Peer-to-Peer Networks

14 Security

Christian Schindelhauer
Technical Faculty
Computer-Networks and Telematics
University of Freiburg

Free-Net

- Ian Clarke, Oskar Sandberg, Brandon Wiley, Theodore Hong, 2000
- Goal
  - peer-to-peer network
  - allows publication, replication, data lookup
  - anonymity of authors and readers
- Files
  - are encoding location independent
    - by encrypted and pseudonymously signed index files
    - author cannot be identified
  - are secured against unauthorized change or deletion
  - are encoded by keys unknown by the storage peer
    - secret keys are stored elsewhere
  - are replicated
    - on the look up path
  - and erased using "Least Recently Used" (LRU) principle

Free-Net

- Network Structure
  - is similar to Gnutella
  - Free-Net is like Gnutella Pareto distributed
- Storing Files
  - Each file can be found, decoded and read using the encoded address string and the signed subspace key
  - Each file is stored together with the information of the index key but without the encoded address string
  - The storage peer cannot read his files
    - unless he tries out all possible keywords (dictionary attack)
- Storing of index files
  - The address string coded by a cryptographic secure hash function leads to the corresponding peer
    - who stores the index data
      - address string
      - and signed subspace key
  - Using this index file the original file can be found

Free-Net

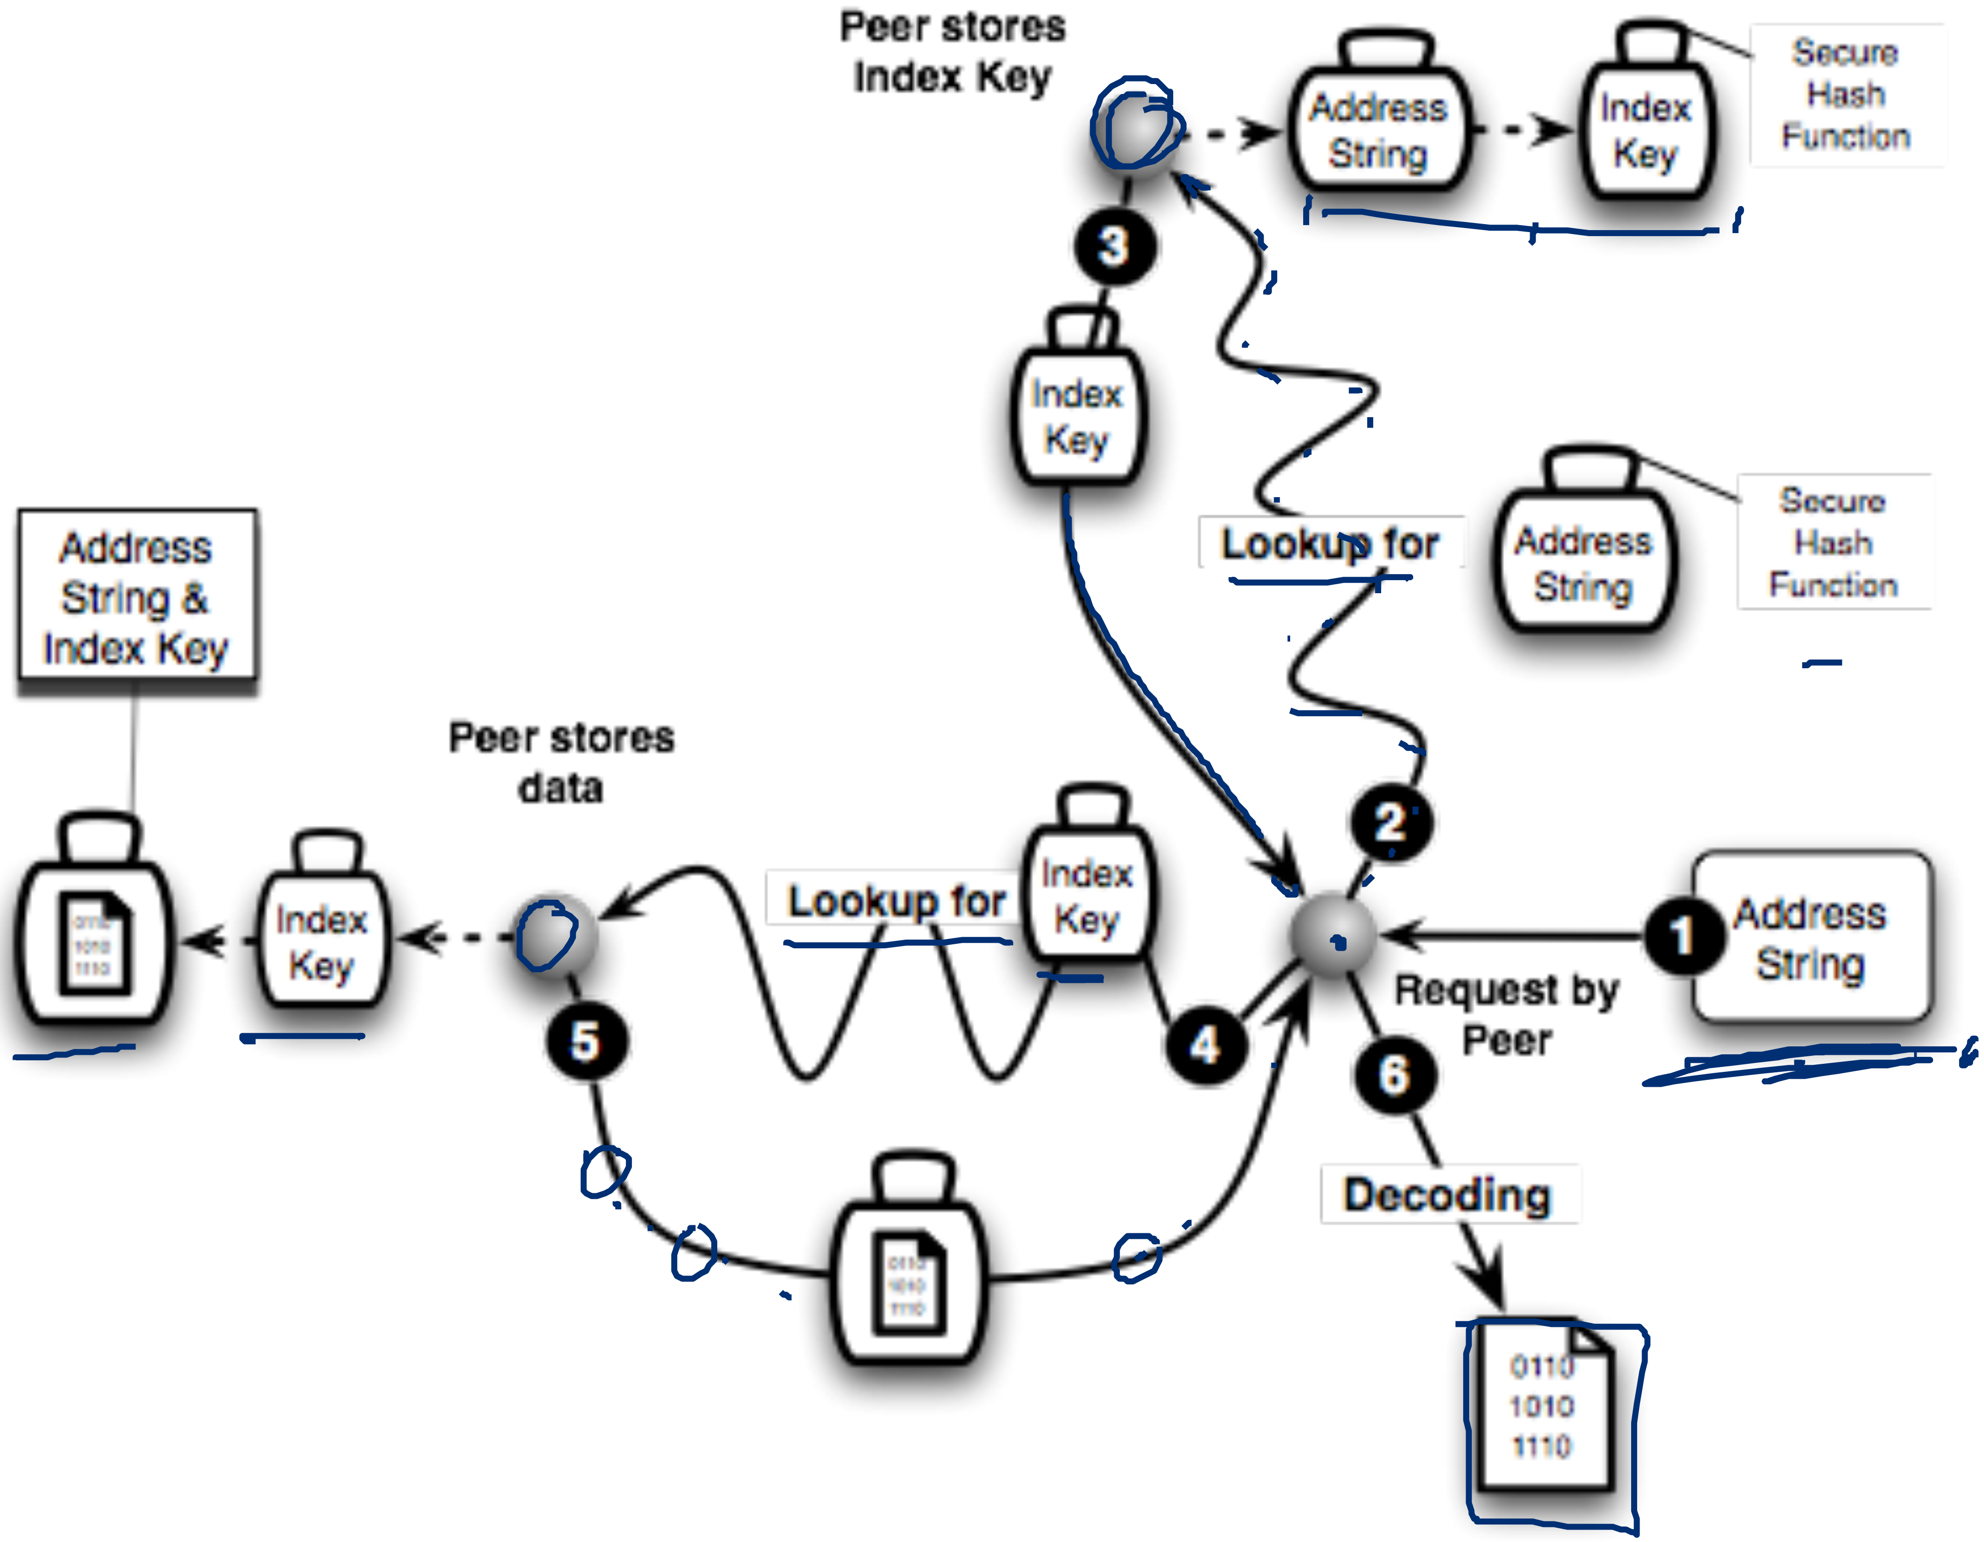

Peer stores
Index Key
Address
String
Index
Key
Secure
Hash
Function
3
Index
Key
Lookup for
Address
String
Secure
Hash
Function
Address
String &
Index Key
Peer stores
data
2
Index
Key
Lookup for
Index
Key
1
Address
String
Request by
Peer
5
4
6
Decoding
0110
1010
1110

Free-Net

- Lookup
  - steepest-ascent hill-climbing
    - lookup is forwarded to the peer whose ID is closest to the search index
  - with TTL field
    - i.e. hop limit
- Files are moved to new peers
  - when the keyword of the file is similar to the neighbor's ID
- New links
  - are created if during a lookup close similarities between peer IDs are discovered

Efficiency of Free-Net

- Network structure of Free-Net is similar to Gnutella
- The lookup time is polynomial on the average

$2^{\log_2 T} = 2^{a \cdot \log_2 n + b}$

$T = n^a \cdot 2^b = O(n^a)$

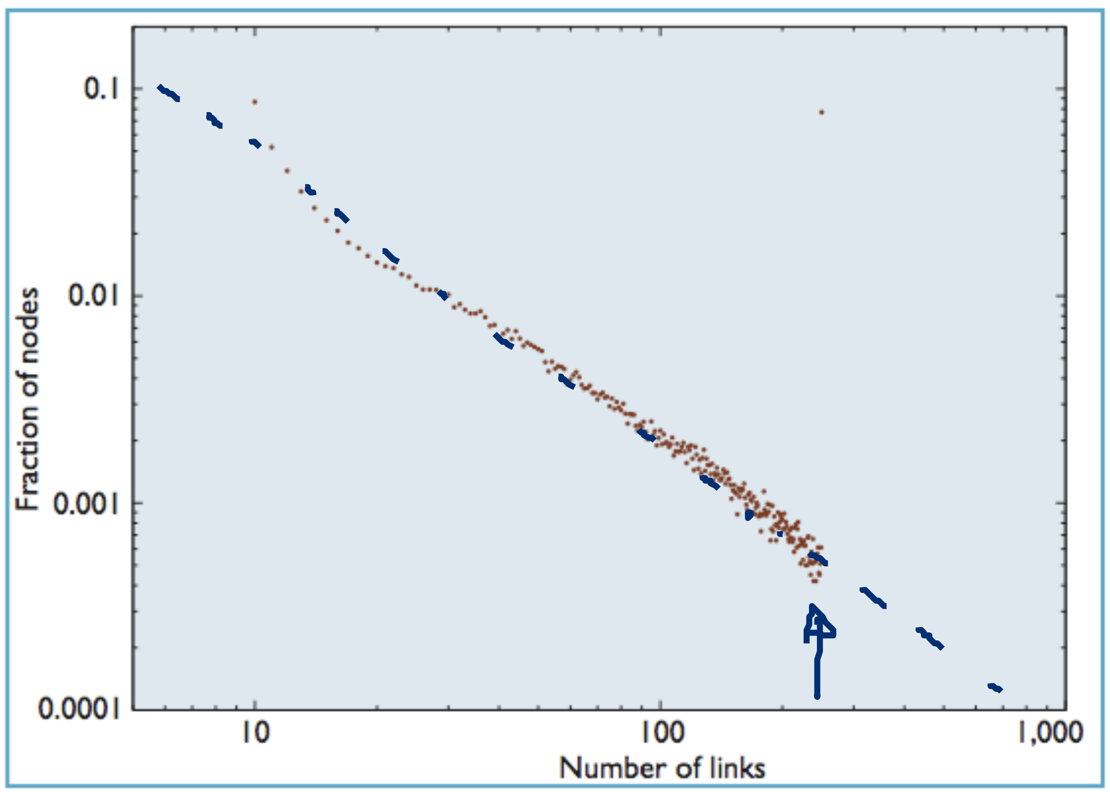

Figure 2. Degree distribution among Freenet nodes. The network shows a close fit to a power-law distribution.

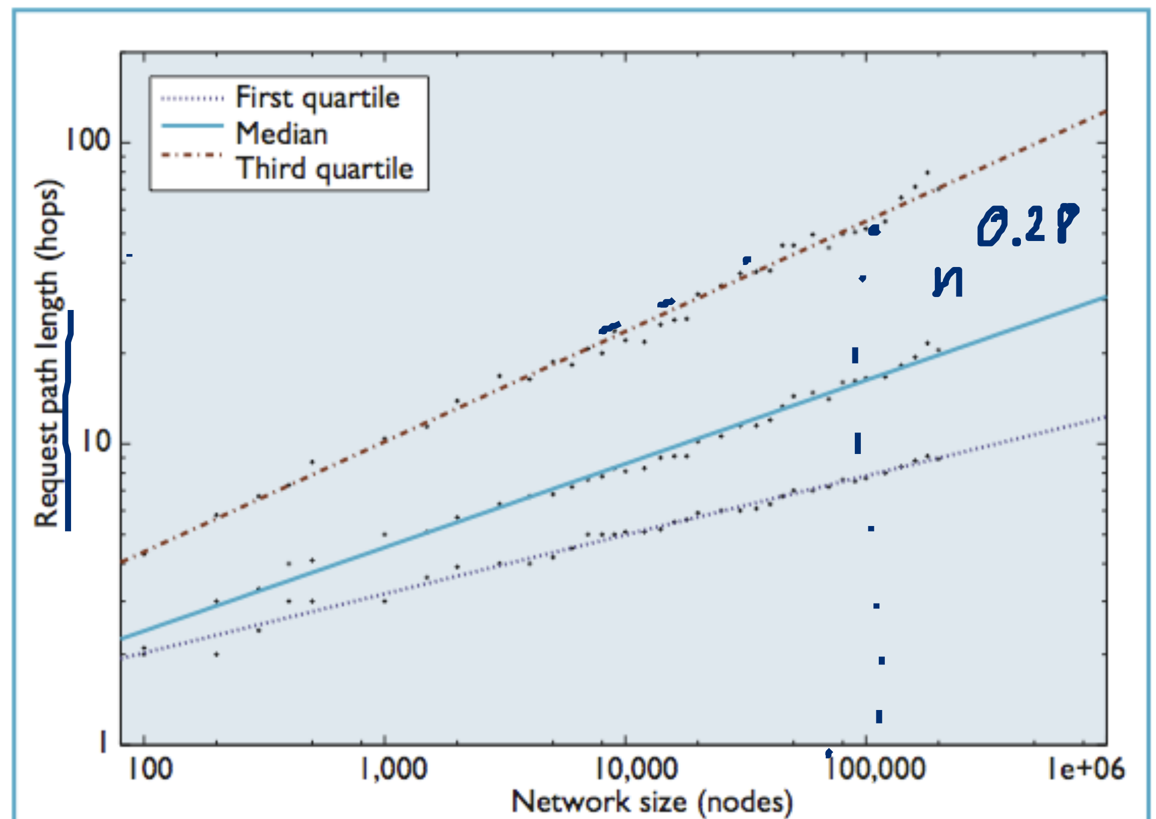

Figure 3. Request path length versus network size. The median path length in the network scales as $N^{0.28}$.

Chord: $O(\log n)$

CAN: $O(n^{1/d})$

$0.28 \sim \frac{1}{4}$

Dark-Net & Friend-to-Friend

- Dark-Net is a private Peer-to-Peer Network
  - Members can trust all other members
  - E.g.
    - friends (in real life)
    - sports club
- Dark-Net control access by
  - secret addresses,
  - secret software,
  - authentication using password, or
  - central authentication
- Example:
  - WASTE
    - P2P-Filesharing up to 50 members
    - by Nullsoft (Gnutella)
  - CSpace
    - using Kademlia

or Bittorrent Sync

CoNe Freiburg

Solutions to the Sybil Attack

- Survey paper by Levine, Shields, Margonin, 2006
- Trusted certification
  - only approach to completely eleminate Sybil attacks
    - according to Douceur
  - relies on centralized authority
- No solution
  - know the problem and deal with the consequences
- Resource testing
  - real world friends
  - test for real hardware or addresses
    - e.g. heterogeneous IP addresses
  - check for storing ability
- Recurring cost and fees
  - give the peers a periodic task to find out whether there is real hardware behind each peer
    - wasteful use of resources
  - charge each peer a fee to join the network
- Trusted devices
  - use special hardware devices which allow to connect to the network

→ Skype
→ central server

Find x s.t.
$h(x) = y$
$x \in \{0,1\}^{100000}$

PAST
- smartcard

UNI FREIBURG

Solutions to the Sybil Attack

- Survey paper by Levine, Shields, Margonin, 2006

- In Mobile Networks
  - use observations of the mobile node
    - e.g. GPS location, neighbor nodes, etc.
- Auditing
  - perform tests on suspicious nodes
  - or reward a peer who proves that it is not a clone peer
- Reputation Systems
  - assign each peer a reputation which grows over the time with each positive fact
  - the reputation indicates that this peer might behave nice in the future
  - Disadvantage:
    - peers might pretend to behave honestly to increase their reputation and change their behavior in certain situations
    - problem of Byzantine behavior

CoNe Freiburg

The Problem of Byzantine Generals

Distributed Systems | Byzanz

- 3 armies prepare to attack a castle
- They are separated and communicate by messengers
- If one army attacks alone, it loses
- If two armies attack, they win
- If nobody attacks the castle is besieged and they win
- One general is a renegade
  - nobody knows who

C

don't attack

A

B ¿

attack

don't

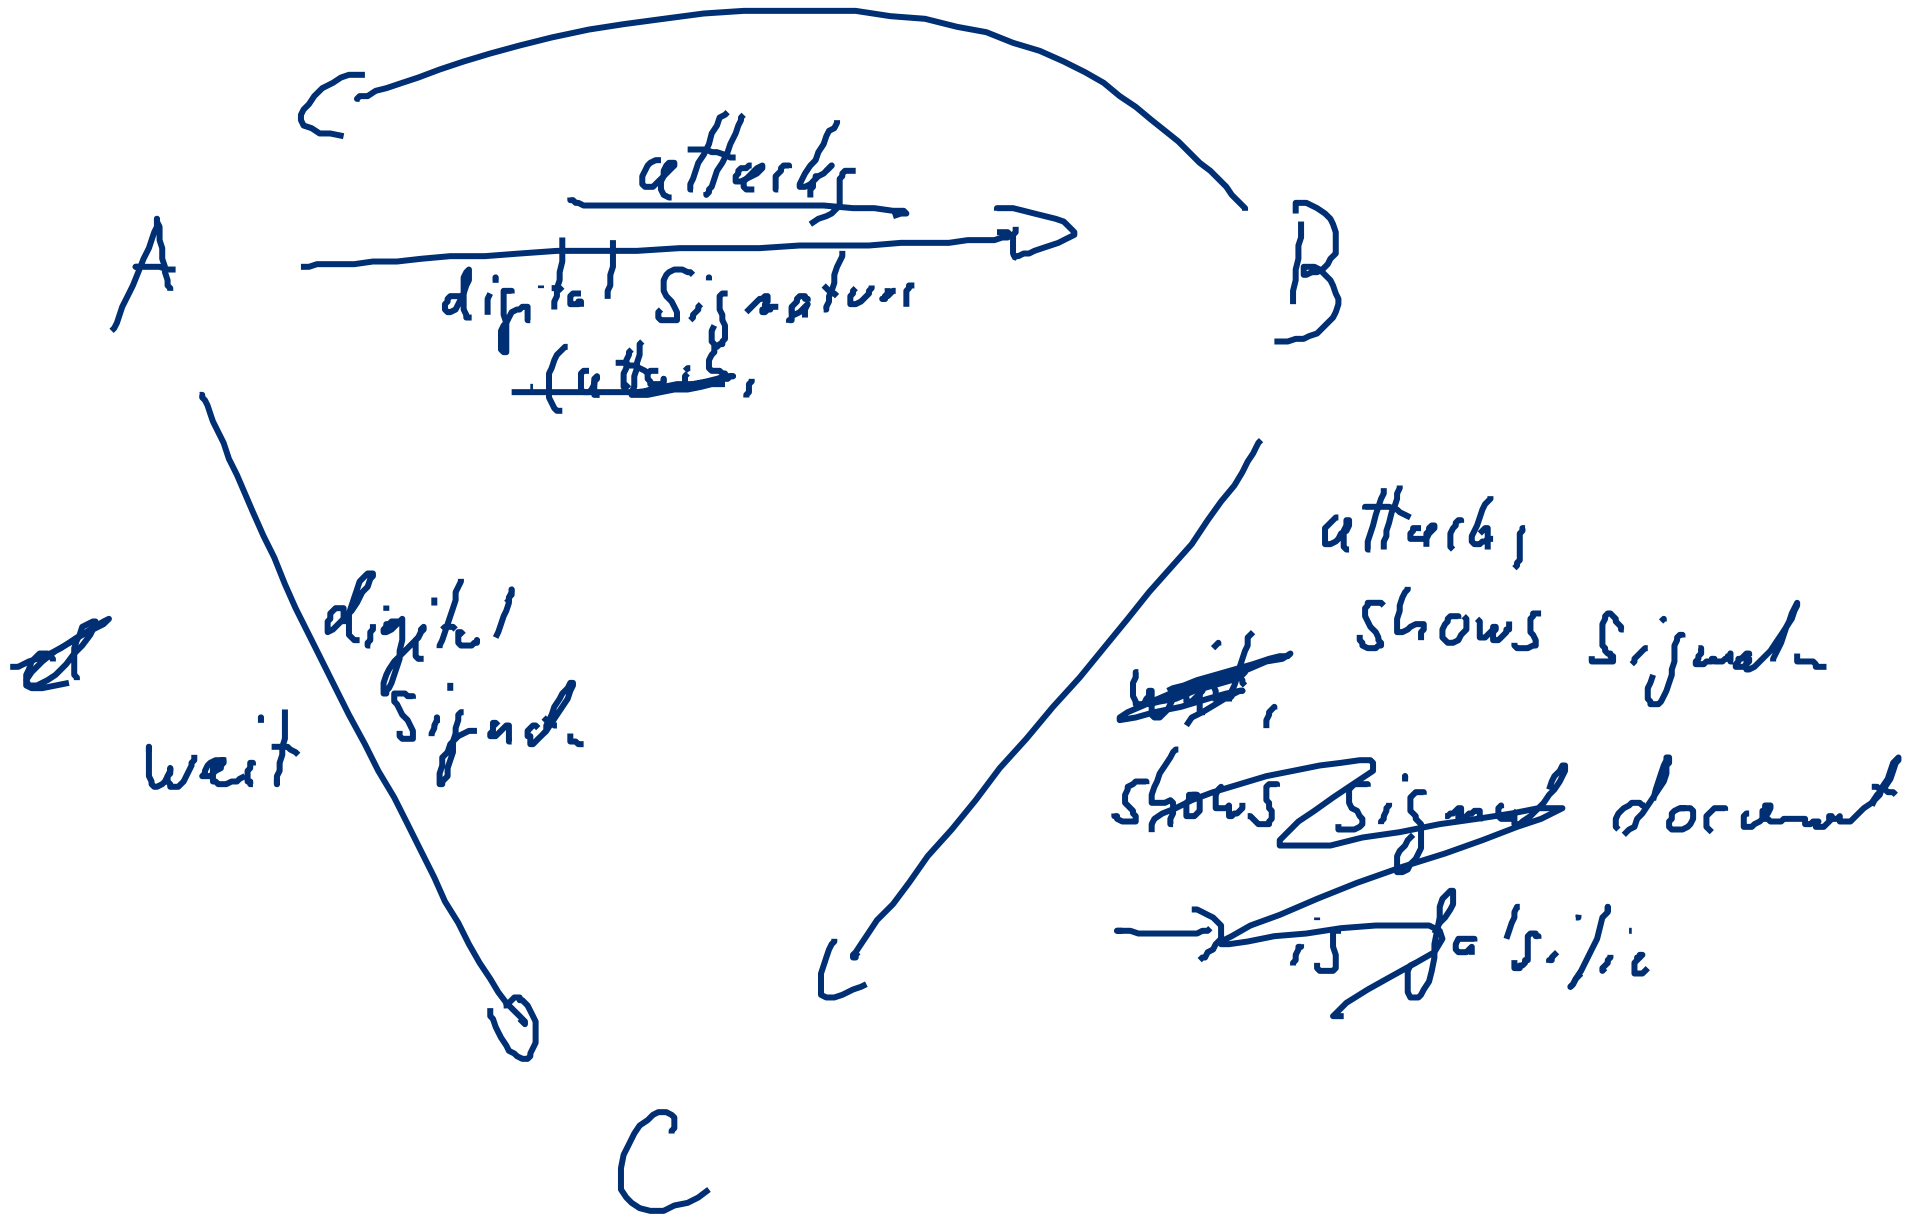
A
B
C
attacks
digital signature
digital signat.
wait
attacks,
shows signature
shows signed document
is falsific

The Problem of Byzantine Generals

- The evil general X tries
  - to convince A to attack
  - to convince B to wait
- A tells B about X‘s command
- B tells B about his version of X‘s command
  - contradiction
- But is A, B, or X lying?

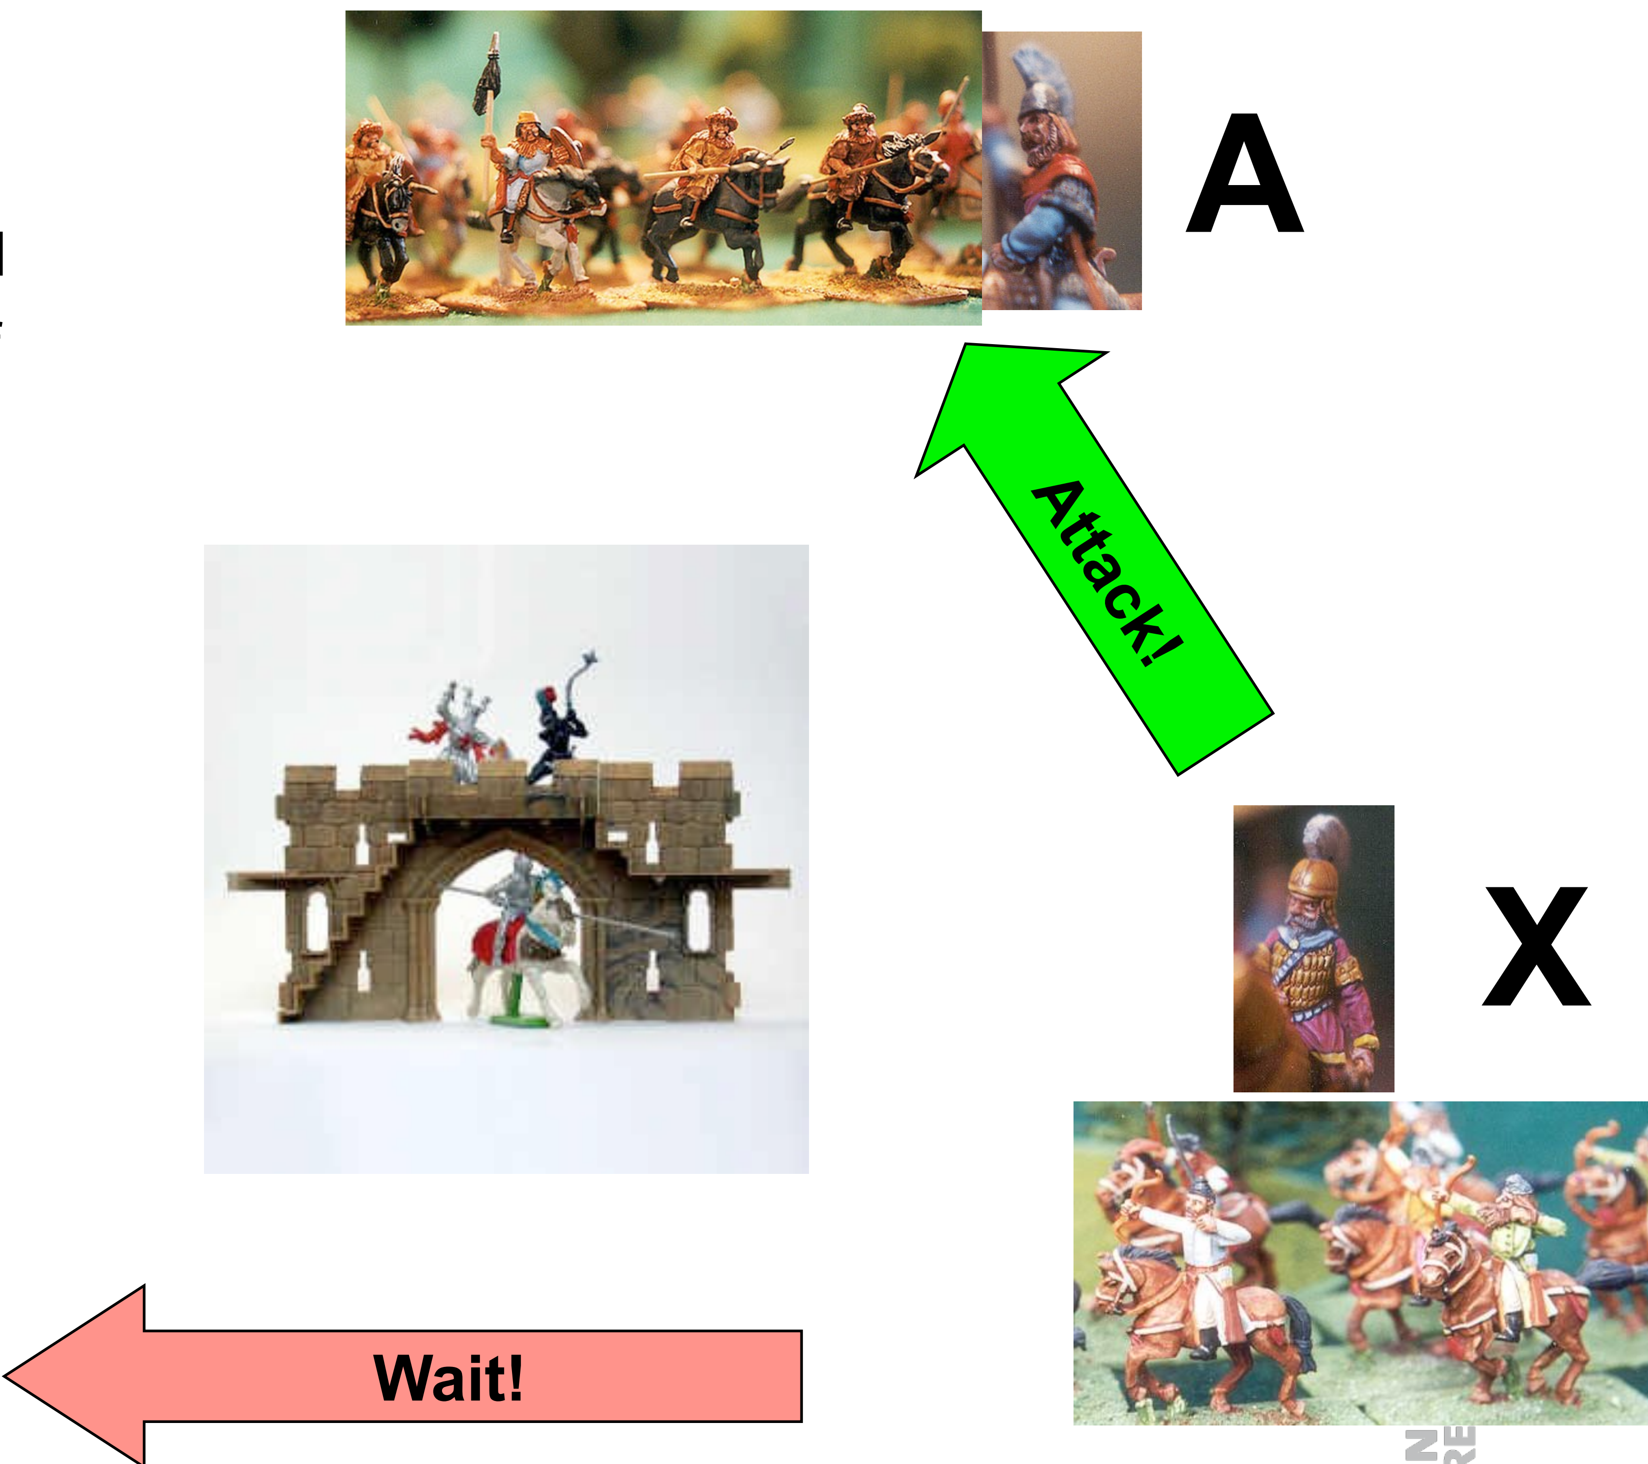

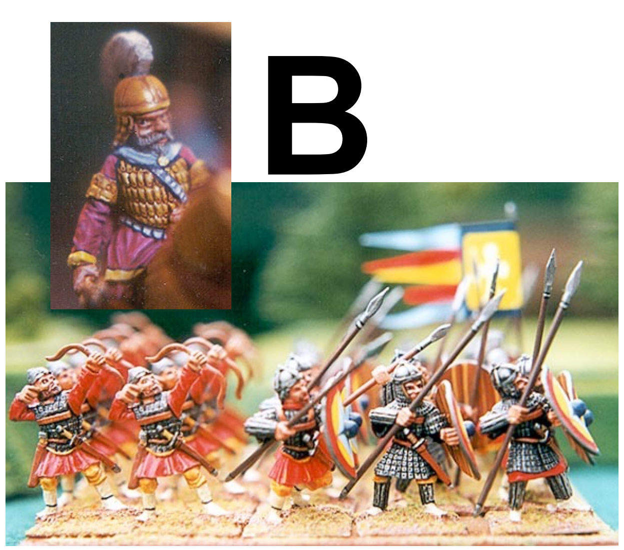

CoNe Freiburg

The Problem of Byzantine Generals

- The evil general X tries
  - to convince A to attack
  - to convince B to wait
- A tells B about X‘s command
- B tells B about his version of X‘s command
  - contradiction
- But is A, B, or X lying?

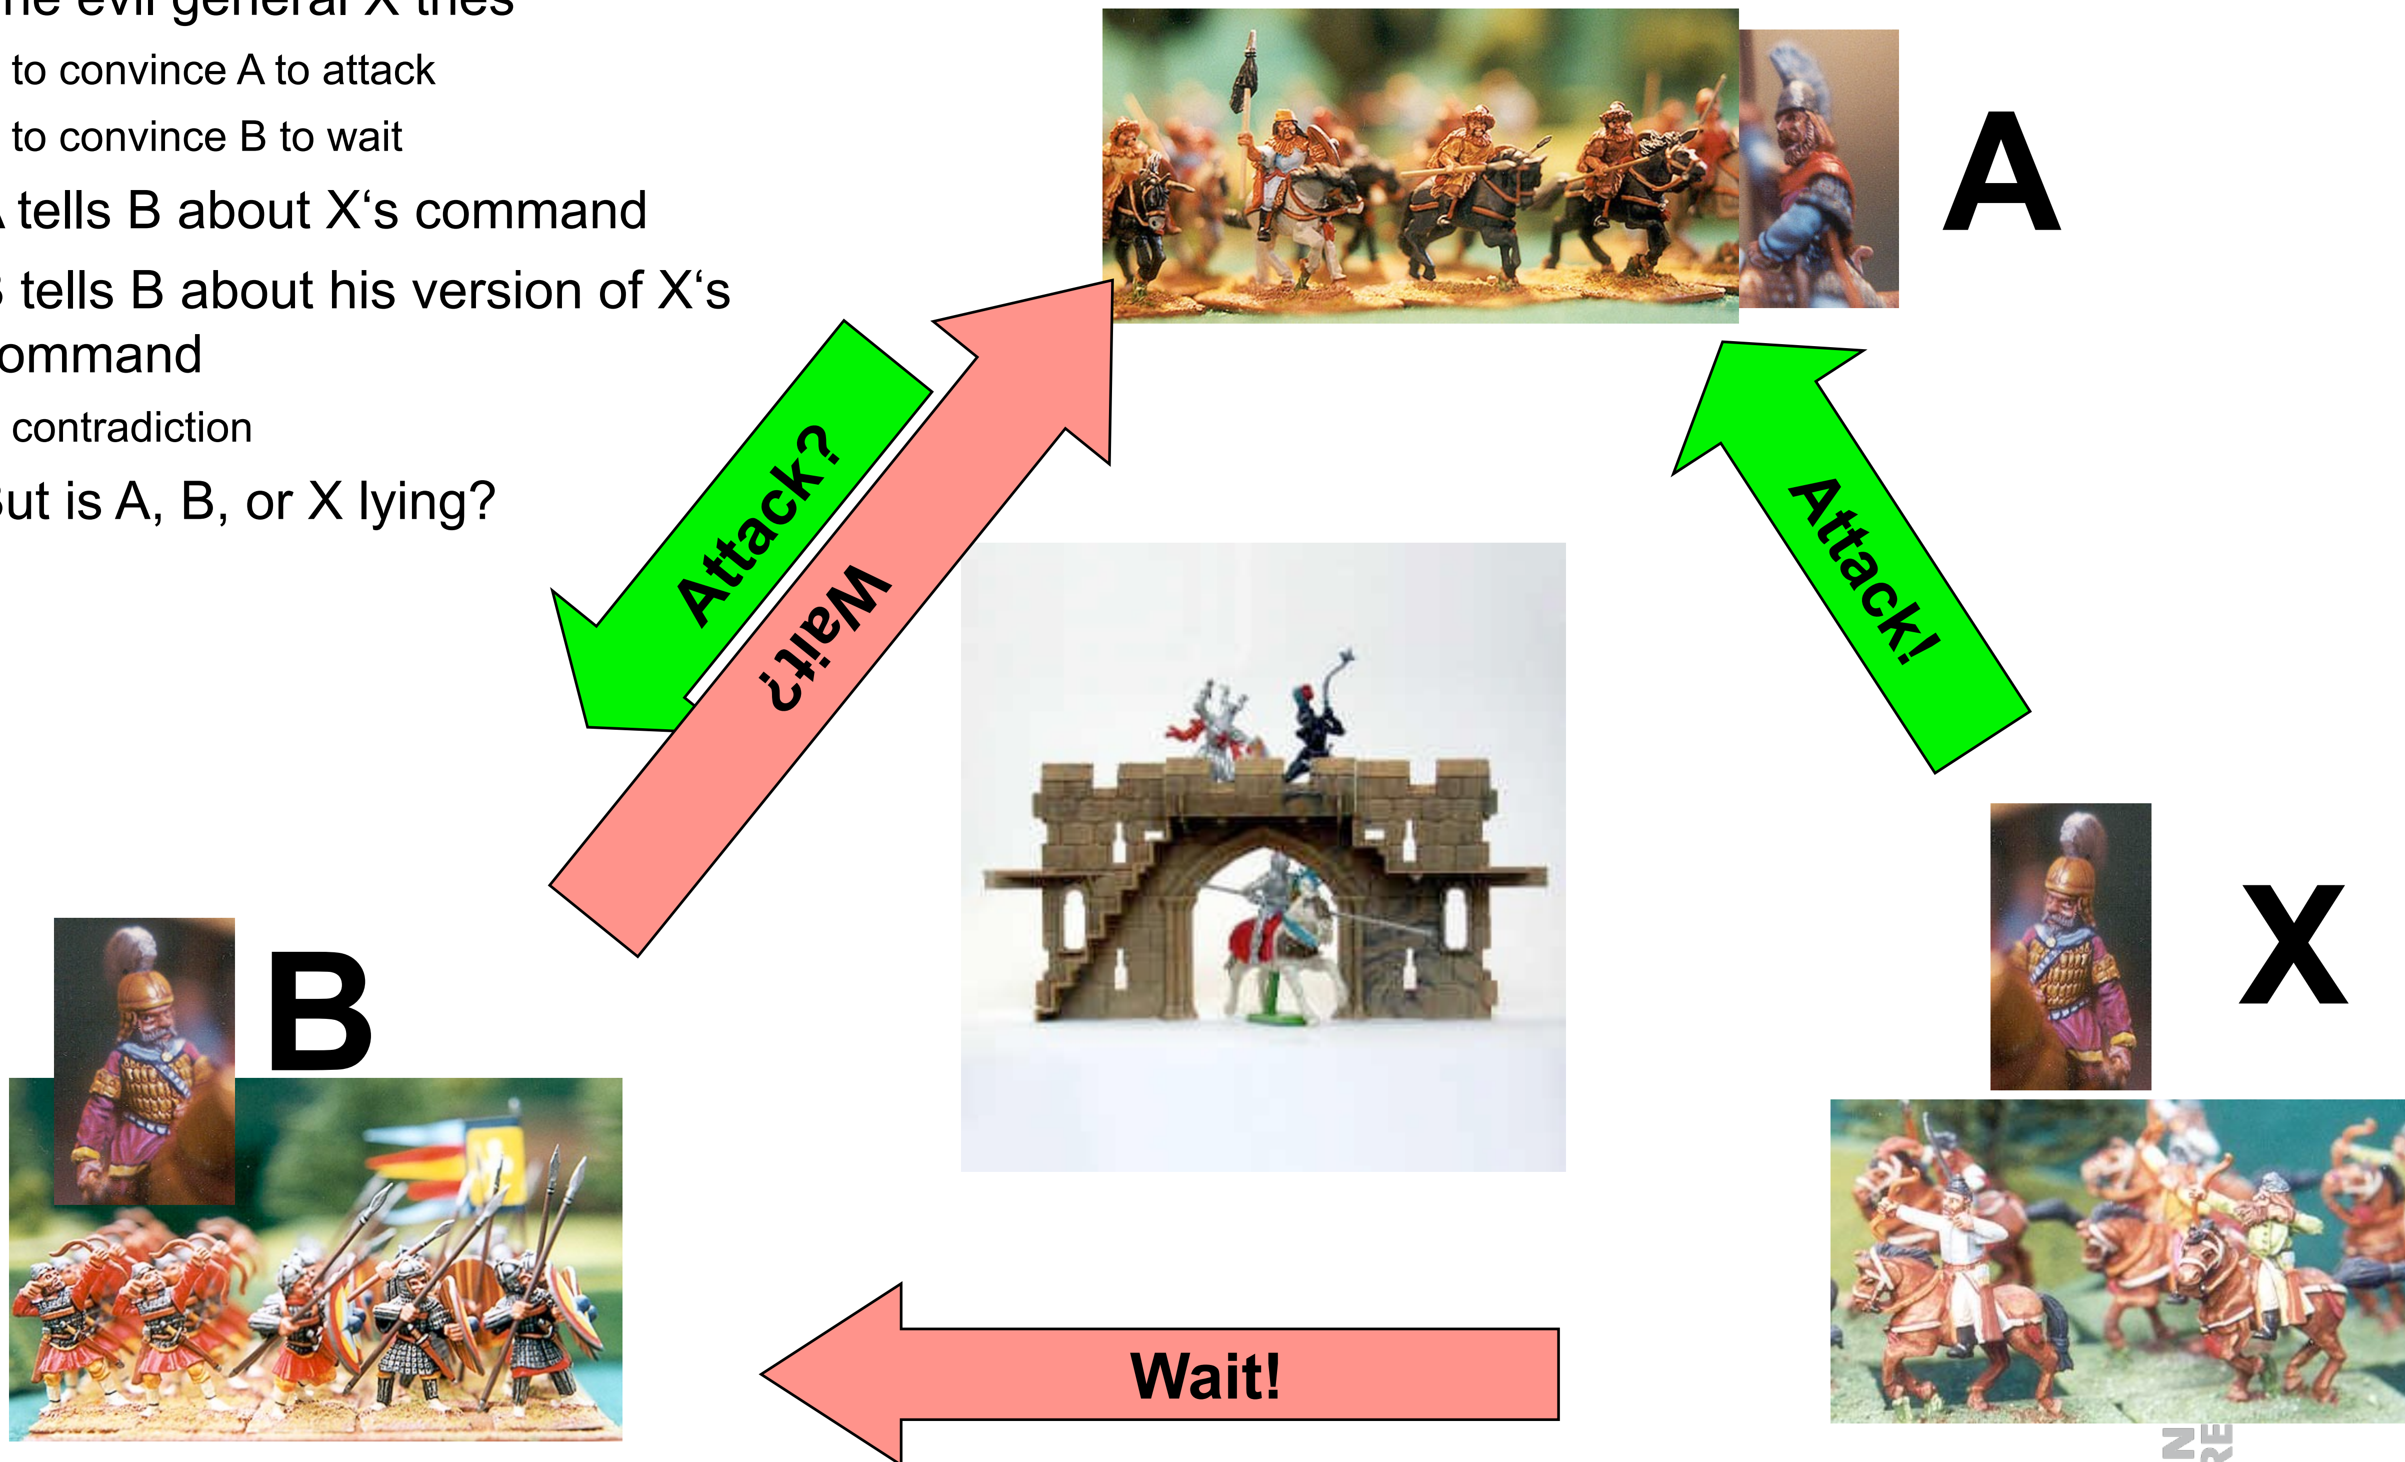

Cryptography

AES

$Encode(m, k_s) = c$

$c \dashrightarrow Decode(c, 1, 2, \ldots) = m$

does it make sense?

$Decode(c, k_s)$

Byzantine Agreement

- Theorem
  - The problem of three byzantine generals cannot be solved (without cryptography)
  - It can be solved for 4 generals
- Consider: 1 general, 3 officers problem
  - If the general is loyal then all loyal officers will obey the command
  - In any case distribute the received commands to all fellow officers
  - What if the general is the renegade?

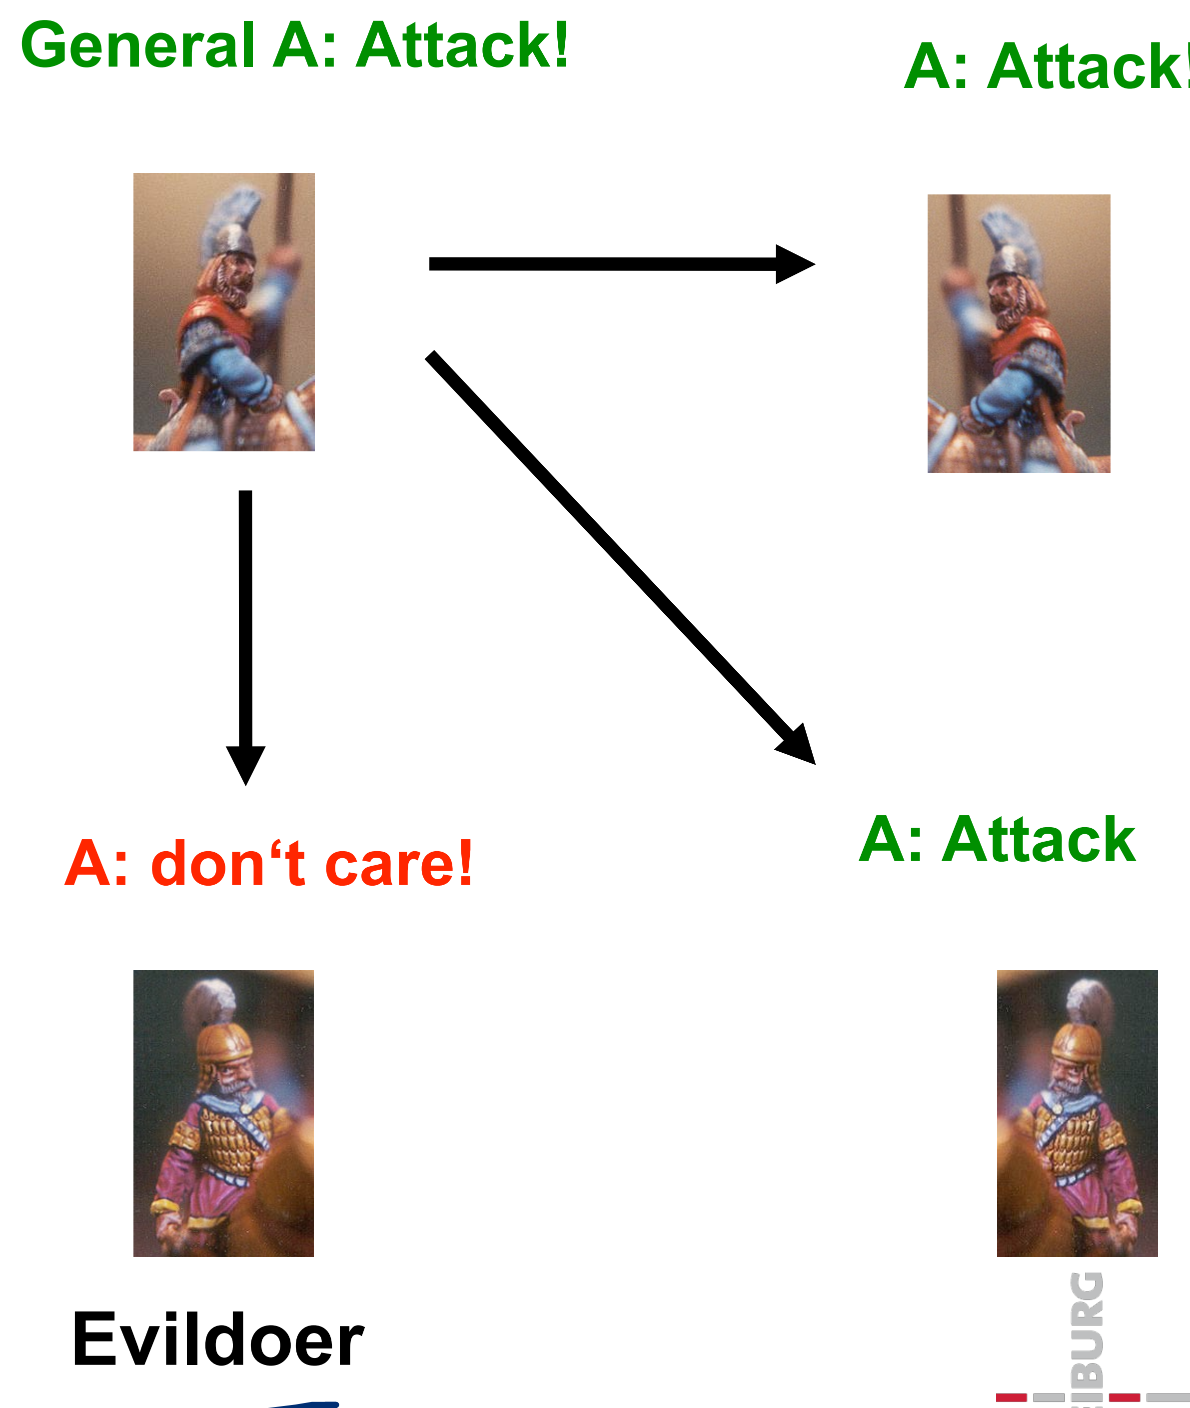

Byzantine Agreement

- Theorem
  - The problem of four byzantine generals can be solved (without cryptography)
- Algorithm
  - General A sends his command to all other generals
    - A sticks to his command if he is honest
  - All other generals forward the received commands to all other generals
  - Every generals computes the majority decision of the received commands and follows this command

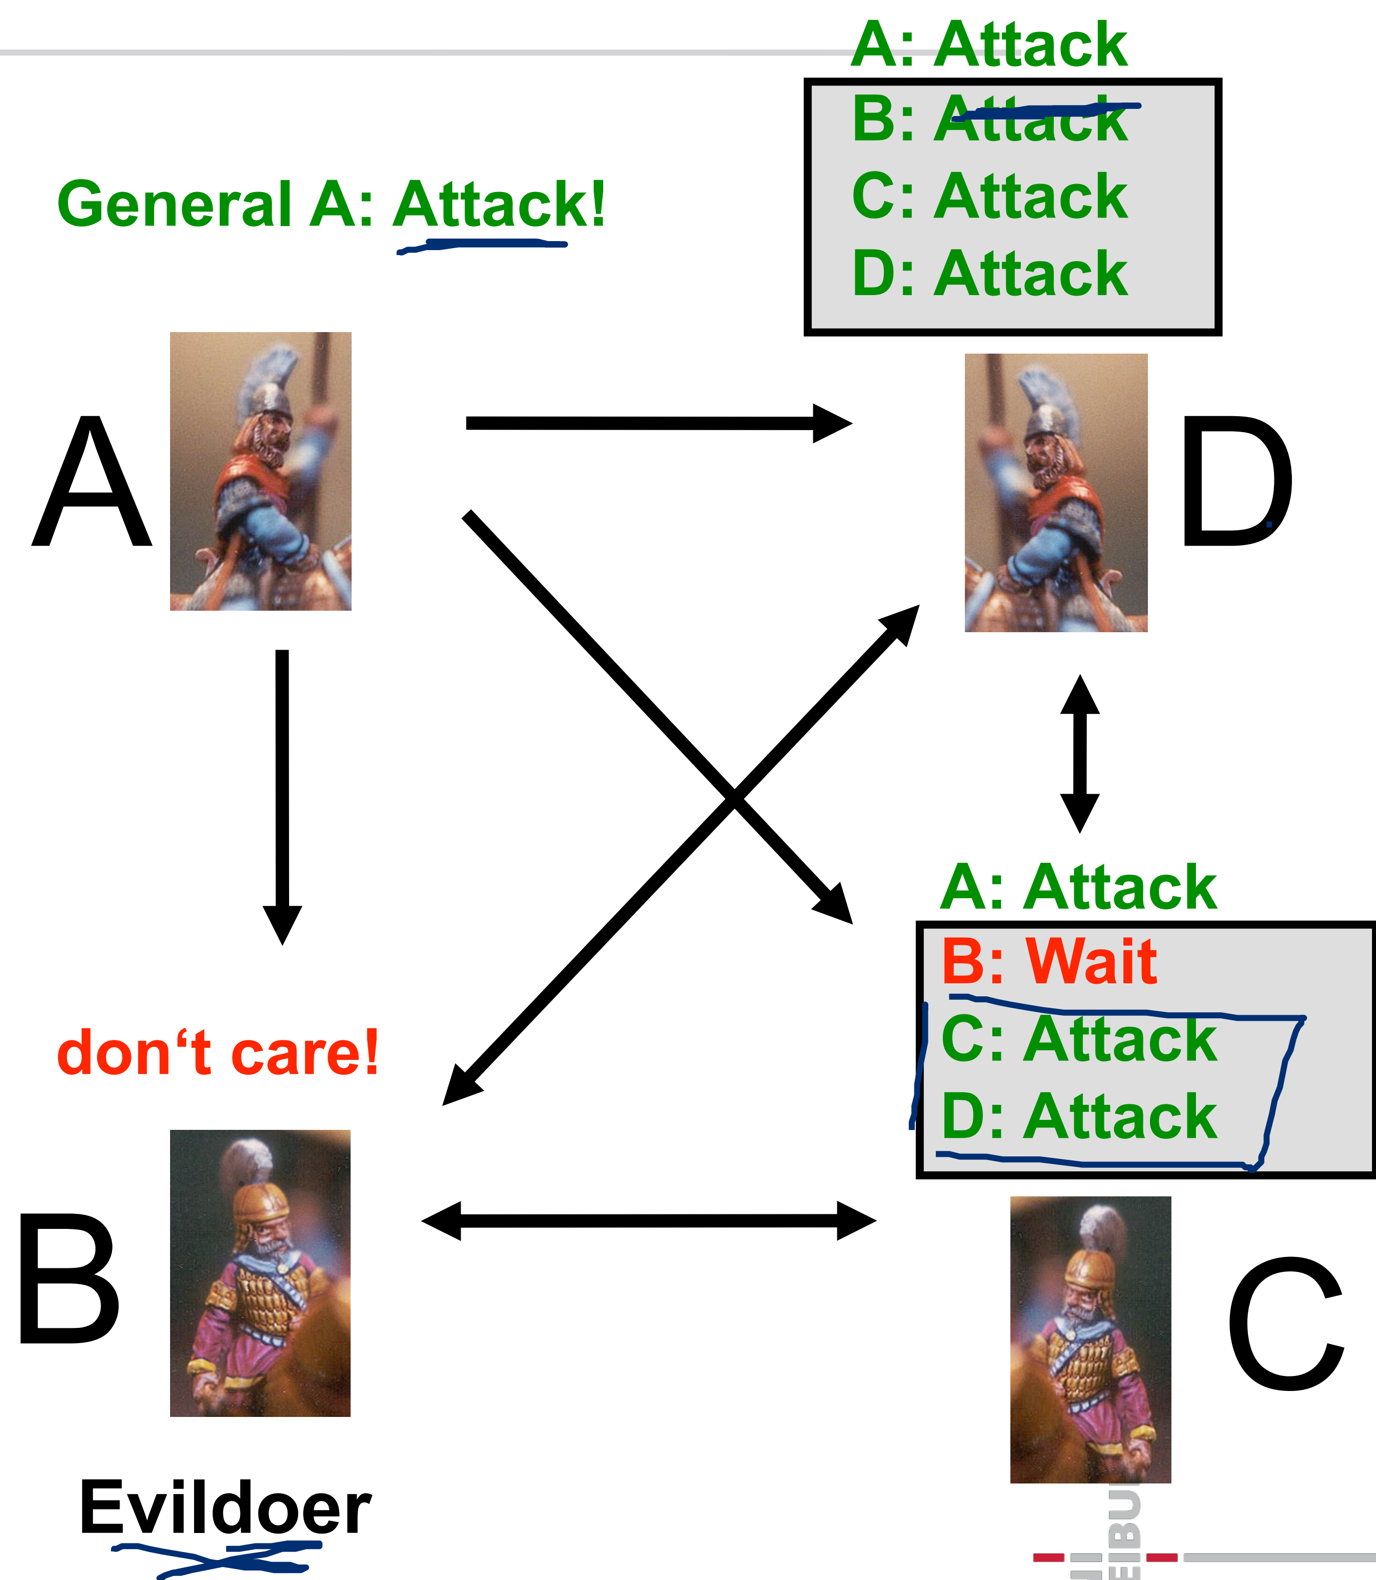

Byzantine Agreement

- Theorem
  - The problem of four byzantine generals can be solved (without cryptography)
- Algorithm
  - General A sends his command to all other generals
    - A sticks to his command if he is honest
  - All other generals forward the received command to all other generals
  - Every generals computes the majority decision of the received commands and follows this command

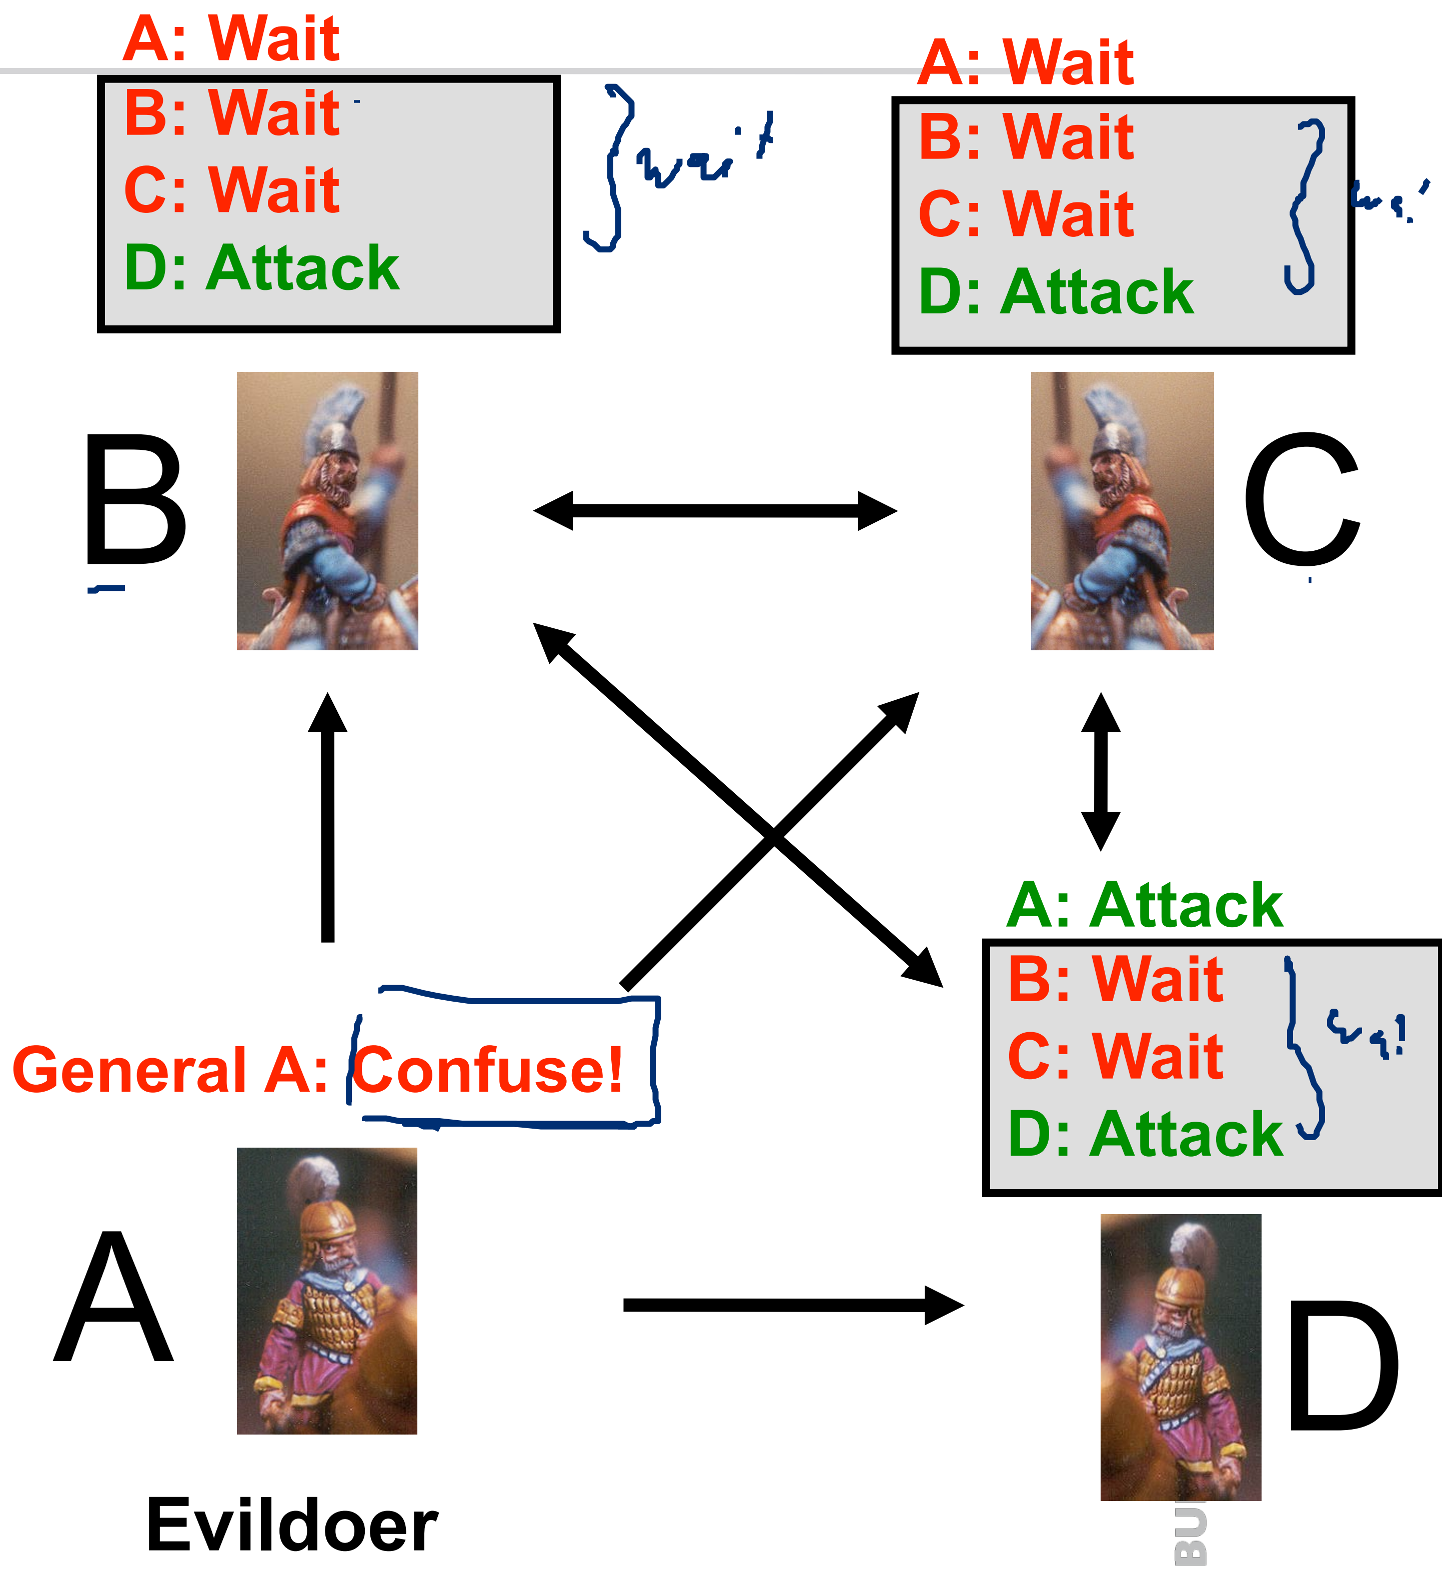

General Solution of Byzantine Agreement

3m + 1

m bad   2m+1 honest

- Theorem
  - If m generals are traitors then 2m+1 generals must be honest to get a Byzantine Agreement
- This bound is sharp if one does not rely on cryptography
- Theorem
  - If a digital signature scheme is working, then an arbitrarily large number of betraying generals can be dealt with
- Solution
  - Every general signs his command
  - All commands are shared together with the signature
  - Inconsistent commands can be detected
  - The evildoer can be exposed

P2P and Byzantine Agreement

- Digital signature can solve the problem of malign peers
- Problem: Number of messages
  - $O(n^2)$ messages in the whole network (for n peers)
- In „Scalable Byzantine Agreement“ von Clifford Scott Lewis und Jared Saia, 2003
  - a scalable algorithm was presented
  - can deal with n/6 evil peers
    - if they do not influence the network structure
  - use only $O(\log n)$ messages per node in the expectation
  - find agreement with high probability

Network of Lewis and Saia

- Butterfly network with clusters of size c log n
  - clusters are bipartite expander graphs
  - Bipartite graph
    - is a graph with disjoint node sets A and B where no edges connect the nodes within A or within B
  - Expander graph
    - A bipartite graph is an expander graph if for each subset X of A the number of neighbors in B is at least c|X| for a fixed constant c>0
    - and vice versa for the subsets in B

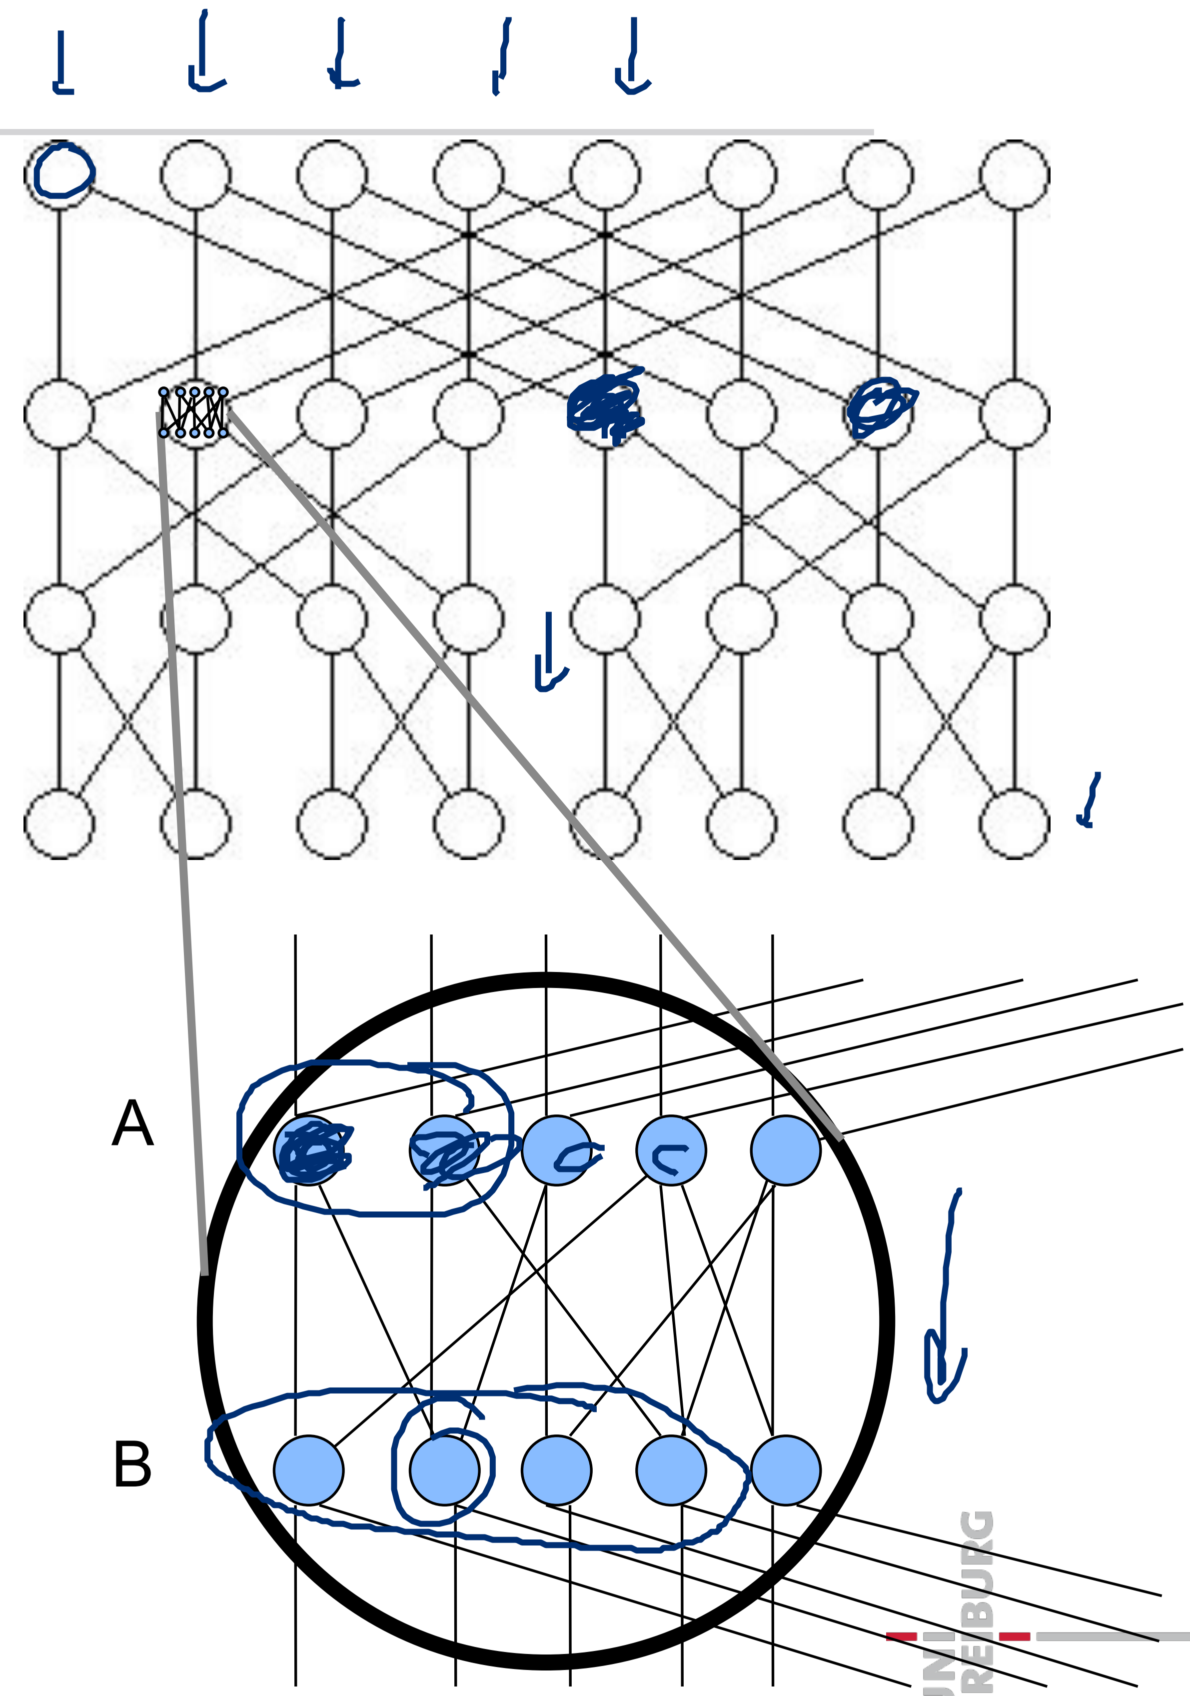

CoNe Freiburg

Discussion

- Advantage
  - Very efficient, robust and simple method
- Disadvantage
  - Strong assumptions
    - The attacker does not know the internal network structure
- If the attacker knows the structure
  - Eclipse attack!

UNI FREIBURG

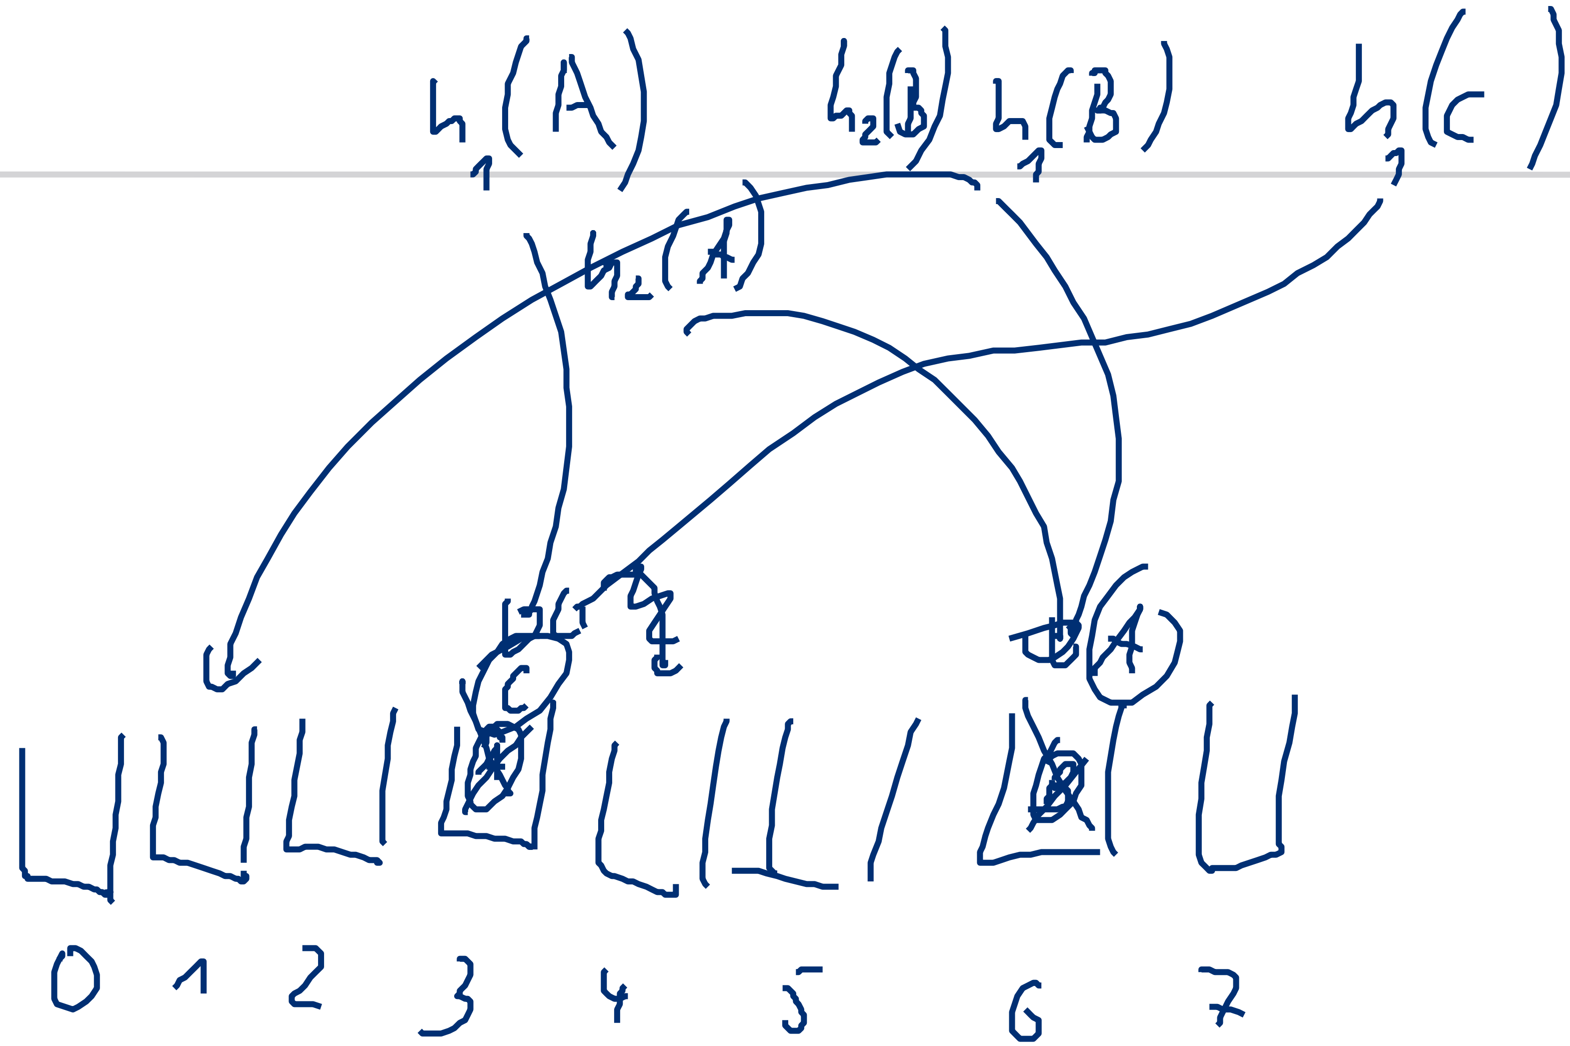
h1(A)
h2(B)
h1(B)
h1(C)
h2(A)
0 1 2 3 4 5 6 7

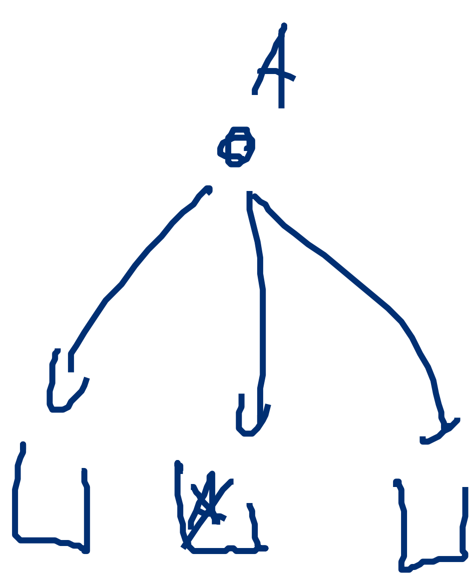
A

CoNe Freiburg

Cuckoo

male

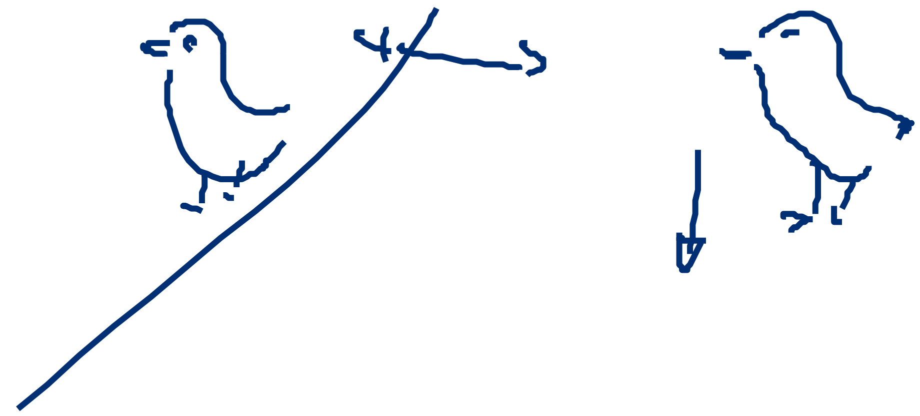

female

complete different bird

UNI FREIBURG

Cuckoo Hashing for Security

- Awerbuch, Scheideler, Towards Scalable and Robust Overlay Networks
- Problem:
  - Rejoin attacks
- Solution:
  - Chord network combined with
  - Cuckoo Hashing
  - Majority condition:
    - honest peers in the neighborhood are in the majority
  - Data is stored with O(log n) copies

Cuckoo Hashing

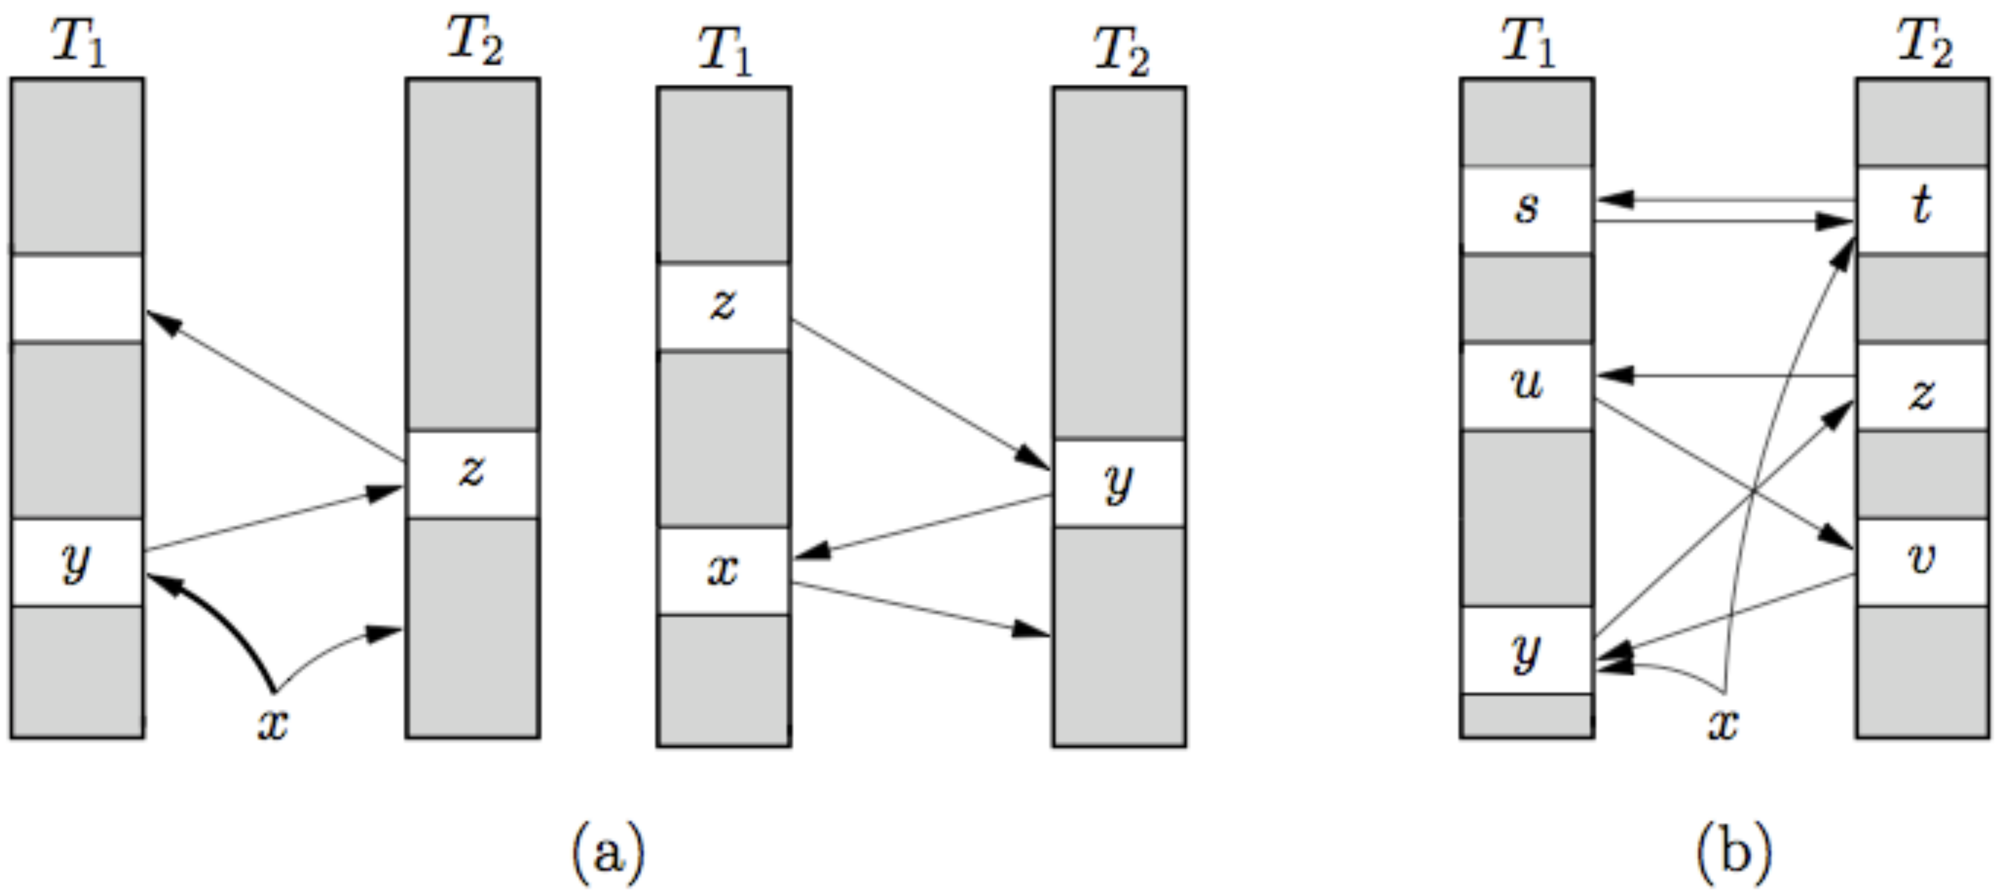

Fig. 1. Examples of CUCKOO HASHING insertion. Arrows show possibilities for moving keys. (a) Key x is successfully inserted by moving keys y and z from one table to the other. (b) Key x cannot be accommodated and a rehash is necessary.

From Cuckoo Hashing
Rasmus Pagh, Flemming Friche Rodler
2004

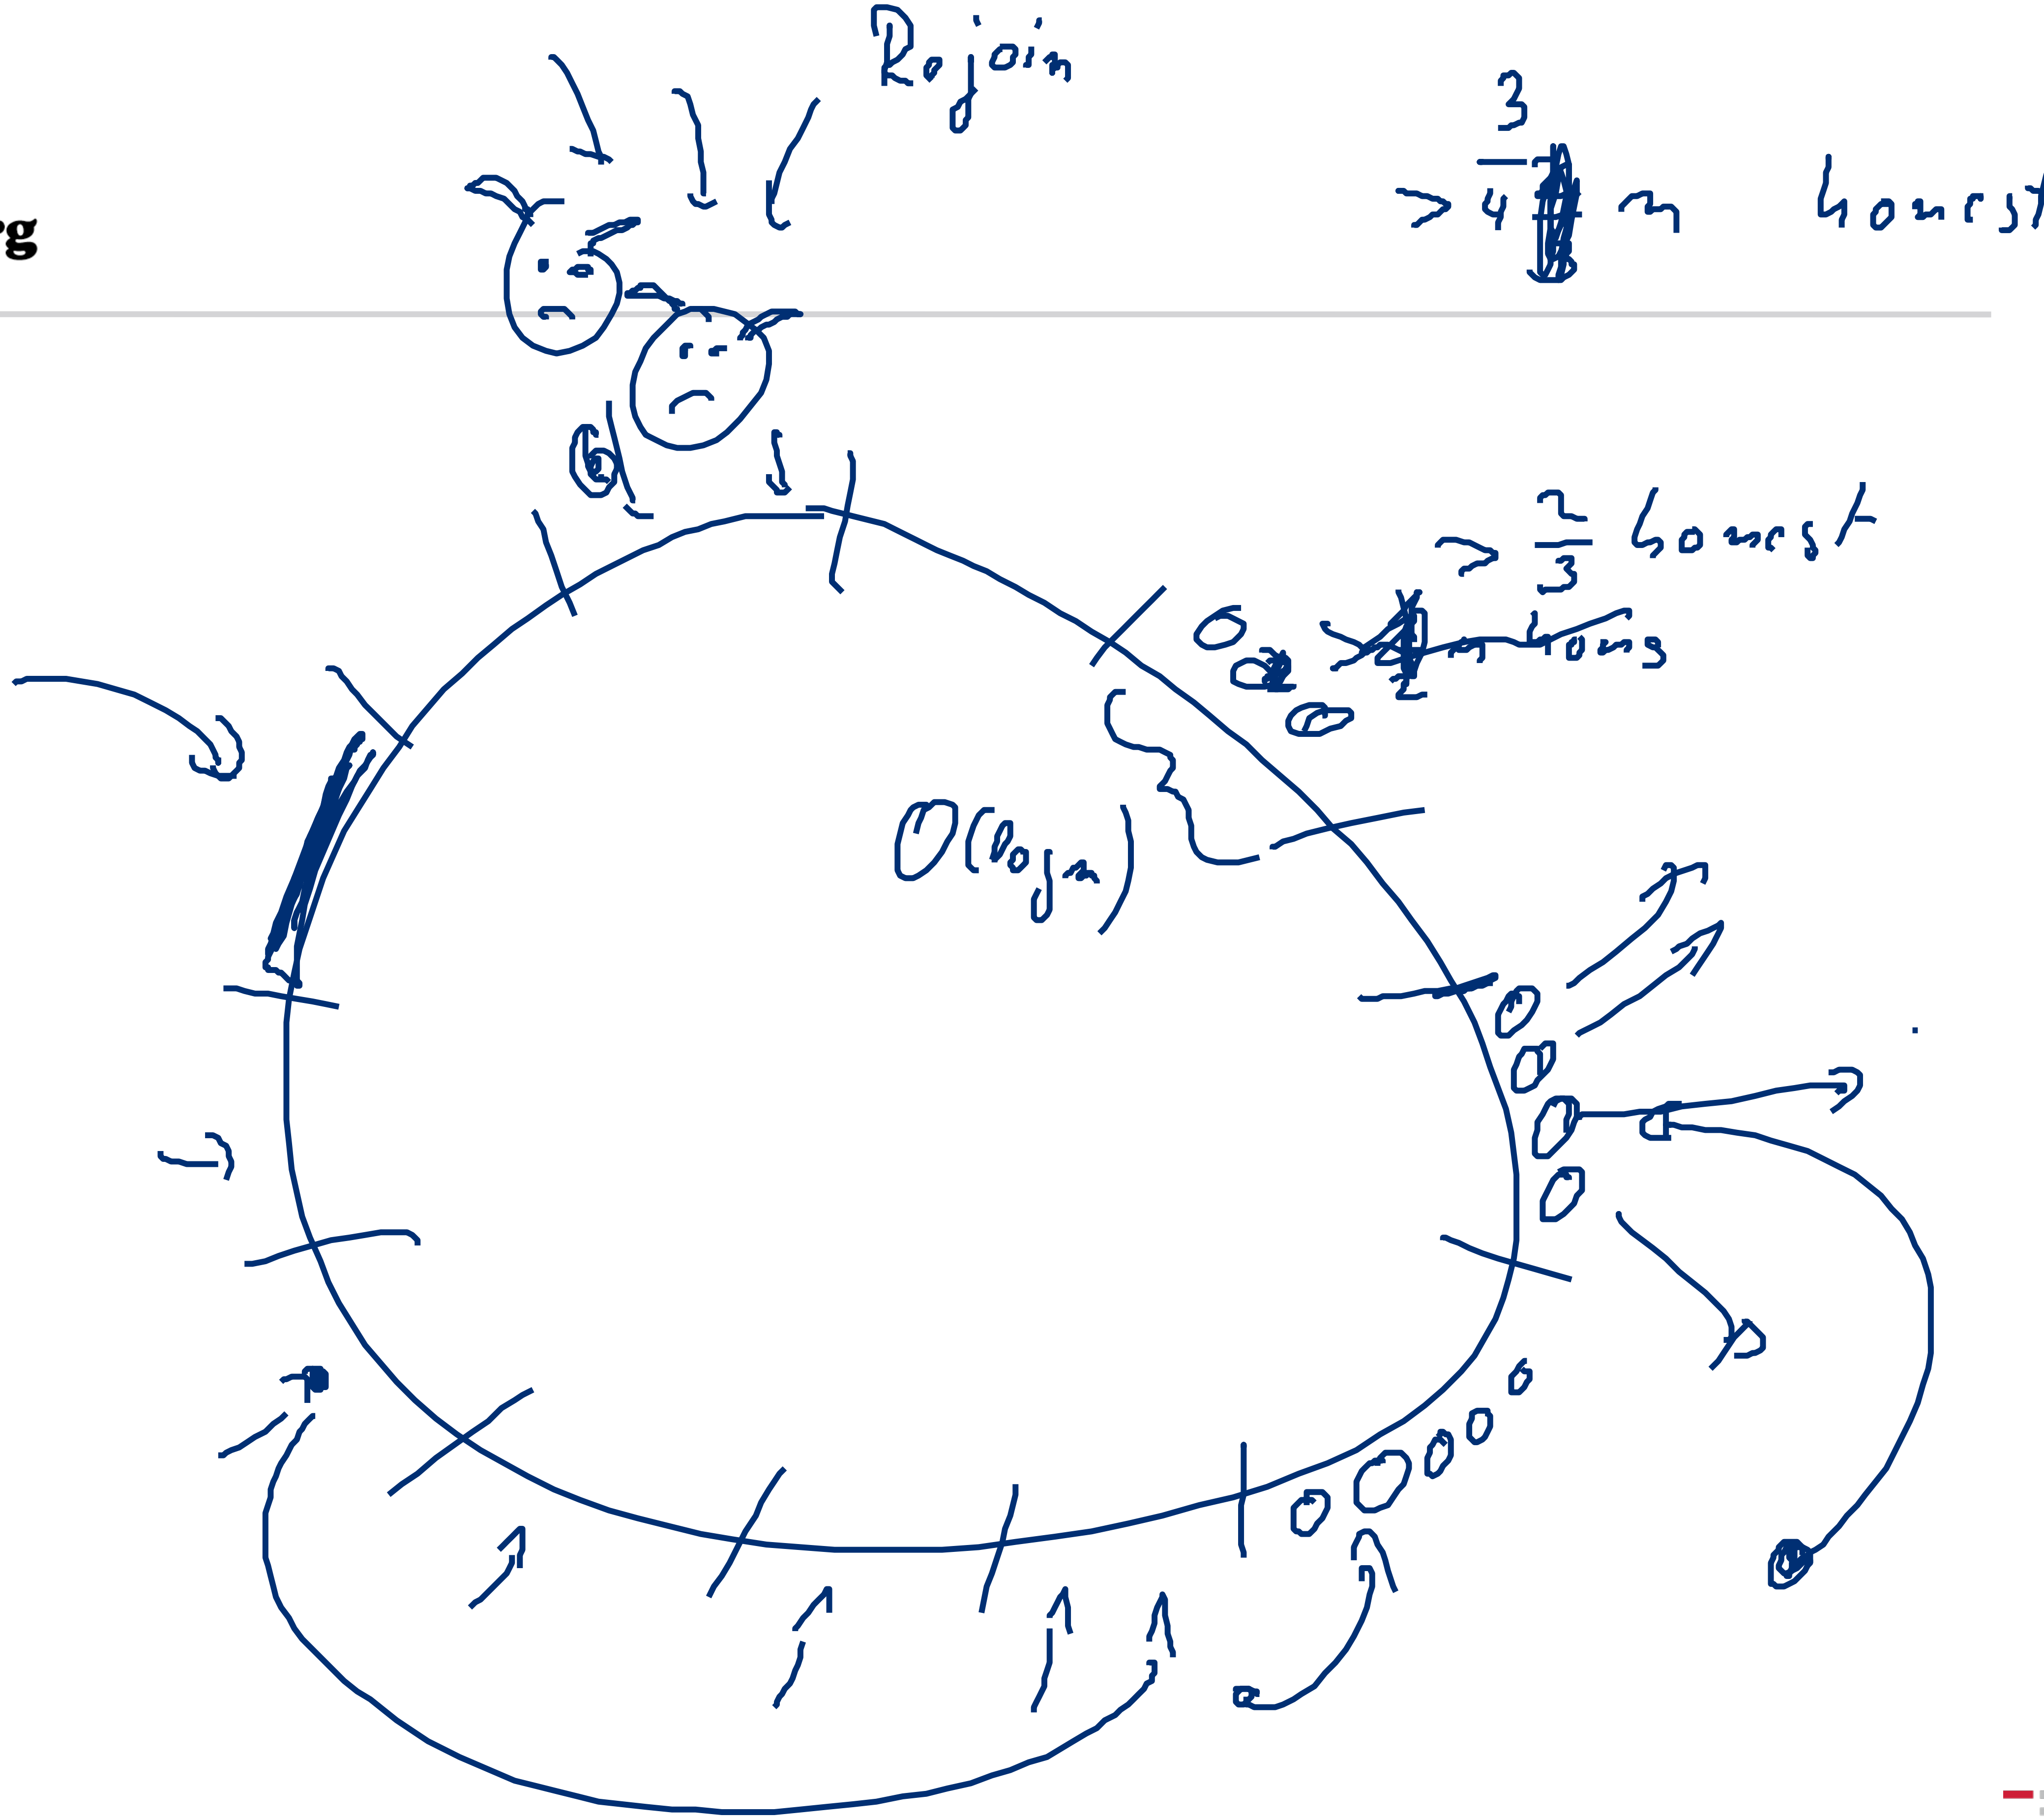

Rejoin
> 3/4 n honest
> 2/3 honest
O(log n)

Efficiency of Cuckoo Hashing

- Theorem
  - Let $\epsilon>0$ then if at most n elements are stored, then Cuckoo Hashing needs a hash space of $2n+\epsilon$.
- Three hash functions increase the load factor from 1/2 to 91%
- Insert
  - needs O(1) steps in the expectation
  - O(log n) with high probability
- Lookup
  - needs two steps

Chord

- Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek and Hari Balakrishnan (2001)
- Distributed Hash Table
  - range $\{0,..,2^m-1\}$
  - for sufficient large m
- for this work the range is seen as [0,1)
- Network
  - ring-wise connections
  - shortcuts with exponential increasing distance

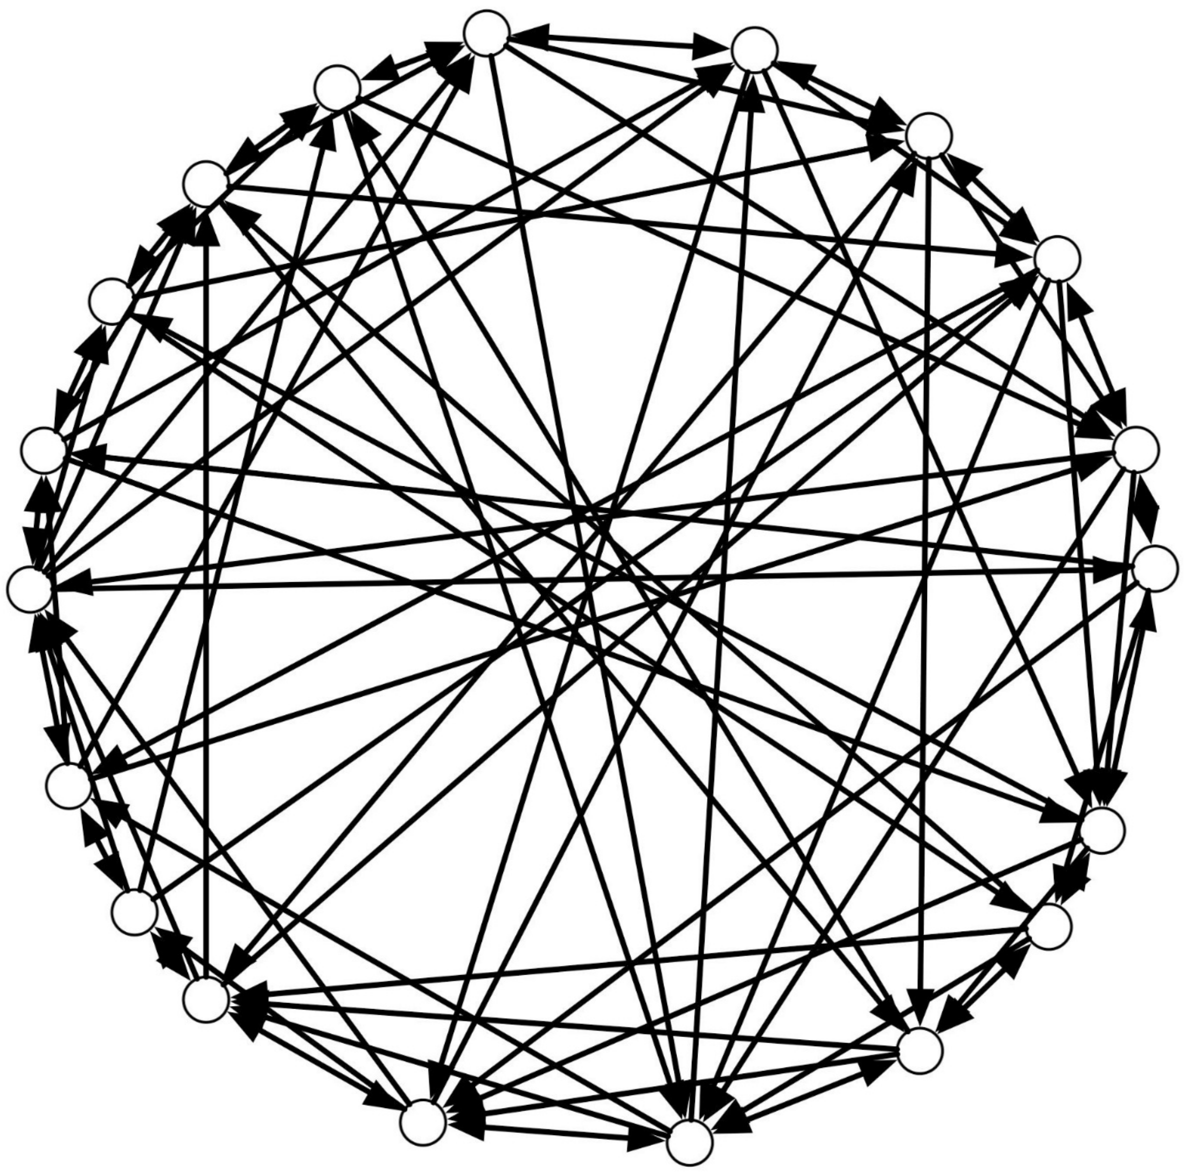

Lookup in Chord

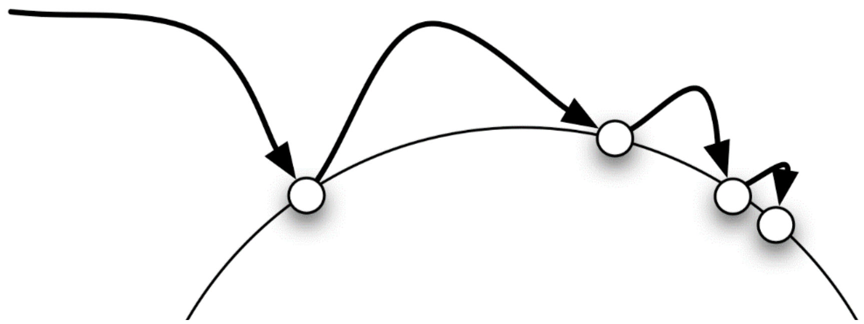

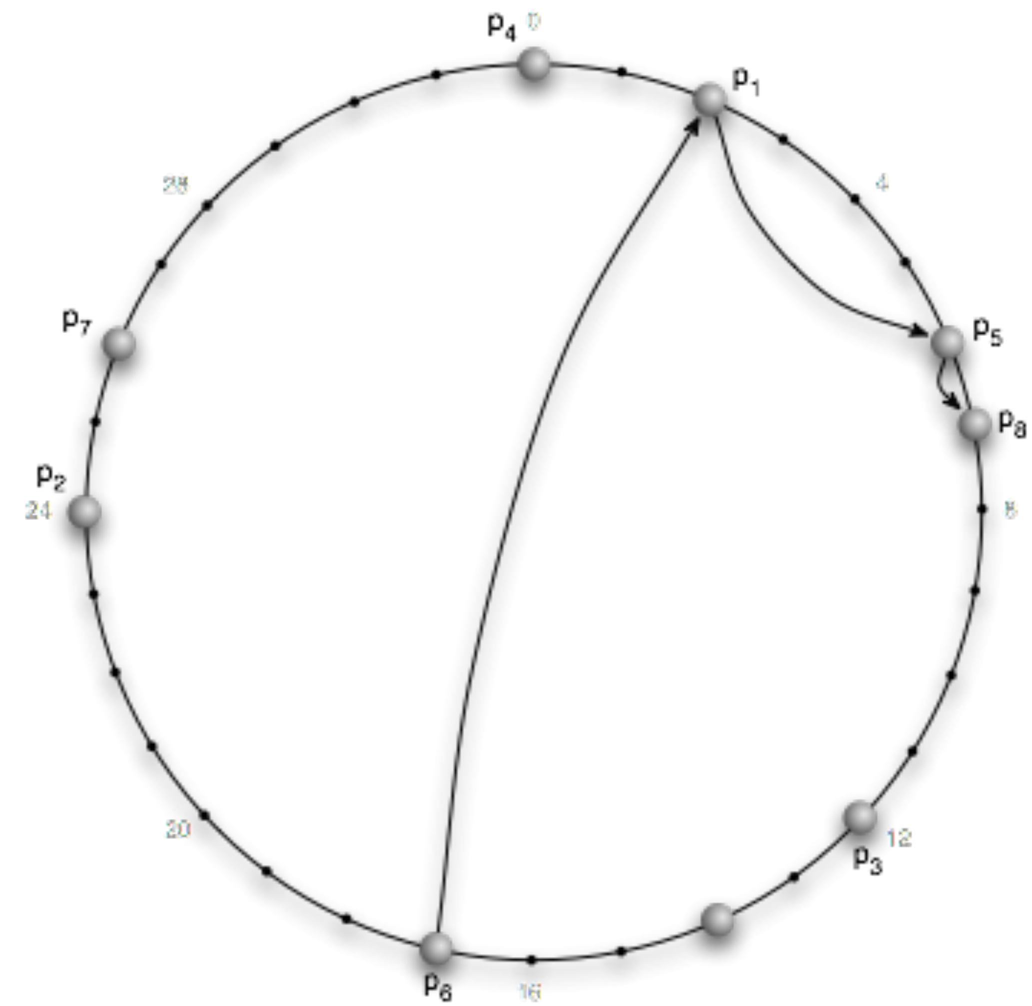

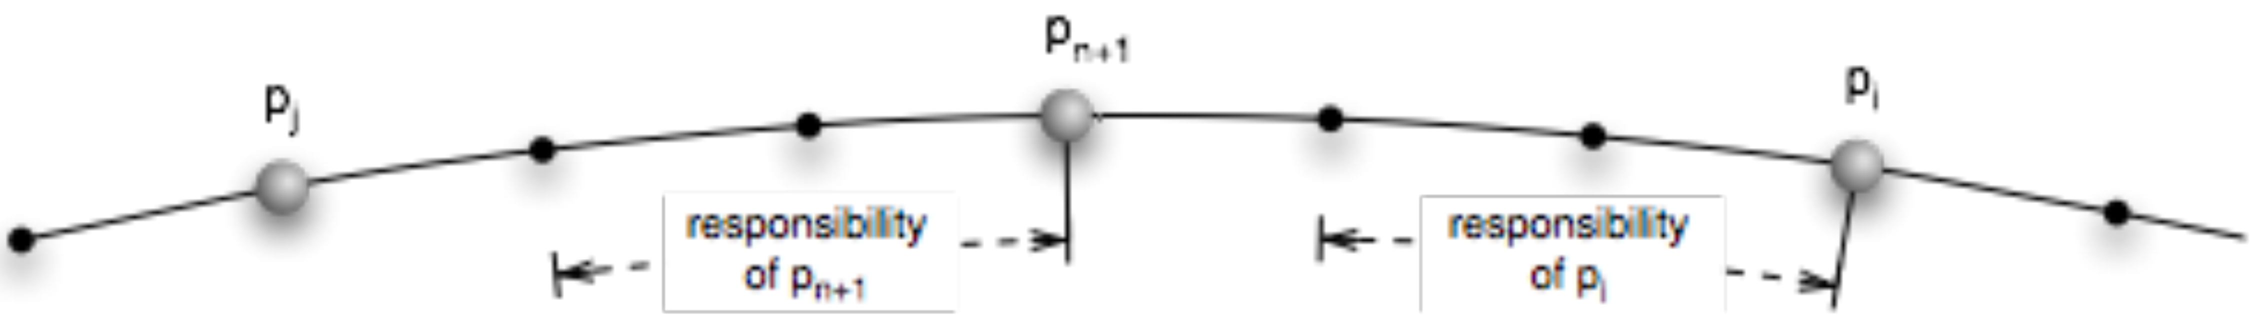

Data Structure of Chord

- For each peer
  - successor link on the ring
  - predecessor link on the ring
  - for all $i \in \{0,..,m-1\}$
    - Finger[i] := the peer following the value $r_V(b+2^i)s$
- For small i the finger entries are the same
  - store only different entries
- Chord
  - needs O(log n) hops for lookup
  - needs $O(\log^2 n)$ messages for inserting and erasing of peers

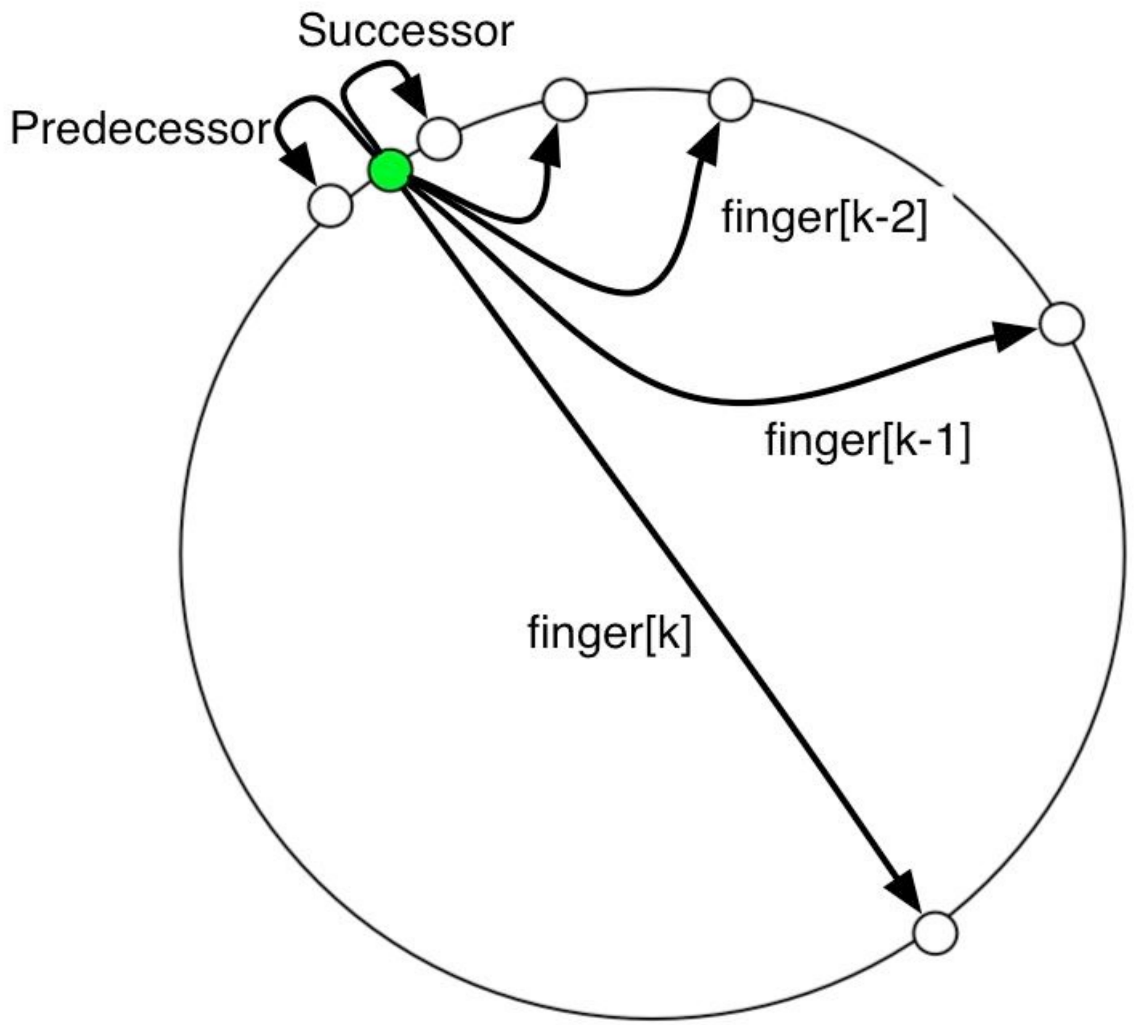

Cuckoo Hashing for Security

- Given n honest peers and ϵ n dishonest peers
- Goal
  - For any adversarial attack the following properties for every interval $I \subseteq [0, 1)$ of size at least (c log n)/n we have
  - Balancing condition
    - I contains $\Theta(|I| \cdot n)$ nodes
  - Majority condition
    - the honest nodes in I are in the majority
- Then all majority decisions of O(log n) nodes give a correct result

Rejoin Attacks

- Secure hash functions for positions in the Chord
  - if one position is used
  - then in an O(log n) neighborhood more than half is honest
  - if more than half of al peers are honest
- Rejoin attacks
  - use a small number of attackers
  - check out new addresses until attackers fall in one interval
  - then this neighborhood can be ruled by the attackers

The Cuckoo Rule for Chord

- Notation
  - a region is an interval of size $1/2^r$ in [0, 1) for some integer r that starts at an integer multiple of $1/2^r$
  - There are exactly 2^r regions
  - A k-region is a region of size (closest from above to) k/n, and for any point $x \in [0, 1)$
  - the k-region $R_k(x)$ is the unique k-region containing x.
- Cuckoo rule
  - If a new node v wants to join the system, pick a random $x \in [0, 1)$.
  - Place v into x and move all nodes in $R_k(x)$ to points in [0, 1) chosen uniformly at random
    - (without replacing any further nodes).
- Theorem
  - For any constants ϵ and k with $\epsilon < 1-1/k$, the cuckoo rule with parameter k satisfies the balancing and majority conditions for a polynomial number of rounds, with high probability, for any adversarial strategy within our model.
  - The inequality $\epsilon < 1 - 1/k$ is sharp

Peer-to-Peer Networks

14 Security

Christian Schindelhauer
Technical Faculty
Computer-Networks and Telematics
University of Freiburg