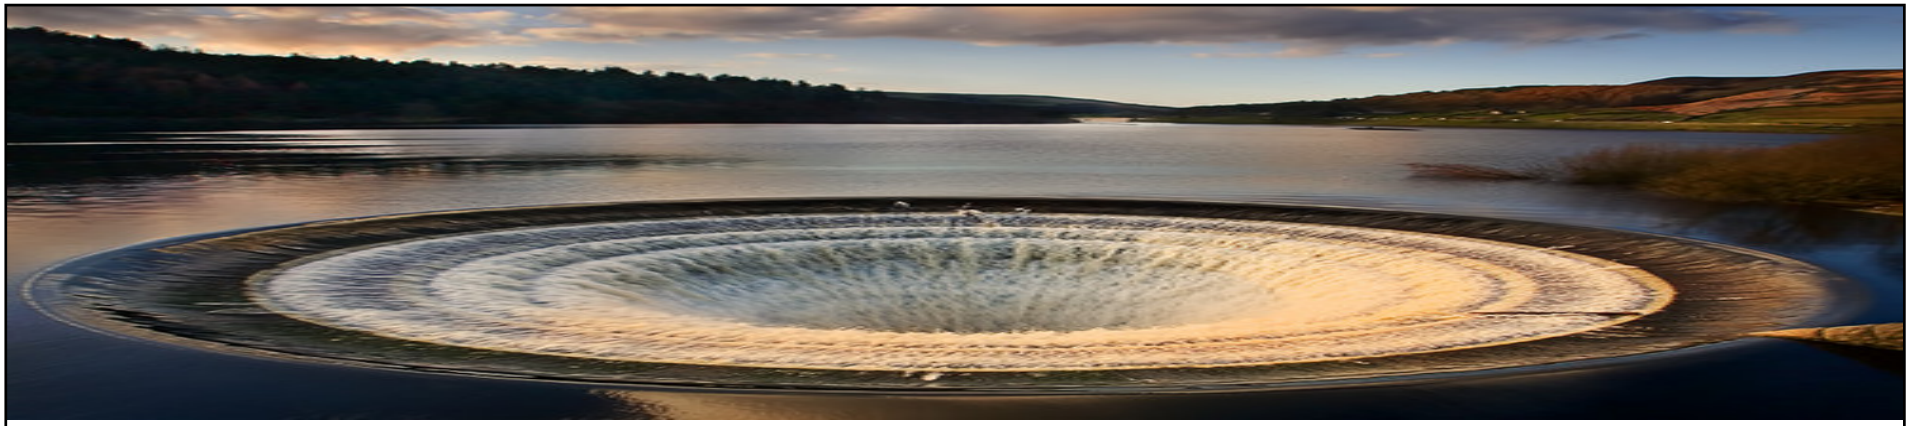

SinkMiner

Mining Botnet Sinkholes for Fun and Profit

Babak Rahbarinia[1], Roberto Perdisci[1,2], Manos Antonakakis[3], David Dagon[2]

[1]University of Georgia
[2]Georgia Tech Information Security Center, [3]Damballa Labs

University of Georgia
Dept. of Computer Science

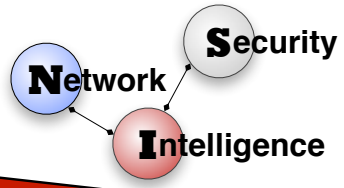

Botnet Sinkholes

- Sinkholes: take over the botnet!
  - disable the botnet
  - Enumerate victims, study C&C protocol, etc.
- Examples
  - *Your Botnet is My Botnet* (ACM CCS 2009)
  - Confiker Working Group – *Lessons Learned*
  - *etc…*

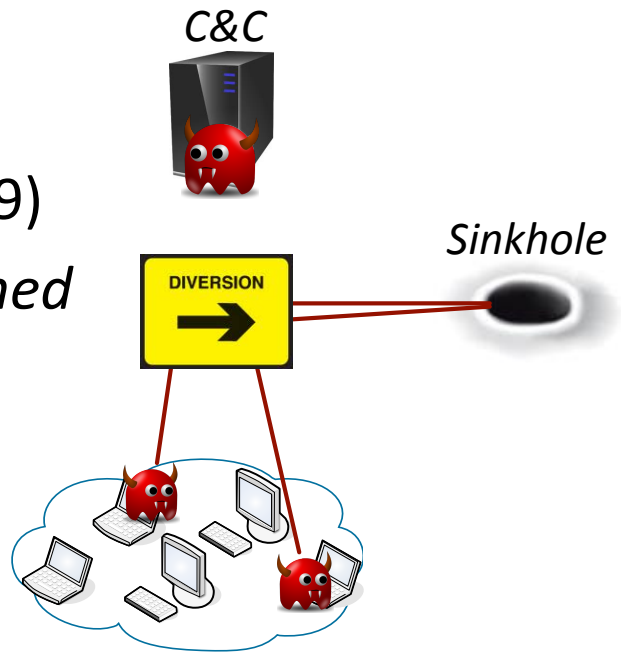

SinkMiner

- Where are the sinkholes?
- Why do you even care?!?
  - measuring *effective* C&C domain lifetime

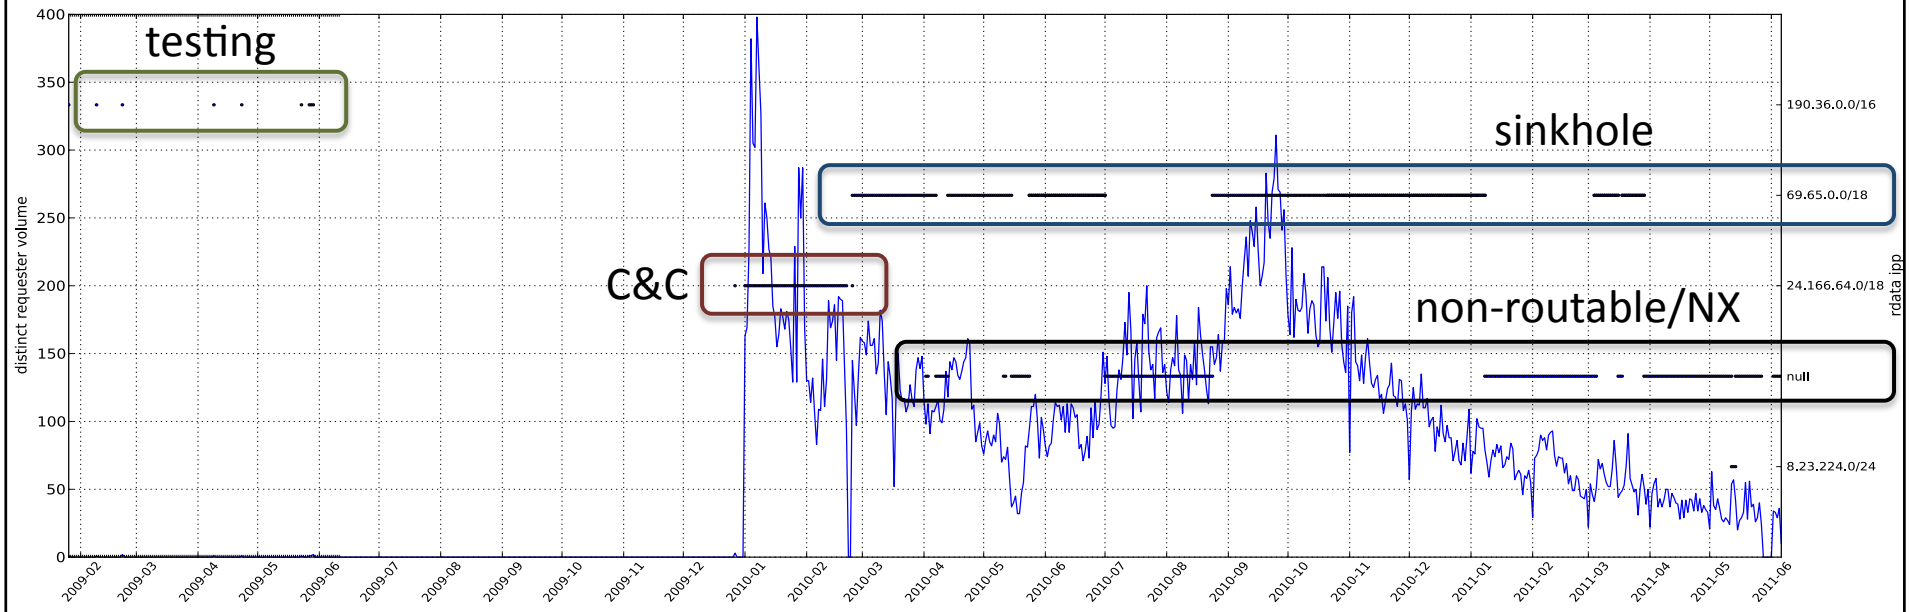

Queries seen from the authoritative NS for the C&C domain

SinkMiner

- Other benign reasons to care
  - Avoid friendly fire

- Not so benign reasons
  - I want your domain blacklist!

Mining Sinkholes

- Surprisingly, sinkholed domains often relocated from a sinkhole IP to another!
  - We thought they would stay put!
- Idea
  - follow the evolution of sinkholed domains

SinkMiner System Overview

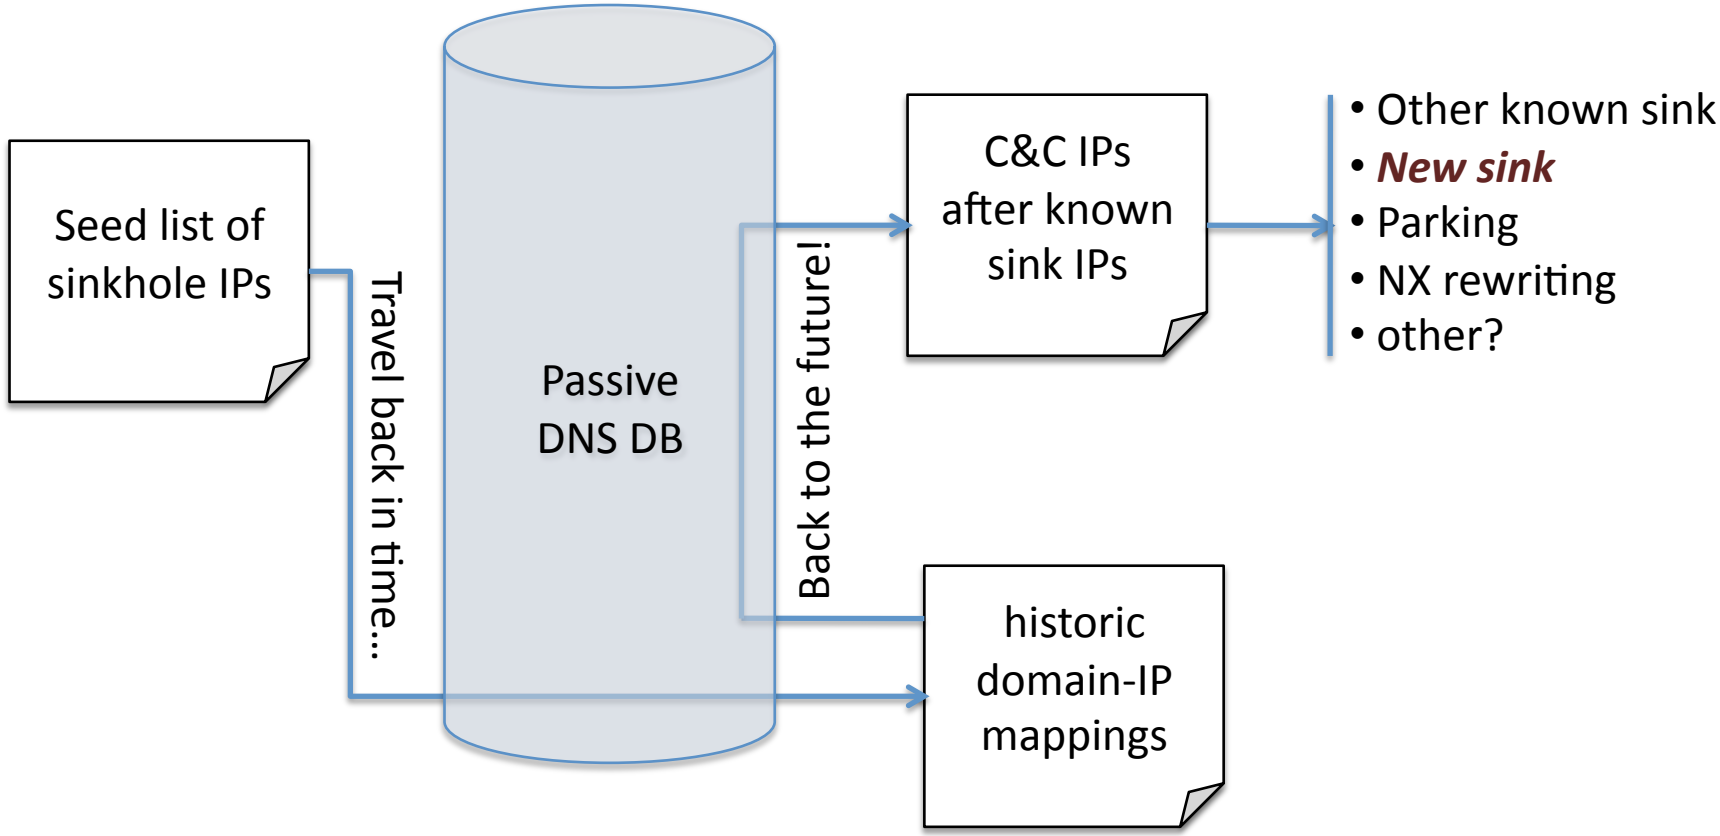

Preliminary Labeling

- Using a set of heuristics:
  - Name server names

IPs resolved by `ns1.sinkhole.ch` are *sinkholes*

- Also, `torpig-sinkhole.org`, `dns3.sinkdns.net`, `sinkhole-00.shadowserver.org`, …

IPs resolved by `nx1.dnspark.net` are *parking*

- Also, `dns1.ns-park.net`, `park1.dns.ws`, `one.parkingservice.com`, …

Preliminary Labeling

- NX-rewriting IPs

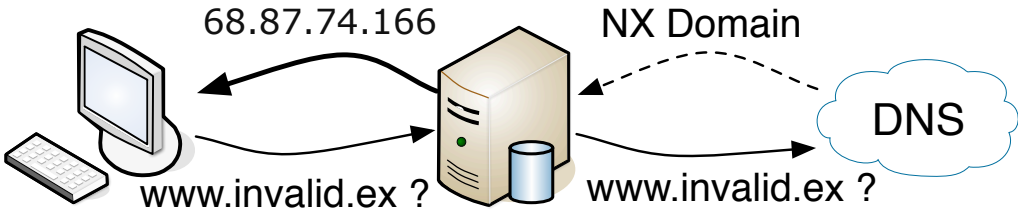

  - IPs that are pointed to by lots of non-existent and/or invalid domain names
  - Very large volumes of domains
    - all “invalid” resolutions from entire networks

Preliminary Labeling

- Volume of domains per IP
  - Observation: same known sinkhole IPs used (in time) to take down lots of domains
  - Heuristics:
    - IPs pointed to by lots of domains are probably sinkholes
    - Especially if they are in the same AS as known sinkholes

Number of domains pointing to known sinkhole IP (cumulative, over time)

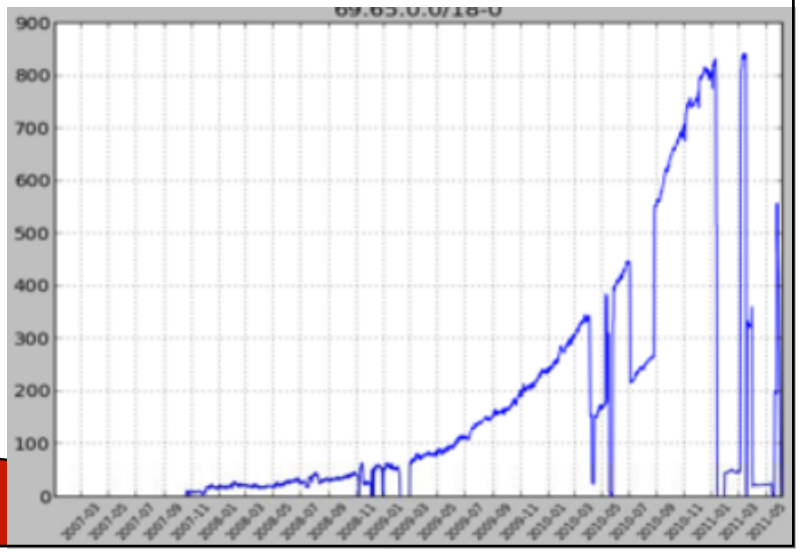

Graph-based Labeling

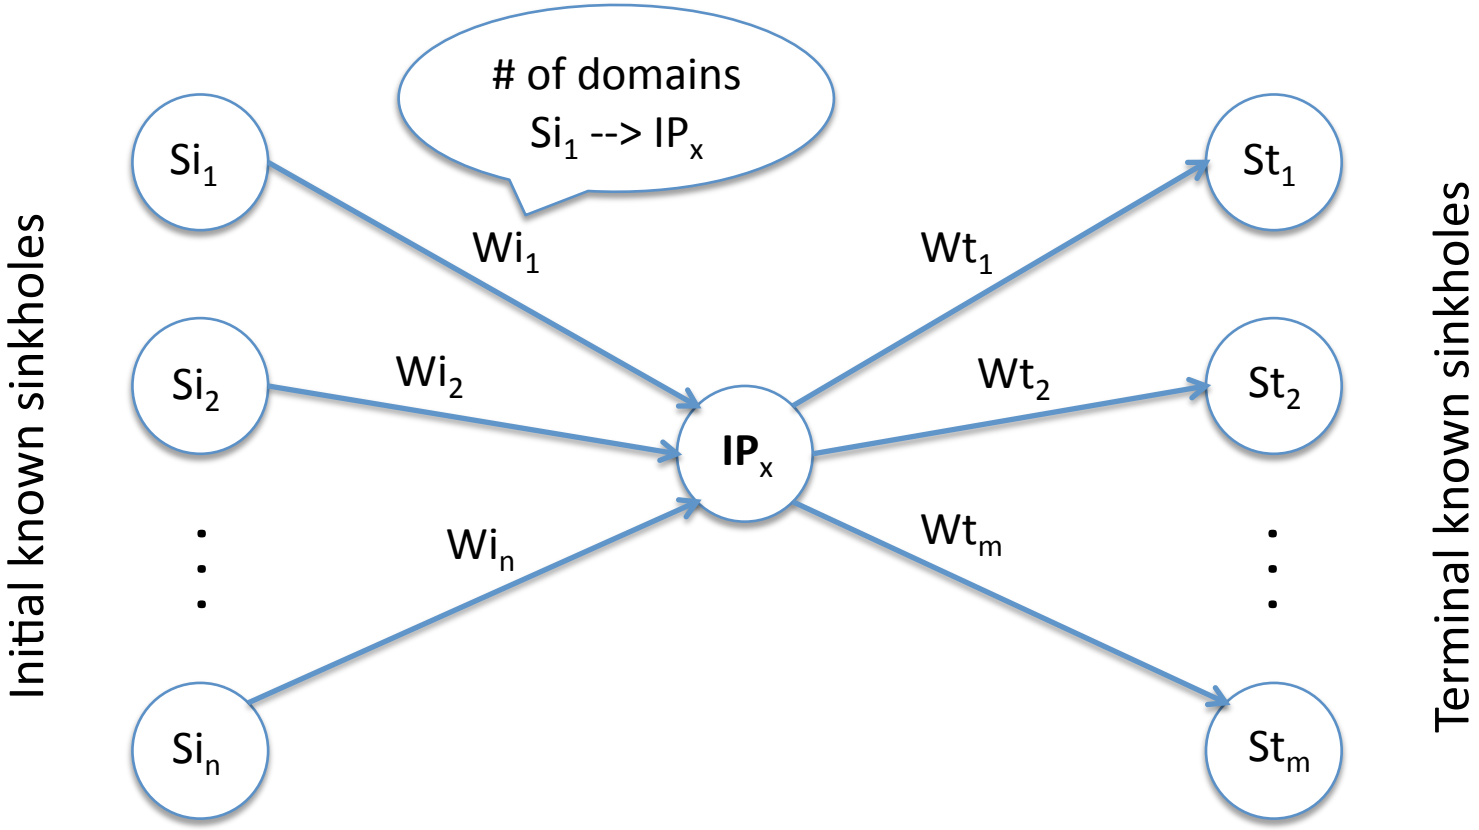

Graph-based Labeling

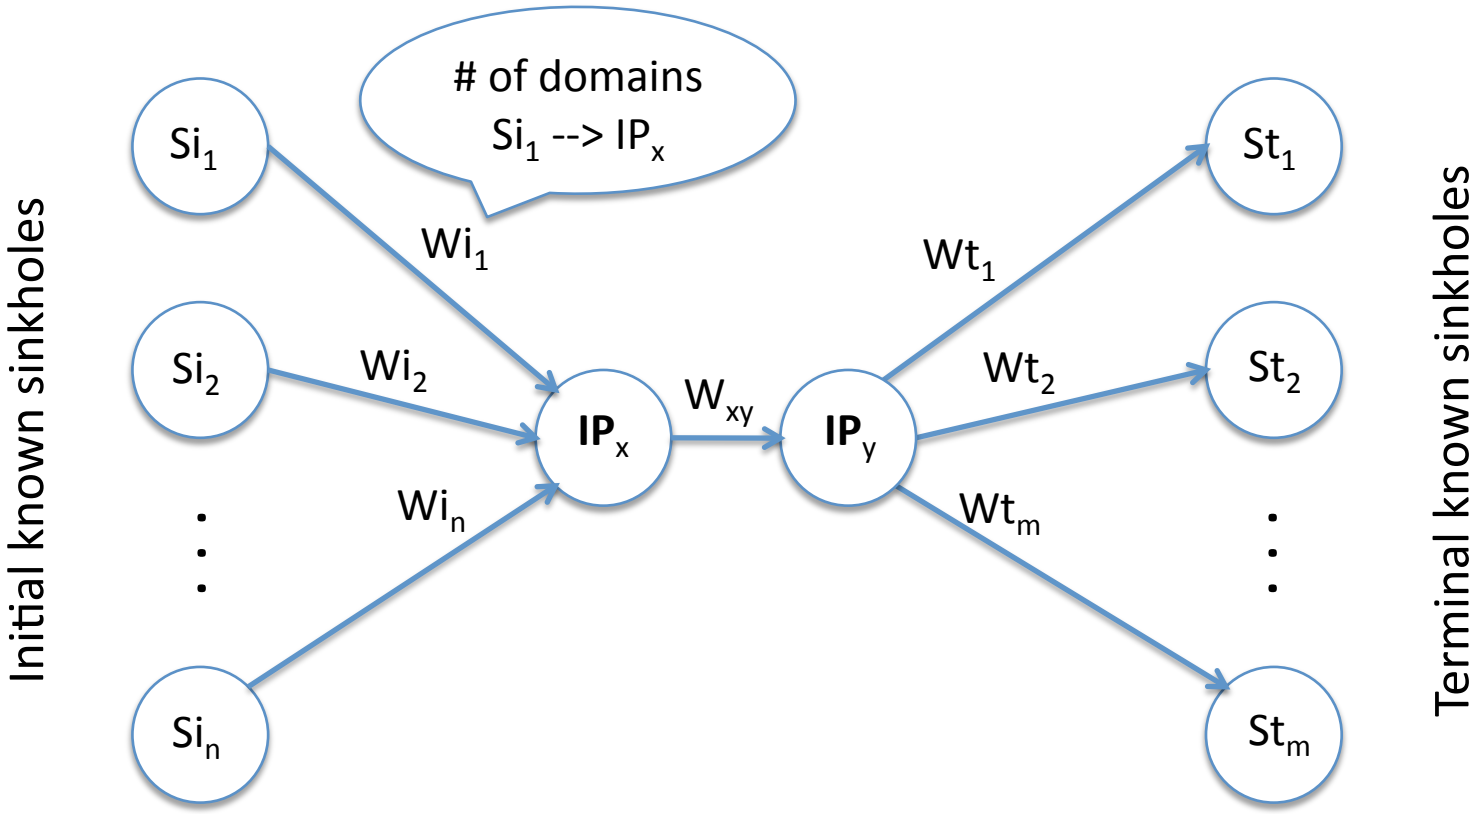

Initial Sinkholes

- 22 known sinkholes (19 different ASes)
  - 2,945,483 sinkholed domains form passive DNS
  - 130,901 after pruning DGA domains

ASN	Organization	Popularity	ASN	Organization	Popularity
14618	AMAZON-AES	46,959	1280	ISC	16,987
8069	MICROSOFT	16,522	2637	GEORGIATECH	15,390
30060	VERISIGN	11,168	15169	GOOGLE	630

- 39% of 130,901 domains changed IP after they were sinkholed
- Potential sinkholes 5,576

Preliminary Labeling

- Name server names
  - 475 parking IPs
  - 15 sinkhole IPs
- 7 NX-rewriting IPs
- Popularity + AS name
  - 23 highly popular sinkholes

Expanded seed list from 22 to 60

Graph-based Labeling

- Graph: >5k nodes and >164k edges
- 49 new sinkholes
  - 12 using 1st GDB query + 37 with 2nd GDB query

Some examples…

IP	ASN	Organization	Popularity
93.170.52.30	44557	DRAGONARA	817,563
216.239.32.21	15169	GOOGLE	535,638
69.25.27.173	10913	INTERNAP	347,902
208.91.197.101	40034	CONFLUENCE	337,539
174.129.212.2	14618	AMAZON	110,381
199.2.137.141	3598	MICROSOFT	1,367

Overall: 87 new likely sinkholes

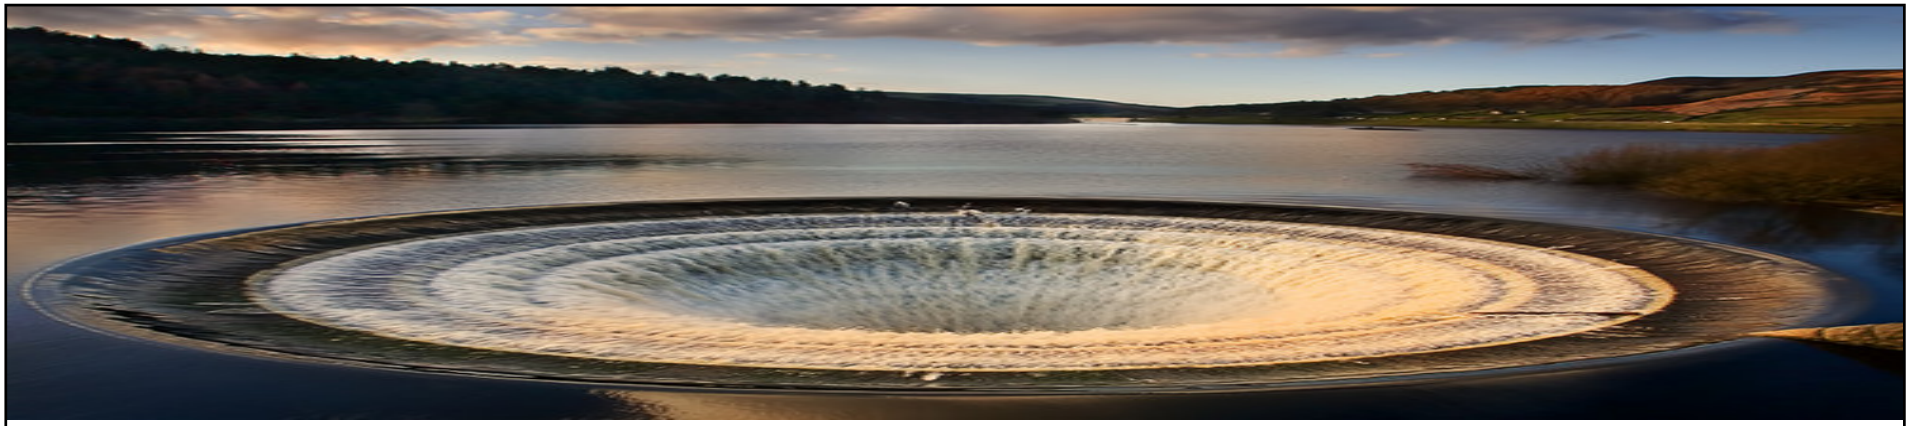

perdisci@cs.uga.edu

University of Georgia
Dept. of Computer Science

Security
Network
Intelligence