A COORDINATED VIEW OF LARGE-SCALE INTERNET EVENTS

Alistair King

alistair@caida.org

Bradley Huffaker, Alberto Dainotti, kc claffy

CAIDA, UCSD

LARGE-SCALE INTERNET EVENTS

(our focus)

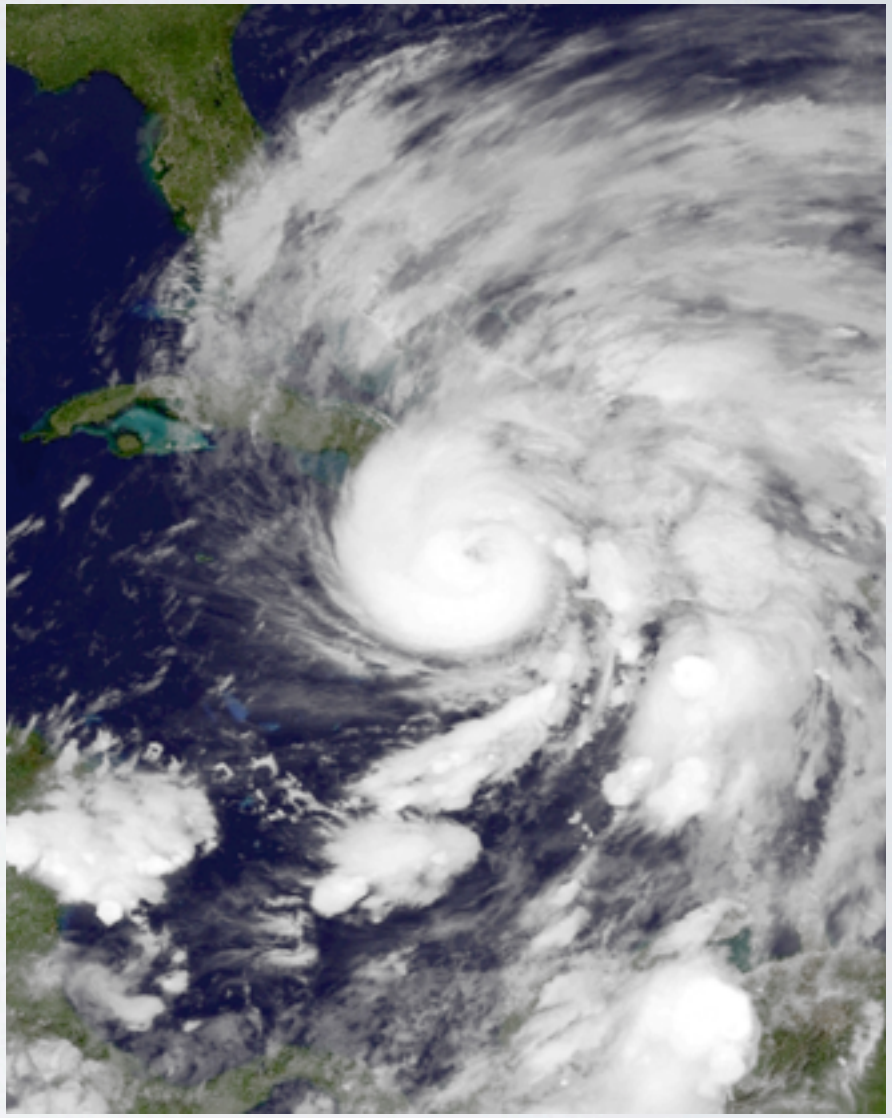

- Events that impact services for a significant section of the Internet
  - Multiple networks/providers
  - Widespread geographic/human impact
- E.g. outages due to Hurricane Sandy; Tohoku Earthquake; malicious scans/attacks; routing hijacks; etc.

EVENT VIEWS

(dimensions)

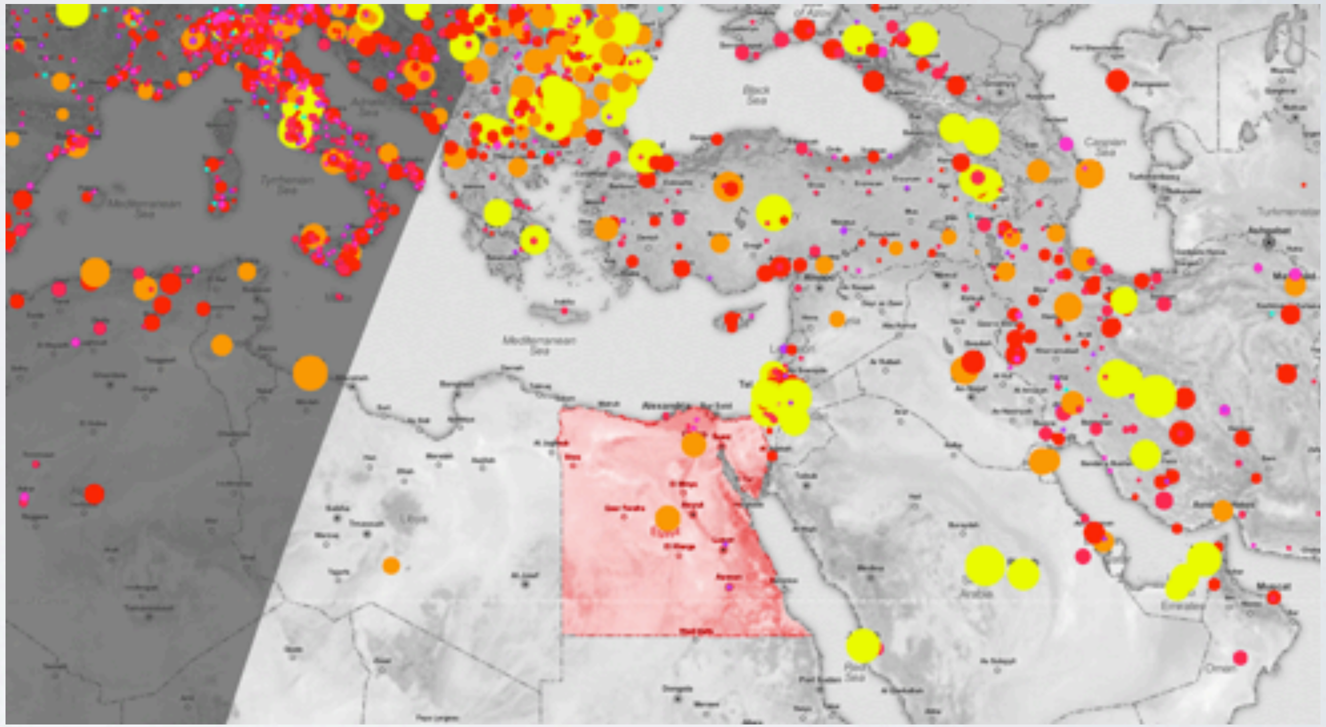

1. **Geographic**
   - City, State, Country, etc
2. **Network Traffic**
   - # packets, # bytes, # sources, etc

- Visualized using the **Cuttlefish** tool

http://www.caida.org/tools/visualization/cuttlefish/

EVENT VIEWS

(dimensions)

3. Internet Address Space

- IP address, Address Ranges, Autonomous Systems, etc
- Visualized using *ipv4-heatmap* tool
- Hilbert space-filling curve

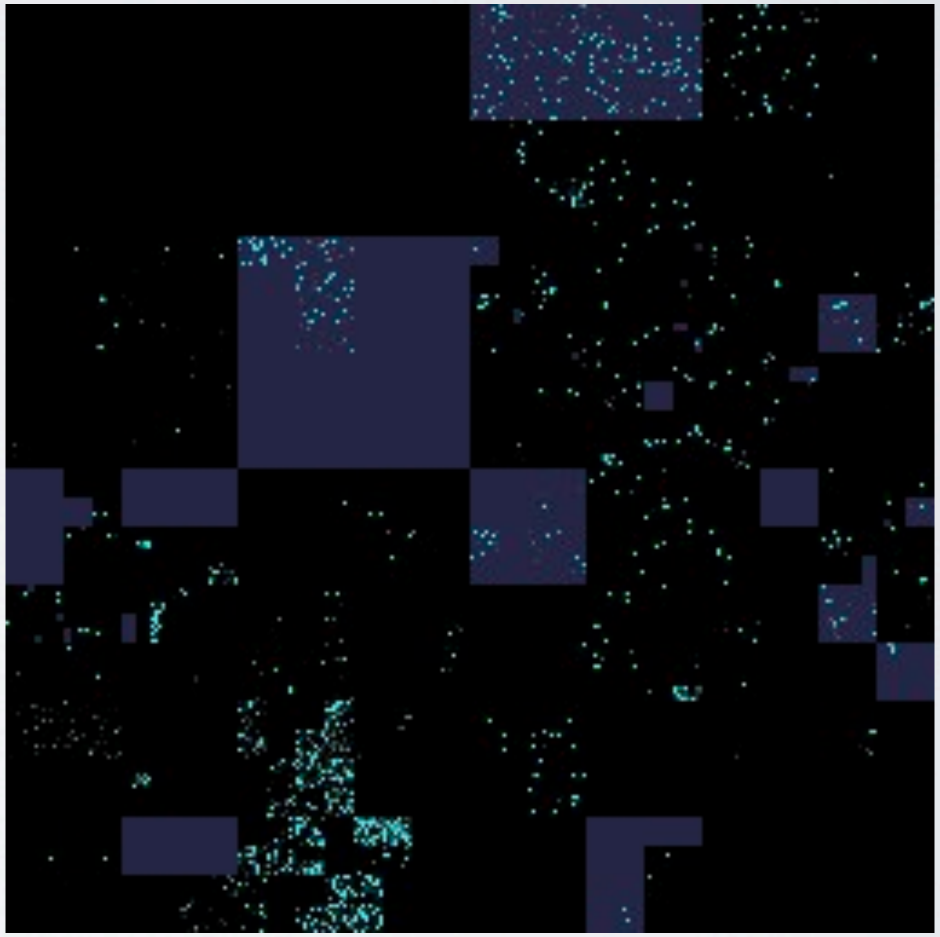

- **all three dimensions evolve over time**

http://maps.measurement-factory.com/software

COORDINATED VIEW

(putting it all together)

- Combine views into a single frame

- Synchronized by time

- Each view augments information shown in others

- *Whole is greater than the sum of the parts*

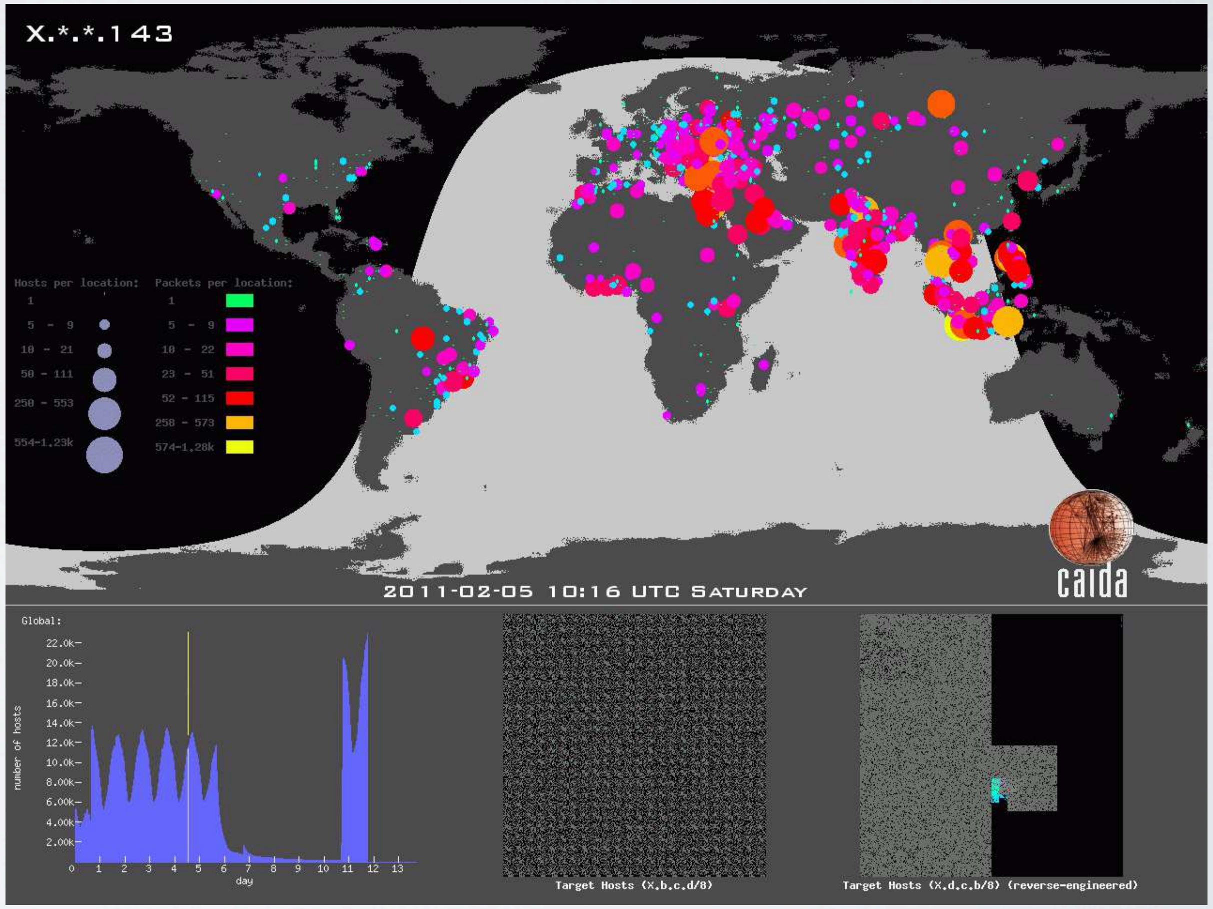

CASE STUDIES

(trying it out)

- Two Case Studies:
  - **The sipscan**
  - **Egypt Internet Blackout**
- Data captured by the UCSD Network Telescope (darknet)
  - Sipscan data available at http://www.caida.org/data/passive/sipscan_dataset.xml
  - Egypt Internet Blackout data will be released as part of an Educational Dataset at the end of 2012

DARKNETS

(or, Network Telescopes)

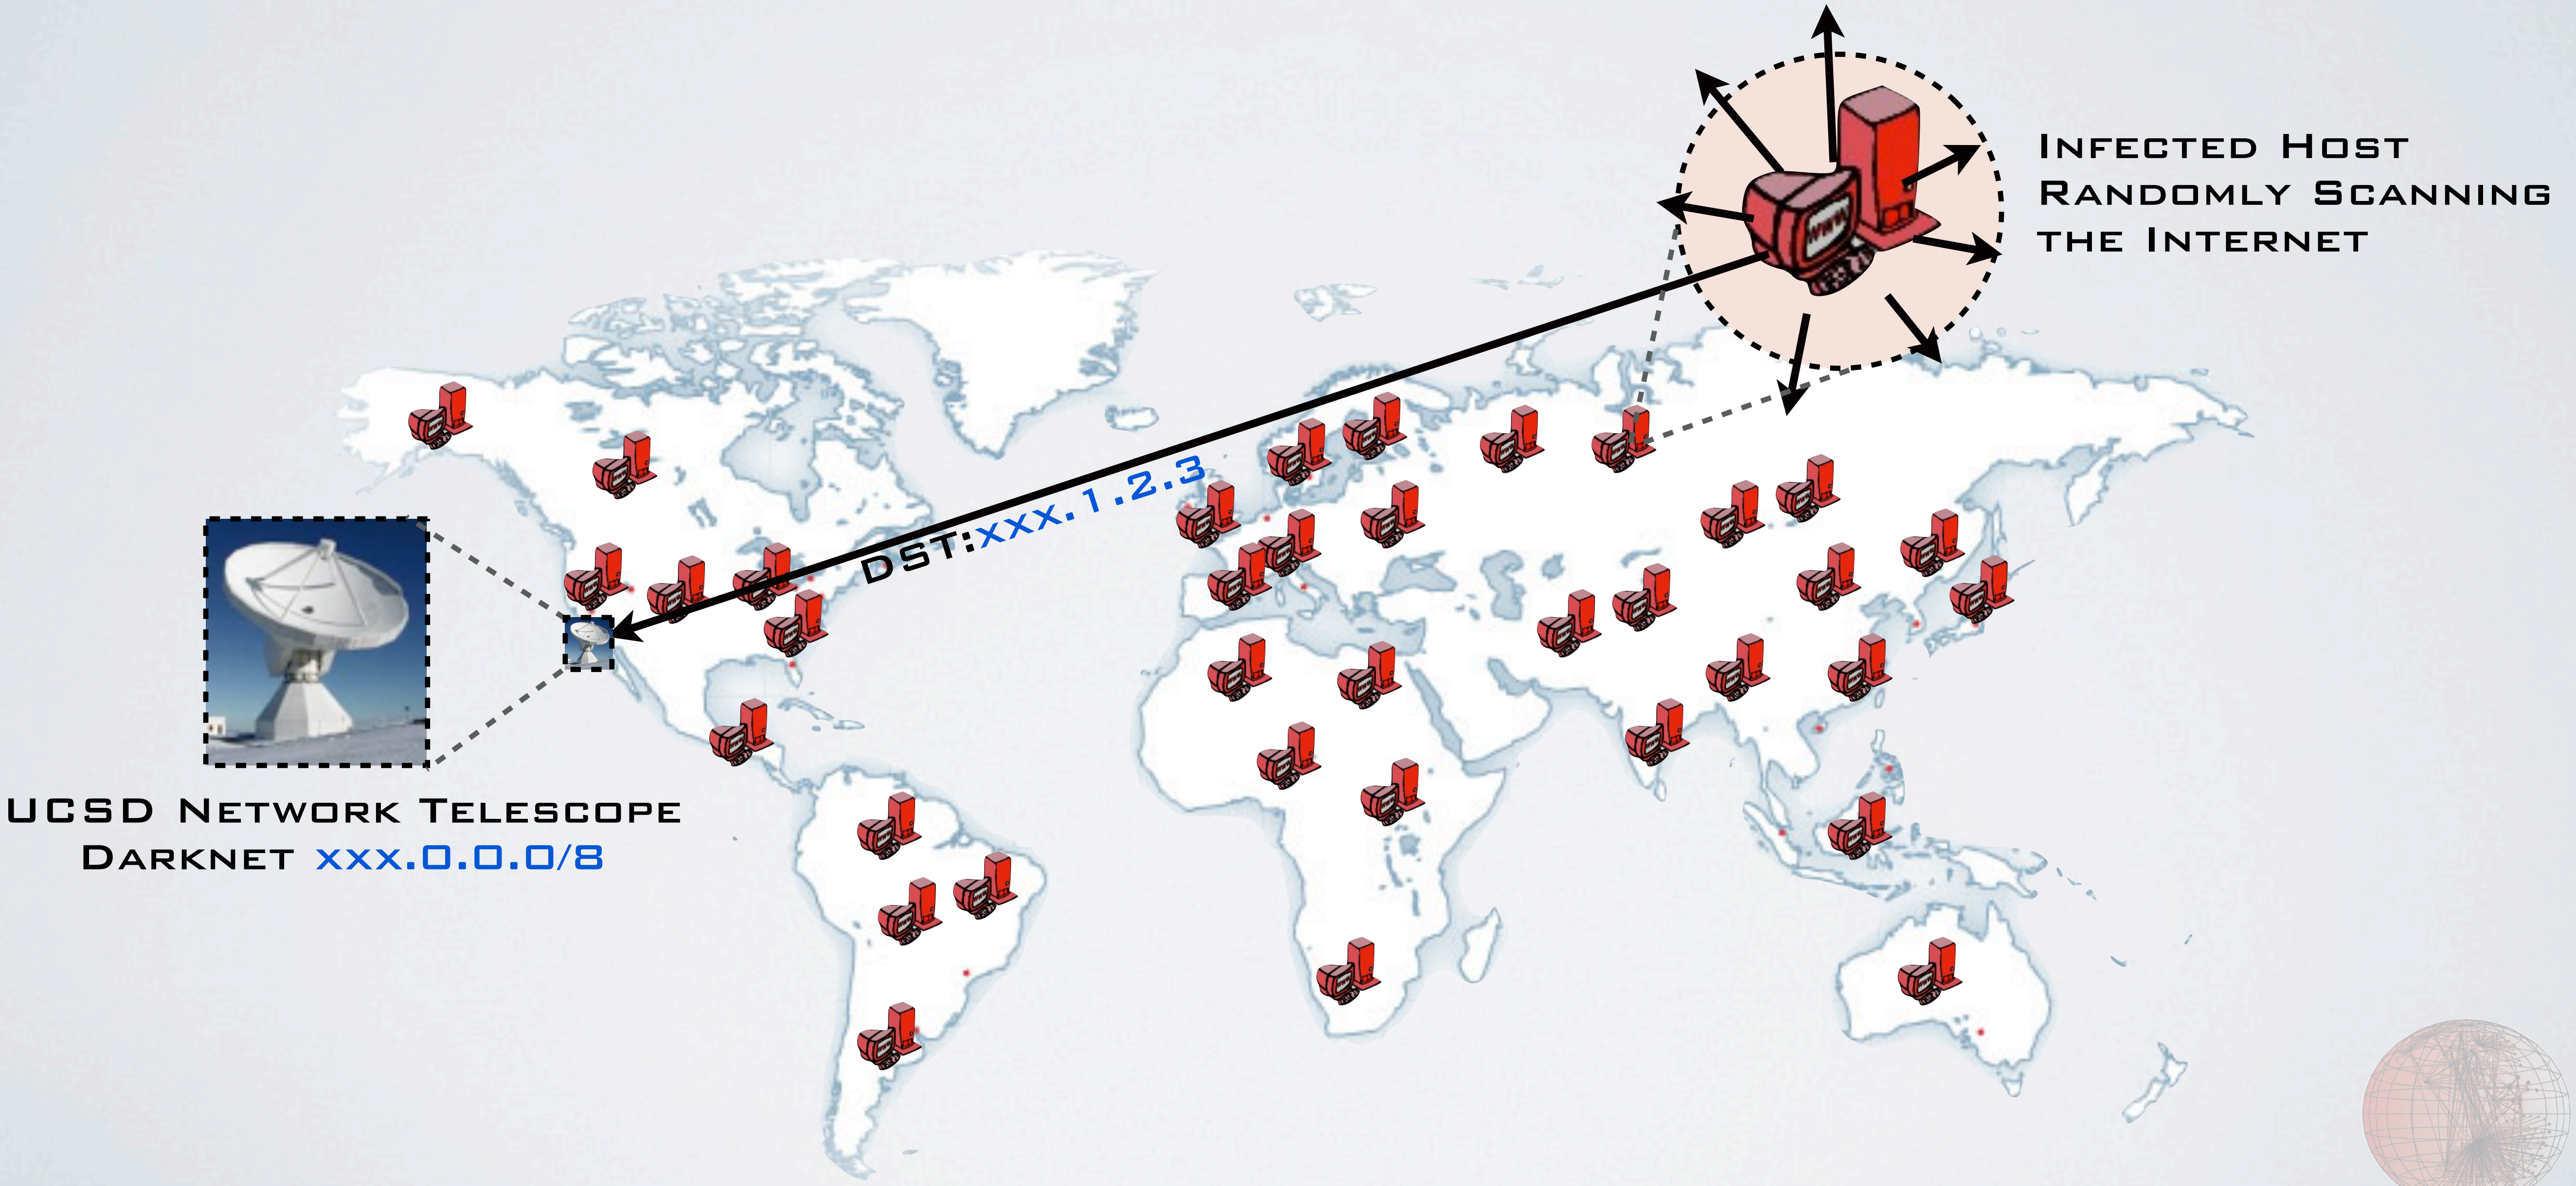

THE SIPSCAN

(a case study)

- "/0" scan from a **botnet**
- February 2011
- **Scanning SIP servers** with a query on UDP port 5060

A. Dainotti, A. King, K. Claffy, F. Papale, A. Pescapè,
"Analysis of a "/0" Stealth Scan from a Botnet",
ACM SIGCOMM Internet Measurement Conference 2012

THE SIPSCAN

(why was it interesting?)

- Covered the **entire IPv4 address space (in 12 days)**
- **Highly Coordinated**
  - Small overlap in targets probed
  - Good coverage
- **Stealthy**
  - Large turnover of geographically distributed bots
  - Reverse byte order increment of target IP

THE SIPSCAN

(why was it interesting?)

- Covered the **entire IPv4 address space (in 12 days)**
- **Highly Coordinated**
  - Small overlap in targets probed
  - Good coverage
- **Stealthy**
  - Large turnover of geographically distributed bots
  - Reverse byte order increment of target IP

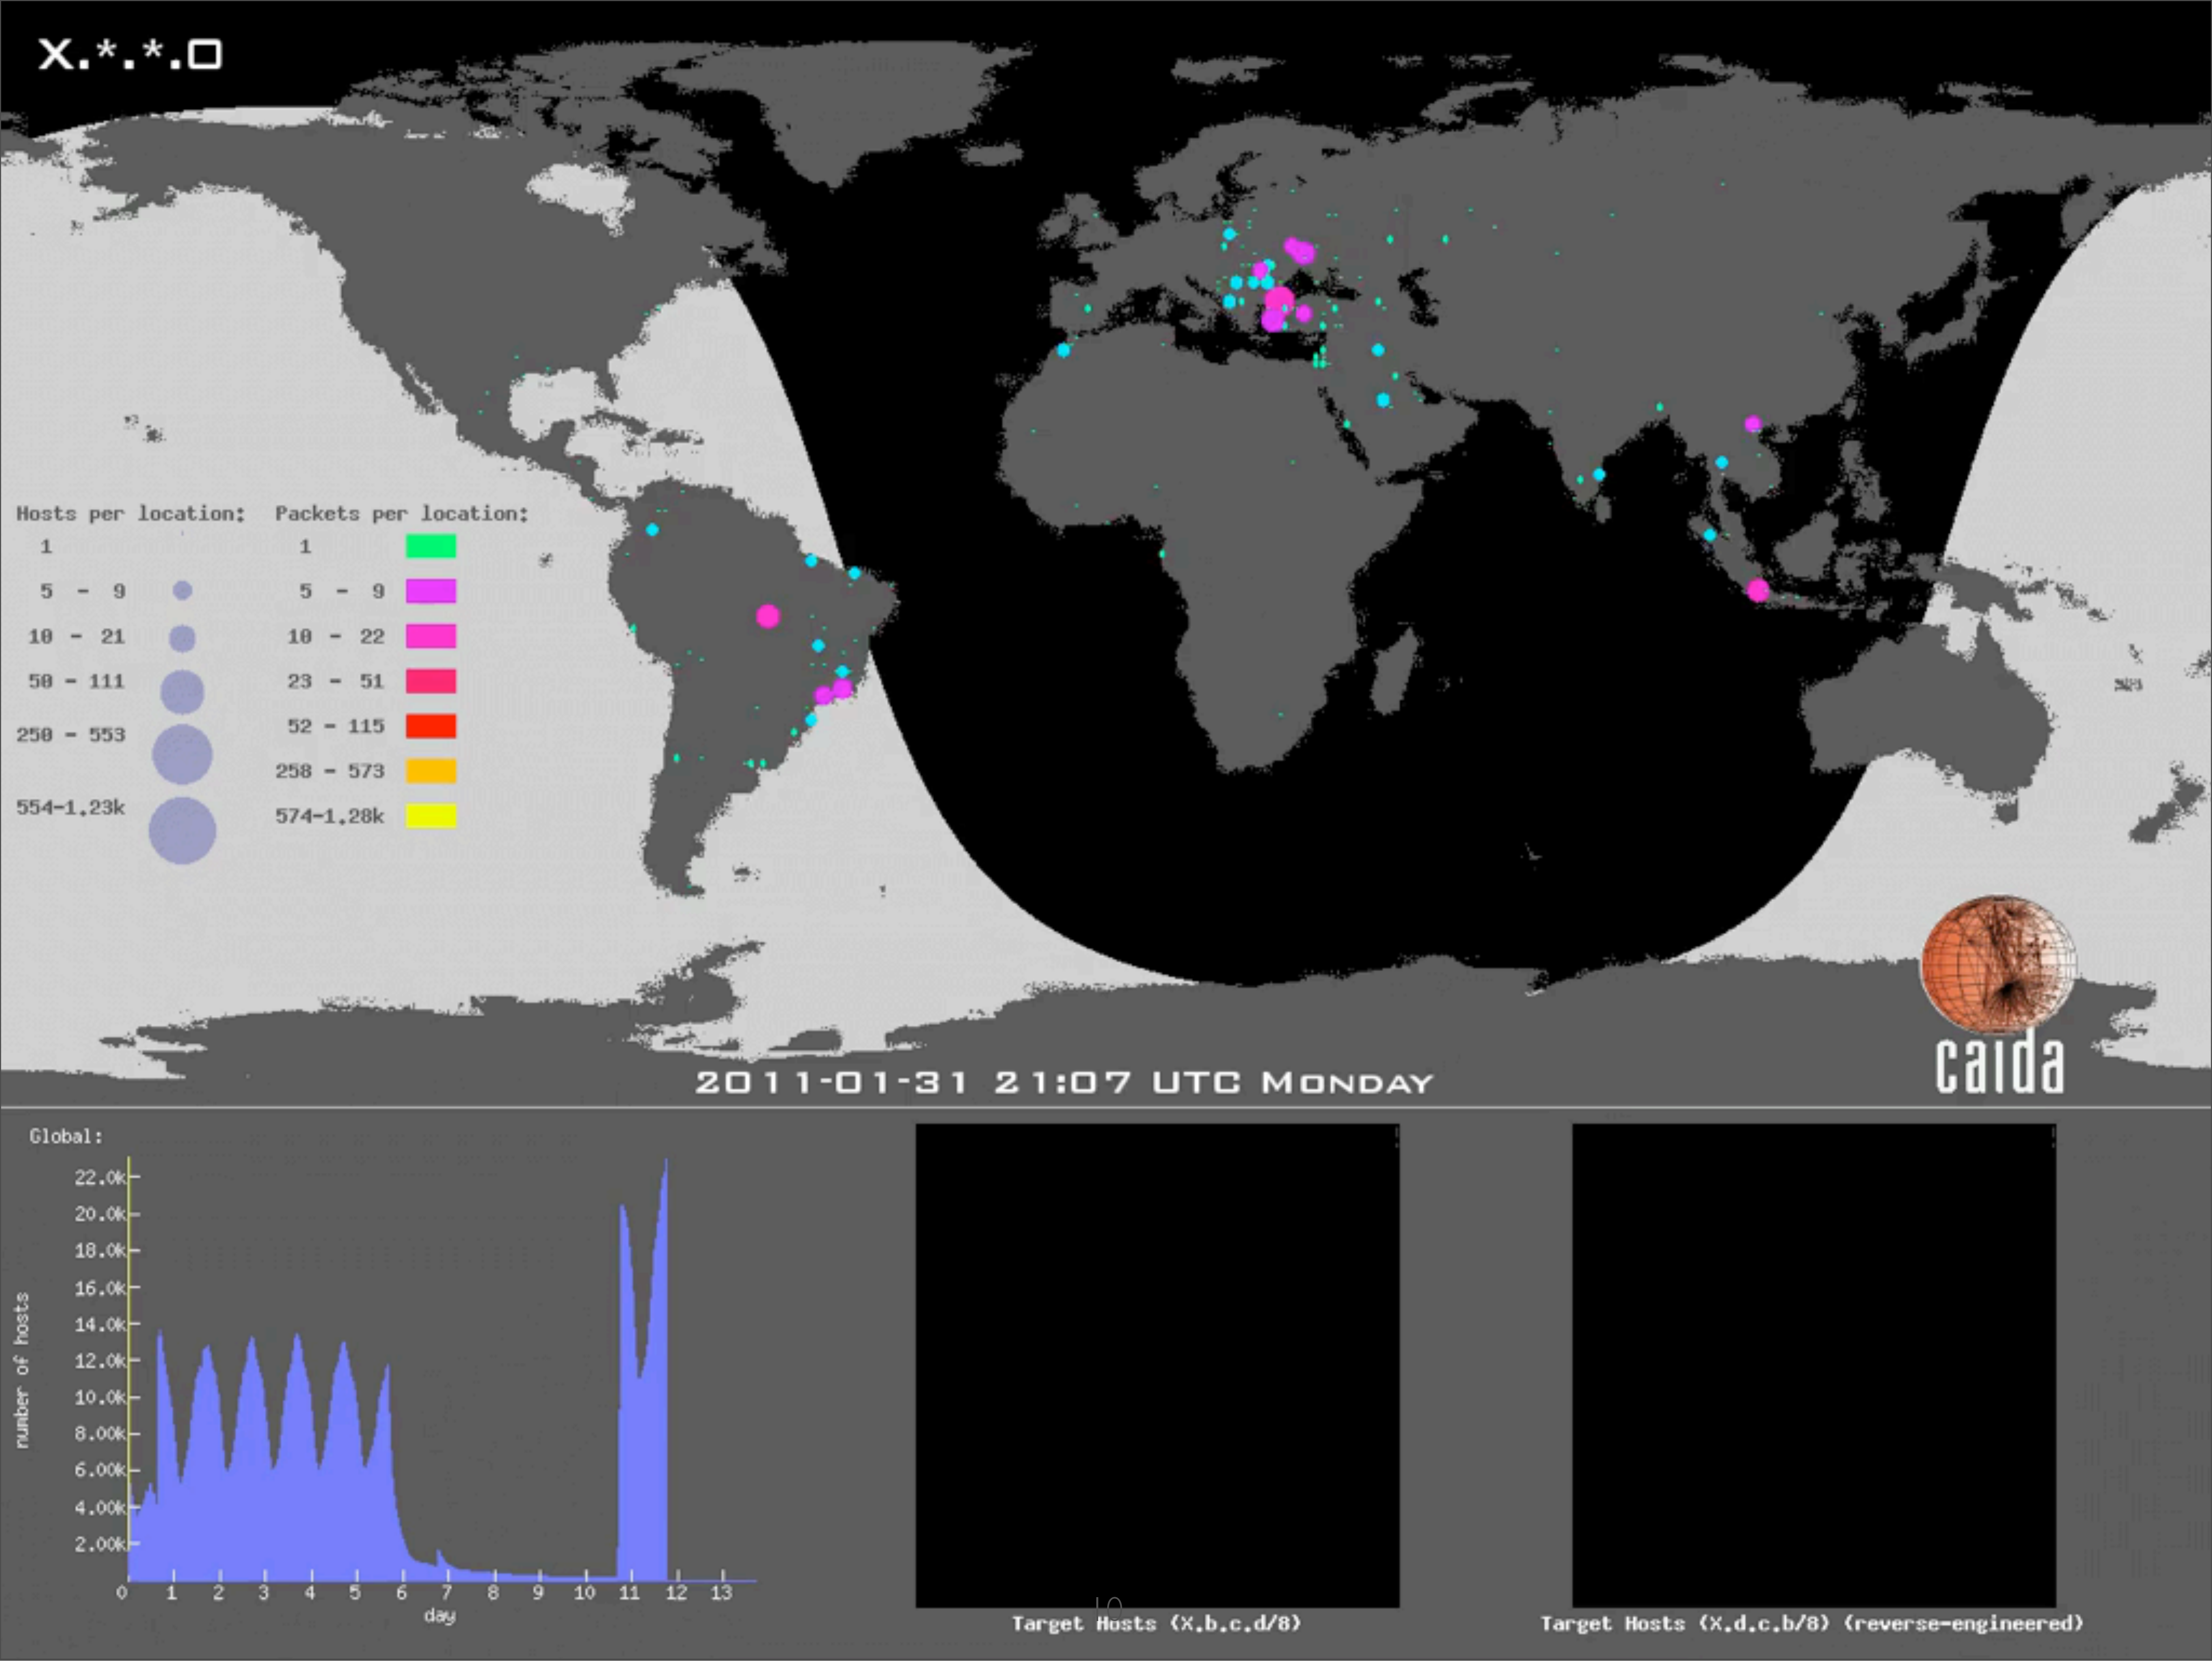
X.*.*.0
Hosts per location:
1
5 - 9
10 - 21
50 - 111
250 - 553
554-1.23k
Packets per location:
1
5 - 9
10 - 22
23 - 51
52 - 115
258 - 573
574-1.28k
2011-01-31 21:07 UTC MONDAY
caida
Global:
number of hosts
22.0k
20.0k
18.0k
16.0k
14.0k
12.0k
10.0k
8.00k
6.00k
4.00k
2.00k
0
1
2
3
4
5
6
7
8
9
10
11
12
13
day
Target Hosts (X.b.c.d/8)
Target Hosts (X.d.c.b/8) (reverse-engineered)

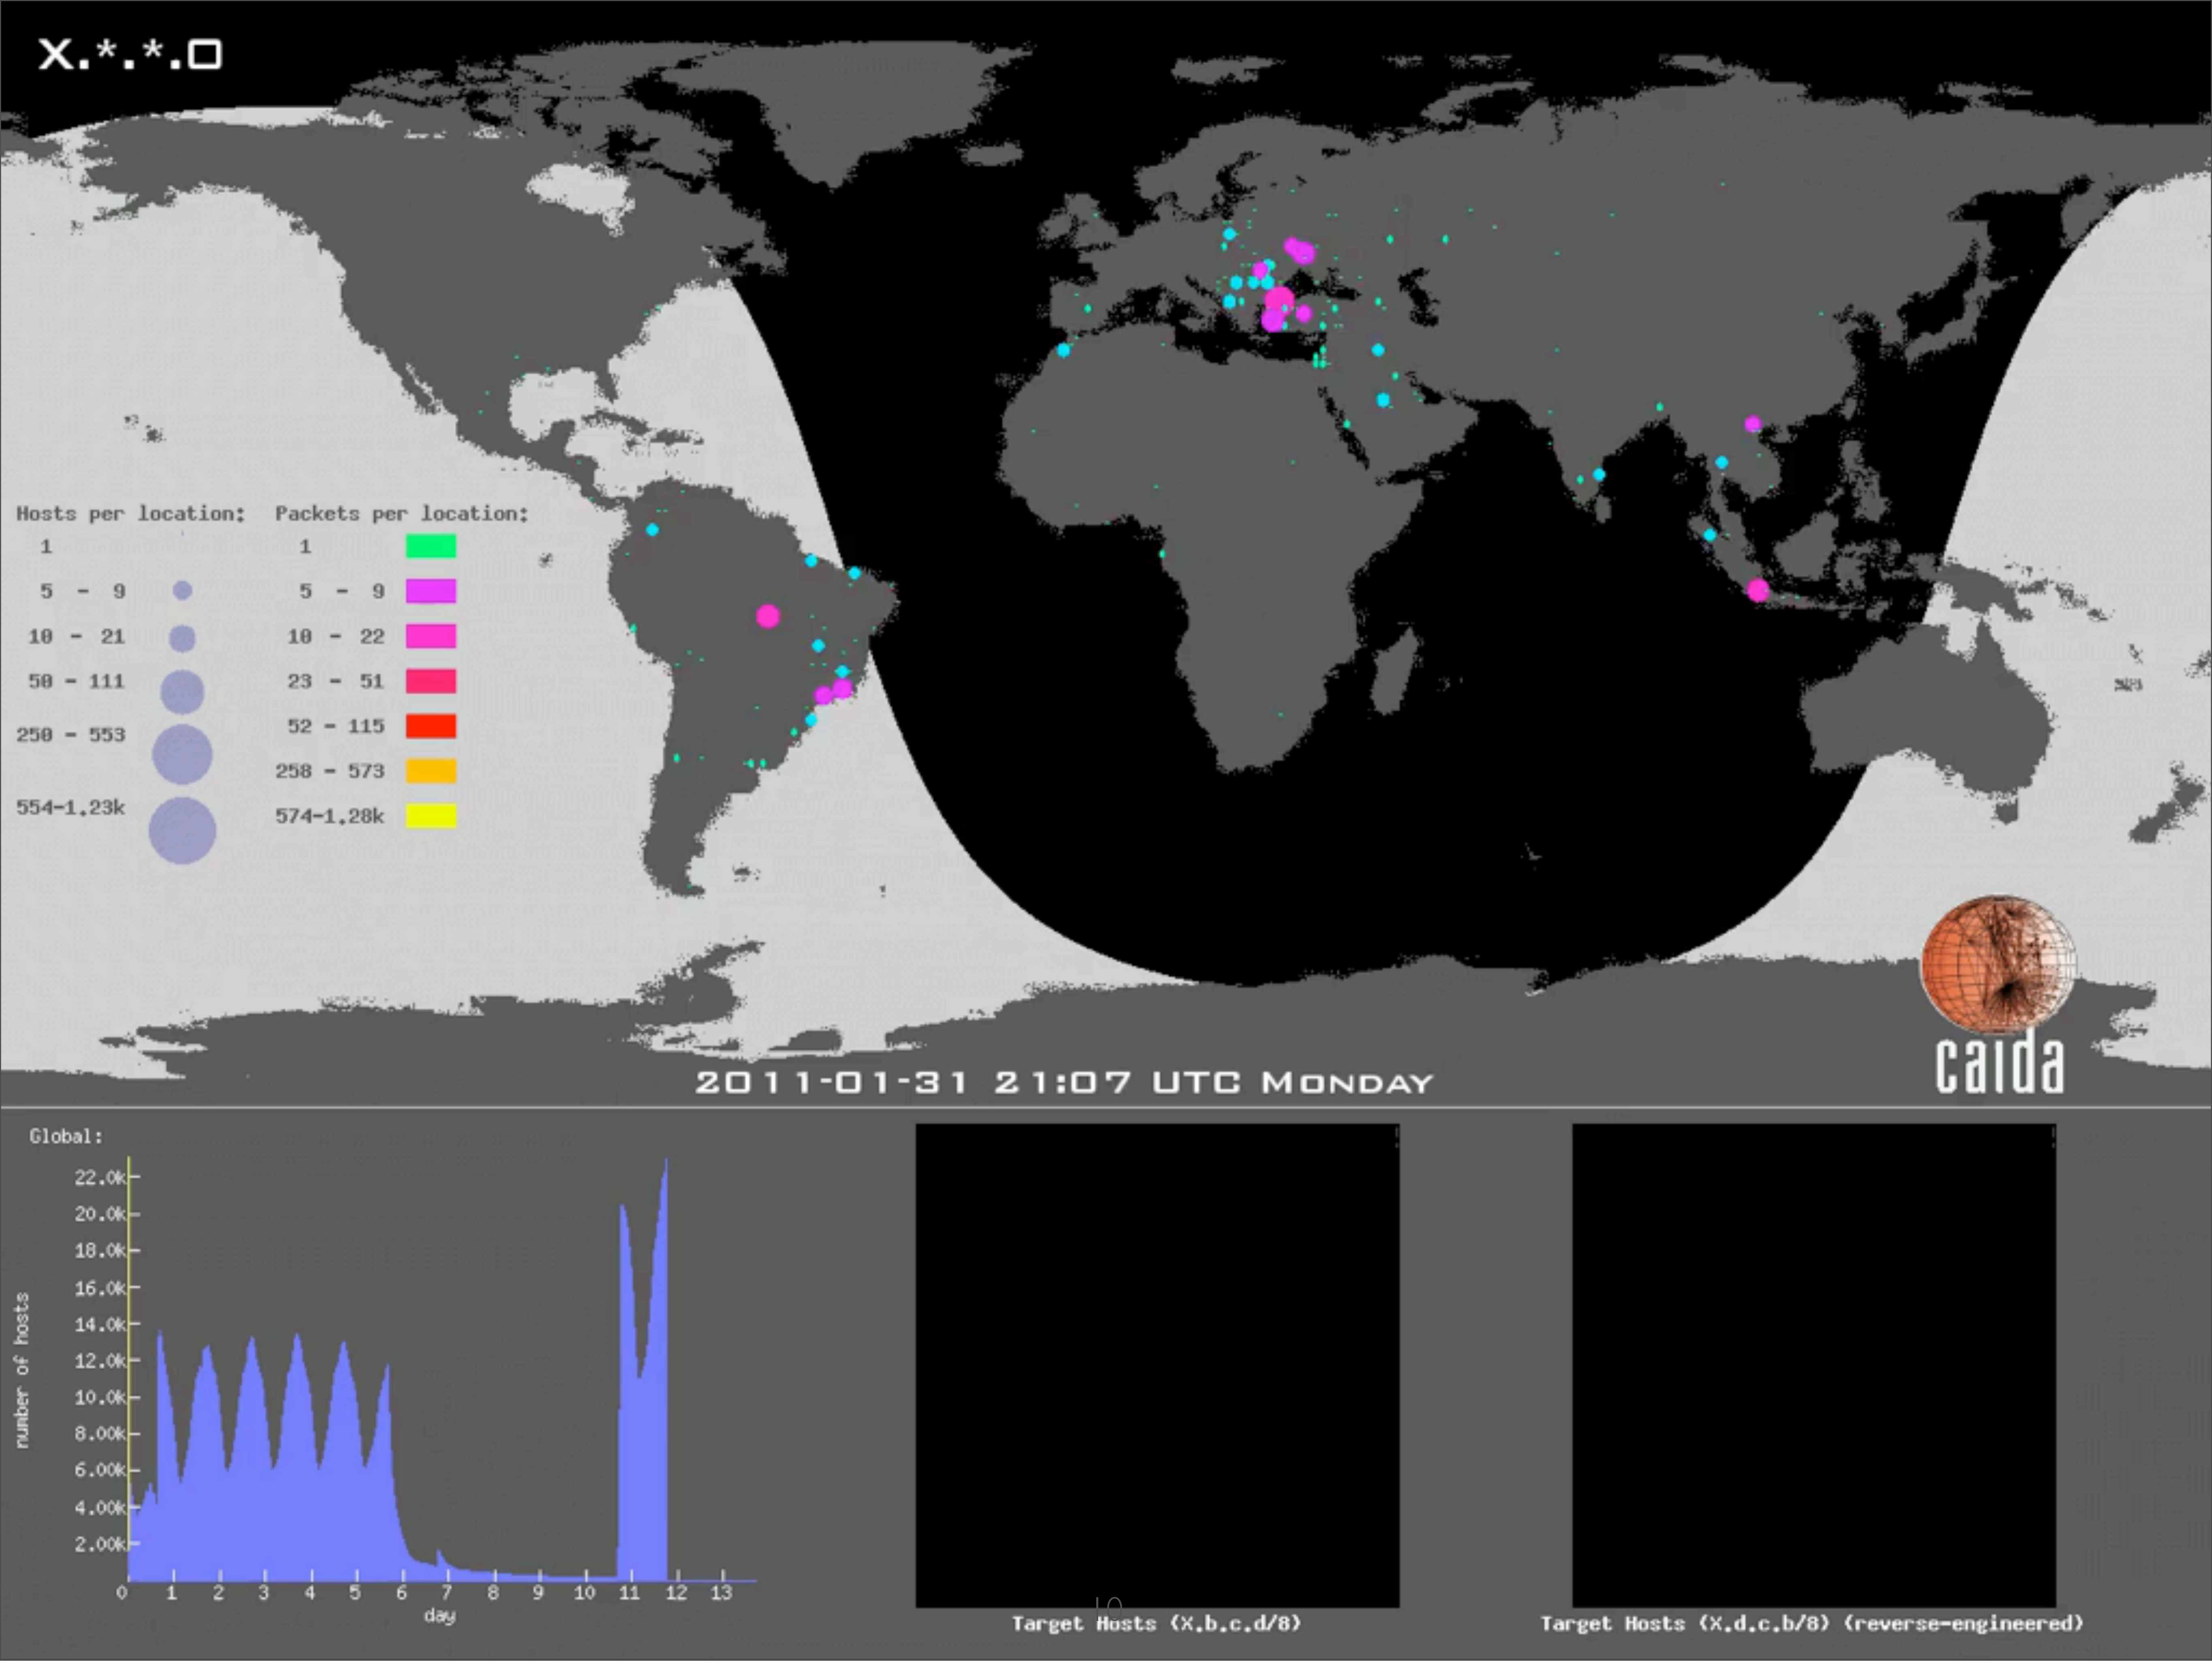X.*.*.0
Hosts per location:
1
5 - 9
10 - 21
50 - 111
250 - 553
554-1.23k
Packets per location:
1
5 - 9
10 - 22
23 - 51
52 - 115
258 - 573
574-1.28k
caida
2011-01-31 21:07 UTC Monday
Global:
22.0k
20.0k
18.0k
16.0k
14.0k
12.0k
10.0k
8.00k
6.00k
4.00k
2.00k
number of hosts
0
1
2
3
4
5
6
7
8
9
10
11
12
13
day
Target Hosts (X.b.c.d/8)
Target Hosts (X.d.c.b/8) (reverse-engineered)

EGYPTIAN INTERNET BLACKOUT

(another example)

- Egyptian government ordered **Internet censorship**
- Most BGP **routes to Egyptian networks withdrawn**
- 5 days beginning January 27 2011

A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescapé.
"Analysis of country-wide internet outages caused by censorship."
ACM SIGCOMM Internet Measurement Conference 2012

EGYPTIAN INTERNET BLACKOUT

(and why it is interesting)

- **Internet access was denied to an entire country**

  ... even to the malware

- Conficker-infected hosts can no longer send packets

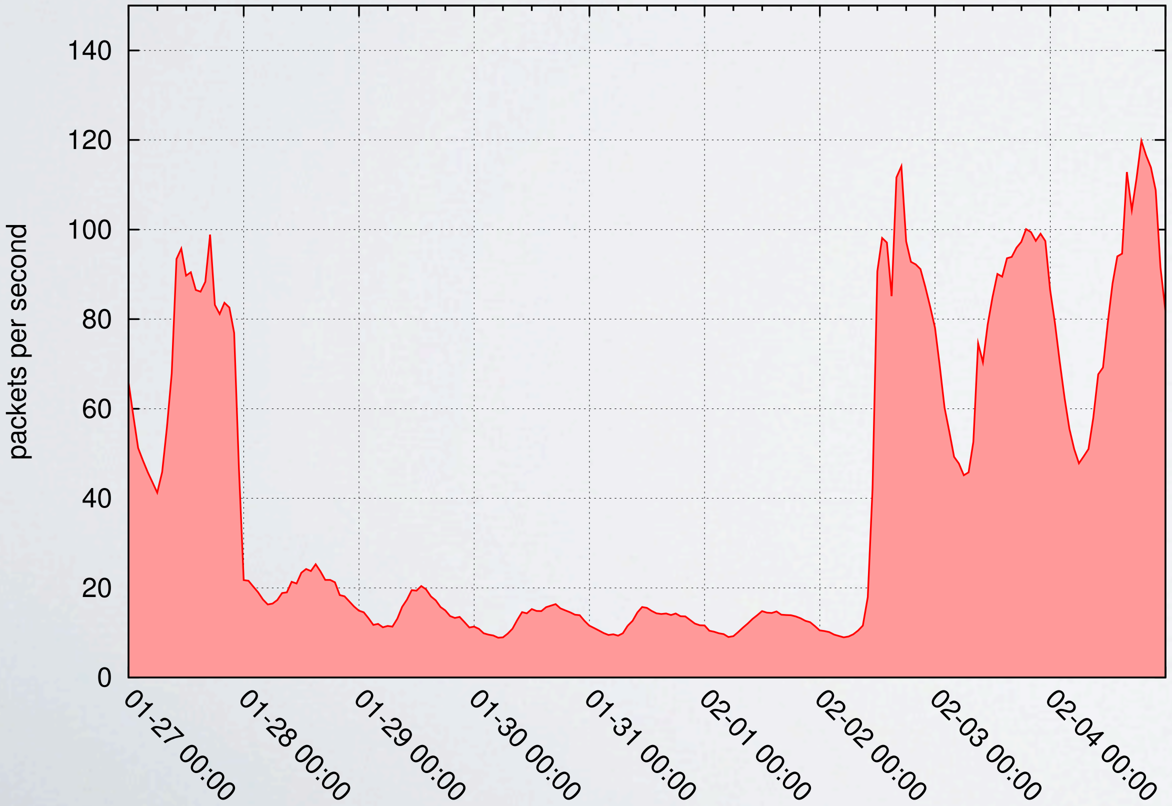

- Drop in packets to TCP port 445 observed by the UCSD Network Telescope

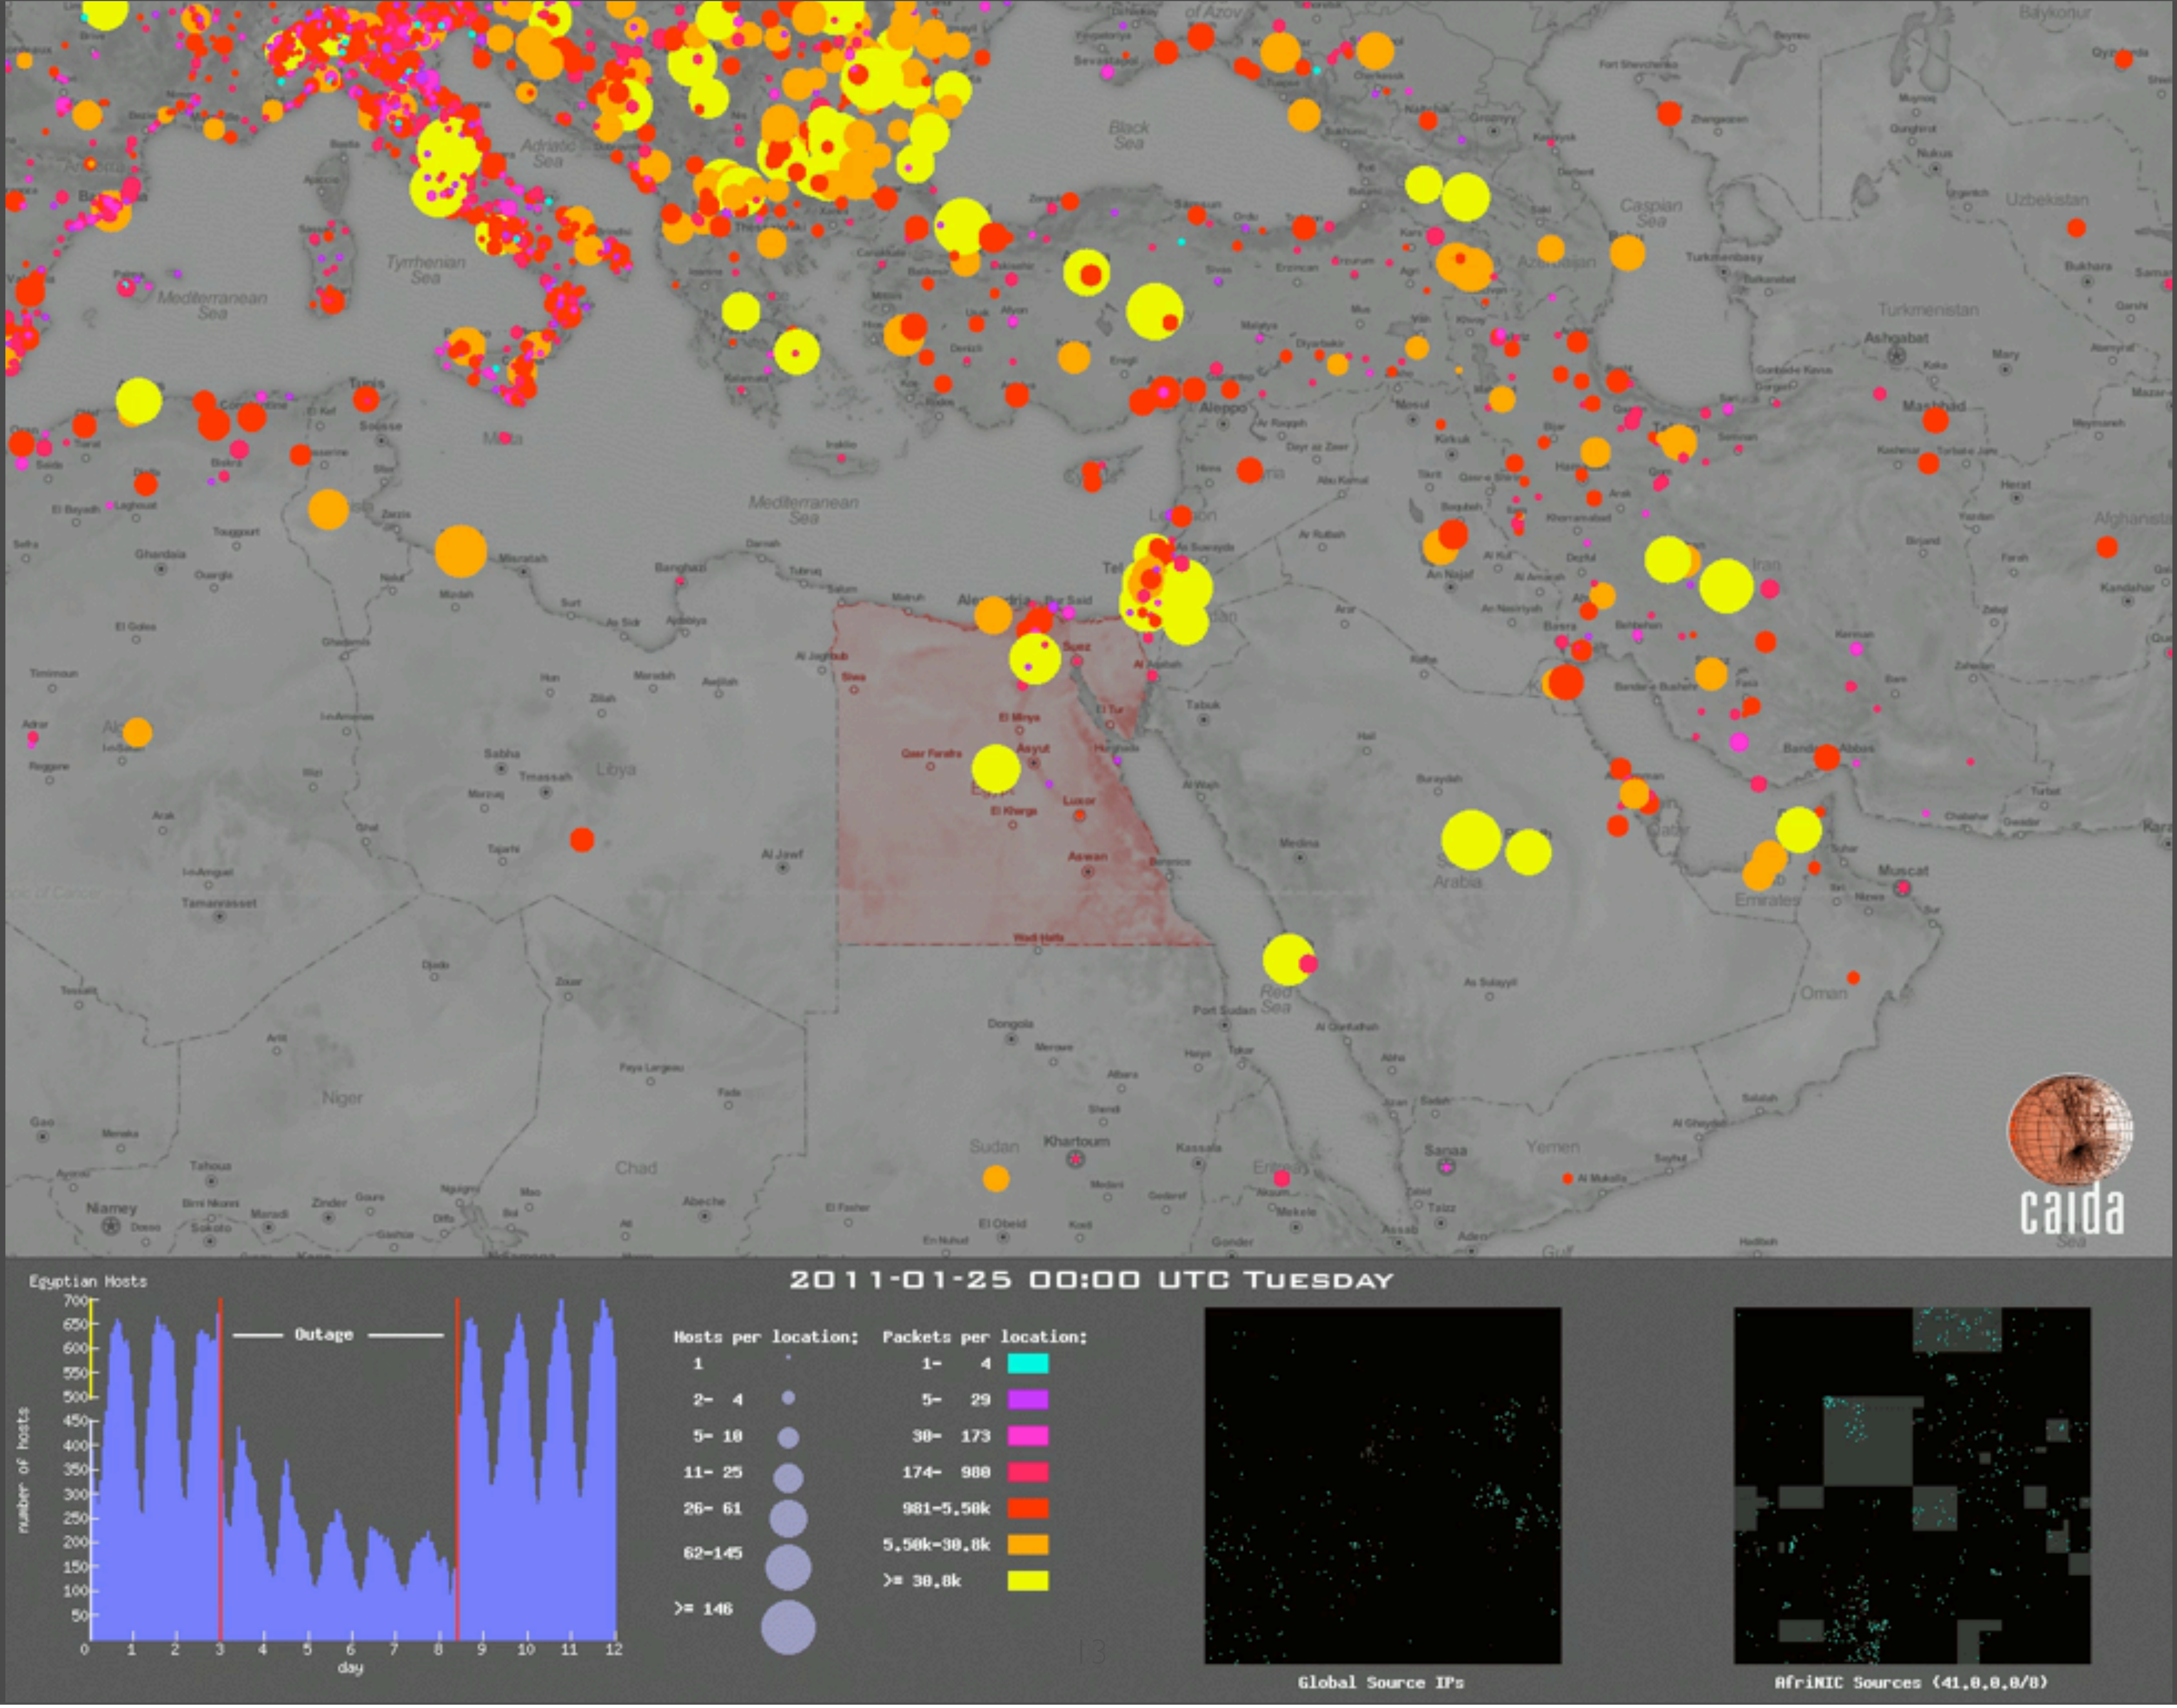
2011-01-25 00:00 UTC TUESDAY
Egyptian Hosts
number of hosts
day
Outage
700
650
600
550
500
450
400
350
300
250
200
150
100
50
0
1
2
3
4
5
6
7
8
9
10
11
12
Hosts per location:
1
2- 4
5- 10
11- 25
26- 61
62-145
>= 146
Packets per location:
1- 4
5- 29
30- 173
174- 980
981-5.50k
5.50k-30.8k
>= 30.8k
Global Source IPs
AfriNIC Sources (41.0.0.0/8)
caida
Mediterranean Sea
Tyrrhenian Sea
Adriatic Sea
Black Sea
Caspian Sea
Sea of Azov
Turkmenistan
Uzbekistan
Afghanistan
Libya
Egypt
Sudan
Chad
Niger
Arabia
Yemen
Oman
Emirates
Red Sea
Iran
Aleppo
Tunis
Khartoum
Sanaa
Muscat
Ashgabat
Mashhad
Baykonur

CONCLUSIONS

- Applied several Information Visualization techniques to large-scale Internet events.

- Used Multiple Coordinated Views to study temporal evolution along different dimensions.

- Potentially allows insights individual views do not

FUTURE WORK

(where are we going with this?)

- Develop additional views/dimensions to include
- Integrate into near-realtime reporting system for Telescope
- Leverage web frameworks (e.g. D3) for interactive viz
- Improve signal to noise ratio by utilizing different geographic aggregation methods (e.g. Voronoi diagrams)

QUESTIONS?

(suggestions?)

- **Animations are available at:** http://www.caida.org/publications/papers/2012/coordinated_view_internet_events/supplemental/