Freenet Darknet Mapping

K.C.N. Halvemaan

University of Amsterdam

System and Network Engineering

Research Project 2 (#86)

July 24, 2017

1. Introduction
2. Research question
3. Freenet
4. Related work
5. Method
   - Experimental setup
   - Traffic detection - step 1: filtering
   - Traffic detection - step 2: comparison to baseline
6. Results
7. Discussion
8. Conclusion
9. Future work
10. References

Introduction

1. Freenet is a distributed semi-structured peer-to-peer file sharing network.
2. First proposed in Clarke [1999], later extended by Clarke et al. [2001] and by Biddle et al. [2002].
3. A censorship resilient membership-concealing overlay network.
4. File sharing, forums, micro blogging, and instant messaging.

Topology

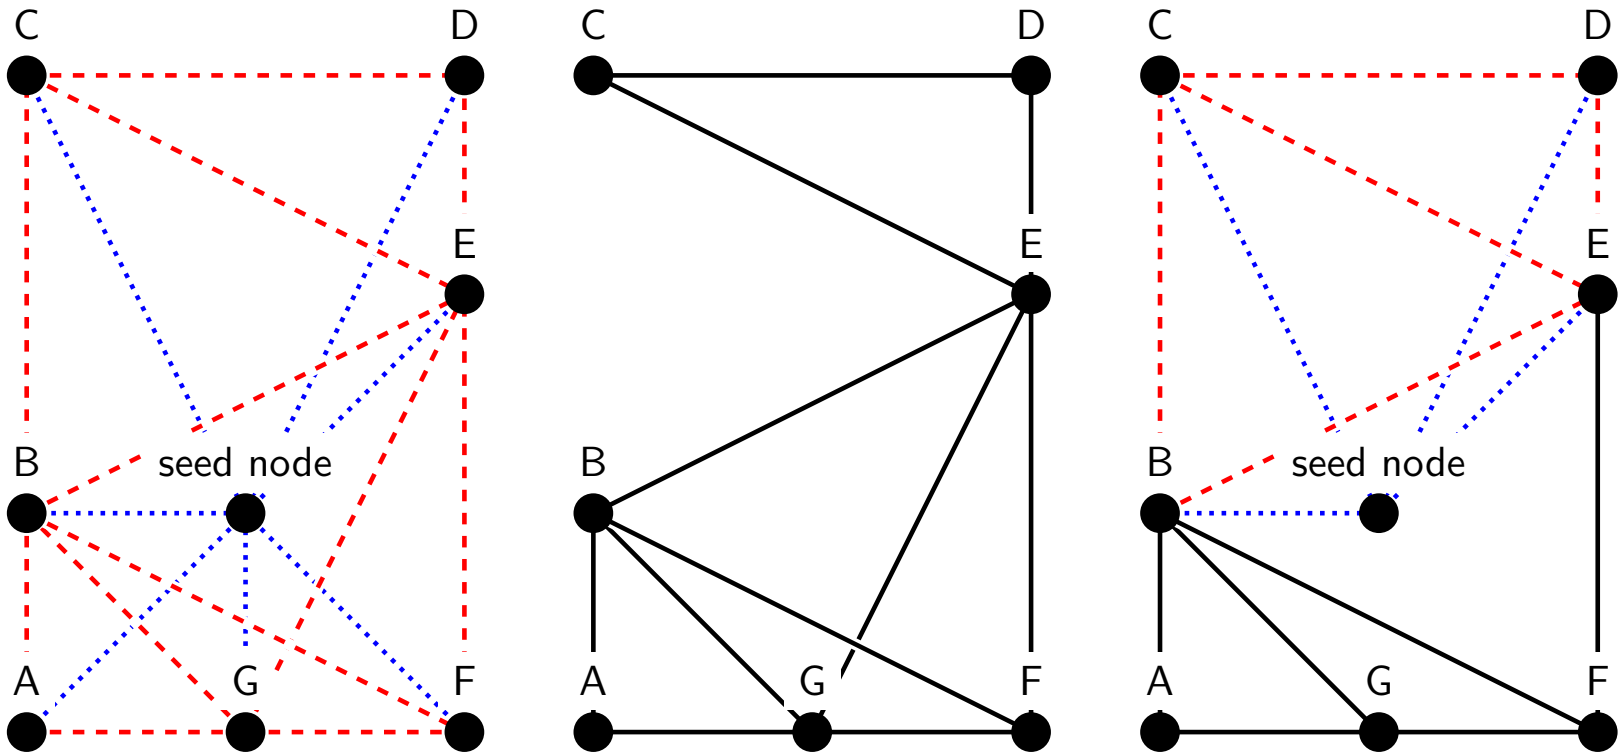

(a) Opennet. (b) Darknet. (c) Hybrid.

Figure: The three possible topologies within Freenet. Solid lines indicate darknet connections, dotted lines are connections to the seed node and dashed lines are connections assigned by a seed node.

Research question

1. Is it possible to discover the IP addresses of nodes participating in a Freenet darknet?

How does Freenet work?

1. Nodes specialising in a part of a distributed hash table.
2. Nodes send messages with a UID to each other via UDP.
3. Routing based on the small-world model by Kleinberg [2000].
4. Files are split into blocks of 32 KiB each.
5. UDP payload is padded to the nearest multiple of 64 with an additional random 0 to 63 bytes.
6. Encrypted with AES in PCFB mode.

Routing

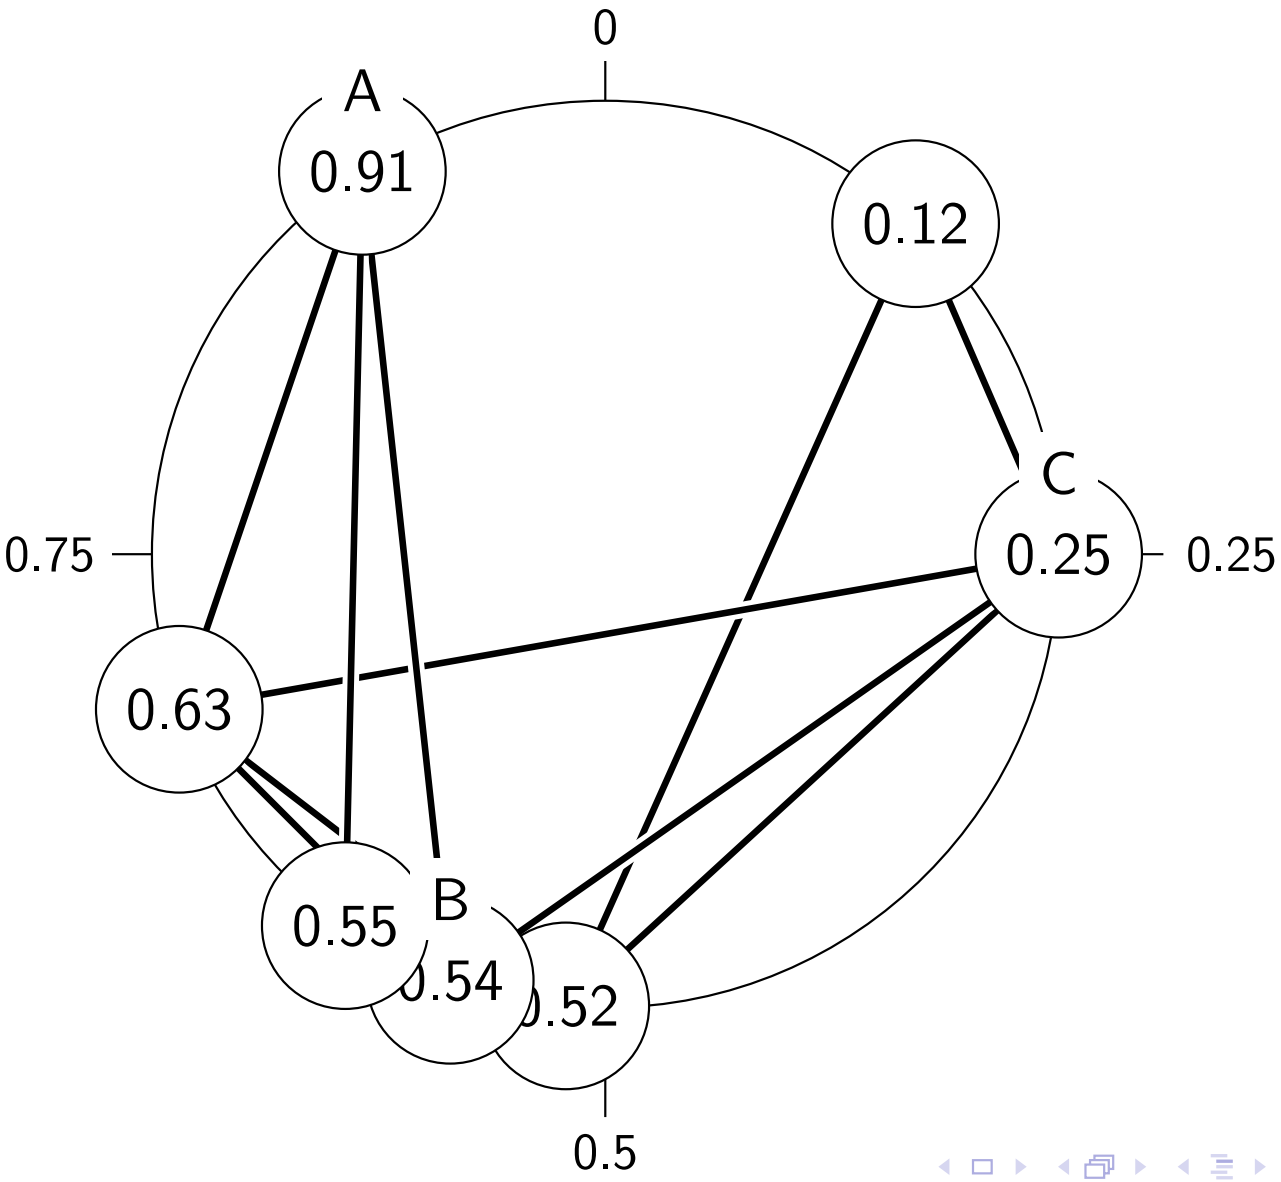

Routing

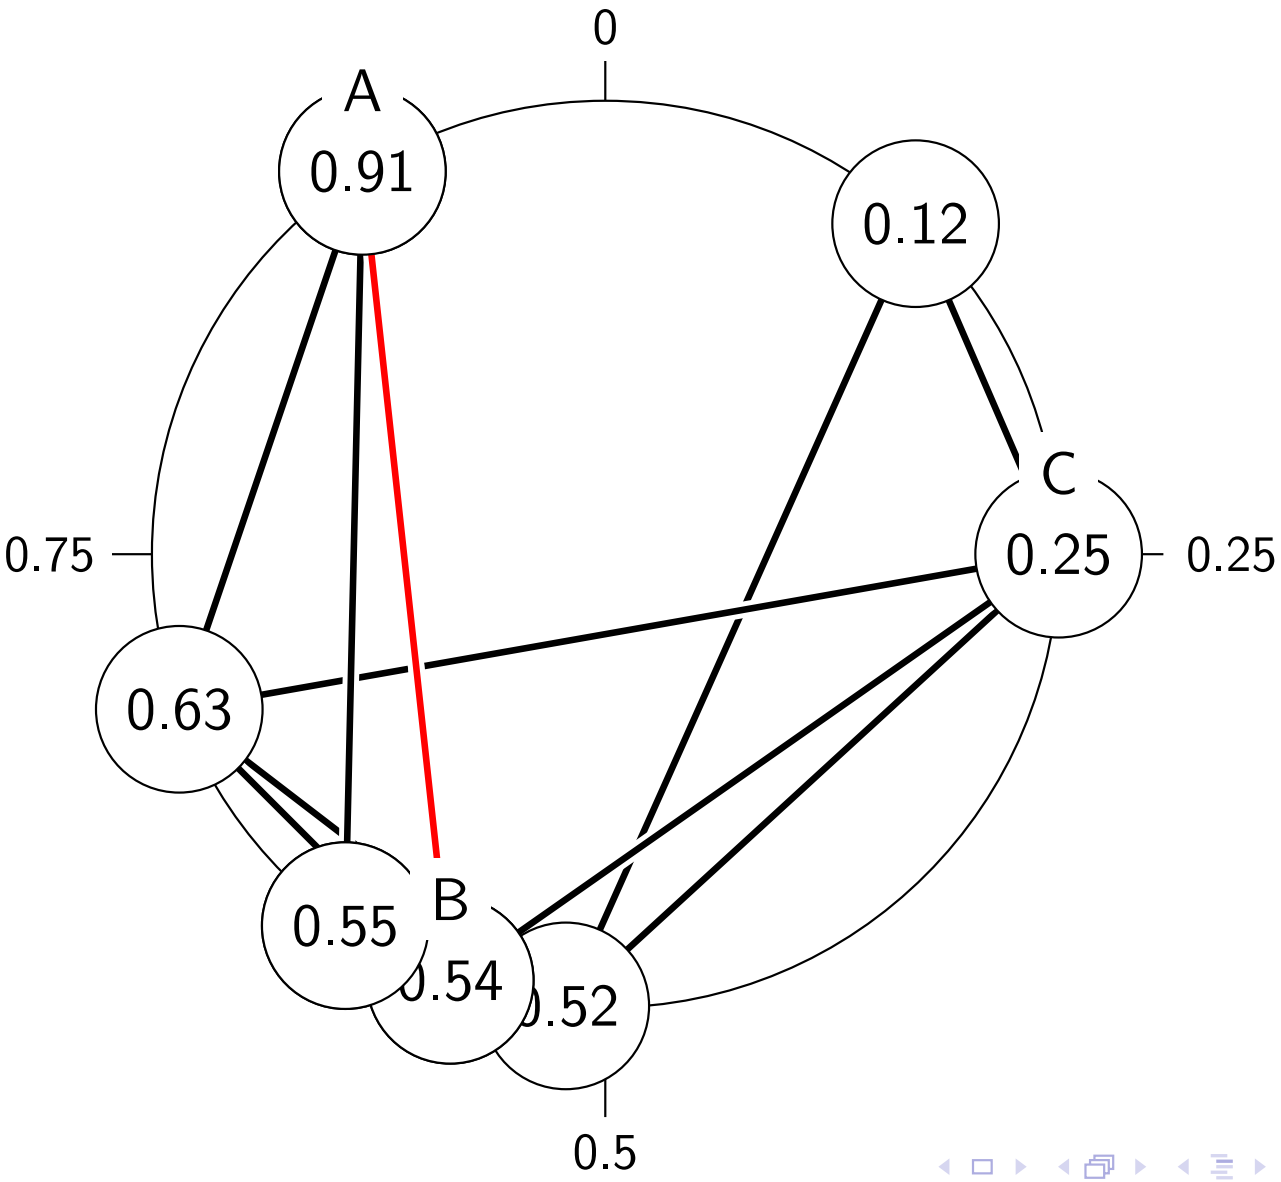

Routing

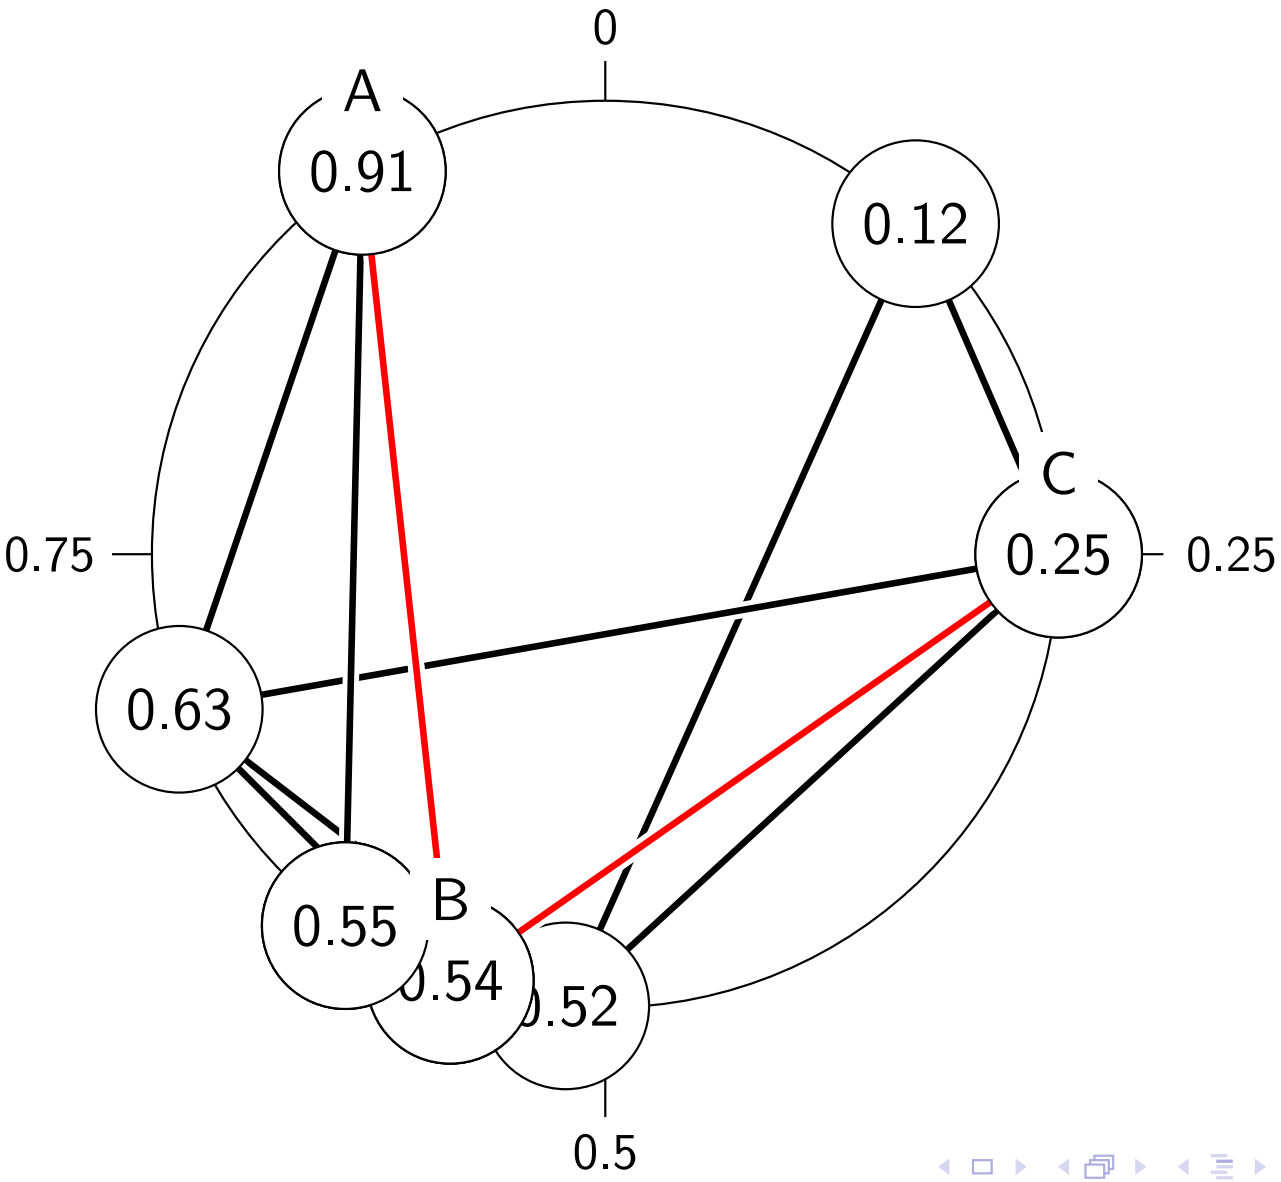

Related work

1. Cramer et al. [2004], Vasserman et al. [2009], and Roos et al. [2014] did monitoring experiments on opennet.
2. DoS "Pitch Black" attack by Evans et al. [2007].
3. Blocking of the FRED by Othman and Kermanian [2008] and the FProxy in Solarwinds.
4. Routing table insertion attack by Baumeister et al. [2012].
5. Message UID traceback attack by Tian et al. [2015] with between 24% and 43% accuracy.

1. Eight Ubuntu 16.04 VMs on a Xen hypervisor, each running a FRED build #1477 (2017-03-09).
2. Physical threat and network threat level to "HIGH".
3. Friend trust level set to "LOW" for all connections.
4. Each node has a degree of at least three.

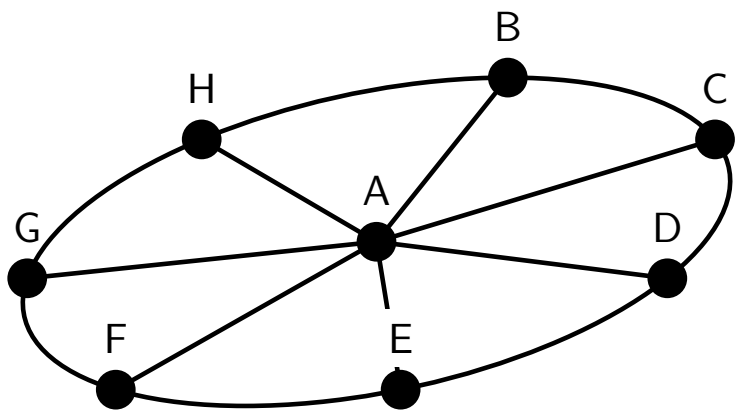

Figure: Topology of the darknet training setup.

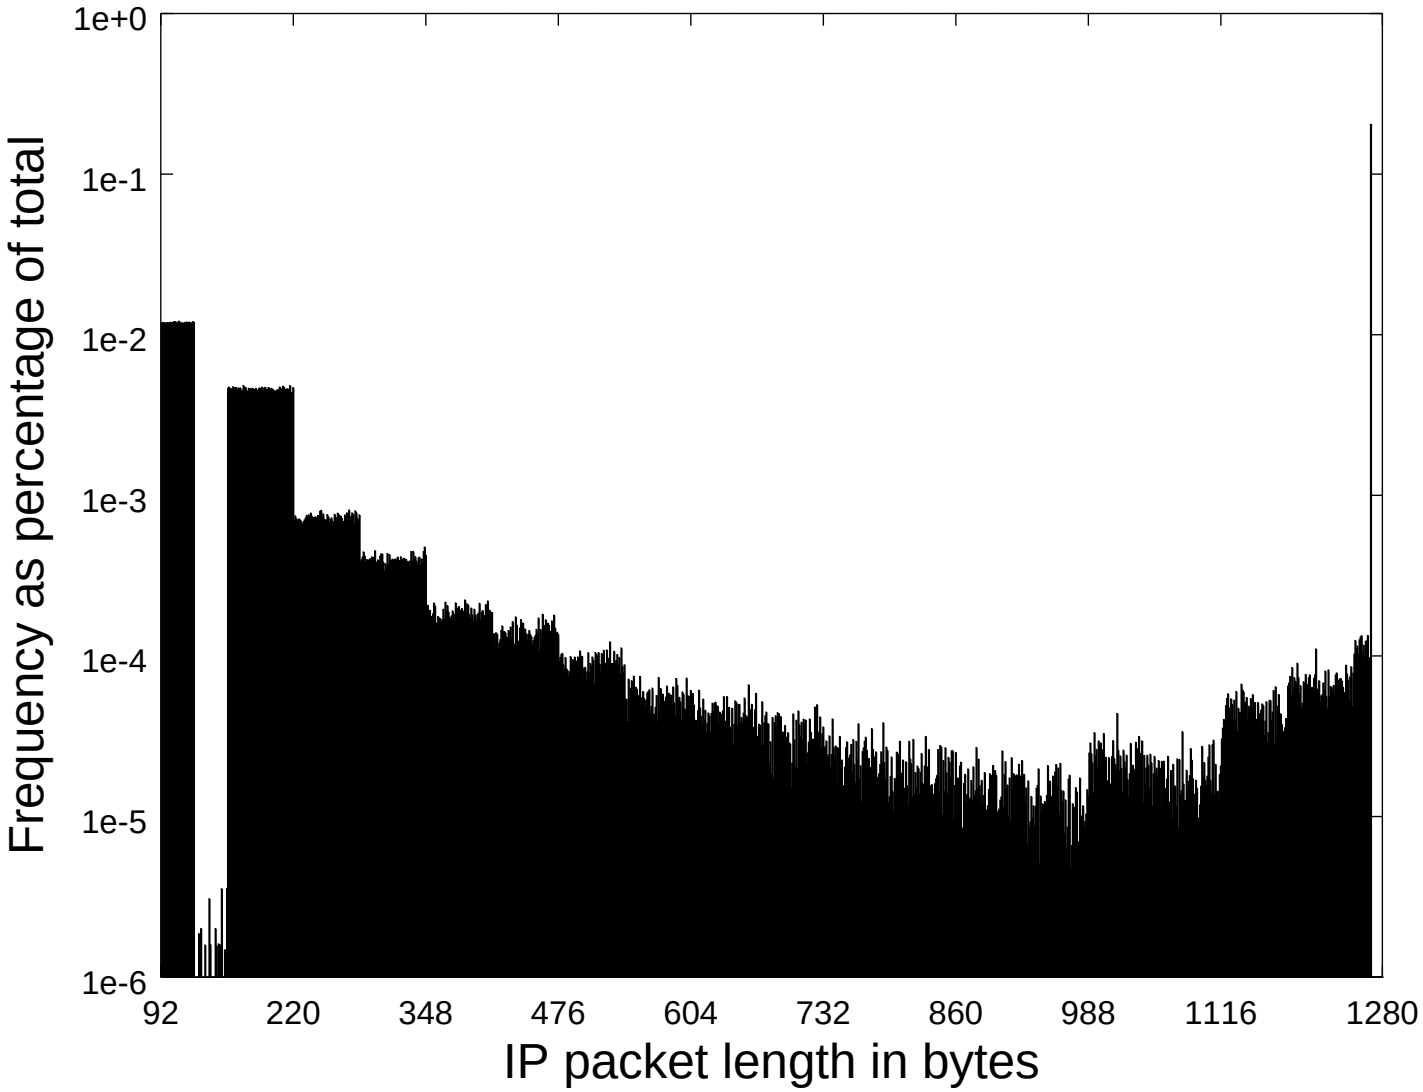
1e+0
1e-1
1e-2
1e-3
1e-4
1e-5
1e-6
Frequency as percentage of total
92
220
348
476
604
732
860
988
1116
1280
IP packet length in bytes

1. Port number between 1024 and 65535.
2. Maximum IP packet length of 1280 bytes.
3. Minimum IP packet length of 92 bytes.
4. Maximum UDP payload of 1232 bytes.
5. Minimum UDP payload of 64 bytes.
6. An IP address receiving packets on the same UDP port from at least three different IP addresses.
7. A socket has to have sent and received at least one packet.

1. A one-class SVM was trained on 5.5 hours of traffic from the test network.
2. As features the normalised packet length frequency per socket were used.
3. Traffic was generated every 10 minutes.
    1. Insert a file with a size between 32 to 320 KiB in each node.
    2. Request the inserted file at a random node.
    3. Request a non-existing file.
4. Check also against some other (P2P) traffic for false positives.

Results - step #1

Table: The number of true positives and false positives in step #1.

Set	**True positives**	**False positives**
darknet 3 hours busy	28 (100%)	0
darknet 3 hours idle	28 (100%)	0
BitTorrent	0	0
OpenArena	0	0
Traceroute	0	0

Results - step #2

Table: The mean score and standard deviation of the 4-fold cross-validation done in step #2.

Set	$\bar{x}$	s
darknet 3 hours busy	43%	17%
darknet 3 hours idle	14%	10%
BitTorrent		
OpenArena		
Traceroute		

Discussion

1. Different accuracy for idle network due to less (re)inserts.
2. Only tested the FRED with default configuration.
3. Small network was tested in a unrealistic setting for a short period of time.

Discussion

1. Different accuracy for idle network due to less (re)inserts.
2. Only tested the FRED with default configuration.
3. Small network was tested in a unrealistic setting for a short period of time.
4. "Making nodes invisible is not easy by any stretch of the imagination and is not something we can or should address before 1.0" [Clarke and Toseland, 2005]
5. The detection method can scale up to ISP or even national level given enough resources.

Conclusion

1. It is possible to identify the IP address of a FRED darknet node based on the network traffic it generates.

Future work

1. Train on a larger and more diverse data set.
2. Apply detection to opennet nodes.
3. Padding payload to a specific size like Tor does.
4. Extract message types based on packet length.
5. Track flow of inserts in the network based on the MTU.
6. Consider implementing the detection method as part of a IDS.

This is the end

Thank you for listening!
Are there any questions?

References I

Todd Baumeister, Yingfei Dong, Zhenhai Duan, and Guanyu Tian. A routing table insertion (rti) attack on freenet. In *Cyber Security (CyberSecurity), 2012 International Conference on*, pages 8–15. IEEE, 2012.

Peter Biddle, Paul England, Marcus Peinado, and Bryan Willman. The darknet and the future of content protection. In *ACM Workshop on Digital Rights Management*, pages 155–176. Springer, 2002.

Ian Clarke. A distributed decentralised information storage and retrieval system. Master's thesis, University of Edinburgh, 1999.

References II

Ian Clarke and Matthew Toseland. Freenethelp.org wiki, 2005. URL `http://www.freenethelp.org/html/AttacksAndWeaknesses.html`. Consulted on 2017-06-21. The page contains an informal discussion on attacks and weaknesses of Freenet. *Toad* is the pseudonym used by Matthew Toseland.

Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Designing Privacy Enhancing Technologies*, pages 46–66. Springer, 2001.

Curt Cramer, Kendy Kutzner, and Thomas Fuhrmann. Bootstrapping locality-aware p2p networks. In *Networks, 2004.(ICON 2004). Proceedings. 12th IEEE International Conference on*, volume 1, pages 357–361. IEEE, 2004.

References III

Nathan S Evans, Chris GauthierDickey, and Christian Grothoff. Routing in the dark: Pitch black. In *Computer Security Applications Conference, 2007. ACSAC 2007. Twenty-Third Annual*, pages 305–314. IEEE, 2007.

Jon Kleinberg. The small-world phenomenon: An algorithmic perspective. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 163–170. ACM, 2000.

Mohamed Othman and Mostafa Nikpour Kermanian. Detecting and preventing peer-to-peer connections by linux iptables. In *Information Technology, 2008. ITSim 2008. International Symposium on*, volume 4, pages 1–6. IEEE, 2008.

References IV

Stefanie Roos, Benjamin Schiller, Stefan Hacker, and Thorsten Strufe. Measuring freenet in the wild: Censorship-resilience under observation. In *International Symposium on Privacy Enhancing Technologies*, pages 263–282. Springer, 2014.

Solarwinds. Solarwinds forum, 2017. URL `https://thwack.solarwinds.com/thread/77015`. Consulted on 2017-06-21.

Guanyu Tian, Zhenhai Duan, Todd Baumeister, and Yingfei Dong. A traceback attack on freenet. *IEEE Transactions on Dependable and Secure Computing*, 2015.

Eugene Vasserman, Rob Jansen, James Tyra, Nicholas Hopper, and Yongdae Kim. Membership-concealing overlay networks. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 390–399. ACM, 2009.