On The Power and Limitations of Detecting Network Filtering via Passive Observation

Passive and Active Measurement Conference 2015

March 19, 2015

Matthew Sargent, Jakub Czyz, Mark Allman, Michael Bailey

CASE SCHOOL
OF ENGINEERING
CASE WESTERN RESERVE
UNIVERSITY

Motivation

- Anecdotally, we know edge network policies exist

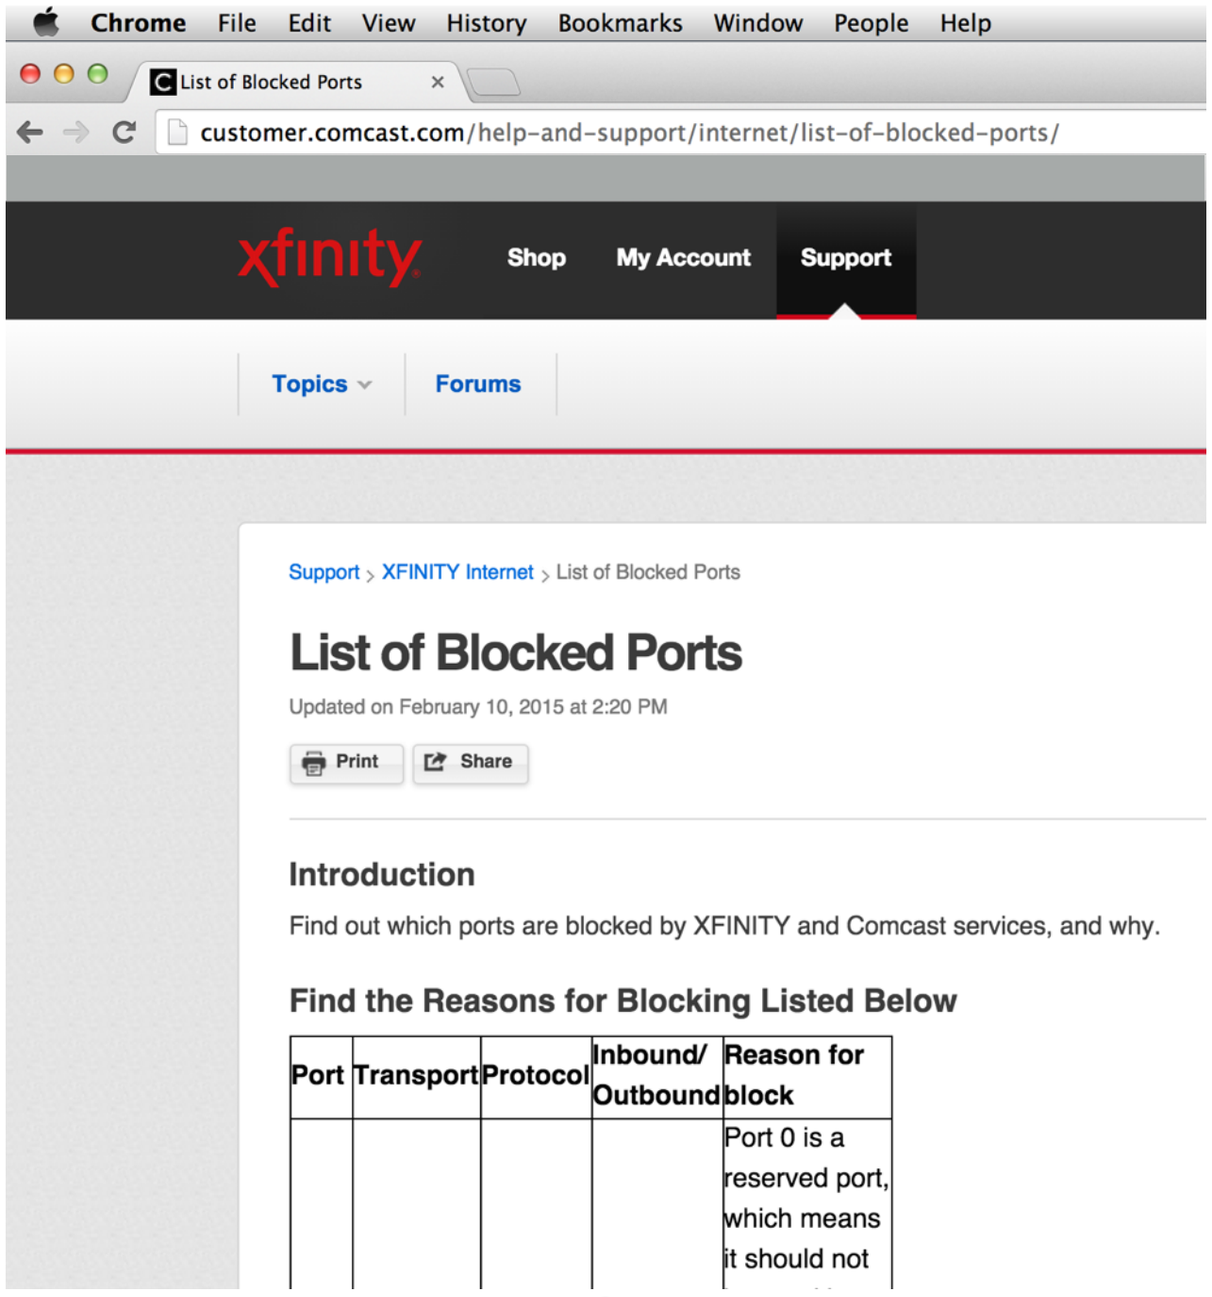

Strategy 1

ISP	Sessions	Country	Blocked (%) Win	SMTP	MSSQL
Comcast	14,765	US	99	8	
RoadRunner	6,321	US			
Verizon	4,341	US	7	21	
SBC	3,363	US	52	74	
Deutsche Telekom	2,694	DE	76		
Cox Cable	2,524	US	93	77	88
Charter Comm.	1,888	US	95	22	36
Qwest	1,502	US	18	6	
BE Un Limited	1,439	UK		49	
BellSouth	1,257	US	59	69	96
Telefonica	1,206	ES		7	
Arcor	1,206	DE	32		
Shaw Cable	1,198	US	5	59	
British Telecom	1,098	UK	10		

Kreibich, Christian, et al. "Netalyzr: illuminating the edge network."
Proceedings of the 10th ACM SIGCOMM conference on Internet measurement. ACM, 2010.

Strategy 2

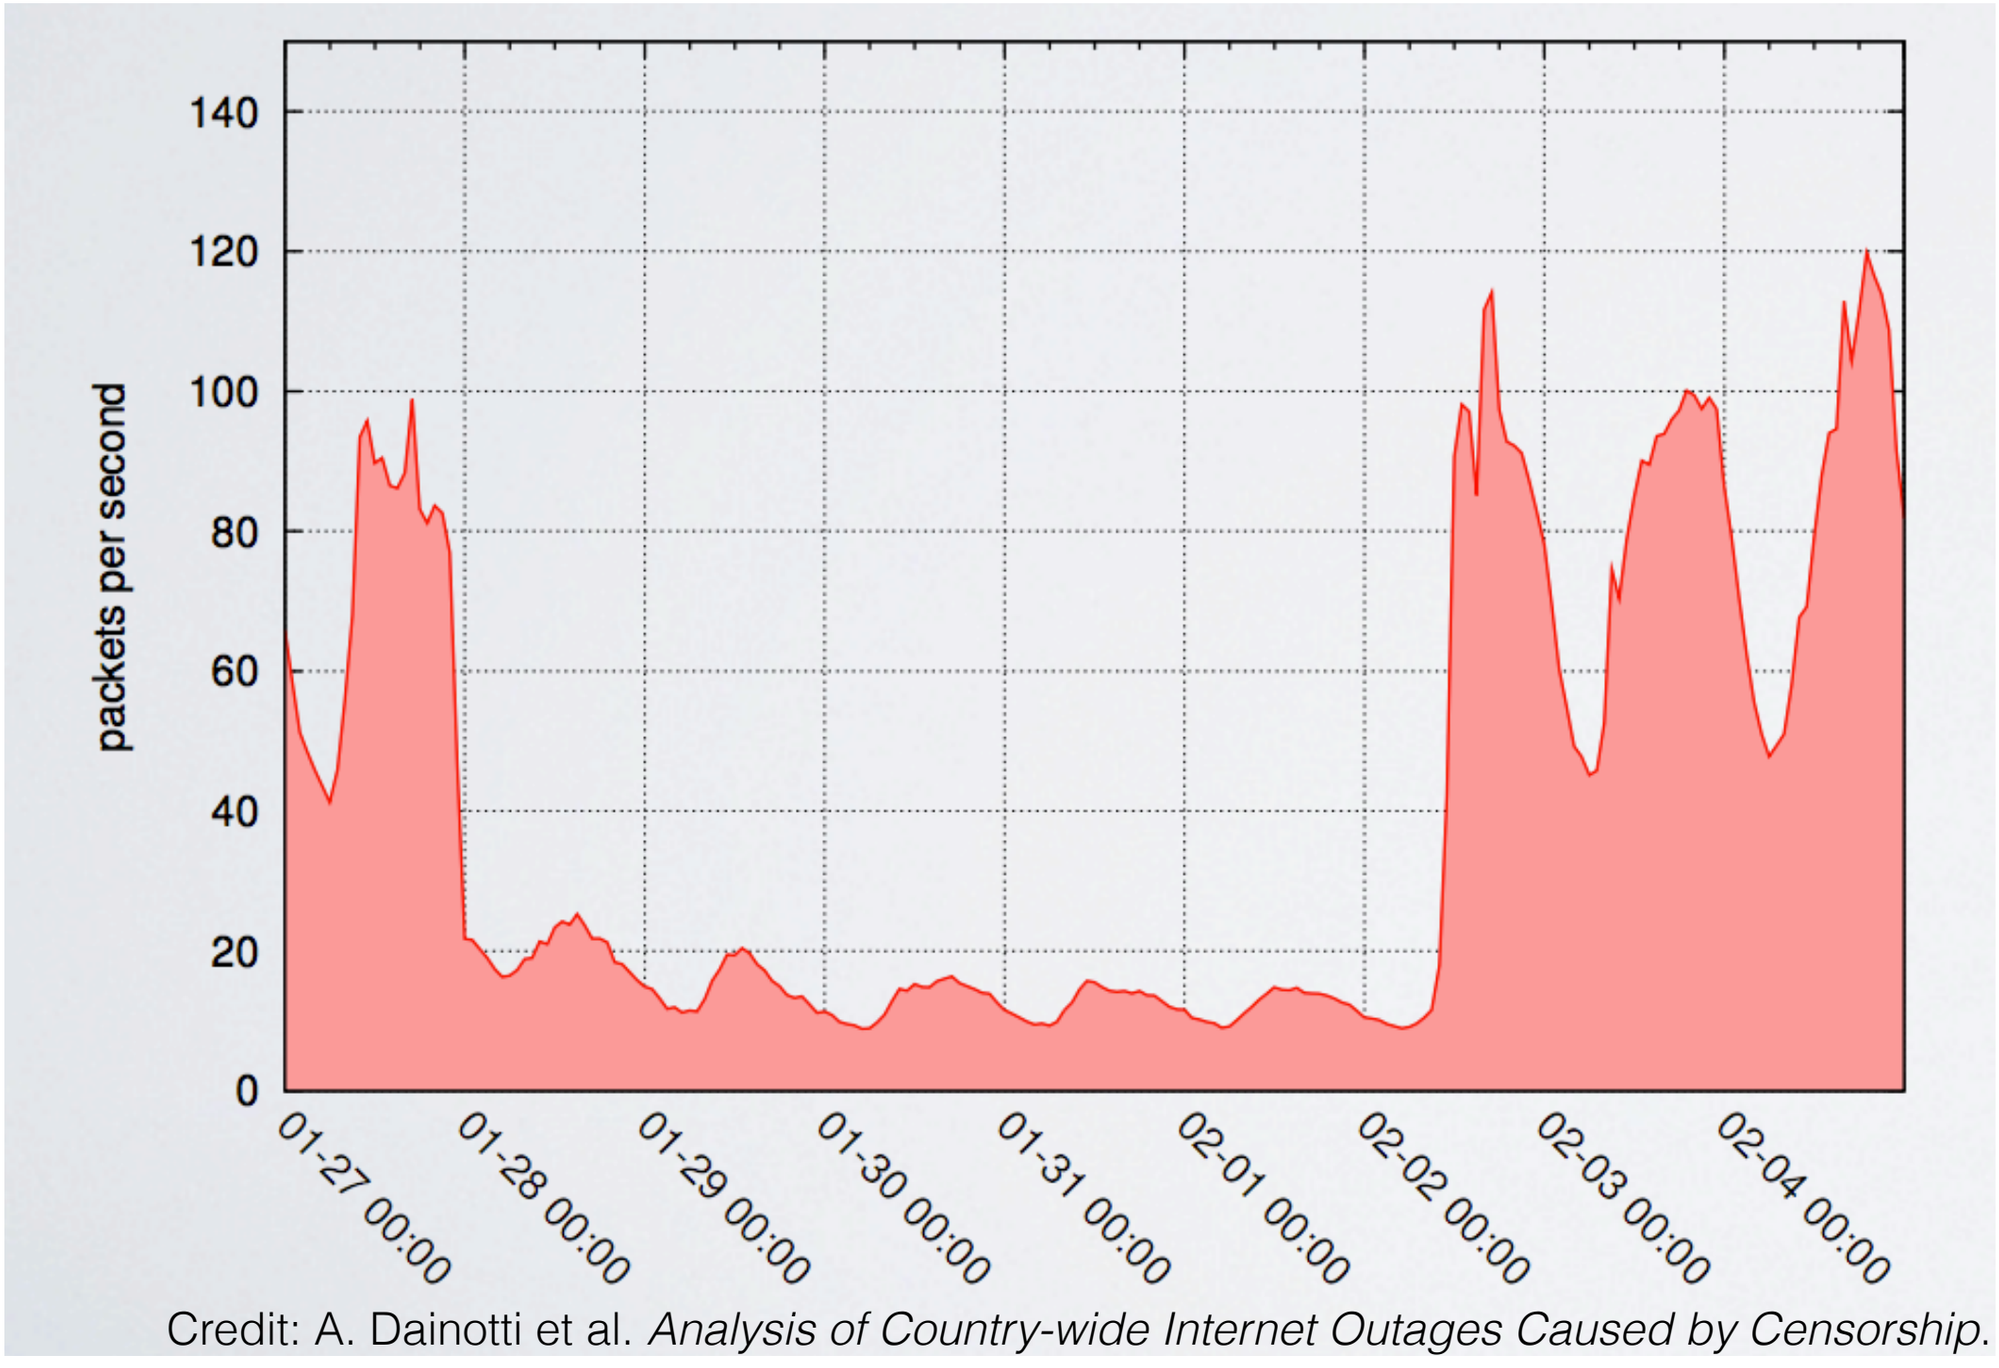

Credit: A. Dainotti et al. *Analysis of Country-wide Internet Outages Caused by Censorship*.

Motivation

- Leveraging darknet space allows us to develop an expectation of seeing certain types of traffic
  - Absence of expected traffic becomes telling

Problem

- We need to see specific types of traffic from many places on the network

- We introduce the concept of *traffic markers*

- We pick an energetic traffic marker, the Conficker worm, as our exemplar

Data

- Two main sources of data:

1. Packet traces from five /8 darknets
   - 2.25% of IPv4 address space
2. Known Conficker host list
   - From Conficker domain sinkhole

Data

Address Block	Packets (billions)	Bytes (trillions)	Rate (Mbps)	Rate (Kpps)	Source /24s (millions)
100/8	22.1	1.7	22.5	36.7	3.1
105/8	17.1	1.1	15.0	28.2	2.1
23/8	16.9	1.8	23.4	28.0	2.6
37/8	21.7	1.5	20.3	35.9	2.4
45/8	18.2	1.3	16.6	30.1	2.3
All	96.1	7.4	97.8	159	4.1

Data Coverage

- Corresponds to 98.8% of IP address space based on routed prefix

- 1.6M of the 4.1M /24s contain Conficker infectees, per the Conficker sinkhole data

Results - /24s

- We first judge /24s where we expect to see Conficker

Expect Conficker?	Observe Conficker?	>=5* known infectees?	Judgement
F	F	-	None
F	T	-	Rare
T	T	-	No Filter
T	F	T	Filtering
T	F	F	None

* Threshold developed in paper

Results - /24s

- Can make judgments on 55% of the /24s containing Conficker infectees
  - 434K do not filter
  - 448K filter
- 747K /24s do not have a enough infectees to form a solid judgment

Results - Routed Prefix

Classification	Amount	Percentage
No Filtering	10,084	13%
Filtering	27,351	35%
Multiple Policies	14,536	18%
Low Signal	22,075	28%
Muddled/No Filtering	5,178	7%

Results - Routed Prefix

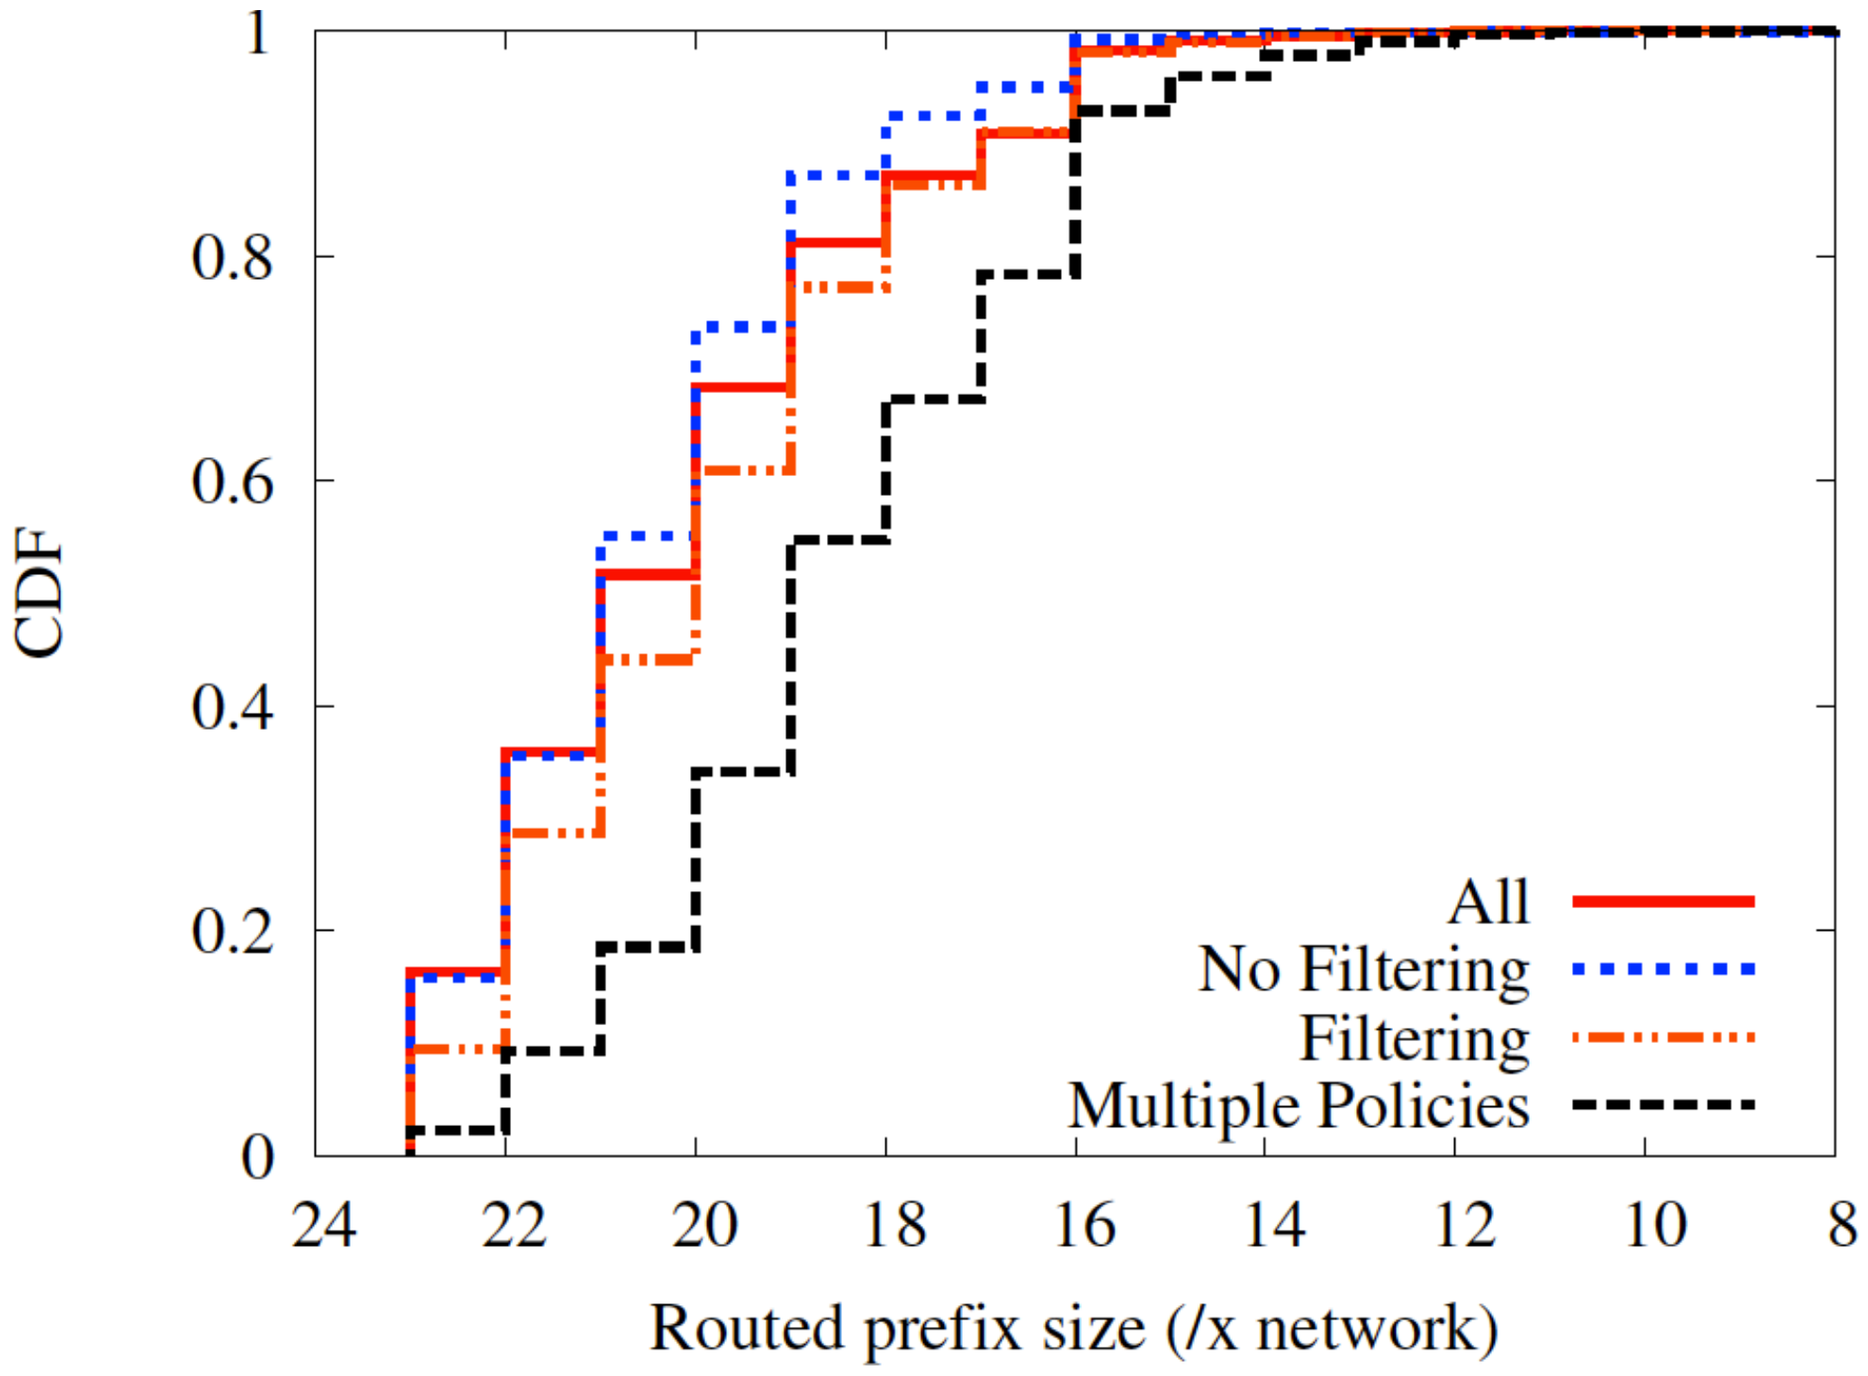

Results - Routed Prefix

- Anecdotal result from Comcast:
  - Can detect and verify TCP/445 filtering on Comcast's /15 network
- Can determine single policy for 699M IP addresses or 28% of the routable addresses.
  - Original Netalyzr study had 130K test runs from 100K IP addresses

Limitations

Darknet	# /24s Receiving SYNs	% /24s w/SYN for TCP/80	TCP/139	TCP/1433	TCP/22
100/8	2.0M	14.2%	1.5%	<1%	<1%
105/8	1.5M	4.0%	1.1%	<1%	<1%
23/8	1.7M	6.2%	1.0%	<1%	<1%
37/8	1.6M	21.6%	1.0%	<1%	<1%
45/8	1.6M	5.6%	1.1%	<1%	<1%
All	3.1M	18.2%	1.3%	<1%	<1%

Conclusions

- Original hypothesis is half-true
  - Traffic markers within darknet data can detect fine-grained network policy
  - Limited to large outbreaks with predictable traffic
- While limited in scope, policy coverage is up to 27x as much as previous work

Questions?

On The Power and Limitations of Detecting Network Filtering via Passive Observation

Passive and Active Measurement Conference 2015

March 19, 2015

Matthew Sargent, Jakub Czyz, Mark Allman, Michael Bailey

CASE SCHOOL
OF ENGINEERING
CASE WESTERN RESERVE
UNIVERSITY