Hardening of P2P networks' stack against the Sybil attack: history, good practices and current state

Thibault Cholez, Claudia Ignat, Victor De Moura Netto

Université de Lorraine, CNRS, Inria, LORIA, Nancy, France

September 30, 2024

Outline

Peer-to-Peer (P2P) networks

Principles:

- Network which links are defined at the application level aka "overlay network"
- Follow its own communication protocol
- Direct service exchange between peers

Quick history:

- Started with Napster (1999), quickly followed by Gnutella (2000), BitTorrent (2002), etc.
- Implement different services (file sharing, blockchains, etc.)
- Prime in 2008/2009 (more than half of Internet traffic)
- P2P network architectures evolved because of dependability and scalability issues

Evolution of P2P network architectures

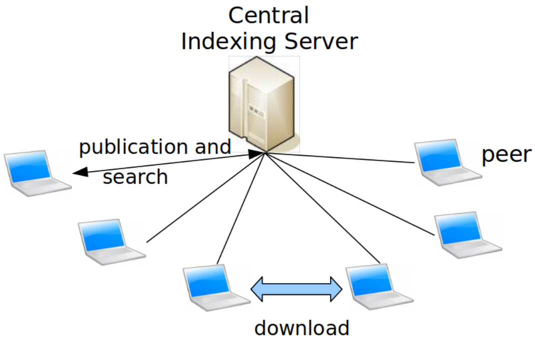

Evolution of P2P network architectures

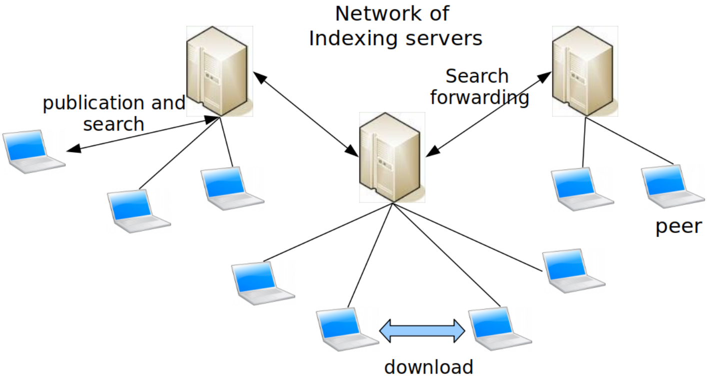

Evolution of P2P network architectures

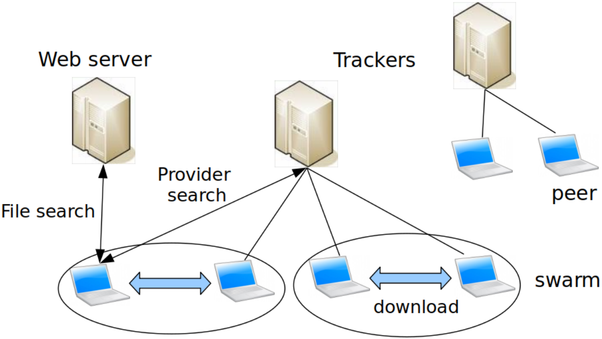

Evolution of P2P network architectures

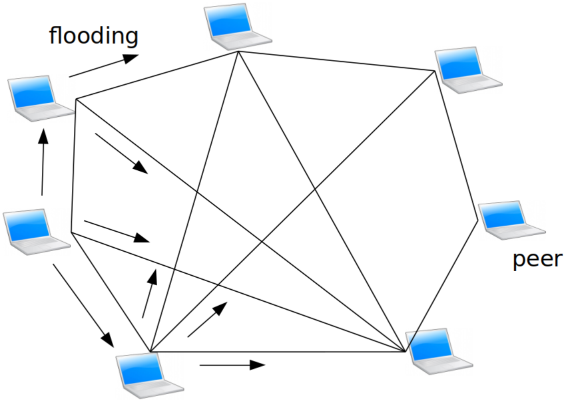

Evolution of P2P network architectures

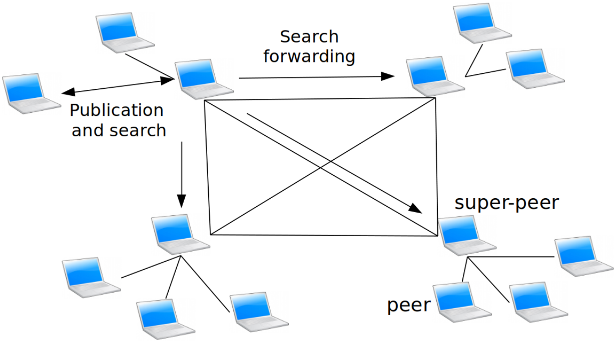

Kademlia Distrbuted Hash Table (DHT) [MM02]

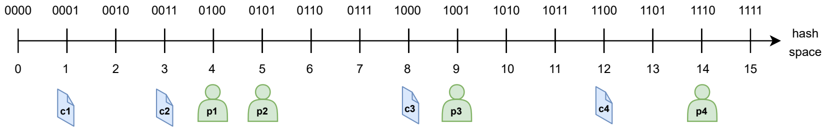

- Distance between IDs is given by a XOR metric
- Peers close to a Key are in charge of it
- What is stored in the DHT? Whatever $< Key, Value >$ pair!
  - PeerID $\rightarrow$ way to contact the peer (IP, port, public key, proxy address, etc.)
  - FileID $\rightarrow$ list of provider peers
  - KeywordID $\rightarrow$ list of corresponding files
- Address space is the size of the hash function output (2^{256})

Kademlia Routing Table Structure

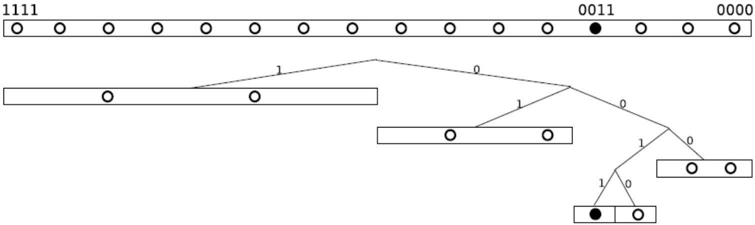

Kademlia Routing Table Structure

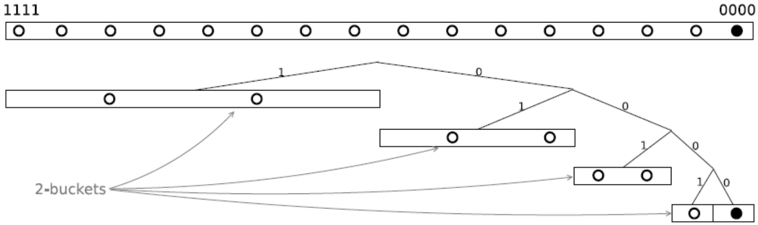

DHT lookup

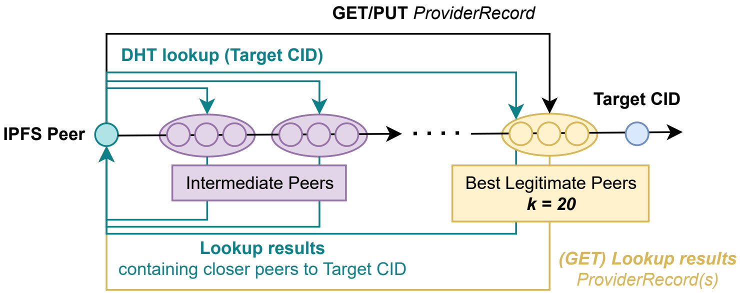

Kademlia Routing Table structure and lookup process ensure a retrieval in $O(\log N)$ jumps (N = Network Size)

Problem statement

DHT Security issues

- Unfortunately DHTs are vulnerable to the Sybil attack
- Sybil attack [Dou02]: one attacker creating many fake identities/peers "Sybils" in the network
- Major threat: very simple to perform, yet very powerful (peer and/or content censorship)

Scope of this talk

- How to perform a Sybil attack? What are the possible consequences?
- How to defend against?
- What is the current state of a recent P2P network, IPFS, regarding the Sybil attack?

Routing Table attacks

Eclipsing a peer

- Attacker fills a peer's routing table with sybils to remove its connections to legitimate peers [CDG+03, SNDW06, WTC+08, PMZ22]
- Disconnect the target to the network
- At a large scale, a well prepared attack can partition the network
- Also works on unstructured P2P networks [MHG18]

Controlling a part of the DHT

- Attacker inserts a massive number of sybils (2^{16}) in peers routing table to take the control over a portion of the DHT ($1/256^{th}$) [SEB07]

Lookup Process attacks

Making a lookup loop indefinitely

- Attacker generates Sybils on the fly when requested during a lookup, each new Sybil being a little closer to the target [KLR09]
- Prevent the lookup to converge on time
- Lookup process reaches a timeout without contacting actual peers holding the data

Controlling a TargetID sourrounded by Sybils

- Objective: place 20 or more Sybils to be the closest to a given Target ID to store all the related data

Lookup Process attacks

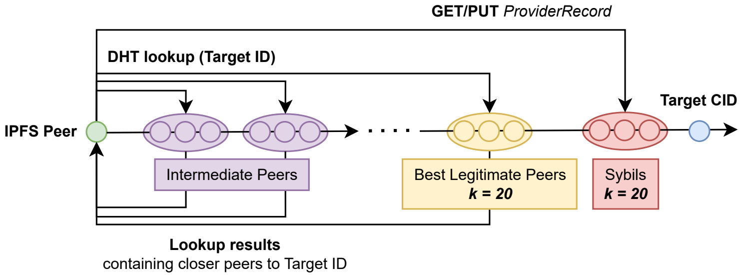

- Monitoring all requests to a targetID [CCF10]
- DoS: attracting all PUT requests but denying GET requests
  → makes a content/peer unreachable [SAK+24]
- Index poisoning [LMSW10]: Sybils return fake results

Rules protecting the routing table [CCF09]

Always check peers' reachability

- Perform an application level three-way handshake before trusting a peer to protect against IP spoofing
- Unresponsive Sybils are discarded
- Prevent the P2P network to send traffic to a DDoS target
  - Blacklisting common ports (53, 80, etc.) also helps

Rules protecting the routing table

Limit the rate of routing table update

- Limit the rate of unsolicited updates to X/min
- Define a timeout to remove oldest contacts
- Prevent an attacker to flood a routing table and to stay

Enforce IP address diversity

- Allow a single peer per IPv4 subnet (/16) to be inserted in a bucket, and X peers per /16 subnet globally
- Attacker must distribute the attack at the network level (botnet)

Rules protecting the lookup process

Enforce IP address diversity

- Allow a single peer per IPv4 subnet (/16) to be considered during a given lookup
- Attacker must distribute the attack at the network level (botnet)

Perform parallel and decorelated lookups

- S/Kademlia proposal [MB07]: run 3 independent parallel lookups (never stepping on a same peer) and not sharing found contacts
- Prevent the attack to succeed as soon as a Sybil is on the path

Rules protecting the lookup process

Check statistical distribution of PeerIDs [CCFD12]

- PeerIDs' distribution should be uniform on the ID space (output of a hash function)
- **CPL** = Common Prefix length between IDs
- Distribution of the CPLs of Peers returned by a lookup depends on the network size

Two steps process

1. Init: estimate current PeerID's distribution with lookups to random IDs
2. For each lookup: Compare the distribution around an ID with the empirical distribution to detect attacks (Sybils insertion create a bias)

Sybil attack detection through PeerIDs distribution

How to compare?

- Challenge: small sample (10 to 20 peers according to the replication factor)
- Most statistical test do not work
- KL-divergence is efficient but needs proper threshold to balance false positives and false negatives (defined empirically)

Kullback-Leibler divergence (G-test):

$$D_{KL}(M \mid T) = \sum_i M(i) \log \frac{M(i)}{T(i)} \tag{1}$$

attack is detected if KL-distance $>$ threshold

Example of PeerIDs distribution after a lookup on IPFS

$NetSize = 13239$				
CPL	**Probability**	**Nodes**	**Nodes (learned)**	**Nodes (attack)**
	...		...	
8	1.3%	0.3	0.3	0
9	34.3%	6.8	6.8	0
10	32.1%	6.4	6.4	0
11	16.2%	3.2	3.2	0
12	8.1%	1.6	1.6	0
13	4.0%	0.8	0.8	20
	...		...	
	$\pm$ 100%	$\pm 20 = k$	$\pm 20 = k$	$20 = k$

Region-based Mitigation – Sridhar et *al* [SAK+24]

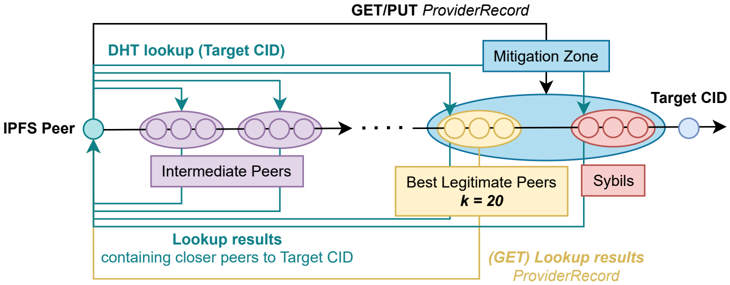

- Send stored value to every peer in a region of ID space defined to contain at least 20 legitimate peers
- During a search, legitimate peers can return the true value
- Alternative countermeasure: discard peers on the most suspicious CPL

InterPlanetary File System (IPFS) [Ben14]

Why is it interesting?

- Modern iteration of P2P system based on Kademlia
- Also implements a second unstructured overlay
- Active community (Protocol Labs), open source
- Main purpose: storage platform for decentralized apps
- P2P network stack became an autonomous project as libp2p [com23]
- Base for other projects: HIVE, DTube, etc.

Publishing/fetching content in IPFS

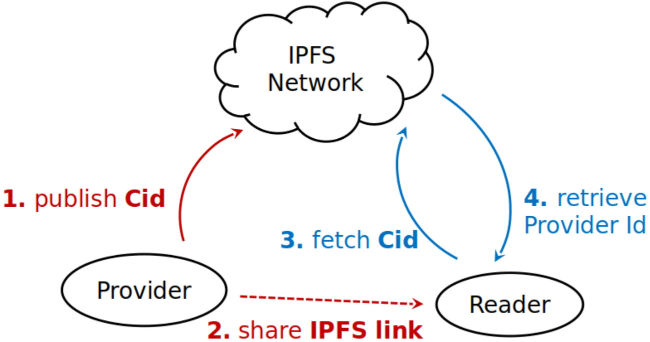

- Providers publish a Document identified by a Content Identifer (Cid) based on the content hash and shared out of band
- A reader interested in a Cid will be directed to the Provider that stores the file identified by the Cid

IPFS Document structure

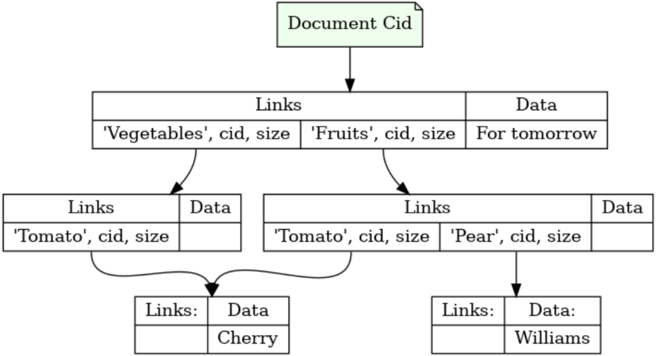

Kademlia DHT for peer and content discovery

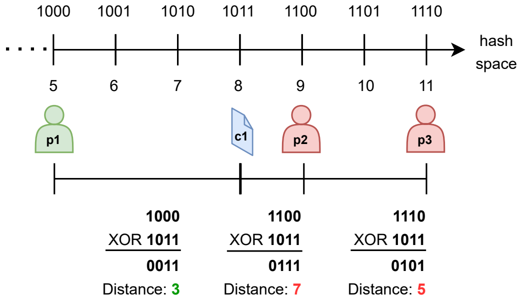

- Peers identified by a PeerID (hash of the public key)
- Distance between identifiers computed by XOR
- Records published on the DHT
    - Provider Record: (PeerID, Cid)
    - Peer Record: (PeerID, Multiaddress), i.e. information to connect to a peer (@IP, port)

Sybil Attack Design

Sybil ID generation

Challenge: PeerIDs are constrained (hash of a cryptographic key), so an attacker must first pre-compute Sybils' PeerID

- IPFS network monitoring with 200 probes during 3 days
- Counted 6,800 PeerIDs and 3,500,000 Cids
- Estimated empirically that placing Sybils at a maximum distance of 2^{230} to a TargetID is close enough to get control of 99.95% of Cids
- Took 1h30 on a 8 cores desktop computer to brute force the 20 Sybil's PeerID
- All generated PeerIDs can be saved for other attacks

Implementation and experimental setup

Implementation of Sybils

- Sybil client is a sightly modified IPFS Kubo client
- Behaves normally except for the target Cid
- Sybils advertise each other during the lookup process

Experiment

- Generate a random "target" file and share it in IPFS with a regular client
- Start Sybils and let them 15 minutes to be connected
- Try to retrieve the file with another regular client

Evaluation

- Attack success is the inability to retrieve the targeted file
- Upon attack failure we investigate how many records were captured by Sybils out of 20

Kubo vers.	Nb sybils	Nb IP@	Nb attack success	Nb Records intercepted in case of failure
19.2	27	27	9/11	19 and 19/20
20	27	27	10/12	17 and 19/20
20	20	1	11/11	-
20	20	1	12/12	-

- Attack is very effective overall
- IP-level distribution is not enforced. Running all Sybils on a single computer achieves 100% attack success
- Still work on latest versions (0.29), but not with the Region-based Mitigation from Sridhar et *al* [SAK+24]

Conclusion

Take away

- Sybil attack has always been a major threat to opened P2P systems based on a DHT
- Basic rules can make the life of the attacker harder
- IPFS did not learn from the past...
- Despite "sota" defense mechanisms, optimized Sybil attacks can still prevent content access in 2/3 attempts

Future work

- Collaboration with HIVE[1] and Inria Alvearium
- Didactic survey of P2P security mechanisms
- Improve defenses against active attacker scenario in IPFS

[1] https://www.hivenet.com/

Thank you for your attention.

Any questions?

- Juan Benet.
  IPFS - content addressed, versioned, P2P file system.
  CoRR, abs/1407.3561:11, 2014.

- Thibault Cholez, Isabelle Chrisment, and Olivier Festor.
  Evaluation of Sybil Attacks Protection Schemes in KAD.
  In *3rd International Conference on Autonomous Infrastructure, Management and Security - AIMS 2009*, volume 5637 of *LNCS*, pages 70–82, Enschede, Netherlands, June 2009. Springer.

- Thibault Cholez, Isabelle Chrisment, and Olivier Festor.
  Monitoring and Controlling Content Access in KAD.
  In *International Conference on Communications - ICC 2010*, pages 1–6, Capetown, South Africa, May 2010. IEEE.

- Thibault Cholez, Isabelle Chrisment, Olivier Festor, and Guillaume Doyen.

Detection and mitigation of localized attacks in a widely deployed P2P network.
Peer-to-Peer Networking and Applications, 6(2):155–174, May 2012.

Miguel Castro, Peter Druschel, Ayalvadi Ganesh, Antony Rowstron, and Dan S. Wallach.
Secure routing for structured peer-to-peer overlay networks.
SIGOPS Oper. Syst. Rev., 36(SI):299–314, December 2003.

IPFS community.
Libp2p IPFS Docs.
`https://docs.ipfs.tech/concepts/libp2p/`, 2023.
Accessed: 20-09-2023.

John R. Douceur.
The sybil attack.

In *Revised Papers from the First International Workshop on Peer-to-Peer Systems*, IPTPS '01, page 251–260, Berlin, Heidelberg, 2002. Springer-Verlag.

Michael Kohnen, Mike Leske, and Erwin P. Rathgeb.
Conducting and optimizing eclipse attacks in the kad peer-to-peer network.
In *NETWORKING 2009*, LNCS, pages 104–116, Berlin, Heidelberg, 2009. Springer.

Thomas Locher, David Mysicka, Stefan Schmid, and Roger Wattenhofer.
Poisoning the kad network.
In *Distributed Computing and Networking*, pages 195–206, Berlin, Heidelberg, 2010. Springer.

Sebastian Mies and Ingmar Baumgart.
S/kademlia: A practicable approach towards secure key-based routing.

In *Parallel and Distributed Systems, International Conference on*, volume 2, pages 1–8, Los Alamitos, CA, USA, dec 2007. IEEE Computer Society.

Yuval Marcus, Ethan Heilman, and Sharon Goldberg.
Low-resource eclipse attacks on ethereum's peer-to-peer network.
Cryptology ePrint Archive, Paper 2018/236, 2018.

Petar Maymounkov and David Mazières.
Kademlia: A peer-to-peer information system based on the xor metric.
In *Revised Papers from the First International Workshop on Peer-to-Peer Systems*, IPTPS '01, page 53–65, Berlin, Heidelberg, 2002. Springer-Verlag.

Bernd Prünster, Alexander Marsalek, and Thomas Zefferer.
Total eclipse of the heart – disrupting the InterPlanetary file system.

In *31st USENIX Security Symposium (USENIX Security 22)*, pages 3735–3752, Boston, MA, August 2022. USENIX Association.

Srivatsan Sridhar, Onur Ascigil, Navin Keizer, François Genon, Sébastien Pierre, Yiannis Psaras, Etienne Rivière, and Michał Król.
Content censorship in the interplanetary file system.
In *31st Annual Network and Distributed System Security Symposium, NDSS 2024*, pages 1–17. The Internet Society, 2024.

Moritz Steiner, Taoufik En-Najjary, and Ernst W. Biersack.
Exploiting KAD: Possible uses and misuses.
SIGCOMM Comput. Commun. Rev., 37(5):65–70, October 2007.

Atul Singh, Tsuen-Wan Ngan, Peter Druschel, and Dan S. Wallach.
Eclipse attacks on overlay networks: Threats and defenses.

In *Proceedings of IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*, pages 1–12. IEEE, April 2006.

Peng Wang, James Tyra, Eric Chan-Tin, Tyson Malchow, Denis Foo Kune, Nicholas Hopper, and Yongdae Kim.
Attacking the kad network.
In *Proceedings of the 4th International Conference on Security and Privacy in Communication Netowrks*, SecureComm '08, pages 1–10, New York, NY, USA, September 2008. ACM.