

Gozar: NAT-friendly Peer Sampling with One-Hop Distributed NAT Traversal

Amir H. Payberah – Jim Dowling – Seif Haridi
{amir,jdowling,seif}@sics.se




Introduction

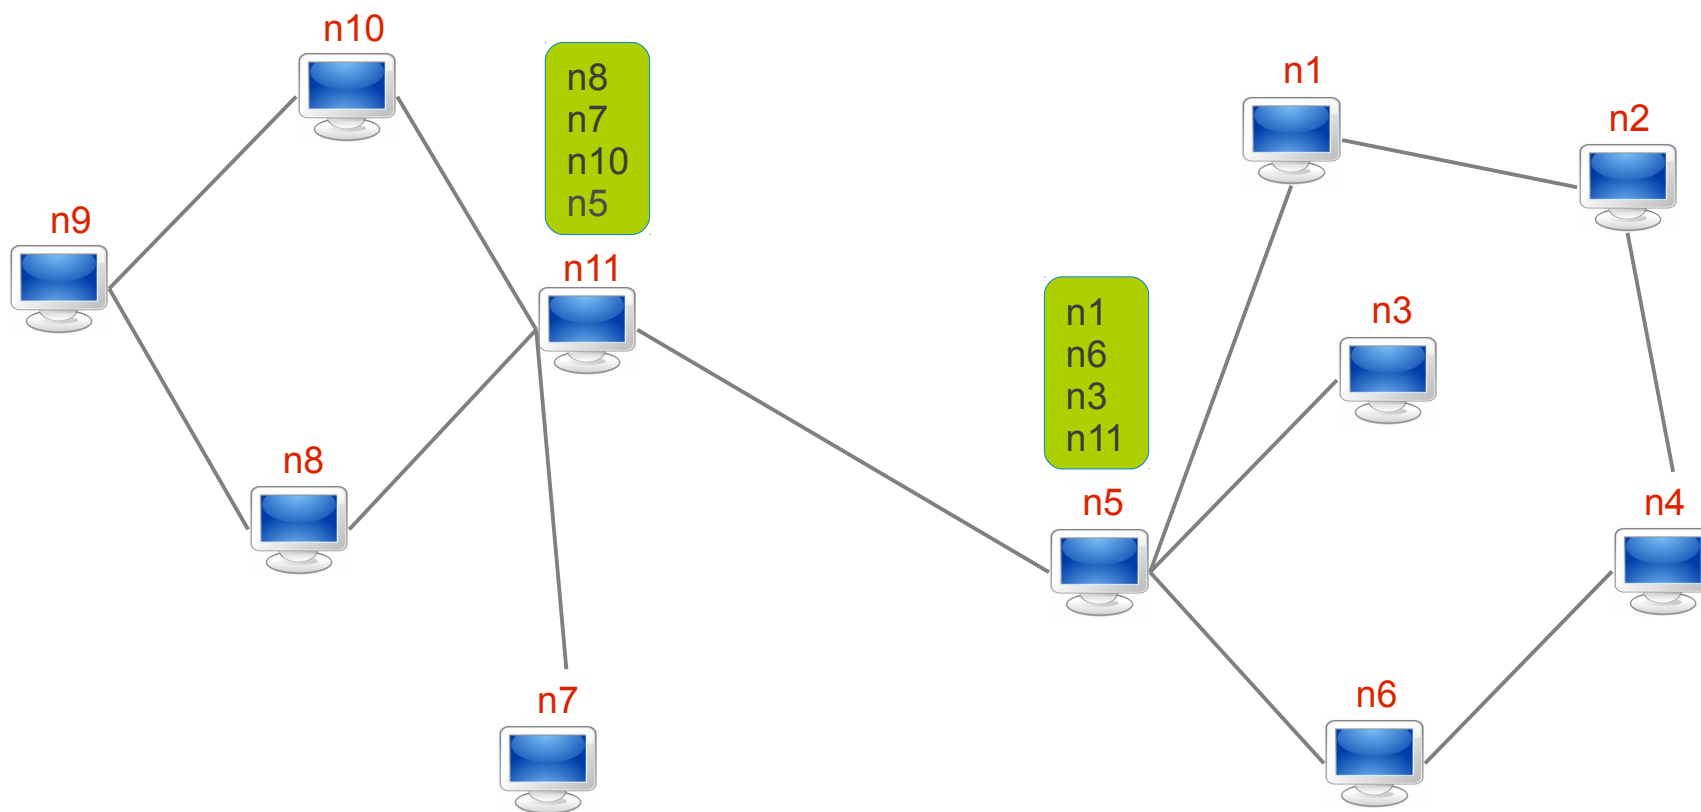
Gossip-based Peer Sampling Service

- Gossip-based peer sampling services (**PSS**) have been widely used in large scale distributed applications.
 - Information dissemination
 - Aggregation
 - Overlay topology management
- A **PSS** provides a node with a **uniform random sample** of live nodes from all nodes in the system (**partial view**).


Gossip-based Peer Sampling Protocol (1/7)




Gossip-based Peer Sampling Protocol (2/7)




Gossip-based Peer Sampling Protocol (3/7)




Gossip-based Peer Sampling Protocol (4/7)




Gossip-based Peer Sampling Protocol (5/7)



Gossip-based Peer Sampling Protocol (6/7)




Gossip-based Peer Sampling Protocol (7/7)



Gossip-based Peer Sampling Design Space

- Peer Selection
- View Exchange
- View Merge



Gossip-based Peer Sampling Design Space

- Peer Selection


- Random
- Tail

- View Exchange

- Push
- Push-Pull


- View Merge

- Blind
- Healer
- Swapper

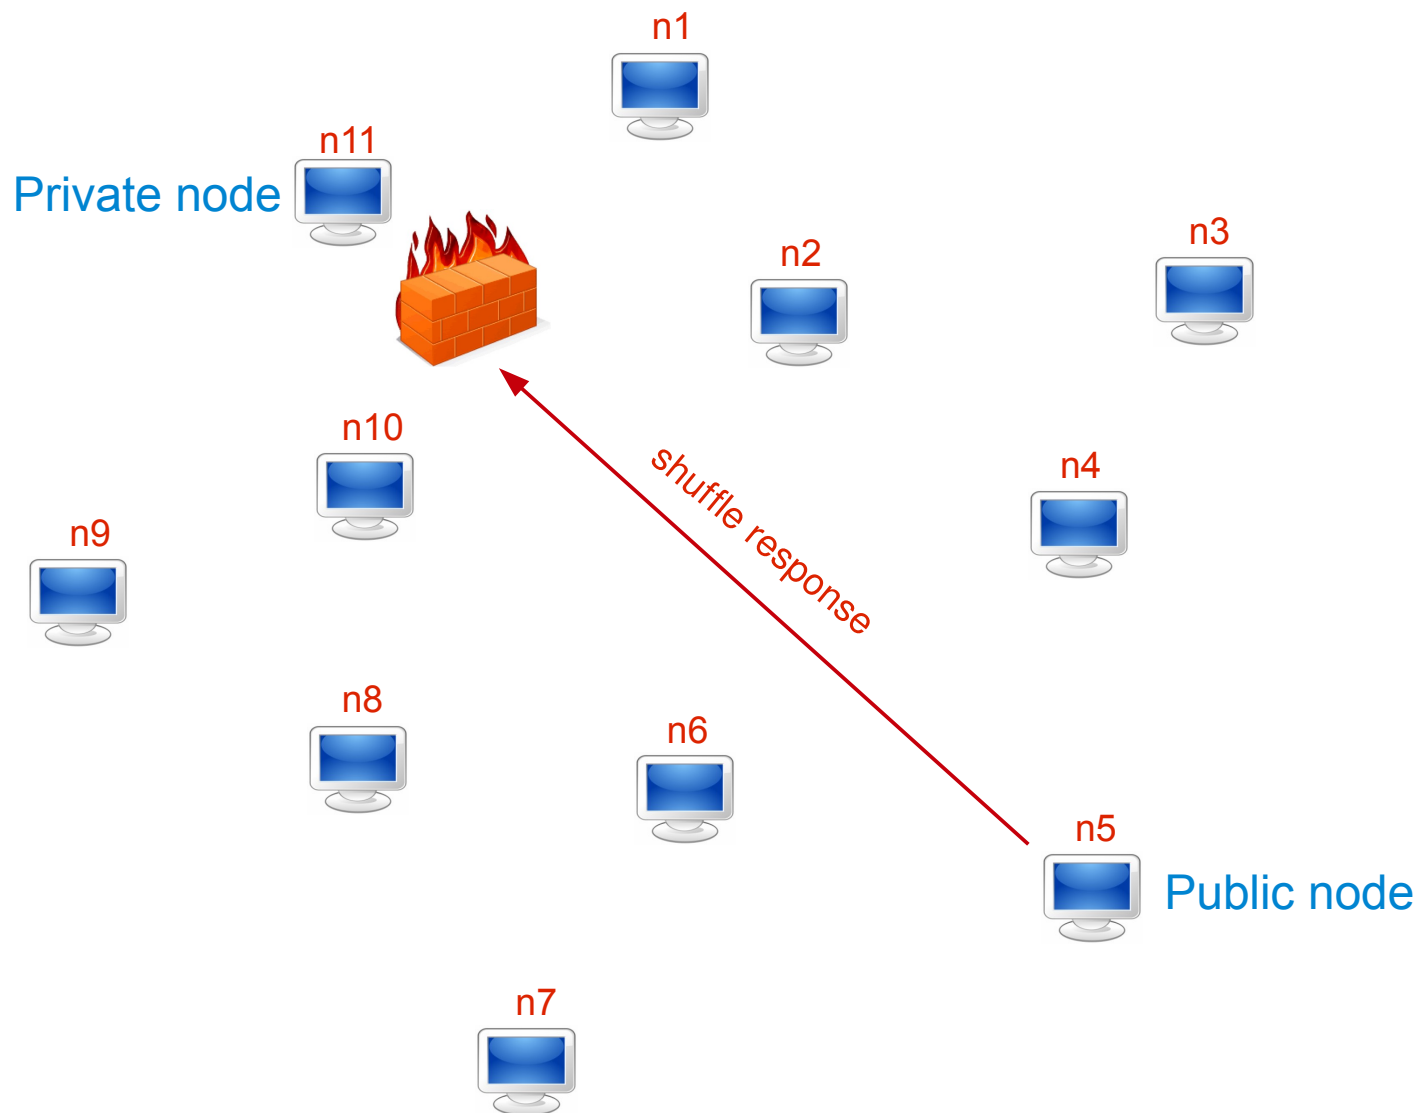


Problem Description


NAT Environments (1/4)




NAT Environments (1/4)



NAT Environments (1/4)




NAT Environments (1/4)



Impact of NATs on PSS' (1/2)


- Size of the biggest cluster for an increasing percentage of NATs.



[A.M.Kermarrec – ICDCS'09]

Impact of NATs on PSS' (2/2)


- Ratio of non-stale references to private nodes.



[A.M.Kermarrec – ICDCS'09]


Solutions for Communicating with Private Nodes (1/3)

- Relay communications to the private node using a **public relay node**.




Solutions for Communicating with Private Nodes (2/3)

- Use a NAT **hole-punching** algorithm to establish a direct connection to the private node using a **public rendezvous node**.



Solutions for Communicating with Private Nodes (3/3)

- Route the request to private nodes using **chains of existing open connections** (of unbounded length).



Research Challenges for a NAT-friendly PSS

- Assuming **distributed relay** and **rendezvous services**, how do nodes discover **which public nodes** act as partners for the private nodes?
- Is **hole-punching** or **relaying** preferable for communicating with private nodes?
 - **How much data** will be sent over the connection and **what are the latency requirements**?
- How **fairly** should the gossiping **load** be balanced over public versus private nodes?

Gozar – NAT friendly Peer Sampling Service

Design Space

- Peer Selection
 - Rand
 - Tail

- View Propagation
 - Push
 - Push-Pull

- View Selection
 - Blind
 - Healer
 - Swapper

Design Space

- Peer Selection

- Rand

- Tail

- View Propagation

- Push

- Push-Pull

- View Selection

- Blind

- Healer

- Swapper

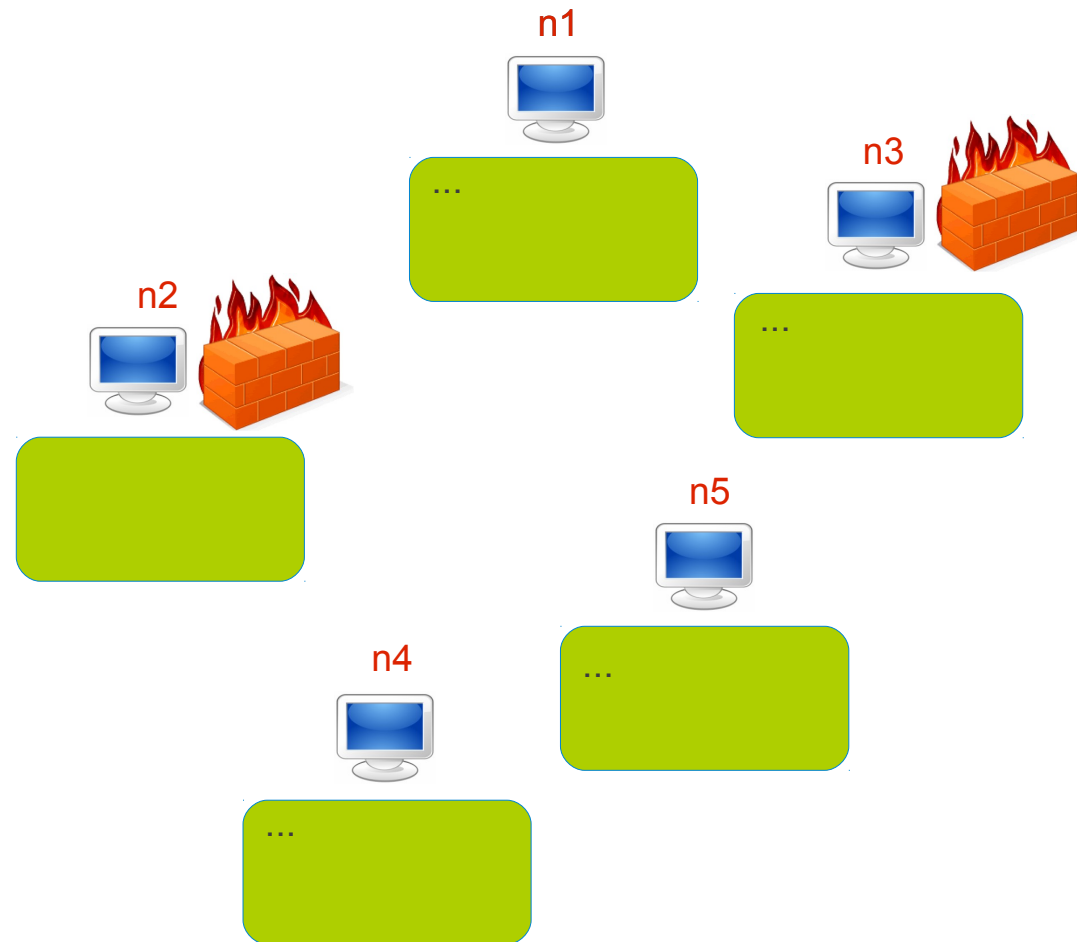
Gozar

NAT Friendliness in Gozar

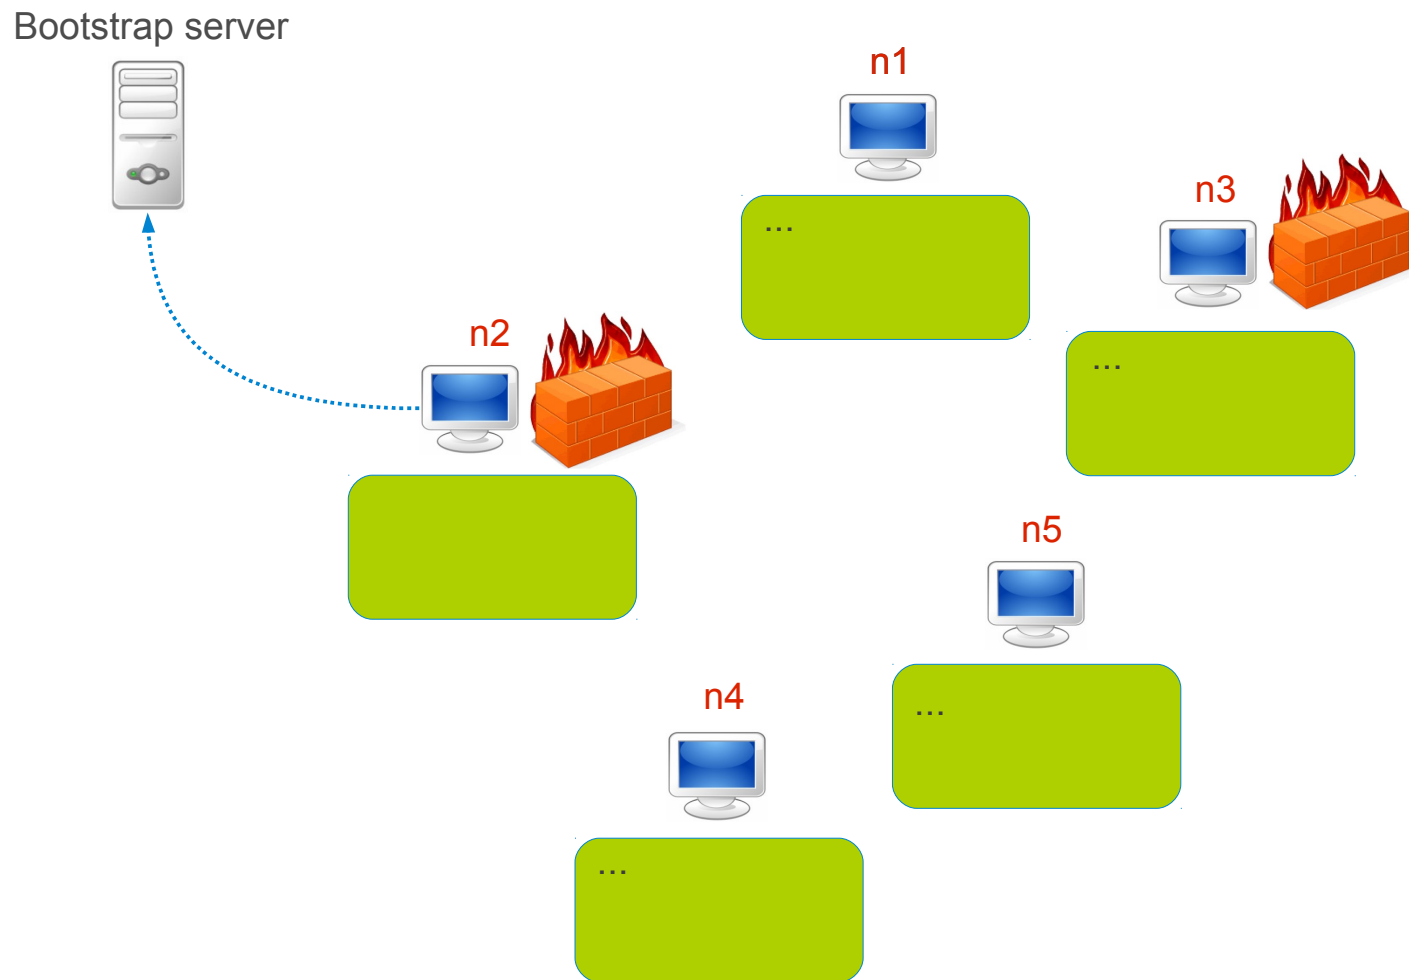
- In Gozar, each **private** node connects to one or more **public** nodes, called **partners** that act as a **relay** or **rendezvous server** on behalf of the private node.
- A **node's descriptor** consists of both its **own address**, its **NAT type**, and its **partners' addresses** at the time of descriptor creation.
- When a node wants to gossip with a private node, it uses the partner addresses in its descriptor to communicate with the private node.

Partnering (1/10)


Bootstrap server



Partnering (2/10)




Partnering (3/10)



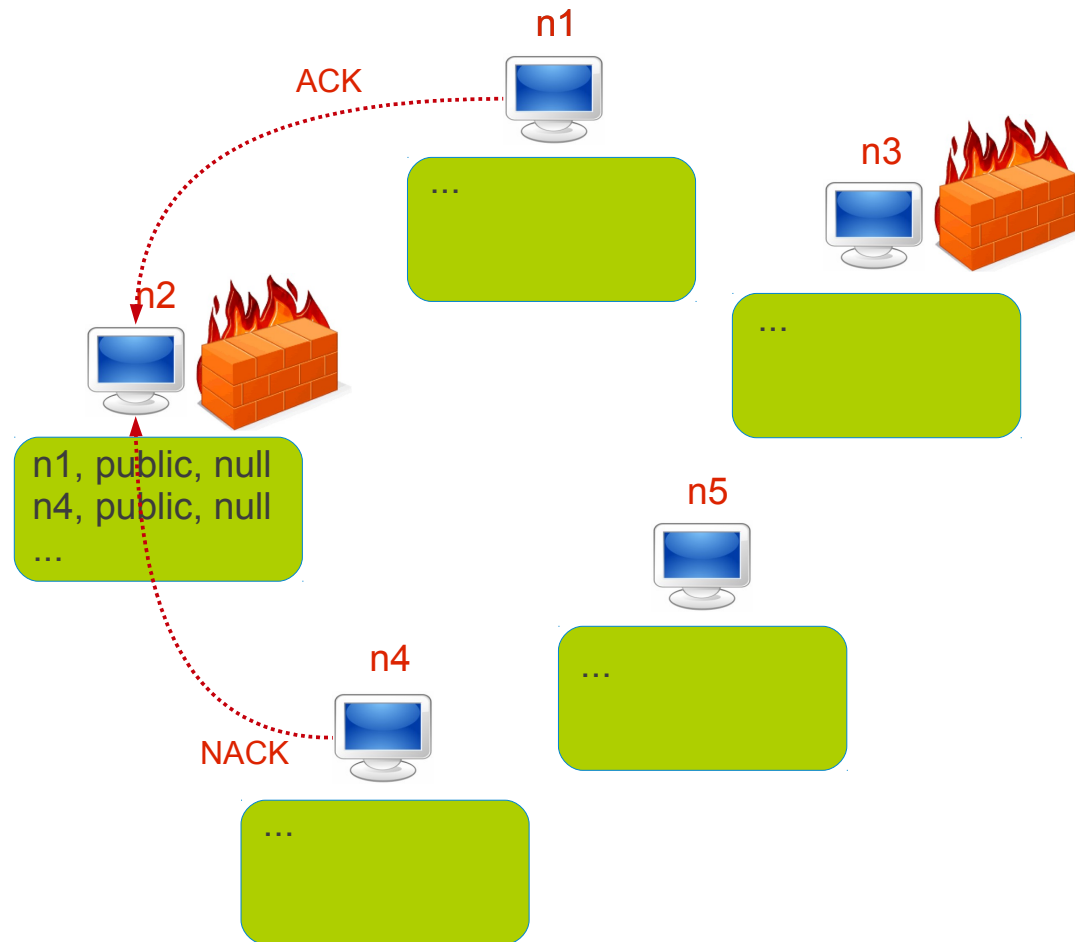
Partnering (4/10)

Bootstrap server




Partnering (5/10)

Bootstrap server

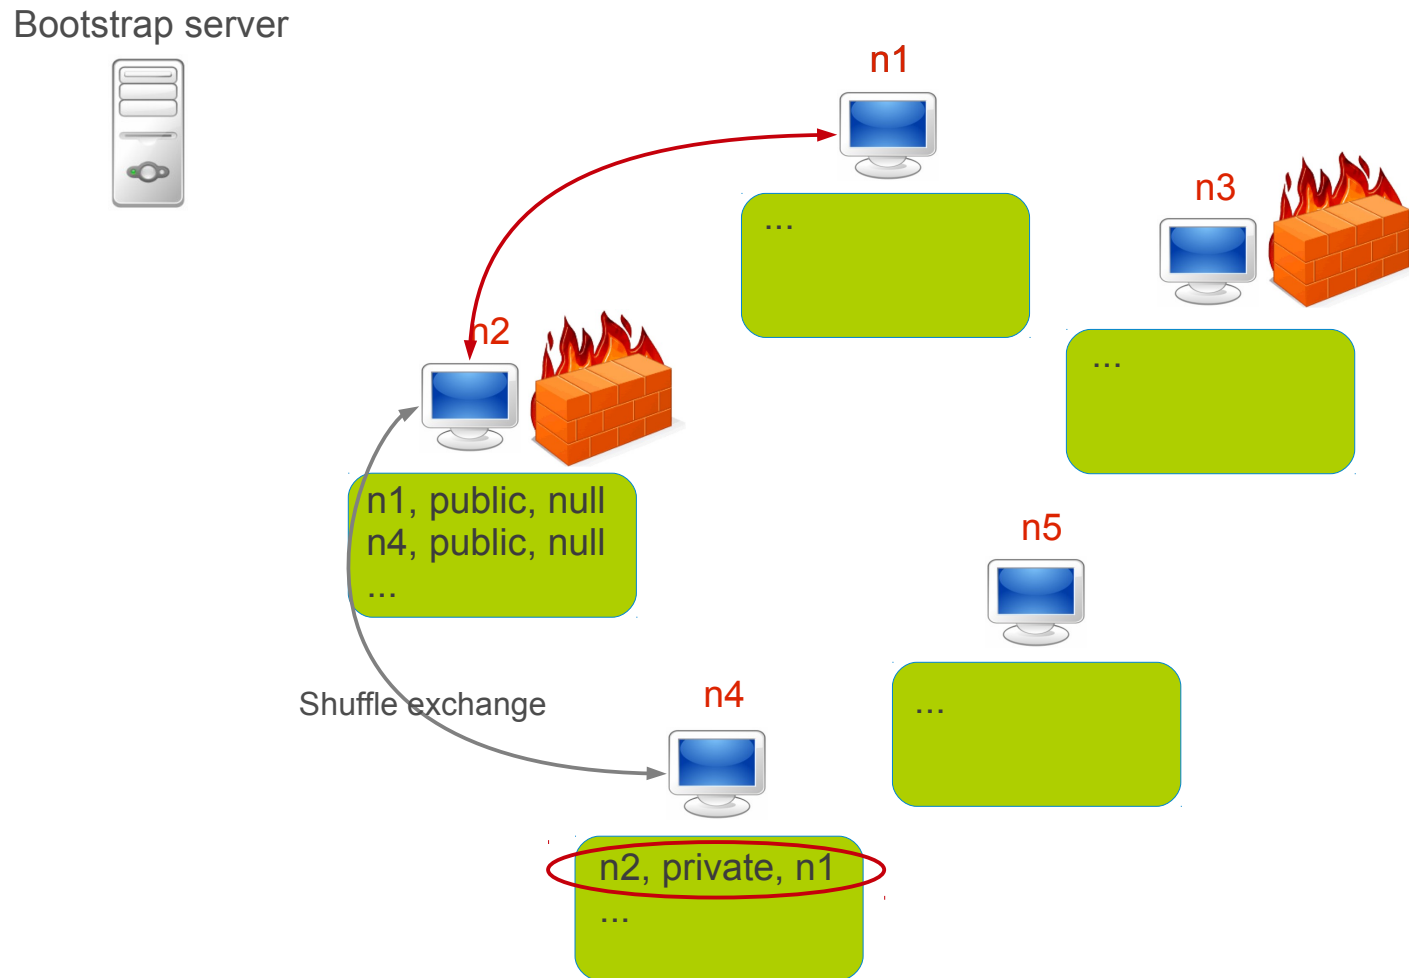


Partnering (6/10)

Bootstrap server

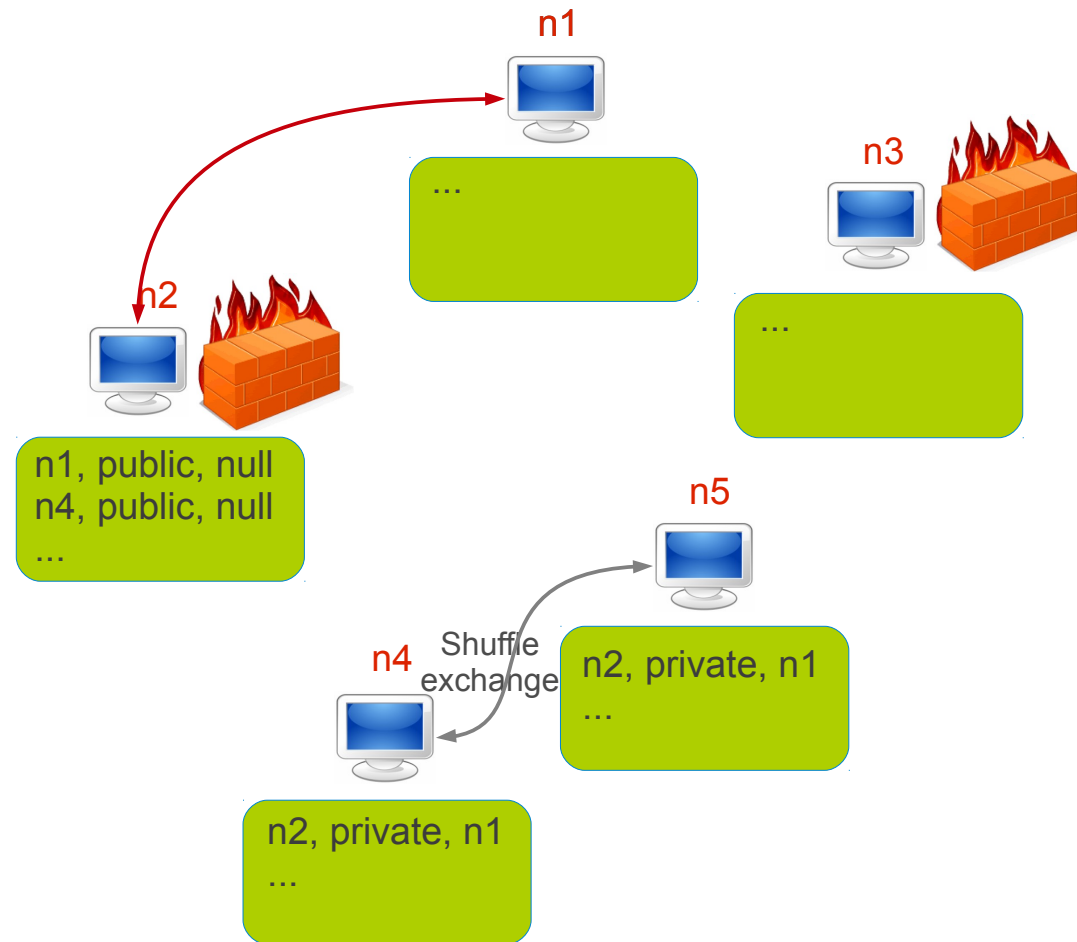


Partnering (7/10)



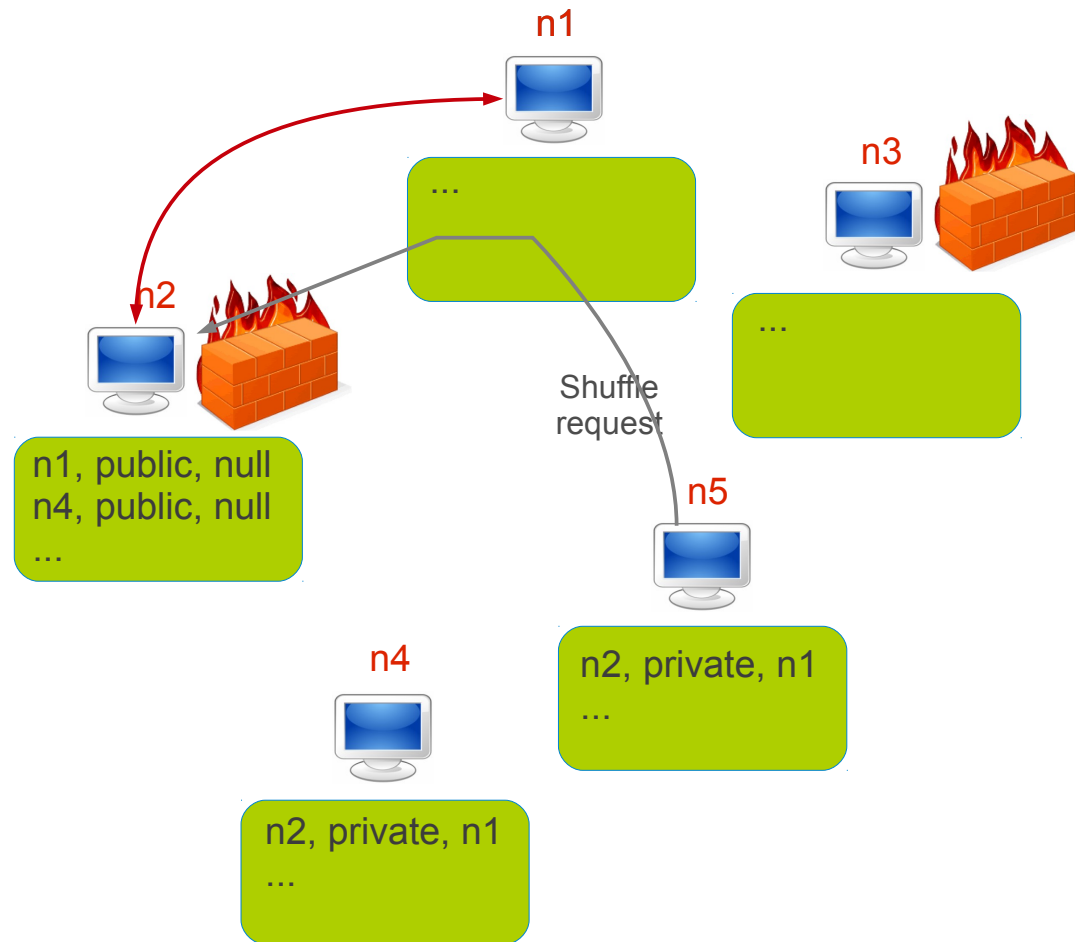
Partnering (8/10)

Bootstrap server



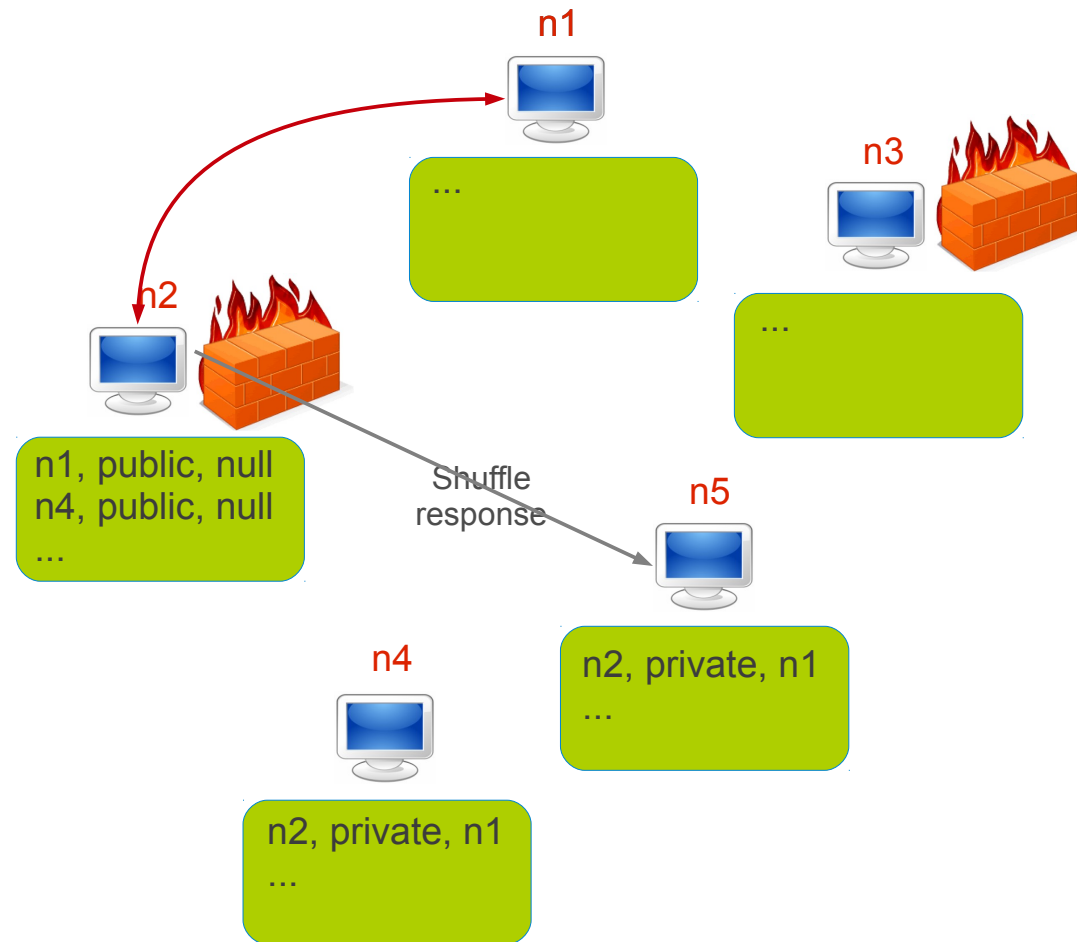
Partnering (9/10)

Bootstrap server



Partnering (10/10)

Bootstrap server



Relaying or Hole Punching?

- Relaying?
 - Lower latency message exchange.
 - Enables lower gossip cycle periods.
 - Necessary in dynamic networks
- Hole punching?
 - Decreases load on public nodes.
 - But not if shuffle messages are small.
- Applications that use the PSS can use partners to hole-punch or relay messages to private nodes.

Experiments

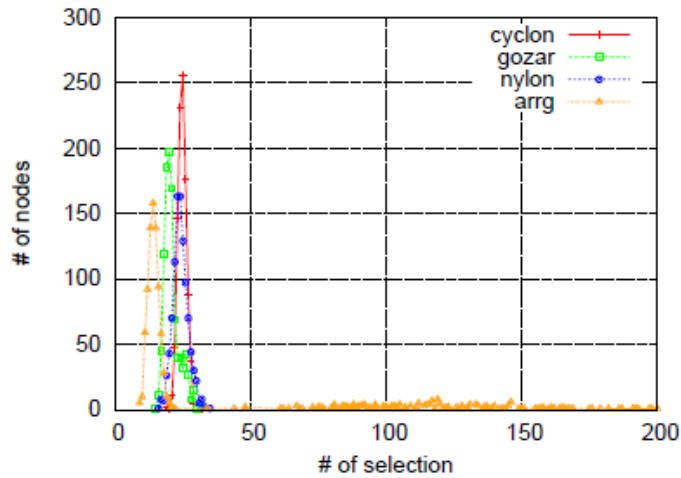
Experiment Setup

- Using the [Kompics](#) as a simulator platform.
- [King dataset](#) is used to model the latencies between nodes.
- [1000](#) nodes, [80%](#) of nodes are private and [20%](#) are public.
- Compare with [Nylon](#) and [ARRG](#).
- [Cyclon](#) is used as a baseline.

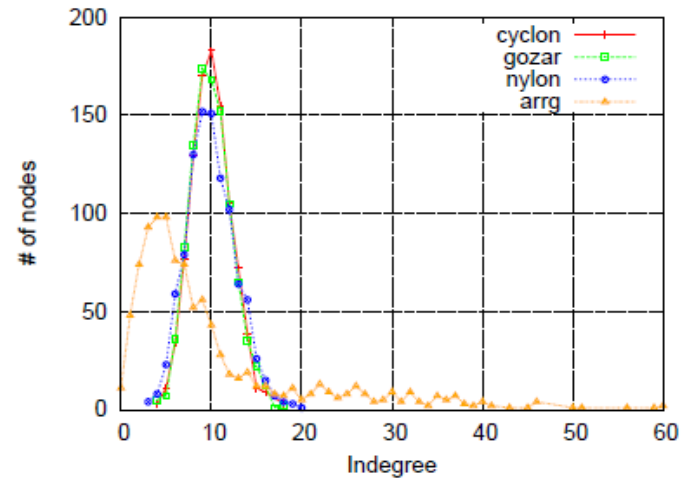
Metrics

- **Randomness** properties:
 - Local randomness
 - In-degree distribution
 - Clustering coefficient
 - Avg. path length
- Protocol **overhead**.
- **Fairness** and **connectivity** in catastrophic failure.

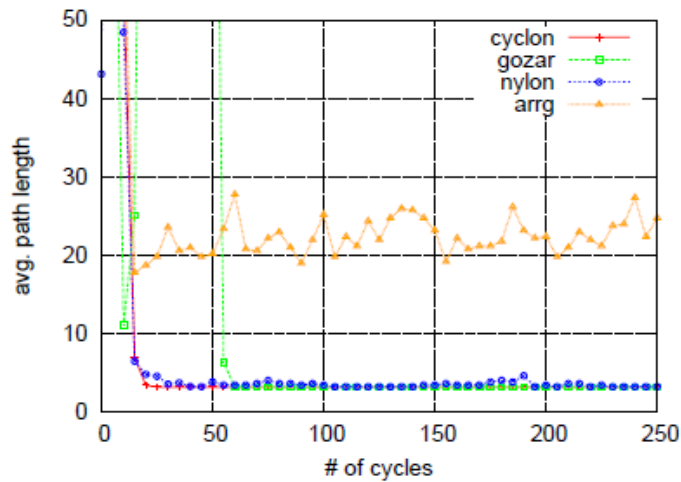
Randomness



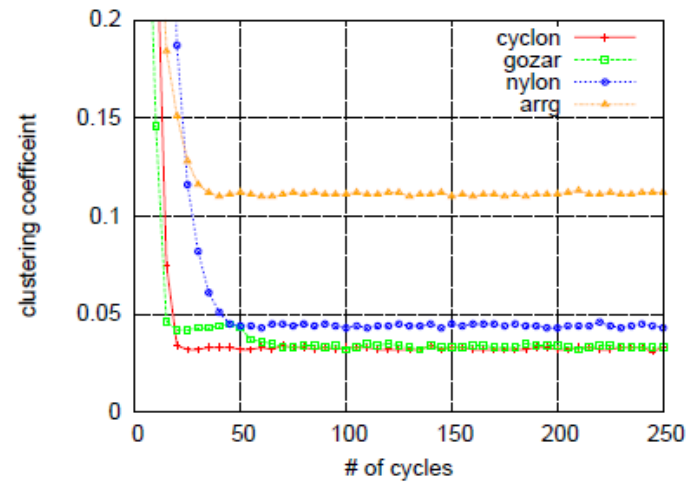
(a) Local randomness.



(b) Indegree distribution.




(c) Average path length.

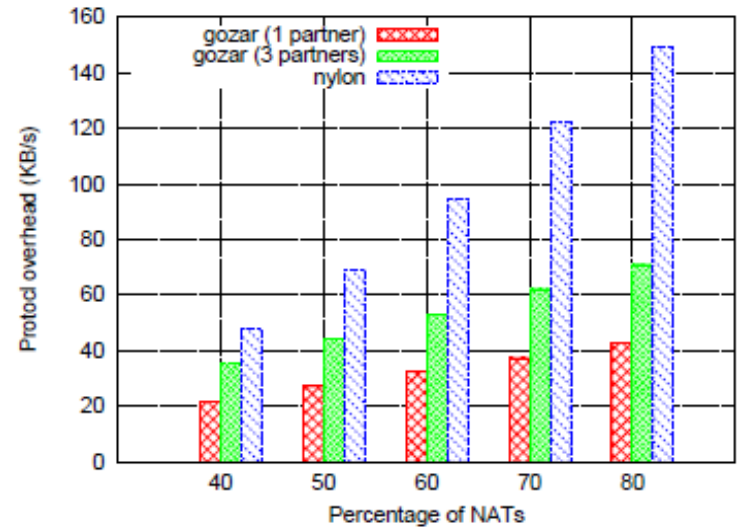


(d) Clustering coefficient.

Protocol Overhead

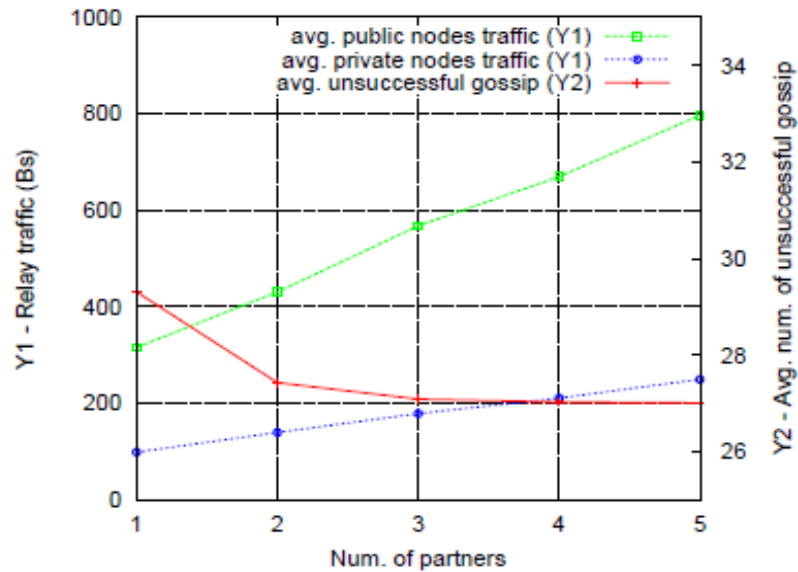


(a) Protocol overhead of Gozar vs. Nylon.




(b) Overhead traffic of Gozar vs. Nylon for varying percentages of private nodes.

Fairness and Connectivity in Failure



(a) Fairness after catastrophic failure: overhead for public and private nodes for varying numbers of parents.



(b) Biggest cluster size after catastrophic failures.

Conclusions

Conclusions

- **Gozar** is a NAT-friendly gossip-based peer sampling service that also provides a distributed NAT traversal service to clients of the PSS.
- **Public** nodes are leveraged to provide both the **relaying** and **hole punching** services.
- Relaying is only used for gossiping to private nodes
 - lower connection latency
 - Enabling a faster gossiping cycle
 - The messages relayed are small

Thank
You