

# Predict, Assess, Risk, Identify Disruptive Internet-scale Network Events (PARIDINE) Kick-off Meeting

April 10, 2018 | Arlington, VA




# IODA-NP: Multi-source Realtime Detection of Macroscopic Internet Connectivity Disruption

Alberto Dainotti | CAIDA, UC San Diego

August 24<sup>th</sup> , 2018

# Team Profile

- **Center for Applied Internet Data Analysis (CAIDA) @ San Diego Supercomputer Center, University of California San Diego**
- **PI:** *Alberto Dainotti, PhD*
- **CoPI:** *Marina Fomenkov, PhD*
- *Alistair King, Rama Padmanabhan, Philipp Winter, Dan Andersen, Paul Hicks, Alex Ma, ...*



caida



# Customer Need

- Timely Detect and Analyze Internet Connectivity Outages
- Focus on: **macroscopic** events, affecting the network **edge**
  - *E.g., a connectivity black-out significantly affecting customers of a large network operator or a large geographic area*
- Context: Cyber attacks, physical attacks, natural disasters, bugs and misconfiguration, government orders, ...
- Application: Public Safety, Situational Awareness, Disaster Recovery, Insurance, Internet Reliability & Performance






# Approach Overview

- IODA: *Internet Outage Detection & Analysis*
  - Started in 2012 with NSF funding
- Approach
  - Combine *active* and *passive* measurements both at the *data plane* and *control plane*
  - Data *aggregation* and event detection per Autonomous System (AS) and Geographic Area
  - Interactive *Visualization*
- IODA-NP: *Next Phase*
  - (i) methodological improvements and evaluation based on rigorous definitions, metrics, ground-truth, cross-validation; (ii) reporting events; (iii) API Framework and Documentation

# An eye-candy moment

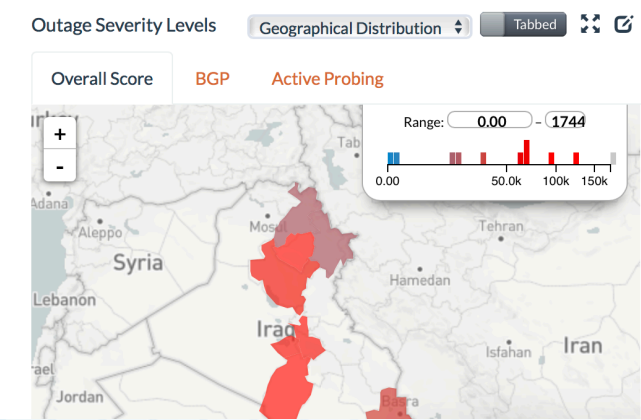


Time	Category	Value 1	Value 2
7:20pm	Probing		
✓ Oct 6th 2016 5:50am	BGP	71,527	71,563
✓ Oct 6th 2016 3:28am	Darknet	28	94
✗ Oct 6th 2016 3:03am	Darknet	22	94
✗ Oct 6th 2016 3:00am	Active Probing	208	949
✗ Oct 6th 2016 2:55am	BGP	70,255	71,563
✗ Oct 6th 2016 2:50am	Active Probing	603	949
✓ Oct 5th 2016 6:20am	Active Probing	791	950

Showing 1 to 15 of 30 entries

[Previous](#) [Next](#)

## Regional Outages for Iraq

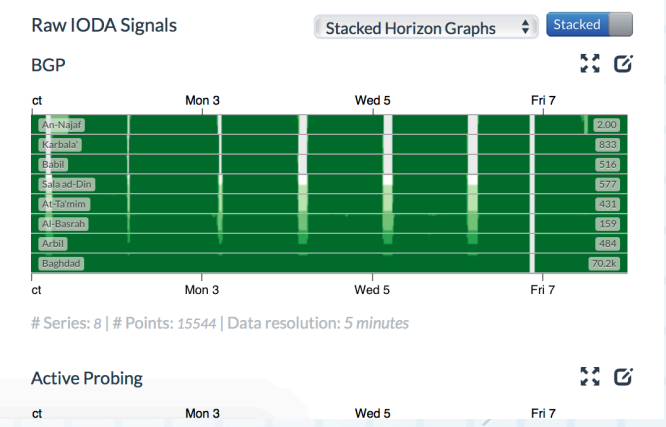


Outage Severity Levels

Show 10 entries

Search:

Region	Overall Score	Active Probing	BGP	Darknet
Baghdad	130M	97.8k	1.33k	
Al-Basrah	51.7M	1.60k	32.3k	
An-Najaf	91.3k		91.3k	
Babil	70.3k		70.3k	
Karbala'	68.9k		68.9k	
Sala ad-Din	62.7k		62.7k	







# Approach (Part 1 - Sources)

- Monitoring the Internet with a combination of **active** and **passive** approaches both at the **data plane** and **control plane**



# IBR (*Passive - Data Plane*)

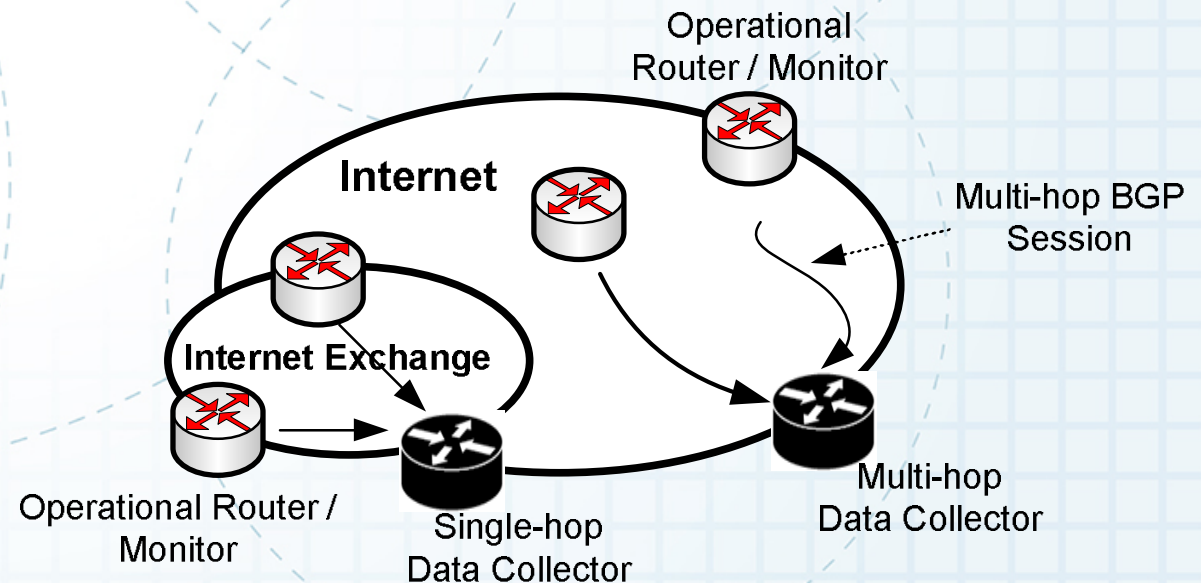
- Internet Background Radiation (IBR) captured by network telescopes





# BGP (*Passive – Control Plane*)

- BGP measurement projects establish peering sessions with ASes to receive their routing tables
  - RouteViews (Univ. Oregon): *371 peers*
  - RIPE RIS (RIPE NCC): *508 peers*
  - **TODO**: sources from CAIDA's BGPStream







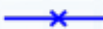
# Active Probing (*Active-Data Plane*)



- ICMP Echo requests
- ISI's **Trinocular** methodology
  - /24 -based probing and inference
- **TODO**: Univ. Maryland's **Thunderping** methodology
  - Per single IP address inference

# Example of Benefit of Multi-Source

Contrasting telescope traffic with BGP measurements **revealed a mix of blocking techniques** that was not publicized by others

The second Libyan outage involved overlapping of **BGP withdrawals** and **packet filtering**


LyStateAS   
IntAS2   
SatAS1 





# Approach (Part 2: Data Aggregation)

- Geography-based Data Aggregation
  - We associate IP addresses, /24 blocks, BGP prefixes with Geographic Coordinates
  - We aggregate post-processed data at Country, State, County level
- AS-level Data Aggregation
  - We associate IP addresses etc. with the operator's AS Number
  - Prefix-to-AS lookups based on BGP data



# Approach (Part 3: Detection)

- For each source type: change point detection on aggregated (i.e., per country, per-state, per-county, per-AS) signals
  - We look for unusual drops
    - Current approach: *naïve* moving-threshold
    - **TODO**: *SARIMA-based detection*
  - **TODO**: (per source type) Link the “drop” to a rigorous definition
- **TODO**: *Detection and Alerting based on fusing data sources*




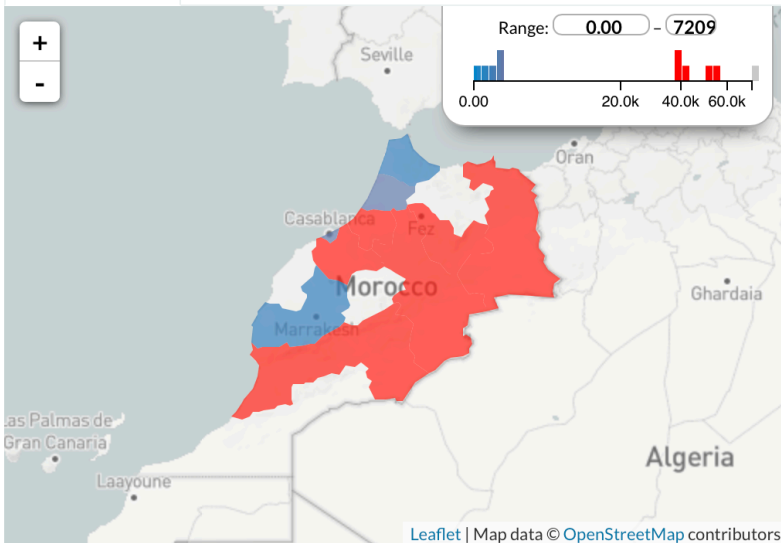
# Approach (Part 4: Interactive Visual Interfaces)

## Regional Outages for Morocco

Outage Severity Levels

Geographical Distribution ▾ Tabbed 🔄 🔗

Overall Score **BGP** **Active Probing**



# Series: 10 | # Points: 14660 | Data resolution: minute

July 19, 2018 5:46pm - July 20, 2018 6:11pm

Outage Severity Levels

Show 10 ▾ entries

Search:

Region	Overall Score	Active Probing	BGP	Darknet
Chaouia - Ouardigha	21.2M	12.4k	1.71k	
Rabat - Salé - Zemmour - Zaer	16.1M	5.71k	2.82k	
Meknès - Tafilalet	15.9M	6.01k	2.65k	
Oriental	11.5M	6.11k	1.89k	
Souss - Massa - Draâ	11.1M	4.56k	2.42k	
Fès - Boulemane	11.0M	5.74k	1.92k	
Gharb - Chrarda - Béni Hssen	929		929	
Grand Casablanca	780		780	
Tanger - Tétouan	449		449	
Marrakech - Tensift - Al Haouz	235		235	


Showing 1 to 10 of 10 entries

[Previous](#) [Next](#)

Raw IODA Signals


Stacked Horizon Graphs ▾ Stacked 🔄 🔗

BGP



# Series: 10 | # Points: 2150 | Data resolution: 5 minutes

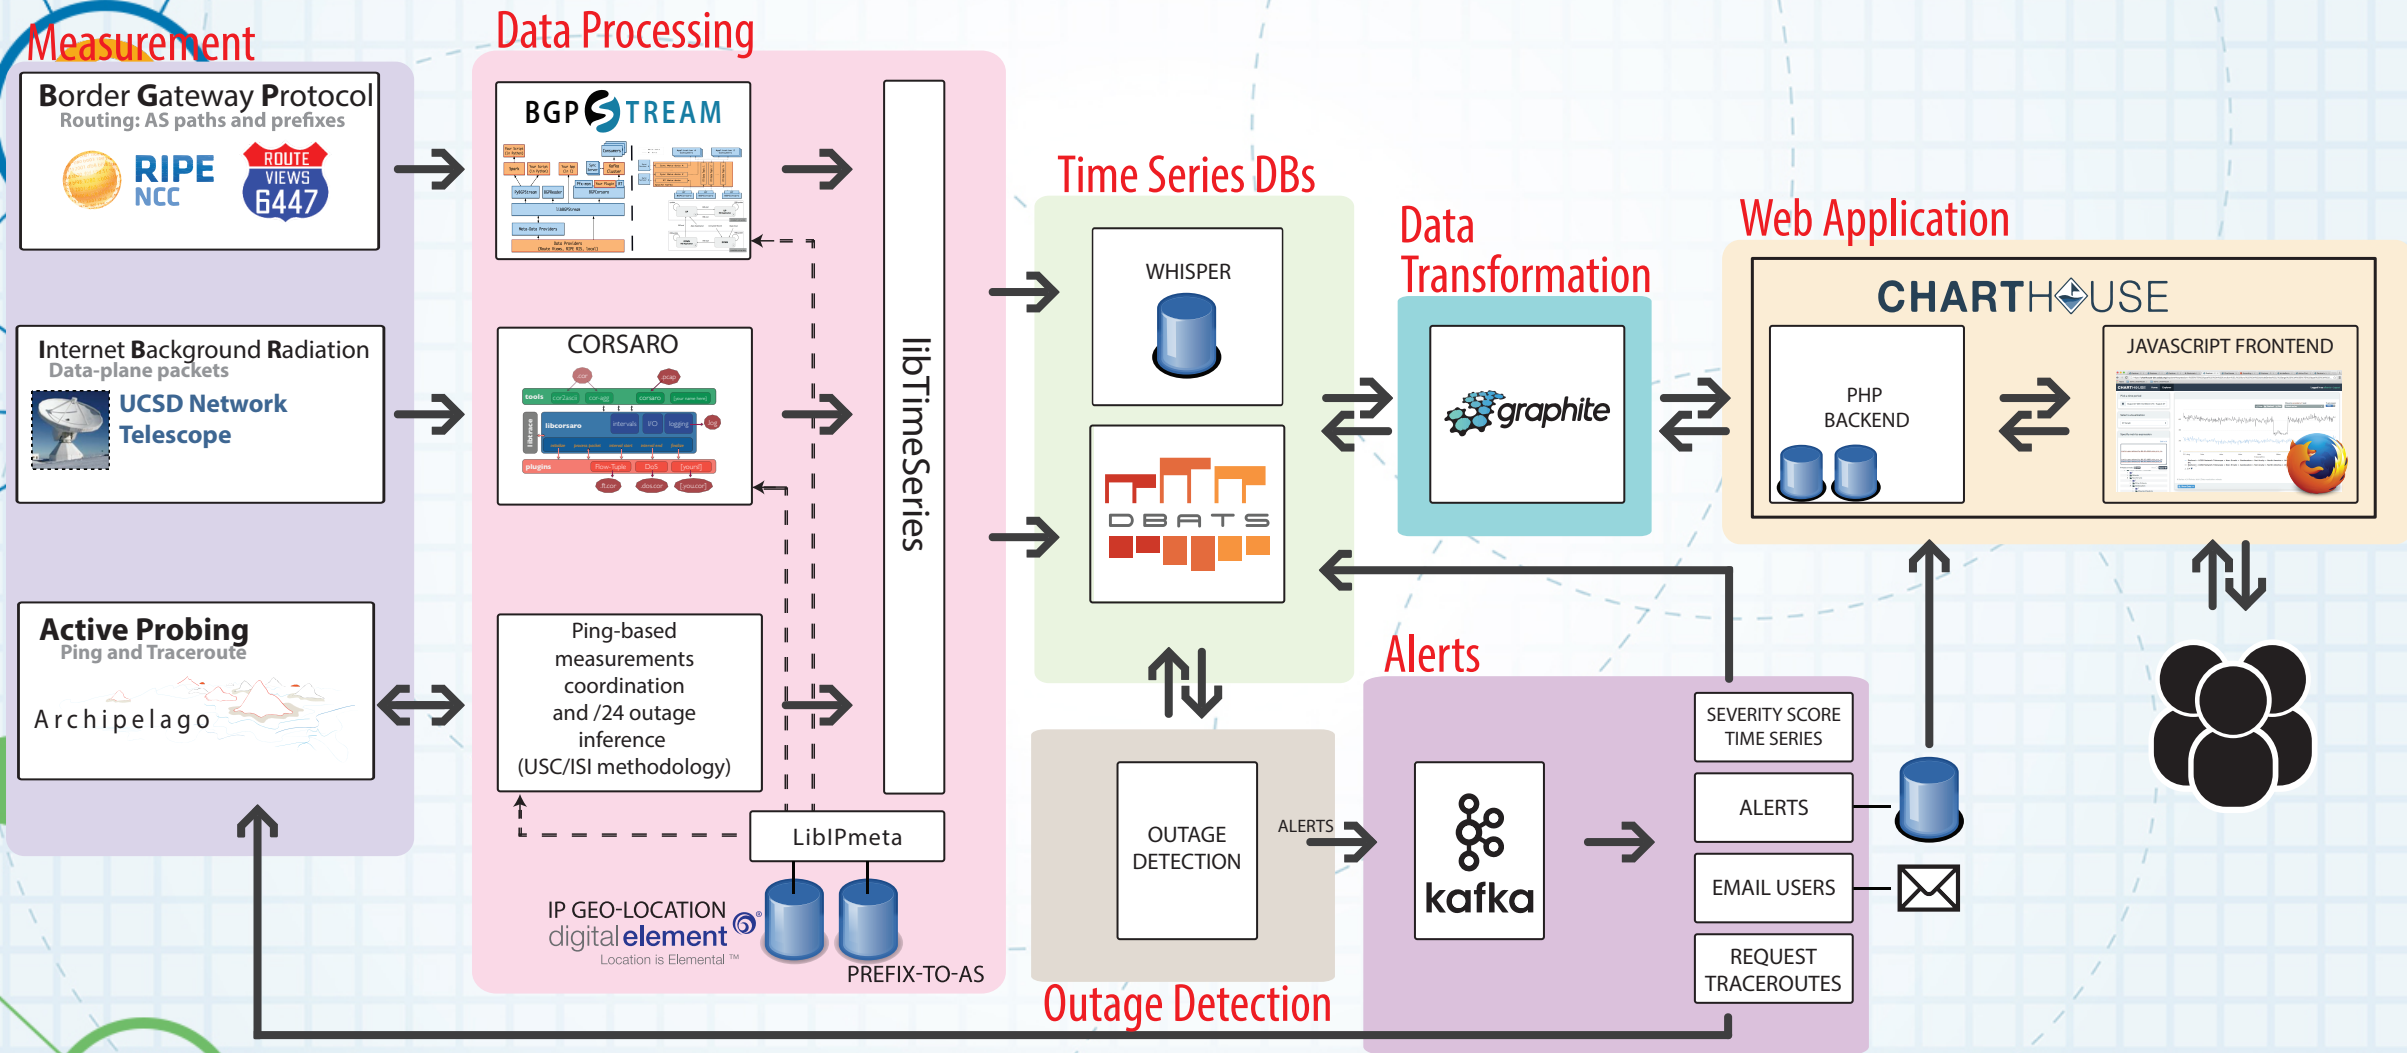
Active Probing



# Series: 6 | # Points: 888 | Data resolution: 10 minutes



# System Overview





# Project Activities + Challenges

- Rigorous definition of targeted event type
  - E.g., 64k related addresses becoming disconnected for more than 5 minutes
  - Investigate different application requirements and intrinsic constraints
- IODA's previous efforts demonstrated the utility of the sources and the approach. However:
  - Need to bridge per-source IODA detection approach with the targeted definition of outage
  - A rigorous evaluation (accuracy, coverage, ...) is missing
  - Current change-point detection generates FP/FNs
  - Need to push to finer geographic granularity (e.g., US counties)
    - E.g., recover filtered out IBR signal, study prefix-geolocation, ...
  - Other data sources can be added
  - The infrastructure needs *reliability* and *latency* improvements



# Project Activities + Challenges

- Focus on US to provide practical insights
  - *Acquire ground truth*
  - *Investigate weather-induced and power outages*
  - *Identify limitations of data sources and approaches in terms of address-block and geographic granularity*
  - *Implement functionalities for US territory and operators*
- Develop and document an API Framework
- Reporting events
  - Already started through the CAIDA blog, a Twitter channel, and cooperating with the KeptOn coalition for politically motivated Internet shutdowns



# Benefits


- Near-realtime detection of macroscopic outages
- Multi-source approach improves:
  - Reliability
  - Coverage
  - Understanding
- Visualization Interface make it intuitive





# Competition


- Oracle's Internet Intelligence Map
  - Focus on country-level
  - Limited interaction/viz functionalities in interface
- ISI / John Heidemann's work
  - IODA uses Trinocular for one data source
  - IODA focuses on per-AS / geographic aggregations
- Akamai
  - State of the Internet reports and some tweets
- Google Transparency report
  - Country-level graphs
- Bgpmon.com
  - BGP only



# Contact Info

<https://ioda.caida.org>  
twitter: [@caida\\_ioda](https://twitter.com/caida_ioda)

**Alberto Dainotti**  
CAIDA, UC San Diego  
[alberto@caida.org](mailto:alberto@caida.org)  
858-534-9249  
Twitter: [@AlbertoDainotti](https://twitter.com/AlbertoDainotti)







# Predict, Assess, Risk, Identify Disruptive Internet-scale Network Events (PARIDINE) Kick-off Meeting

April 10, 2018 | Arlington, VA