Raquel Hill

Hijacking Network Traffic: Early Detection of Routing Anomalies

COLLABORATORS

Pablo Moriano, L. Jean Camp

INDIANA UNIVERSITY
SCHOOL OF INFORMATICS, COMPUTING, AND ENGINEERING

Macroscopic Internet Outages

- Large-scale disruption (“outage”, “leak”, “hijack”)
- Potential causes: natural disasters, cyber attacks, physical attacks, terrorism, war, bugs and misconfigurations, government order, …

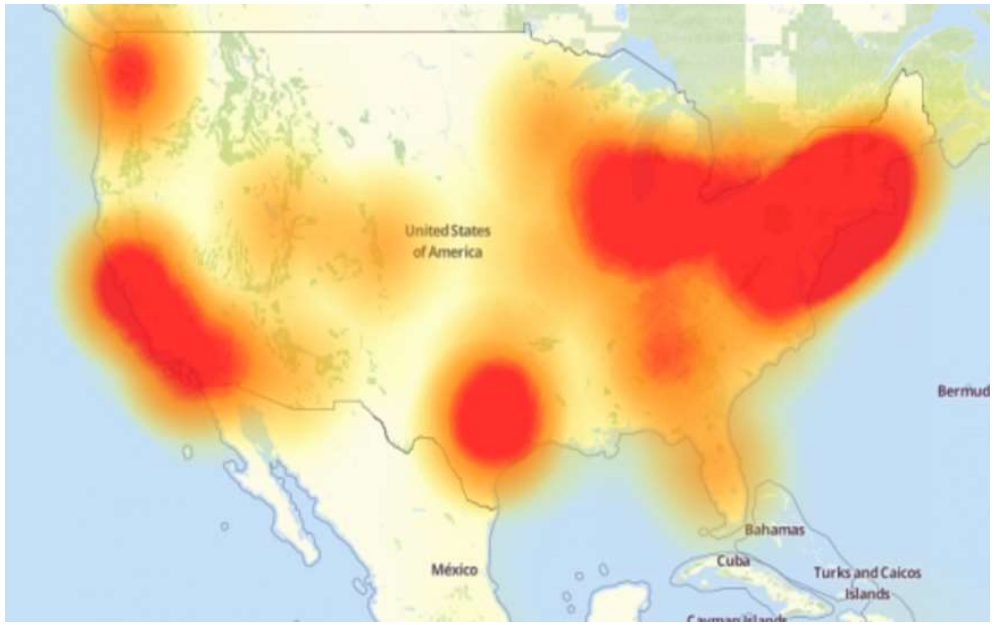

Cost of Outages

CLOUD

5-minute outage costs Google $545,000 in revenue

DYLAN TWENEY @DYLAN20 AUGUST 16, 2013 4:06 PM

How Much Will Today's Internet Outage Cost?

Some companies lose tens of thousands of dollars for every *minute* of a DDoS attack.

ADRIENNE LAFRANCE | OCT 21, 2016 | TECHNOLOGY

[Popular Destinations rerouted to Russia]

*Posted by Andree Toonk - December 12, 2017 - Hijack***0**

Early this morning (UTC) our systems detected a suspicious event where many prefixes for high profile destinations were being announced by an unused Russian Autonomous System. Starting at 04:43 (UTC) 80 prefixes normally announced by organizations such Google, Apple, Facebook, Microsoft, Twitch, NTT Communications and Riot Games were now detected in the global BGP routing […]

REPORT

Internet shutdowns cost countries $2.4 billion last year

Darrell M. West · 6 October 2016

Layered View of the Internet

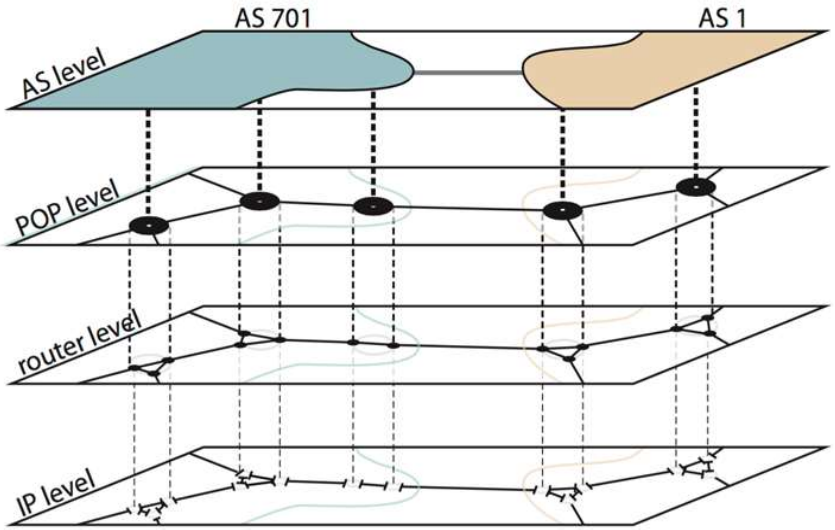

B. Huffaker and Y. Hyun. Interactive Access to Internet Topology Data. CAIDA. 2016.

Border Gateway Protocol (BGP)

- Central nervous system of the Internet
- Determines how data is routed from one AS to another

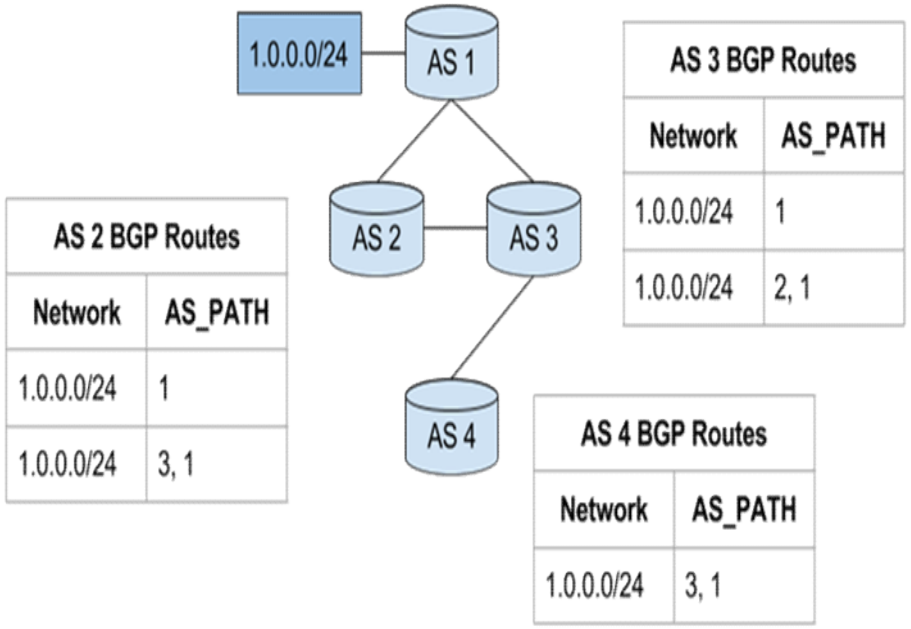

Routing Anomalies: Prefix Hijacking

- BGP (or 'prefix') Hijacking
  - A BGP participant intentionally sends false connectivity information
- Route Leak
  - An unintentional misconfiguration that causes a BGP participant to send incorrect connectivity information
- Both result in traffic flowing incorrectly

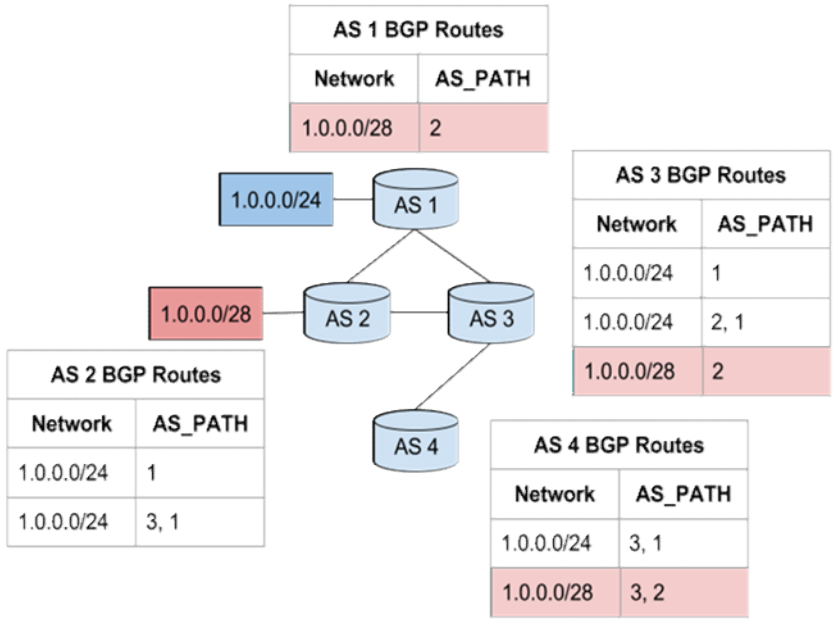

Internet Traffic Misdirection

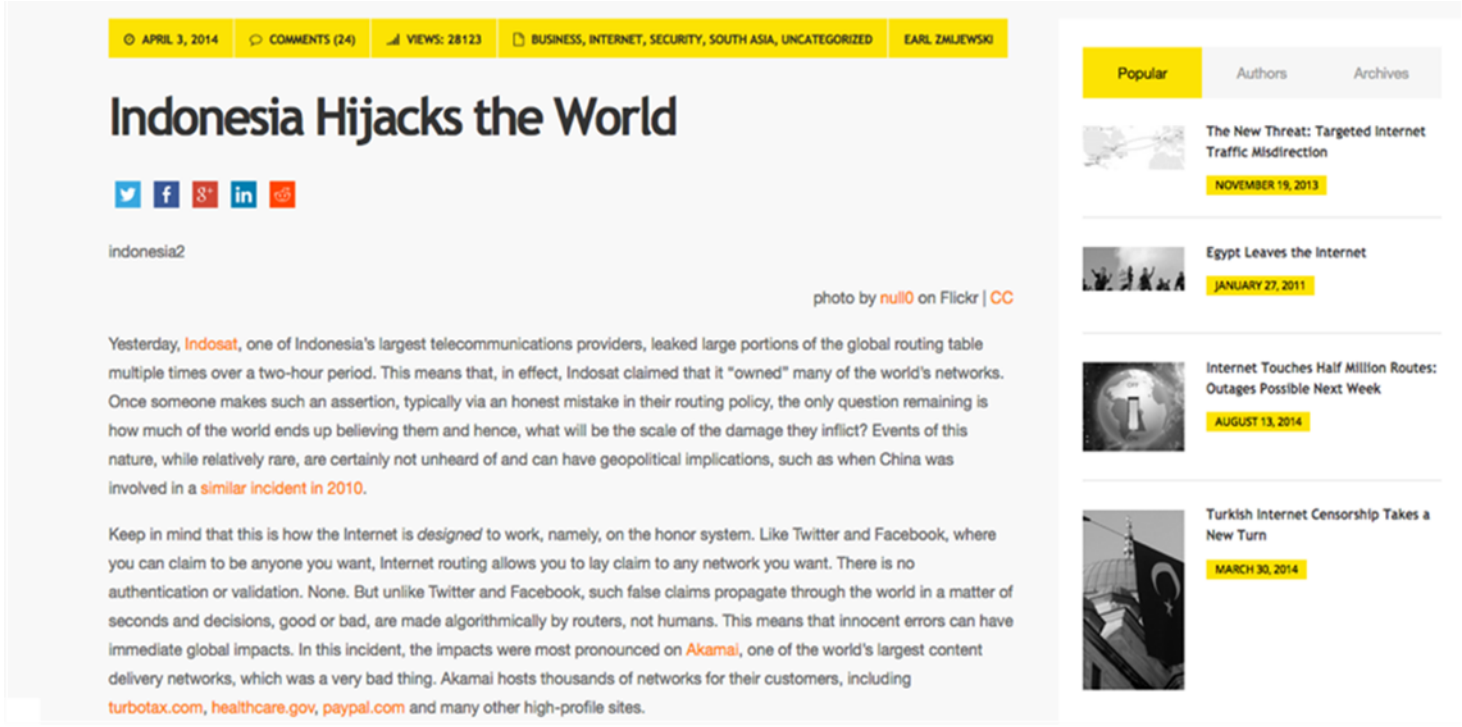

APRIL 3, 2014 | COMMENTS (24) | VIEWS: 28123 | BUSINESS, INTERNET, SECURITY, SOUTH ASIA, UNCATEGORIZED | EARL ZMIJEWSKI

Indonesia Hijacks the World

indonesia2

photo by null0 on Flickr | CC

Yesterday, Indosat, one of Indonesia's largest telecommunications providers, leaked large portions of the global routing table multiple times over a two-hour period. This means that, in effect, Indosat claimed that it "owned" many of the world's networks. Once someone makes such an assertion, typically via an honest mistake in their routing policy, the only question remaining is how much of the world ends up believing them and hence, what will be the scale of the damage they inflict? Events of this nature, while relatively rare, are certainly not unheard of and can have geopolitical implications, such as when China was involved in a similar incident in 2010.

Keep in mind that this is how the Internet is *designed* to work, namely, on the honor system. Like Twitter and Facebook, where you can claim to be anyone you want, Internet routing allows you to lay claim to any network you want. There is no authentication or validation. None. But unlike Twitter and Facebook, such false claims propagate through the world in a matter of seconds and decisions, good or bad, are made algorithmically by routers, not humans. This means that innocent errors can have immediate global impacts. In this incident, the impacts were most pronounced on Akamai, one of the world's largest content delivery networks, which was a very bad thing. Akamai hosts thousands of networks for their customers, including turbotax.com, healthcare.gov, paypal.com and many other high-profile sites.

Popular | Authors | Archives

The New Threat: Targeted Internet Traffic Misdirection
NOVEMBER 19, 2013

Egypt Leaves the Internet
JANUARY 27, 2011

Internet Touches Half Million Routes: Outages Possible Next Week
AUGUST 13, 2014

Turkish Internet Censorship Takes a New Turn
MARCH 30, 2014

Prior Work: Control Plane

- Control Plan Approaches:
  - Monitor BGP updates or routing tables from a distributed set of BGP monitors [Khare et al 2012, Lad et al 2006, Sermpezis et 2018].
  - Look for inconsistencies in the origin of prefixes announced by ASes or unexpected path changes.
  - Limitations
    - Tend to report a large number of false positives
    - Detection typically occurs after ASes have accepted a malicious or misconfigured route

Prior Work: Data Plane

- Use application such as ping and traceroute to detect routing anomalies [Zhang et al 2010, Zheng et al 2007]
- Monitor the reachability of routes from the victim to detect anomalies
- Higher detection accuracy
- Approaches do not scale well since they require a considerable number of active measurements

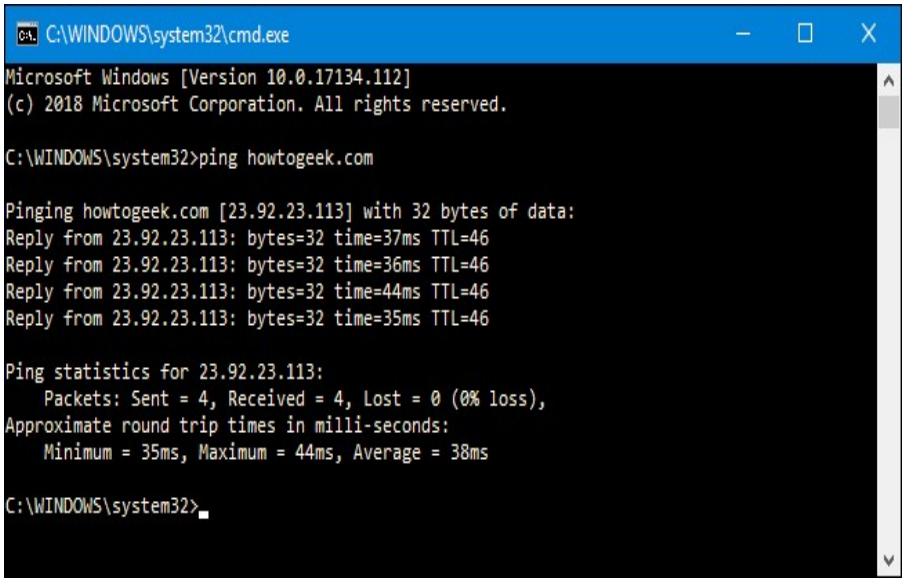

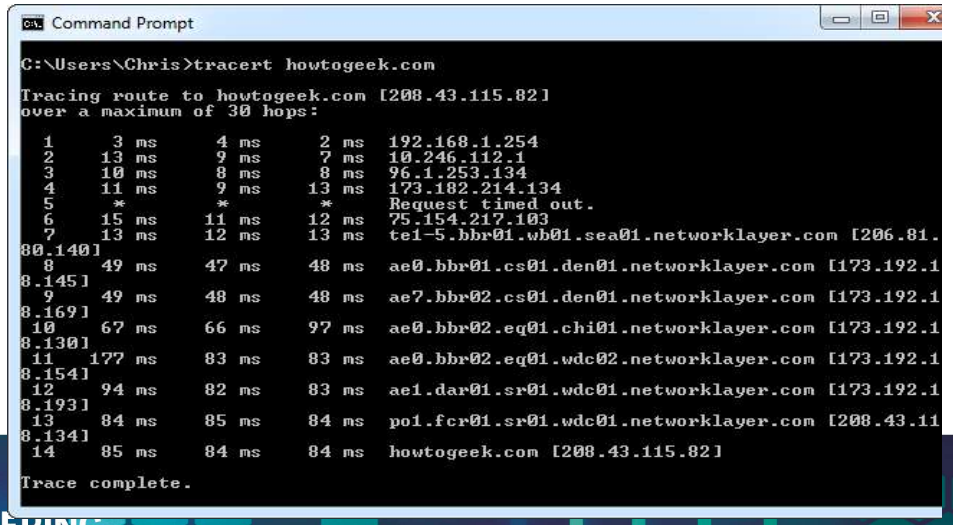

Our Approach

- Goal: Identify routing anomalies as they emerge, several hours before the anomalous event is detected by state of the art approaches
- Use control plane data collected by network monitoring tools (Routeviews, BGPStream)
- Approach relies on the key observation that anomalous BGP announcements are made/sent in bursts
  - Burstiness refers to the tendency of certain events to occur in groups of relatively high frequency, i.e., short inter-arrival time intervals, followed by periods of relatively infrequent events
- We characterize bursty announcements through statistical analysis of inter-arrival times
- We conduct a case-based systematic analysis of the changes in inter-arrival times that are associated with three well-known anomalous events

Measuring BGP

BGPSTREAM

An open-source software framework for live and historical BGP data analysis, supporting *scientific research, operational monitoring,* and *post-event analysis.*

Get BGPStream

Powerful Tools & APIs

Quickly inspect raw BGP data from the command-line, develop Python apps, or build complex systems using a C/C++ API, etc. Designed to run anywhere, from laptops to clusters.

Learn about the components »

Seamless & Live Data Access

Give BGPStream a time range and it will automatically acquire and stream the right data to you. Enable realtime monitoring by changing a single parameter.

See available data »

Tutorials & Docs

Documentation includes software and API reference manuals as well as tutorials with fully-running code samples.

Get started »

Collector Distribution

Name	Location	Country
route-views2	Eugene, OR	US
route-views3	Eugene, OR	US
route-views4	Eugene, OR	US
route-views6	Eugene, OR	US
route-views.eqix	Ashburn, VA	US
route-views.isc	Palo Alto, CA	US
route-views.kixp	Nairobi	KE
route-views.jinx	Johannesburg	ZA
route-views.linx	London	GB
route-views.texatl	Atlanta, GA	US
route-views.wide	Tokyo	JP
route-views.sydney	Sydney	AU
route-views.saopaulo	Sao Paulo	BR
route-views.nwax	Portland, OR	US
route-views.perth	Perth	AU
route-views.sg	Singapore	SG
route-views.sfmix	San Francisco, CA	US
route-views.soxrs	Belgrade	RS
route-views.chicago	Chicago, IL	US

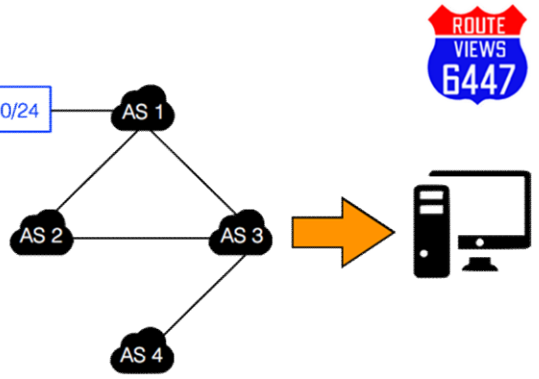

Measuring Burstiness

- $X_{A\to B} = \{X_{A\to B}(t)\}$, $t = 0, 1, \ldots, N$ be a time series of time-stamped announcements sent by AS A and received by collector B.
- Let $\tau_{A\to B}$ be a random variable that represents the time interval between consecutive announcements so that $\tau_{A\to B}$ takes values in $\{X_{A\to B}(1) - X_{A\to B}(0), X_{A\to B}(2) - X_{A\to B}(1), \ldots, X_{A\to B}(N) - X_{A\to B}(N-1)\}$

Burstiness cont.

- The inter-arrival time distribution $P(\tau_{A\to B})$ can be characterized by a burstiness factor defined by

$$B = \frac{\sigma-\mu}{\sigma+\mu}.$$

- Here σ and μ denote the standard deviation and mean of the inter-arrival time distribution.

- Note that ***B*** = −1 for σ = 0, which means regular time intervals
- ***B***= 0 for σ = μ in the case of random time intervals. Finally, it
- has a value of 1 for σ → ∞ and a finite μ in the case of a highly bursty time series of announcements.

Case Study: Idonesia

Indosat- An Indonesian ISP hijacks the world. On April 2, 2014, starting at 18:26 UTC, Indosat (one of the largest telecommunications providers in Indonesia) announced more than 320, 000 IP prefixes belonging to other networks. Indosat announced roughly two-thirds of the entire Internet address space. A large fraction of the hijacked prefixes belonged to Akamai, which is one of the larger Content Delivery Networks. This incident lasted approximately for 2.9 hours until 21:15 UTC. Traffic continued to be delivered; however, the path of the traffic was significantly altered.

Case Study: Malaysia

Global collateral damage of the Telecom Malaysia leak. On June 12, 2015, starting at 08:43 UTC, Telecom Malaysia announced about 179, 000 IP prefixes to Level 3 (the largest crossing AS). Level 3 accepted these announcements and then propagated the routes to their peers and customers around the word. Because Telecom Malaysia is a customer of Level 3, the routes announced by Telecom Malaysia were identified as a preferred delivery route for Level 3. This event caused a significant packet loss and Internet service degradation around the world. Level 3 suffered a significant loss of connectivity from the Asia pacific region and the rest of the world. Note this was a leak, so the data were not delivered after being transmitted to Telecom Malaysia. This incident lasted approximately 2.7 hours. At around 10:40 UTC there were slowly observed improvements, and by 11:15 UTC the errors in the Routing Information Base (RIB) began to be resolved.

Case Study: India

Large scale BGP hijack in India. On November 6, 2015, starting at 05:52 UTC, Bharti Airtel Ltd., claimed the ownership of about 16, 123 IP prefixes. These addresses corresponded to more than two thousand unique ASes. This event became widespread because two large ASes (e.g.,Cogent Communications and GlobeNet Cabos Submarinos S.A.) accepted and propagated these routes to their peers and customers. Legitimate owners of the prefixes included Akamai, Tata Communications, and Apple Inc. This incident lasted approximately 8.9 hours until 14:40 UTC.

Results: Indonesia

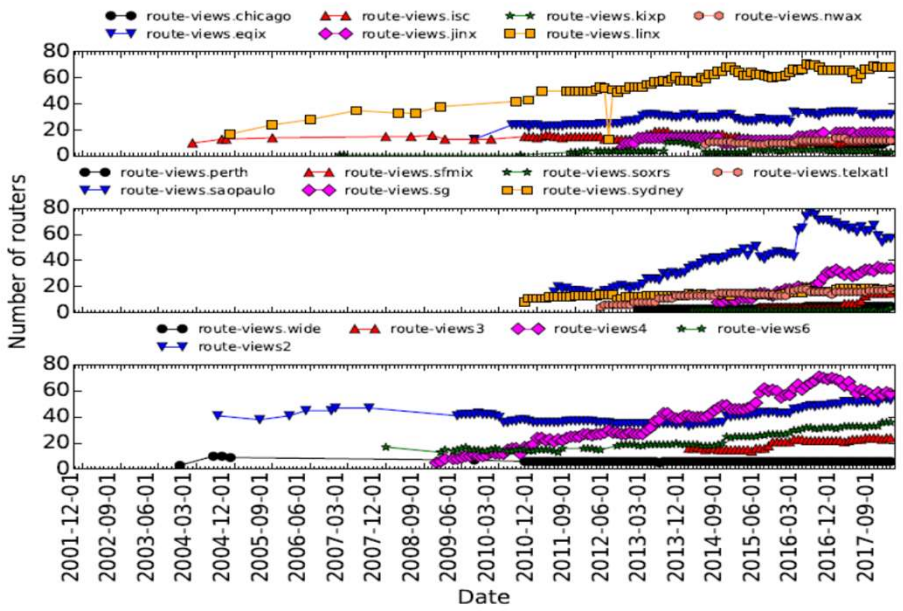

Figure 1: Time series of the number of routers peering with collectors. Collectors are ordered in alphabetical order. Major ticks correspond to nine-month intervals while minor ticks correspond to one-month intervals.

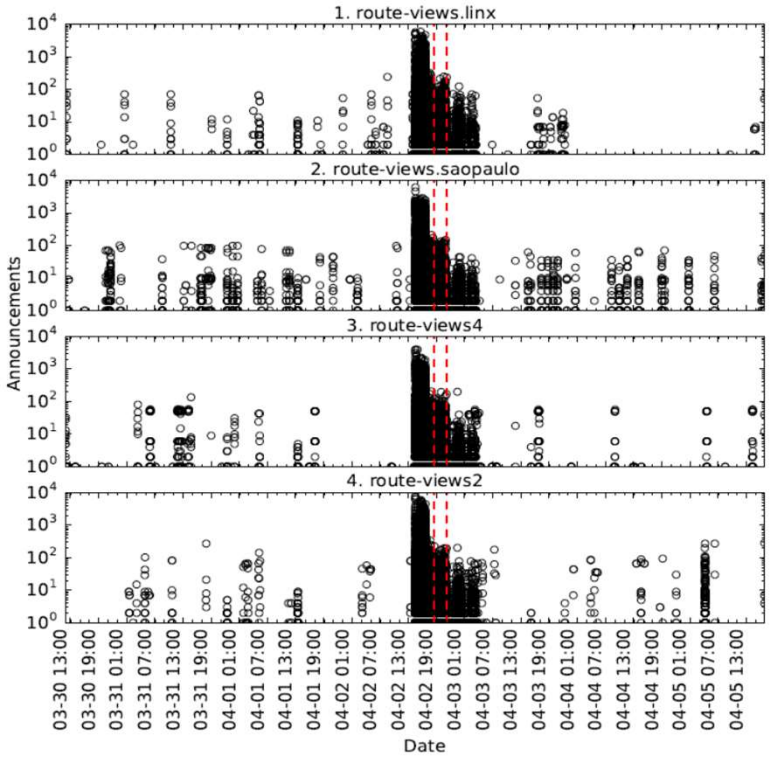

Figure 2: Time series of the number of announcements from AS 4761 that collectors received before, during, and after the Indosat incident in 2014 for the top four collectors. Major ticks correspond to six-hour intervals while minor ticks correspond to two-hour intervals.

Results: Distribution of Burstiness and Number of Announcements

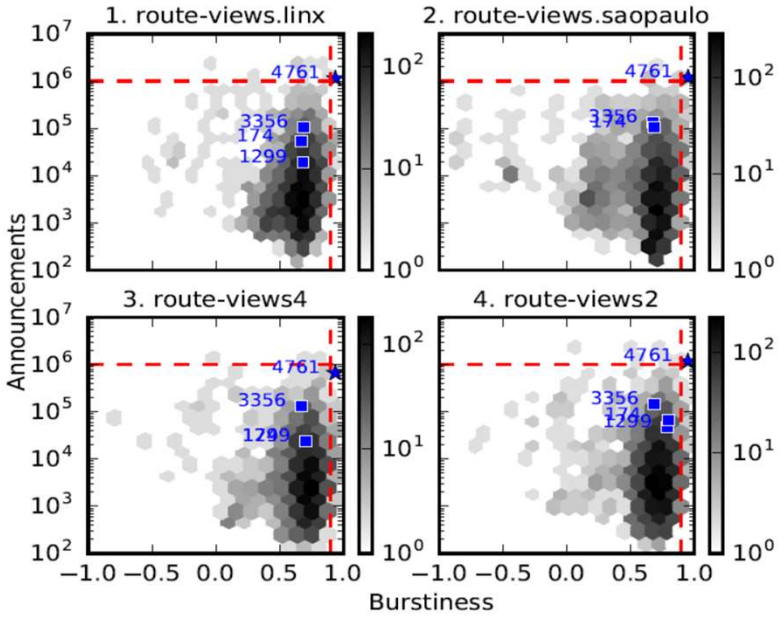

Figure 5: Joint distribution based on the the burstiness (horizontal axis) and number of announcements (vertical axis) during one day interval around the Indosat incident.

Comparison with the Null Case

- Test the null hypothesis: ASes send announcements in a bursty manner even during times where no malicious event was detected.

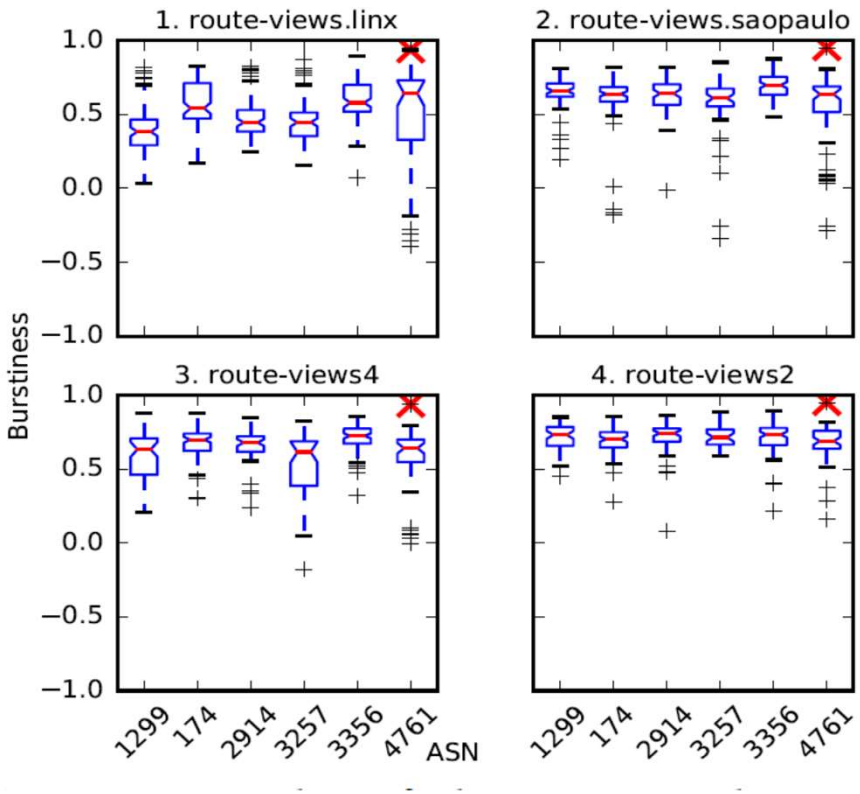

Detection Method

- Measure used in the context of intrusion detection
- $Q_{A->B}$ – number of announcements sent from AS A to collector B; exponentially weighted
- r is the decay factor (1/300); subsequent announcements sent in less than 300 seconds
- $\Delta = X_{A->B(t)} - X_{A->B(t-1)}$; inter-arrival time of consecutive announcements

$$Q_{A\to B}(t) = 1 + 2^{-r\Delta} Q_{A\to B}(t-1). \qquad (1)$$

Detection: Indosat

- Observations that are more than two standard deviations from the moving average are labeled with a star,
- Collector, route-views.linx has the largest number of feeders and detects a deviation from the normal inter-arrival time 3 hrs 43 min 2 seconds before the earliest detection of the event.

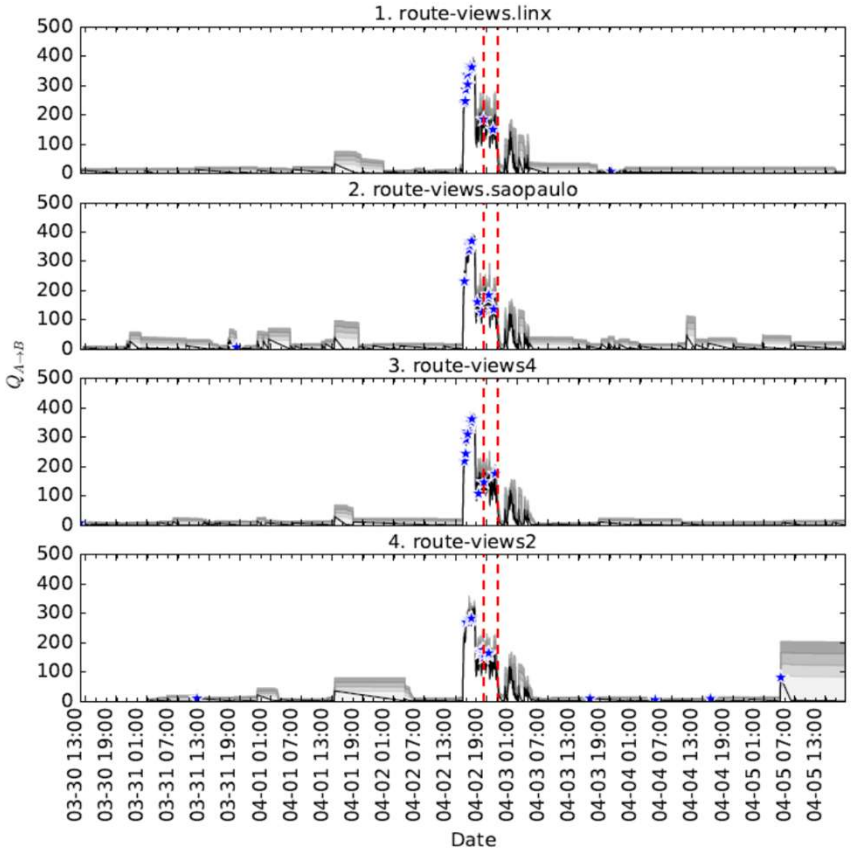

Figure 11: $Q_{4761\to B}$ **time series for the Indosat incident.**

Summary

- Distributed views of the Internet at the AS-level can be used for early detection of large-scale Internet disruptions despite their distributed and incomplete nature.
- The event is visible as anomalies in the inter-arrival time of BGP announcements.
- A view that is constructed from a smaller number of feeders is not as robust and changes in similarity may have a larger range and persist after the event.
- This method has potential for early detection of large-scale control-plane anomalies; enabling quicker mitigation.


THANK YOU
ralhill@indiana.edu
www.cs.indiana.edu/~ralhill
INDIANA UNIVERSITY
SCHOOL OF INFORMATICS, COMPUTING, AND ENGINEERING