Attacking the Kad network-real world evaluation and high fidelity simulation using DVN

Peng Wang, James Tyra, Eric Chan-Tin, Tyson Malchow, Denis Foo Kune, Nicholas Hopper and Yongdae Kim

Wiley Security and Communication Networks 2008

KAIST

SysSec Lab

Minjung Kim

Introduction

P2P Systems

- How to find the desired information?
  - Centralized structured: Napster
  - Decentralized unstructured: Gnutella

P2P Systems

- How to find the desired information?
  - Centralized structured: Napster
  - Decentralized unstructured: Gnutella
  - Decentralized structured: Distributed Hash Table

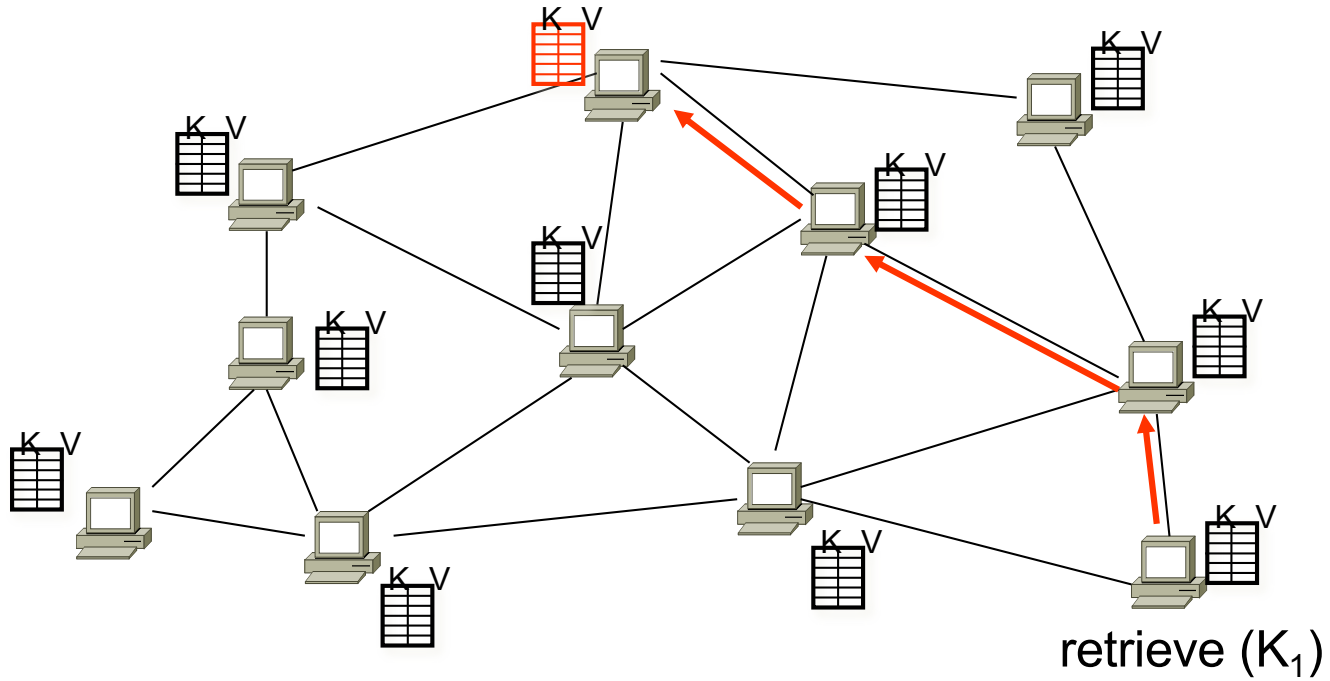

DHT: Terminologies

- Every node has a unique ID: *nodeID*
- Every object has a unique ID: *key*
- Keys and nodeIDs are logically arranged on a *ring* (*ID space*)
- A data object is stored at its *root(key)* and several *replica roots*
  - Closest nodeID to the key (or successor of k)

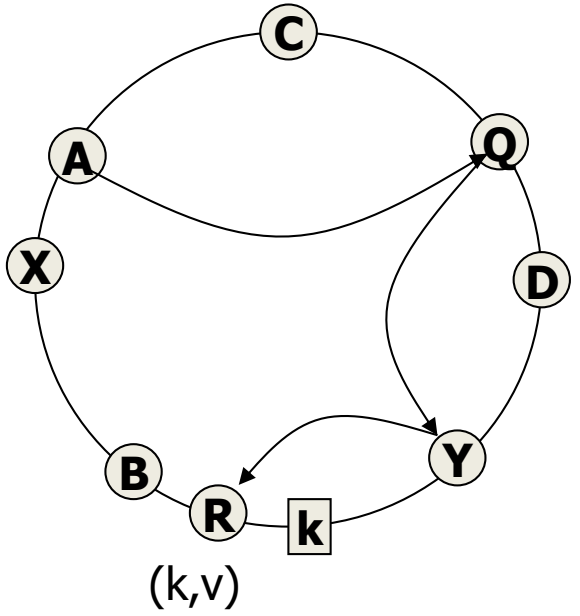

- *Range:* the set of keys that a node is responsible for
- Routing table size: O(log(N))
- Routing delay: O(log(N)) hops

Main Questions?

❖ Any P2P system is used for finding desired information

❖ Questions

- Routing attacks on DHT? What does it mean?
- Is the most popular DHT secure against routing attacks?
- What are the resources?
- How efficient is the attack?
- How to fix it?

Background

Target P2P System

- Kad
  - A peer-to-peer DHT based on Kademlia
- Kad Network
  - BitTorrent
  - Overlay built using eD2K series clients
    - eMule, aMule, MLDonkey
    - Over 1 million nodes, many more firewalled users
  - BT series clients
    - Overlay on Azureus
    - Overlay on Mainline and BitComet

Kademlia Protocol

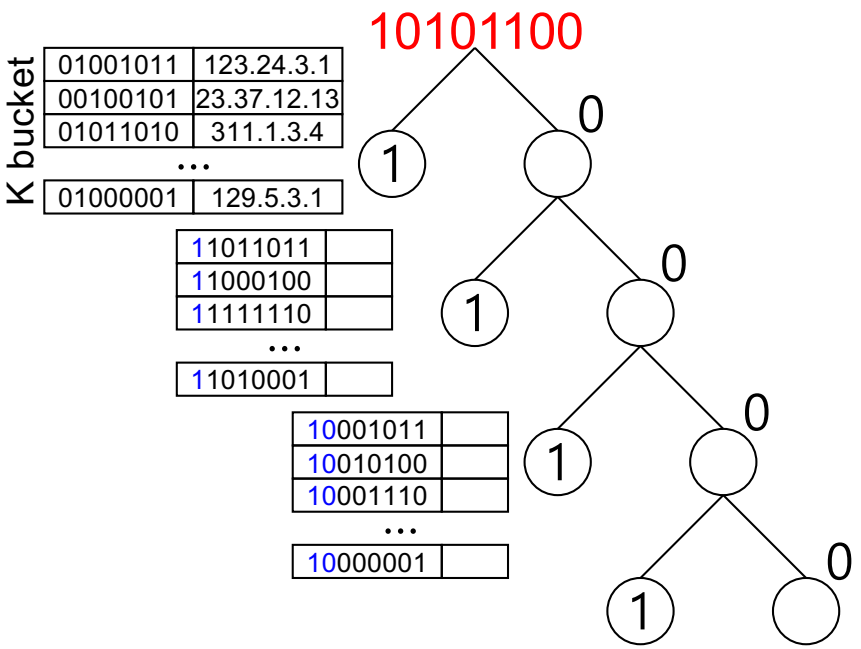

- An entry in k-bucket in i-th level shares at least i-bit prefix with the nodeID
- d(X, Y) = X XOR Y

  ex) 10101100 XOR 11001011

  = 01100111
- Add new contact if
  - k-bucket is not full

Kademlia Protocol

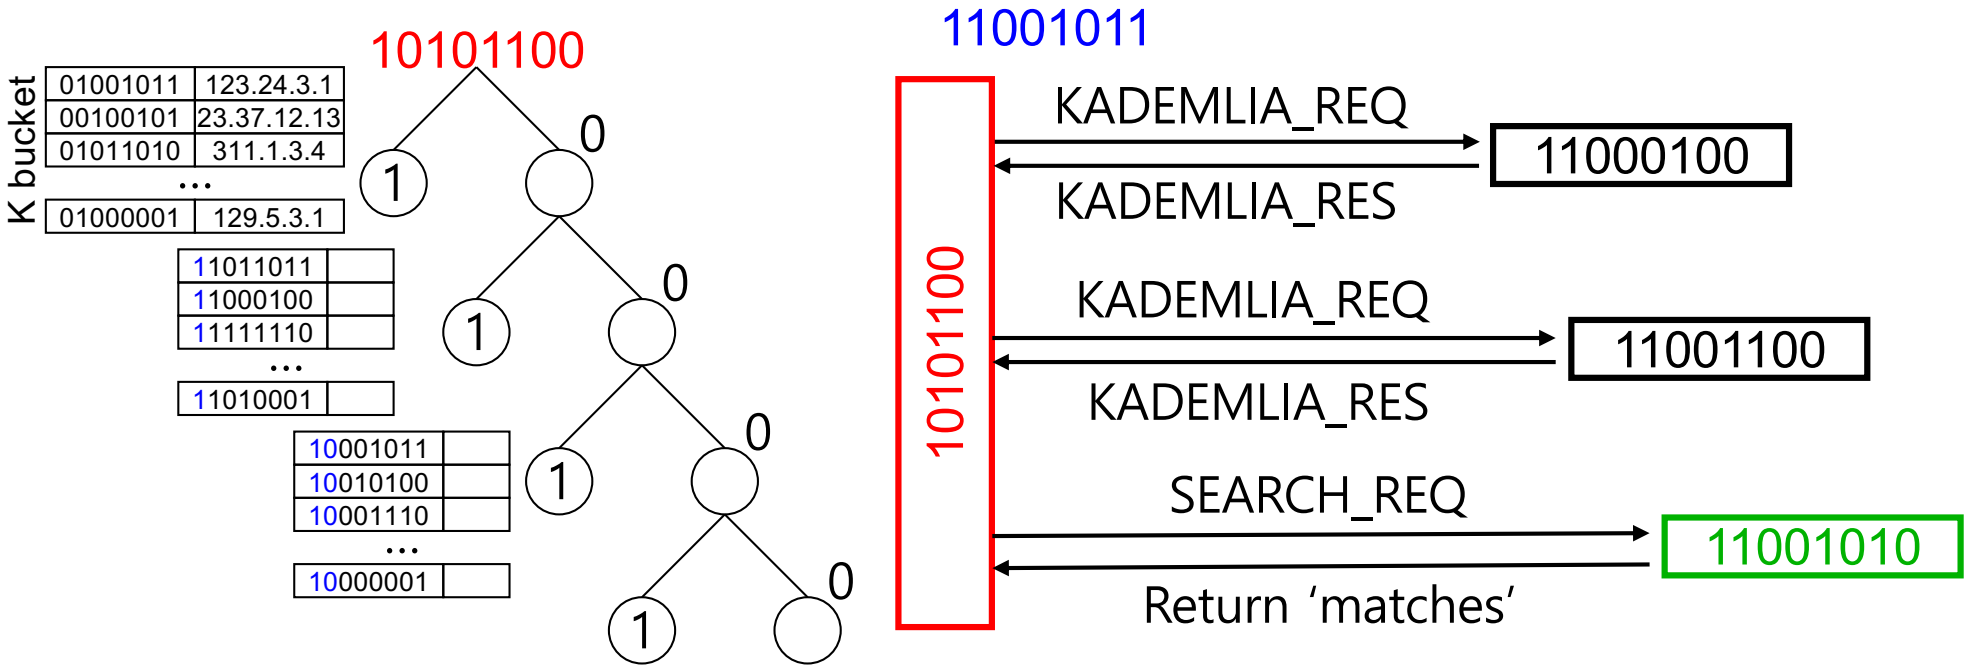

- Iterative, parallel, prefix-matching routing
- Replica roots: k closest nodes

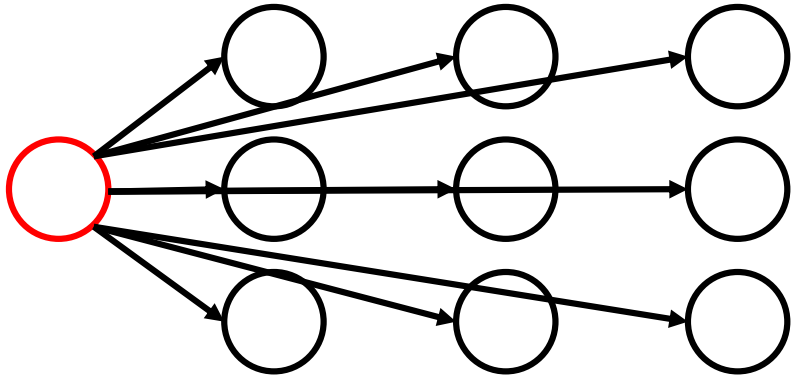

Kad Protocol

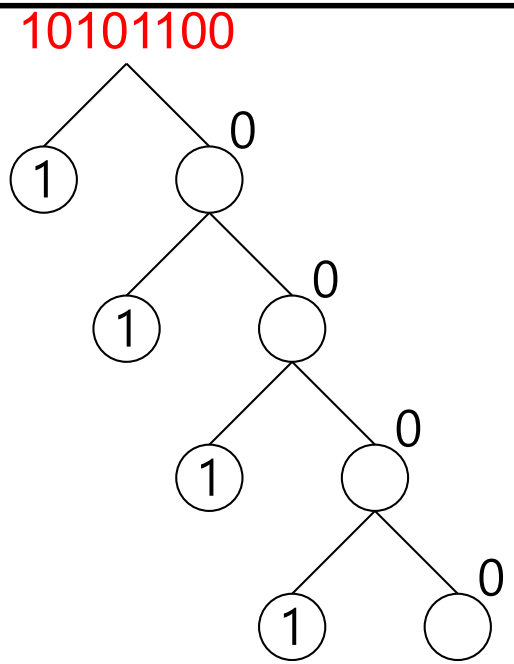

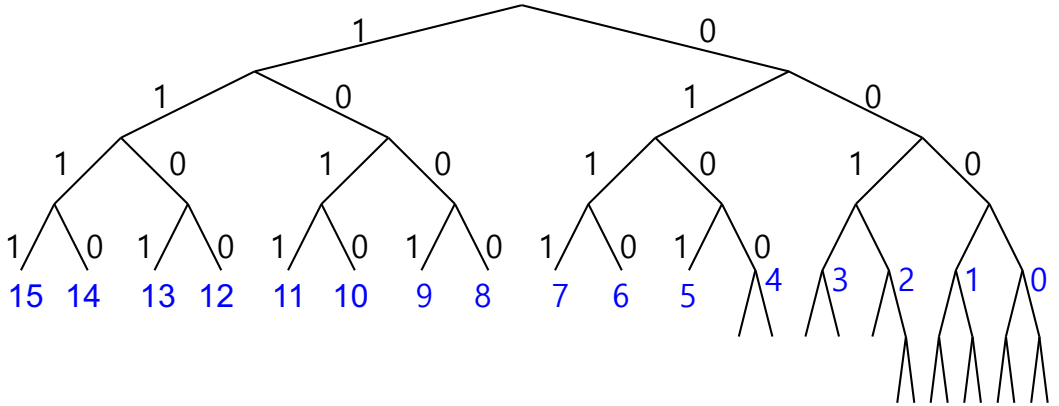

- Wide routing table ➔ short routing path
- No restriction on nodeID
- Replica root: $|r, k| < \delta$

- K bucket in i-th level covers $1/2^i$ ID space
- K buckets with index [0,4] can be split if new contact is added to full bucket

Vulnerabilities of Kad

- No admission control, no verifiable binding
  - An attacker can launch a Sybil attack by generating an arbitrary number of IDs
- Eclipse Attack
  - Stay long enough: Kad prefers long-lived contact
  - (ID, IP) update: Kad client will update IP for a given ID without any verification
- Termination
  - Query terminates when A receives 300 matches.
- Timeout
  - When M returns many contacts close to K, A contacts only those nodes and timeouts.

Attacking the Kad Network

Attack Model

- Attack goal

  Degrade the service of the Kad network, by causing a significant fraction of all keyword as well as node searches to fail.

- Attacker

  - Attacker controls only end-system
  - Does not require corruption or misrouting of IP-layer packet between honest nodes
  - Attacker's primary cost is in bandwidth, and it has enough computational and storage resources

Actual Attack

- Preparation phase
  - Backpointer Hijacking: honest A, attacker M
    - Learns A's Routing Table by sending appropriate queries

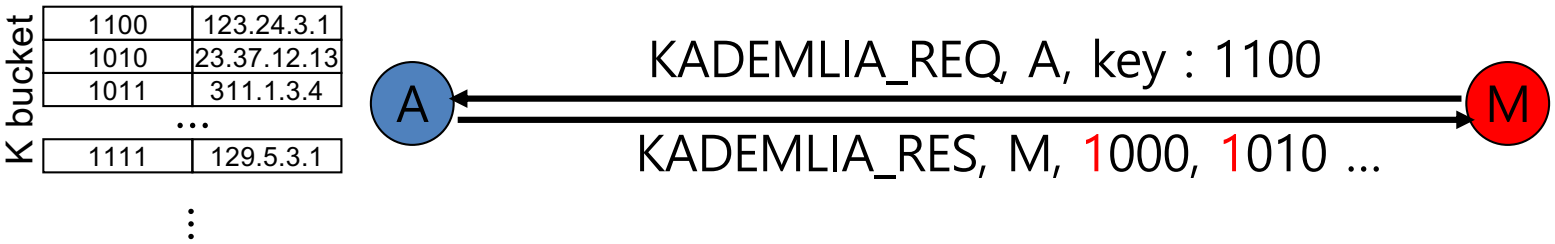

    - Then, change routing table by sending the following message.

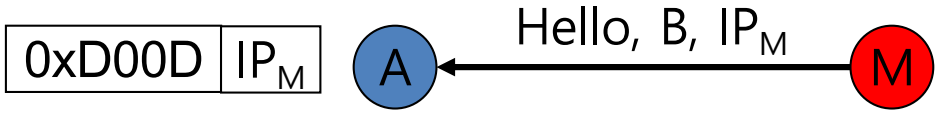

Actual Attack

- Execution Phase
  - Termination condition

    Keyword terminates when the querier A receives more than 300 keyword matches in response to SEARCH_REQ messages

    => malicious node sends a list of 300 bogus matches in reponse
  - Timeout condition

    No reply for 25 seconds, it will stop sending message.

    => Provide many non-existing contacts

Attack Evaluation

Summary of Estimated Cost

- Assumption
  - Total 1M nodes
  - 860 routing table entries
  - 100 Mbps network link
- Preparation phase cost
  - 41.2GB bandwidth to hijack 30% of routing table
  - Takes 55 minutes with 100 Mbps link
- Execution phase cost
  - 100 Mbps link is sufficient to stop 65% of WHOLE query messages.

Large Scale PlanetLab Experiment

- 11,303 ~ 16,105 Kad nodes running on ~500 PlanetLab machines

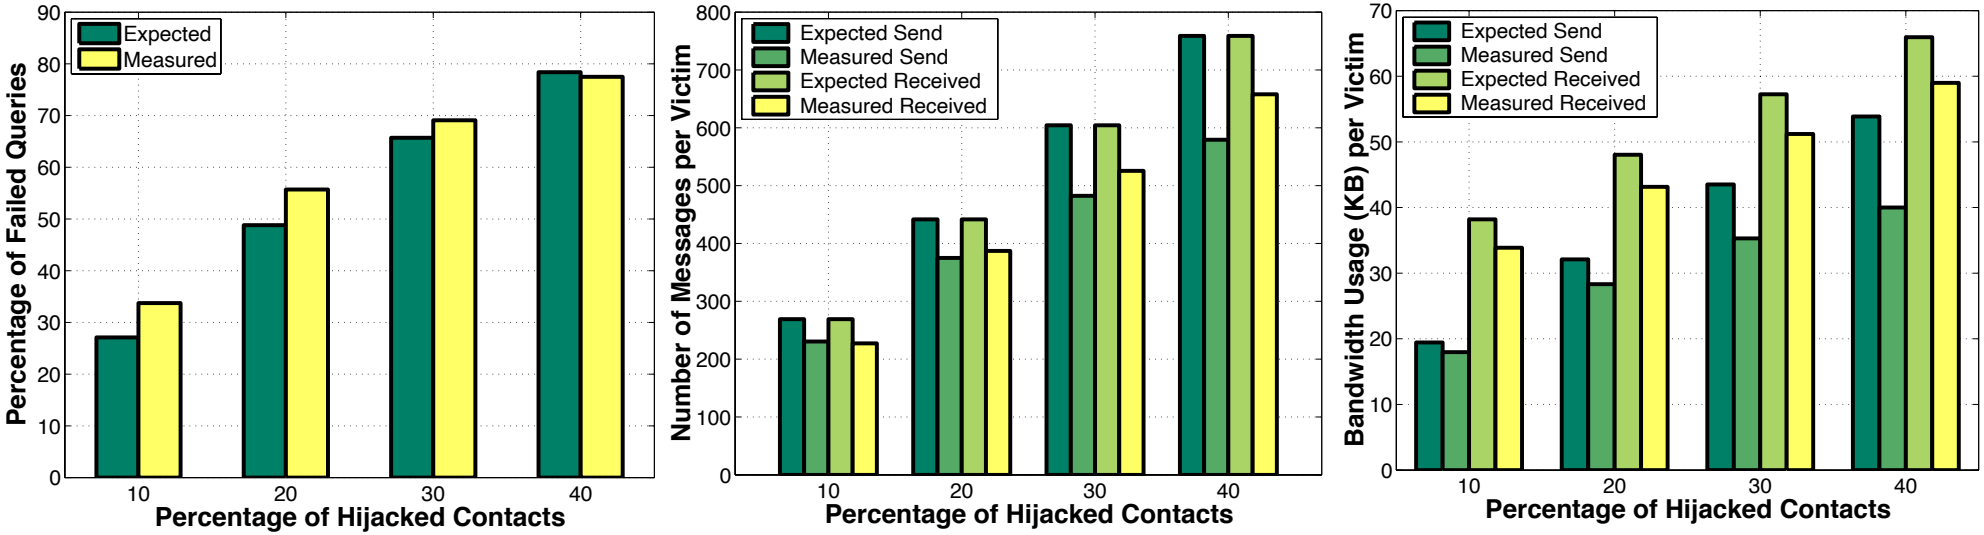

- Comparison between expected and measured
  - keyword query failures
  - Number of messages used to attack one node
  - Bandwidth usage

Large Scale Simulation

- 50,000 nodes and 50 attackers with DVN
- Focus on control plane (routing process)

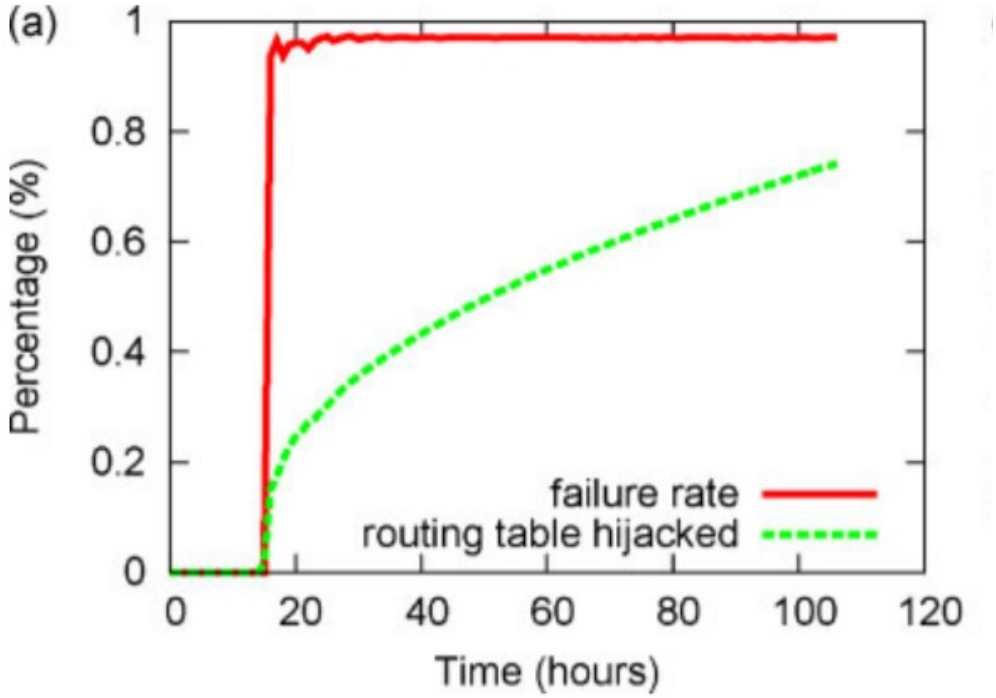

Screen Shots

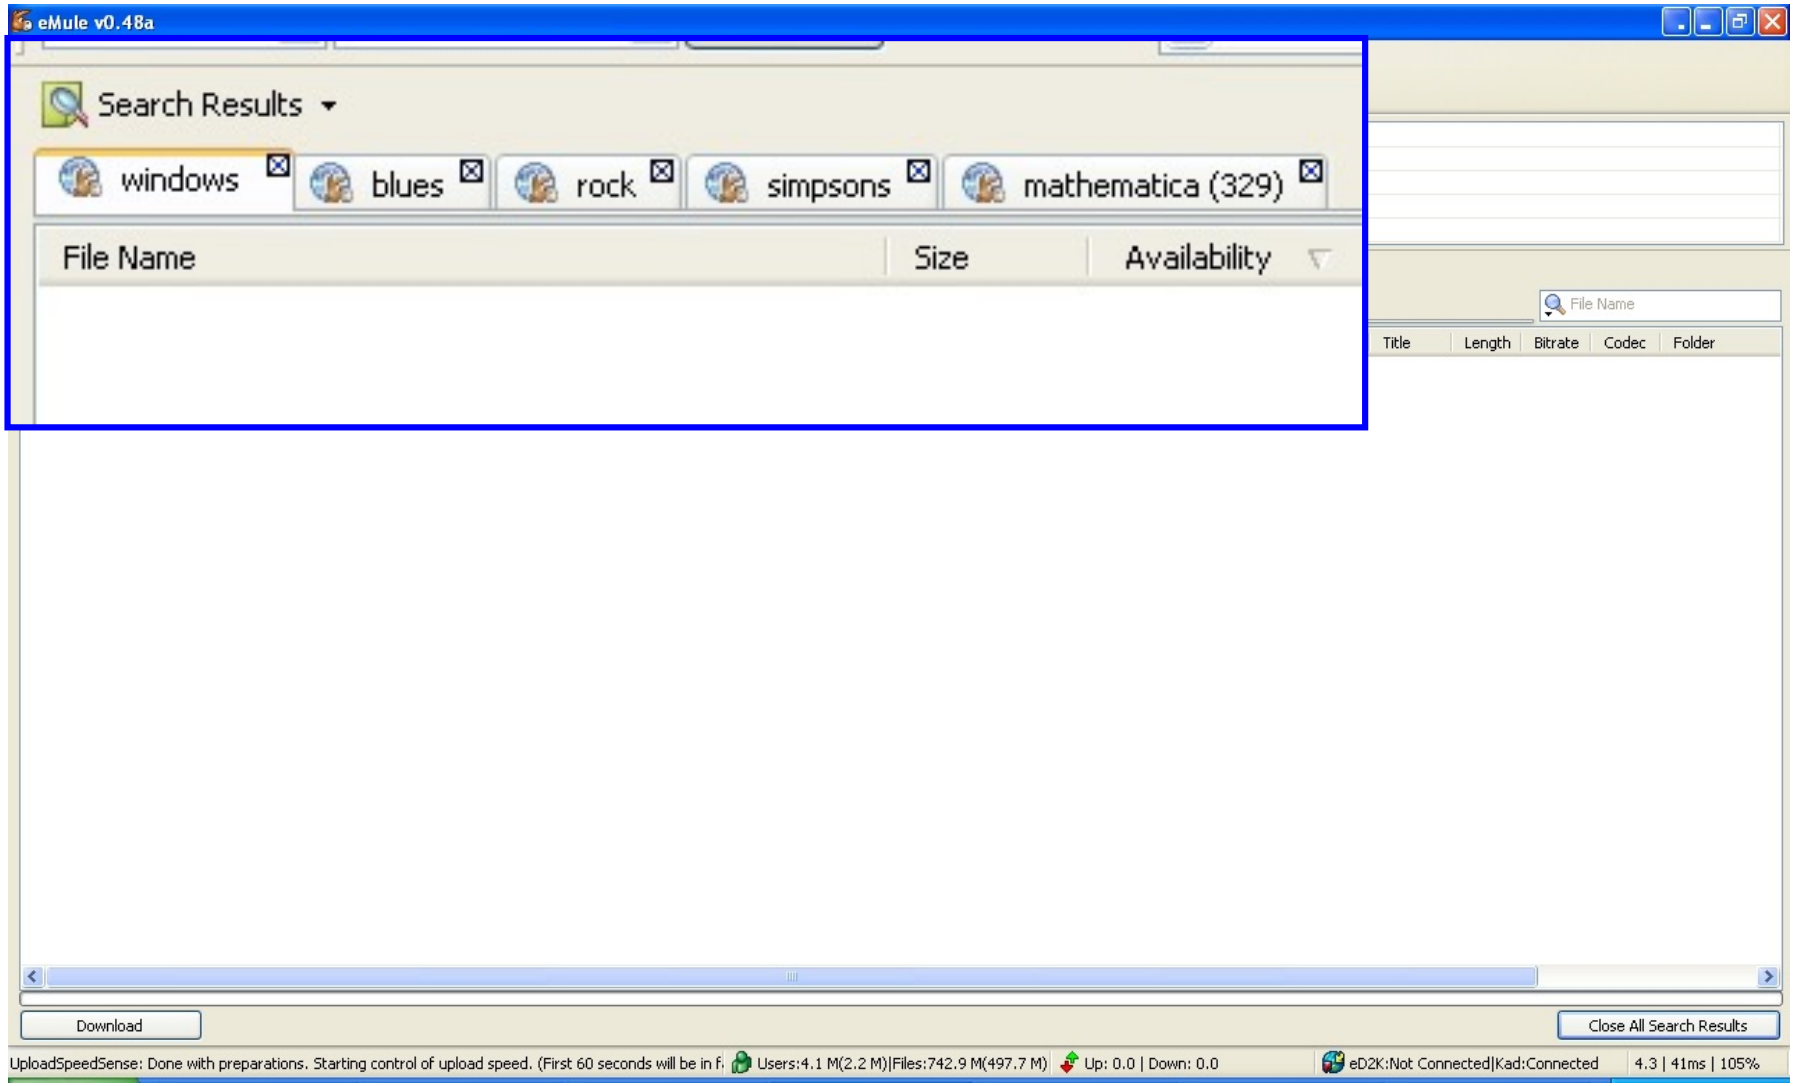

Reflection Attack

- Fill node A's routing table with A itself.

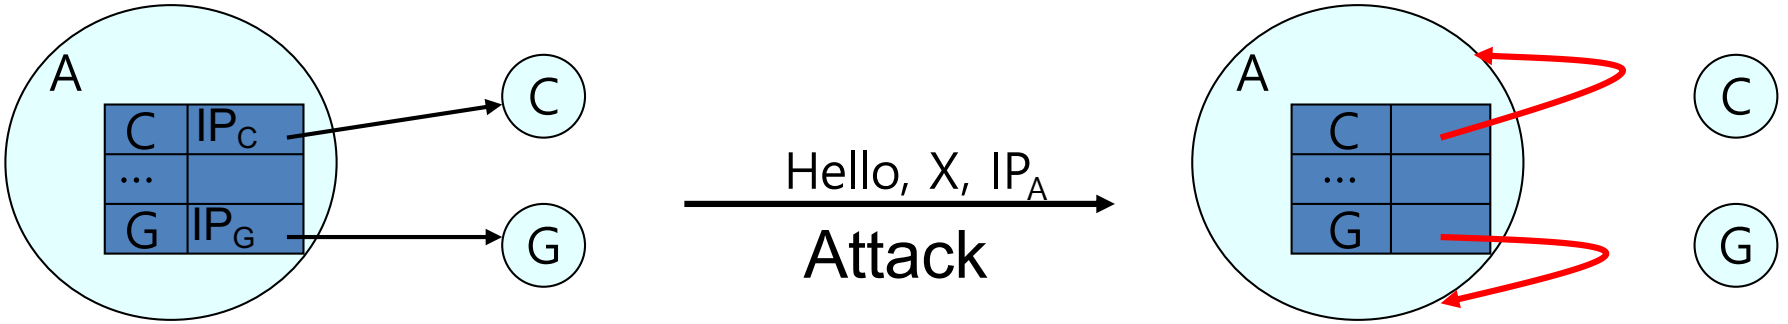

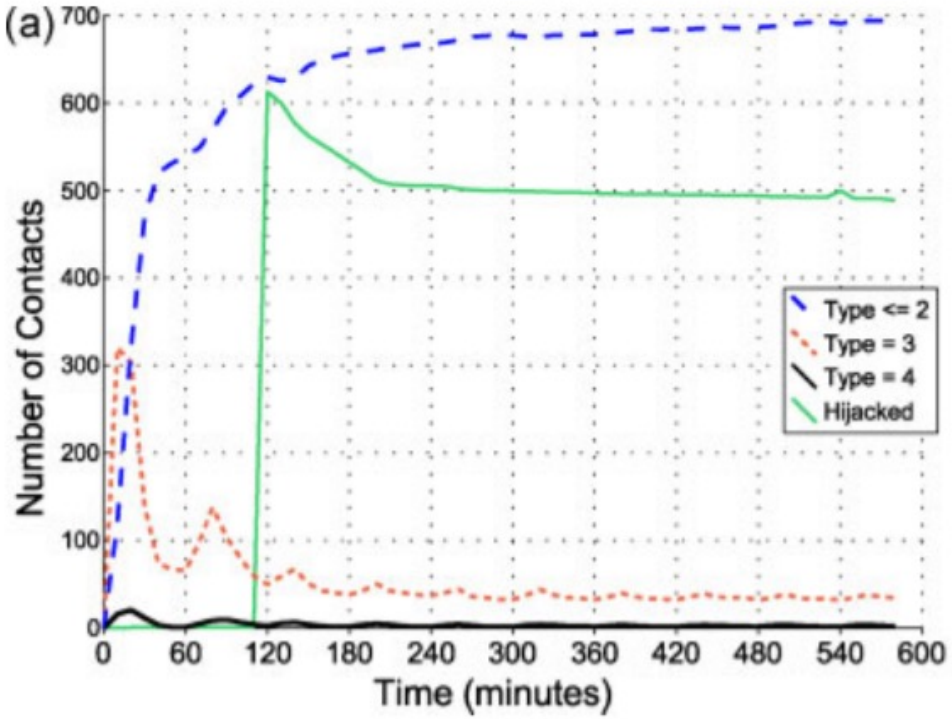

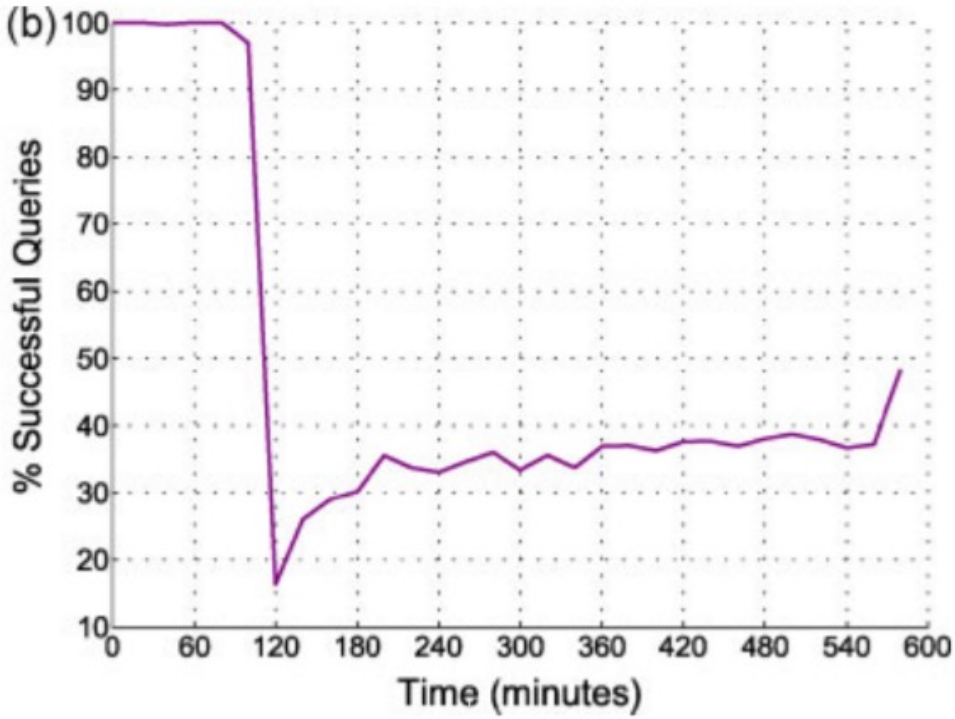

Mitigation

❖ Identity Authentication

Method	Secure	Persistent ID	Incremental deployable
Verify the liveness of old IP	No	Yes	Yes
Drop Hello with new IP	Yes	No	Yes
ID=hash(IP)	Yes	No	No
ID=hash(Public Key)	Yes	Yes	No

❖ Routing Corruption

- 3 parallel lookups but they are not independent

backpointers	Current method	Independent parallel routes
40%	98% fail	45% fail
10%	59.5% fail	1.7% fail

Then

```
----------------------
- Jun, 27. 2008 -
----------------------
.: Several changes were made to Kad in order to defy routing attacks
researched by University of Minnesota guys [Peng Wang, James Tyra, Eric
Chan-Tin, Tyson Malchow, Denis Foo Kune, Nicholas Hopper, Yongdae Kim], in
particular:
.: Kad contacts will only be able to update themself in others routing tables if
they provide the proper key (supported by 0.49a+ nodes) in order to make it
impossible to hijack them
.: Kad uses now a three-way-handshake (or for older version a similar check)
for new contacts, making sure they do not use a spoofed IP
.: Unverified contacts are not used for routing tasks and a marked with a
special icon in the GUI
```

Related Work

- Sit, E., & Morris, R. (2002, March). **Security considerations for peer-to-peer distributed hash tables.** In *International Workshop on Peer-to-Peer Systems* (pp. 261-269). Springer, Berlin, Heidelberg.

- Castro, M., Druschel, P., Ganesh, A., Rowstron, A., & Wallach, D. S. (2002). **Secure routing for structured peer-to-peer overlay networks.** *ACM SIGOPS Operating Systems Review*, *36*(SI), 299-314.

- Fiat, A., Saia, J., & Young, M. (2005, October). **Making chord robust to byzantine attacks.** In *European Symposium on Algorithms* (pp. 803-814). Springer, Berlin, Heidelberg.

Work After This Work

- Cholez, T., Chrisment, I., & Festor, O. (2009, June). **Evaluation of sybil attacks protection schemes in kad**. In *IFIP International Conference on Autonomous Infrastructure, Management and Security* (pp. 70-82). Springer, Berlin, Heidelberg.

- Yu, J., Fang, C., Xu, J., Chang, E. C., & Li, Z. (2009, September**). ID repetition in Kad.** In *Peer-to-Peer Computing, 2009. P2P'09. IEEE Ninth International Conference on* (pp. 111-120). IEEE.

- Fantacci, R., Maccari, L., Rosi, M., Chisci, L., Aiello, L. M., & Milanesio, M. (2009, June**). Avoiding eclipse attacks on kad/kademlia: an identity based approach**. In *Communications, 2009. ICC'09. IEEE International Conference on* (pp. 1-5). IEEE.

- Apostolaki, M., Zohar, A., & Vanbever, L. (2017, May). **Hijacking bitcoin: Routing attacks on cryptocurrencies.** In *Security and Privacy (SP), 2017 IEEE Symposium on* (pp. 375-392). IEEE.

- Blond, S. L., Manils, P., Abdelberi, C., Kaafar, M. A. D., Castelluccia, C., Legout, A., & Dabbous, W. (2011**). One bad apple spoils the bunch: exploiting P2P applications to trace and profile Tor users.** *arXiv preprint arXiv:1103.1518*.

Future Work

- Ethereum uses a variant of Kademlia protocol in node discovery.
- Make nodeID with its public key
- Try to connect with the closest node as a peer

  → Same problem in here?

- Why no verification mechanism in the first place?

Conclusion

- Deny service to a large portion of the Kad network with only 100Mbps of bandwidth
- The attack was successful and efficient.
- This attack is more efficient than currently known attacks such as Sybil and Index Poisoning
- Introduce new simulator, DVN.

FAQ

- **정현식**
  - **How Kad is different from the original Kademlia specification? If it's n ot that different, there are some blockchains using Kademlia to disco ver nodes. Was there any effective attack on those blockchains?**
- **Tuan**
  - **I think these vulnerabilities can be exploited to make more serious att acks, which can affect users worldwide. One example: malicious node s response A with malicous IDs, that IDs contain malwares or ransom ewares, users don't have any sense to know they are malwares and ac cess that data.**
- 고우영
  - This DVN simulator seems very powerful, Can you introduce recent simul ator?
- 김성중
  - Is there a similar attack on the recently proposed P2P network?
- 이태화
  - Many cryptocurrency uses p2p services. Are there same problems?

Thank you!!!