, NY, USA, 2004. ACM.

[16] Shivani Deshpande, Marina Thottan, Tin Kam Ho, and Biplab Sikdar. An Online Mechanism for BGP Instability Detection and Analysis. *IEEE Trans. Comput.*, 58(11):1470–1484, November 2009.

[17] M. Lad, R. Oliveira, D. Massey, and L. Zhang. Inferring the Origin of Routing Changes using Link Weights. In *Proceedings of IEEE International Conference on Network Protocols*, 2007.

[18] Masafumi Watari, Atsuo Tachibana, and Shigehiro Ano. Inferring the origin of routing changes based on preferred path changes. In *Proceedings of the 12th international conference on Passive and active measurement*, PAM'11, pages 163–172, Berlin, Heidelberg, 2011. Springer-Verlag.

[19] J. Wu, Z. Mao, J. Rexford, and J. Wang. Finding a needle in a haystack: pinpointing significant BGP routing changes in an IP network. In *Proceedings of Usenix Symposium on Networked Systems Design & Implementation*, pages 1–14, 2005.

[20] Jian Zhang. Learning-Based Anomaly Detection in BGP Updates. In *In MineNet 05: Proceeding of the 2005 ACM SIGCOMM workshop on Mining network data*, pages 219–220, 2005.

[21] Yiyi Huang, Nick Feamster, Anukool Lakhina, and Jun (Jim) Xu. Diagnosing network disruptions with network-wide analysis. In *In Sigmetrics*, pages 61–72, 2007.

[22] University of Oregon Route Views Project. `http://www.routeviews.org/`.

[23] RIPE Routing Information Service. `http://www.ripe.net/ris/`.

[24] Internet Topology Collection. `http://irl.cs.ucla.edu/topology/`.

[25] Y. Hyun, B. Huffaker, D. Andersen, E. Aben, M. Luckie, kc claffy, and C. Shannon. The IPv4 Routed 24 AS Links Dataset. `http://www.caida.org/data/active/ipv4_routed_topology_aslinks_dataset.xml`.

[26] H. Haddadi, D. Fay, S. Uhlig, A. Moore, R. Mortier, and A. Jamakovic. Mixing biases: Structural changes in the AS topology evolution. In *Proceedings of the Second International Workshop on Traffic Measurements and Analysis (TMA)*, pages 32–45, April 2010.

[27] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed internet routing convergence. *IEEE/ACM Transactions on Networking*, 9(3):293–306, June 2001.

[28] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271 (Draft Standard), January 2006. Updated by RFCs 6286, 6608.

[29] T.G. Griffin and B.J. Premore. An experimental analysis of BGP convergence time. In *Proceedings of IEEE International Conference on Network Protocols*, pages 53–61, November 2001.

[30] S. Deshpande and B. Sikdar. On the impact of route processing and MRAI timers on BGP convergence times. In *Proceedings of IEEE Global Communication Conference*, pages 1147–1151, December 2004.

[31] P. Jakma. Revised default values for the BGP 'Minimum Route Advertisement Interval'. Internet draft, draft-jakma-mrai-02.txt, work in progress, November 2008.

[32] B. Prakash, N. Valler, D. Andersen, M. Faloutsos, and C. Faloutsos. BGP-lens: patterns and anomalies in internet routing updates. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1315–1324, June 2009.

[33] A. Lambert, M.-O. Buob, and S. Uhlig. Improving internet-wide routing protocols convergence with MRPC timers. In *Proceedings of ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, pages 325–336, 2009.

[34] Sara Burkle. BGP convergence analysis. Master's thesis, Universitat des Saarlandes, 2003. Done on a theme of Prof. Anja Feldmann, Ph.D. Supervisor Olaf Maennel.

[35] B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the Internet AS-Level Topology. *ACM Computer Communication Review Magazine*, 35(1):53–61, January 2005.

[36] R. Oliveira, B. Zhang, D. Pei, and L. Zhang. Quantifying path exploration in the internet. *IEEE/ACM Transactions on Networking*, 17:445–458, April 2009.

[37] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang. The (In)completeness of the Observed Internet AS-Level Structure. *IEEE/ACM Transactions on Networking*, 18:109–122, 2010.

[38] R. Teixeira, S. Uhlig, and C. Diot. BGP route propagation between neighboring domains. In *Proceedings of the Passive and Active Measurement Conference*, pages 11–21, 2007.

[39] A. Fabrikant, U. Ali Syed, and J. Rexford. There's something about MRAI: Timing diversity can exponentially worsen BGP convergence. In *Proceedings of IEEE International Conference on Computer Communications*, pages 2975–2983, April 2011.

[40] C. Labovitz, G. Malan, and F. Jahanian. Origins of Internet Routing Instability. In *Proceedings of IEEE International Conference on Computer Communications*, pages 218–226, 1999.

[41] B. Zhang R. Oliveira, R. Izhak-Ratzin and L. Zhang. Measurement of Highly Active Prefixes in BGP. In *Proceedings of IEEE Global Communication Conference*, 2005.

[42] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: Mapped? In *Proceedings of ACM Internet Measurement Conference*, pages 336–349, 2009.

[43] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a large european IXP. *SIGCOMM Computer Communication Review*, 42(4):163–174, October 2012.

Appendix A. Updates in correlated spikes and connections between monitors

From the table 5, the number of correlated spikes depends on number of neighbours and connections between neighbours. But if we look at the absolute number of updates in correlated spikes, there is no such dependence.

It happens because of differences between correlated spikes received by different collectors from different monitors. We noticed that the average number of updates in correlated spike differs from the average number of updates over all spikes. And correlation between average number of updates over all spikes and average number of updates in correlated spike defines which portion of updates belongs to correlated spikes. See table A.7 for details.

Collector	Total spikes	Total updates	Total correlated spikes	Total updates in correlated spikes	Updates per spike	Updates per correlated spike	Proportion	Updates in correlated spikes, 30s interval
Route-views4	101734	1371562	92454	1018277	13.48	11.01	0.82	74%
Wide	25092	464883	23439	438122	18.53	18.69	1.01	94%
Eqix	88938	1857771	83058	1669158	20.89	20.10	0.96	90%
Isc	107649	2849866	101762	2078749	26.47	20.43	0.77	73%
Linx	308440	6868316	295005	6310863	22.27	21.39	0.96	92%
Route-views2	579456	10128729	566169	8188484	17.48	14.46	0.83	81%

Table A.7: Proportion between average number of updates in correlated spikes and number of updates over all spikes. 30s time interval, 99% same prefix updates in two spikes to be concerned correlated.

Number of updates per spike (column 6) in the table A.7 was calculated by dividing of total number of updates (column 3) on total number of spikes (column 2). Number of updates per correlated spike (column 7) was calculated by the same way (using columns 5 and 4).

Column "Proportion" shows ratio of average number of updates in correlated spikes (column 7) to average number of updates over all spikes (column 6). In other words, how different is the average size of correlated spikes from average size of all spikes. If this ratio is more than 1, correlated spikes are smaller than the global average and vice versa.

Percent of updates in spikes concerned correlated (the last column) depends on this ratio.

Appendix B. Classification algorithm details

This appendix provides details on spike classification methods that were used in Sections 4 and 5.

Appendix B.1. Classification of spikes into single and correlated

Let us first introduce a few notations. Let x be a spike which we want to classify as corresponding to a local or global routing event.

- $SameAS\_updates_x$**:** Total number of BGP updates in correlated spikes received from the same AS as the one that received spike x.
- $OtherAS\_updates_x$**:** Total number of BGP updates in the correlated spikes received from ASs other than the AS that received spike x.
- $Max\_SameAS\_updates_x$**:** Number of BGP updates in largest correlated spike received from the same AS as the one that received spike x.
- $Max\_otherAS\_updates_x$**:** Number of BGP updates in largest correlated spike received from an AS other than the AS that received spike x.

Now, we consider that a spike x is single (caused by a local event), if it satisfies the following condition:

$$OtherAS\_updates_x < threshold^{12} \times SameAS\_updates_x. \quad \text{(B.1)}$$

This condition implies that the most of correlated updates were received from the same AS as the spike x. If this condition is not true, it is likely that the corresponding routing event is not local, implying that correlated updates are likely to have been propagated in the Internet.

In addition, we also require that:

$$Max\_otherAS\_updates_x < threshold \times Max\_SameAS\_updates_x. \quad \text{(B.2)}$$

This condition imposes that no significant correlated spike were received from AS different from that sent spike x.

[12]The results in Section 4 correspond to the 0.1 value of the threshold. If we change it to 0.25, the fraction of updates in single spikes increases from 19.42% to 22,43%

If many spikes correlated with x were received from different ASs, we consider that the spike is correlated. Furthermore, if any large enough correlated spike is received from another AS, we also consider that the spike is correlated.

Only if both conditions are satisfied[13], we consider that a given spike is single and might be related to a local routing event.

The presented methodology allows to precisely identify amount of correlation among spikes received by different collectors. To prove it, we have compared our results from Section 4 with the amount of correlated prefixes found among all spikes. The results for the last are shown on the Figure B.13.

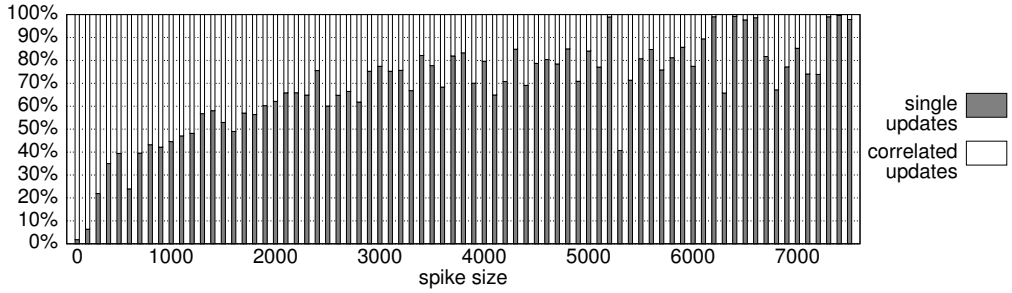

Figure B.13: Fraction of correlated and non-correlated BGP updates on a per-bin basis, 120s time interval.

The distribution of correlated updates is similar to the one on the Figure 6. Moreover, the number of correlated prefix updates found among all updates form 80.32% from total, which is also similar to the result of 80.58%, calculated using the developed methodology in the Section 4.

Appendix B.2. Estimating the extent of correlated spikes

To take into account the AS-level visibility in the classification of single and correlated spikes, we need to shortly revisit the threshold used in condition B.1. Imagine that a group of correlated spikes have been received from 2 ASs and have been observed by 4 vantage points. 3 of the vantage points have received the correlated spikes from one of the 2 AS, and the fourth vantage point from the second AS. In such a situation, if the threshold from

[13]In Section 4 we also require that correlated spike contains more than *threshold* of prefixes, marked as correlated with prefixes from other monitors. Otherwise we classify it as single, even if these conditions are not satisfied

condition B.1 is set to less than 0.25, the considered spike might be classified as single, as it is likely that 3 vantage points receiving the correlated spikes from one AS will collect 75% of the updates and only 25% of the updates will be collected by the fourth vantage point. Therefore, the threshold for condition B.1 should be dependent on the visibility we have from the observation points. For this, we introduce two more parameters:

- $Monitor\_pairs_x$: number of $\langle monitor, collector \rangle$ pairs that observed spikes correlated with the considered spike x.
- $Same\_AS\_pairs_x$: number of $\langle monitor, collector \rangle$ pairs that observed spikes correlated with spike x and belong to the same AS as the AS which observed spike x.

We replace the threshold from condition B.1 by the following:

$$threshold_{visibility} = 1/2 \times \frac{Monitor\_pairs_x - Same\_AS\_pairs_x}{Monitor\_pairs_x}. \quad \text{(B.3)}$$

$threshold_{visibility}$ adapts the threshold from condition B.1 by removing the impact of multiple observation points in the same AS as the one of the considered spike.

If both conditions B.1 and B.2[14] are satisfied, we consider the spike as single and classify it into one of the three groups mentioned earlier, depending on our visibility of the connectivity between the ASs from which the spike was received. If one of the conditions is not satisfied, we consider the spike as correlated and compute the maximum AS hop distance between the ASs that observed the correlated spike.

Not all ASs that send correlated updates should be included into the maximum AS hop distance computation. Indeed, imagine that the group of correlated spikes were received by multiple observation points from multiple ASs, involving $10,000$ BGP prefixes overall. If one of the observation points received from one AS BGP updates for only 50 of these prefixes, we may want to ignore the BGP updates seen for the corresponding $\langle observed\ AS, collector \rangle$ pair, as they might very well not be related to the underlying events that

[14] After we have modified the threshold in the condition B.1, we consider the condition B.2 only as supplementary. For this case, it was decided to use the constant threshold, equal to 0.33.

created the correlated spikes. We want to have a high enough confidence that our correlated spikes really belong to global BGP propagation and not some randomly occurring BGP updates that only remotely look like correlated spikes.

Before we explain in more details which pairs we include in the distance computation, let us define the following variables:

- $Update\_Sum(x)$**:** number of all BGP updates in all correlated spikes related to spike x.
- $Max\_Spike(x)$**:** the largest spike in number of BGP updates, in the group of correlated spikes related to spike x, on a $\langle monitor, collector \rangle$ granularity.
- $Max\_SameAS\_updates(x)$**:** the largest spike in number of BGP updates, in the group of correlated spikes related to spike x, received by any $\langle monitor, collector \rangle$ pair where the monitored AS is the same as the one that sent spike x.

Now, we consider that the BGP updates seen by a given pair $\langle as, r \rangle$ should be considered in the computation of the maximum distance for spike x if the following condition holds:

$$Update\_Sum_{<as,r>}(x) > threshold_{pairs} \times Update\_Sum(x). \tag{B.4}$$

where $threshold_{pairs} = 1/2 * 1/Monitor\_pairs_x$ and $Update\_Sum_{\langle as,r \rangle}(x)$ is the restriction of $Update\_Sum(x)$ where only the BGP updates of pair $\langle as, r \rangle$ are considered.

Condition B.4 implies that the amount of BGP updates received in correlated spikes by pair $\langle as, r \rangle$ is large enough compared to the total BGP updates involved in all correlated spikes. If condition B.4 is not satisfied, we also check if

$$Max\_Spike_{\langle as,r \rangle}(x) > 0.33 \times Max\_SameAS\_updates(x) \tag{B.5}$$

where $Max\_Spike_{\langle as,r \rangle}(x)$ is the restriction of $Max\_Spike(x)$ to pair $\langle as, r \rangle$. Condition B.5 checks if a given pair has seen a large enough number of BGP updates compared to the largest spike received from the AS that sent spike x.

If conditions B.4 and B.5 are not satisfied by a given pair, we skip it for the maximum distance computation. We therefore include only $\langle observation\ AS$,

collector⟩ pairs, where the route collector has received a sufficient amount of BGP updates in comparison to other pairs or relative to the total amount of BGP updates involved in the correlated spikes.