

# Connectivity Properties of Mainline BitTorrent DHT Nodes

Raul Jimenez, Flutra Osmani, Björn Knutsson

KTH, Sweden




ROYAL INSTITUTE  
OF TECHNOLOGY




# Background

- P2P-Next
  - Large EU project
  - Content distribution platform on top of BitTorrent
- Kademlia DHT (Mainline implementation)
  - Used as distributed BitTorrent tracker
  - Iterative routing algorithm

# Kademlia: lookup



# Kademlia: lookup



20% dead nodes in  
95% of the lookups

One minute lookups


[Crosby-07]

Non-transitive  
connectivity and  
DHTs


[Freedman-05]

# Connectivity Properties


## Reciprocity



## Transitivity




## Persistence



# Experiment

- Nodes' connectivity properties
- 3,6 million unique nodes in 24 hours



N triggers experiment


Up to 5 queries per vantage point

1 minute timeout

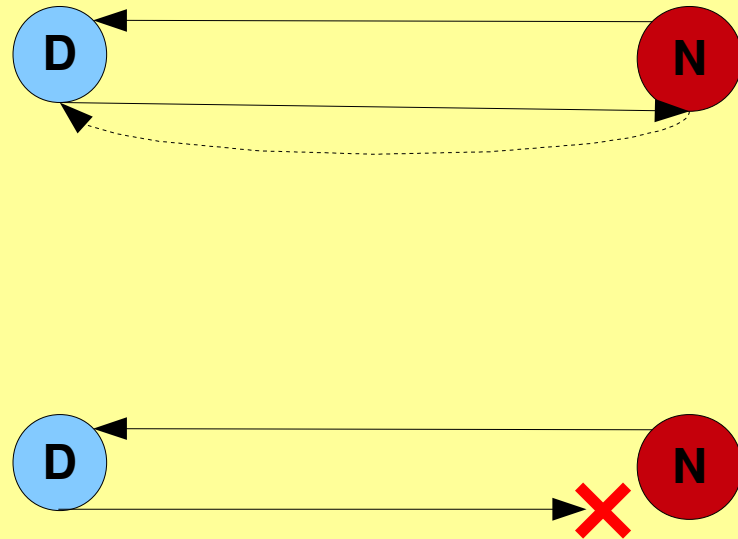
Repeat after 5 minutes

# Reciprocal Connectivity

## Reciprocity




## Experiment





# Reciprocal Connectivity

## Reciprocity




## Experiment






# Transitive Connectivity

## Transitivity




## Experiment




# Transitive Connectivity

## Transitivity




## Experiment



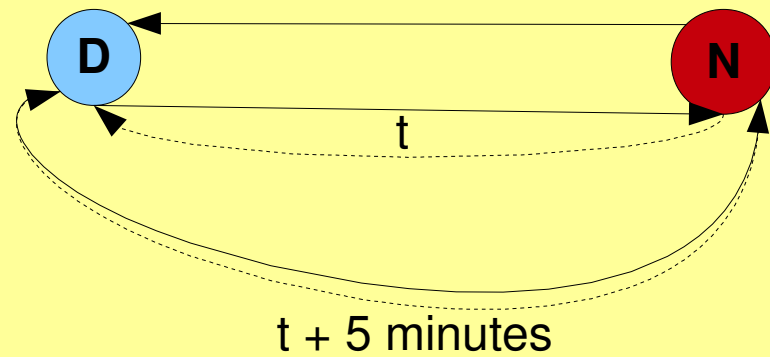
**40 %**

# Persistent Connectivity

## Persistence




## Experiment




# Persistent Connectivity

## Persistence




## Experiment



**44 %**

# Experiment Results



- Clear overlap: transitivity & persistence
- Just 1/3 of the nodes show “good” connectivity

# Discussion

- Possible causes
  - Firewalls → Reciprocity
  - NATs → Transitivity and/or persistence
  - DHT implementation → All
- Potential effects on the DHT performance
  - Higher (apparent) churn
  - Broken lookup routes

# Conclusion

- Understand the underlying connectivity issues
  - Explicit definition of connectivity properties
  - Large scale experiment on a real-world DHT
- Use this knowledge to achieve faster lookups

Thank you!




# Possible Causes

t	t + 5 min	%	Obs
		10.6	Firewall
		31.3	Port restricted cone & Symmetric NAT
		2.7	Full cone NAT & Real churn
		35.5	Open Internet

# Experiments

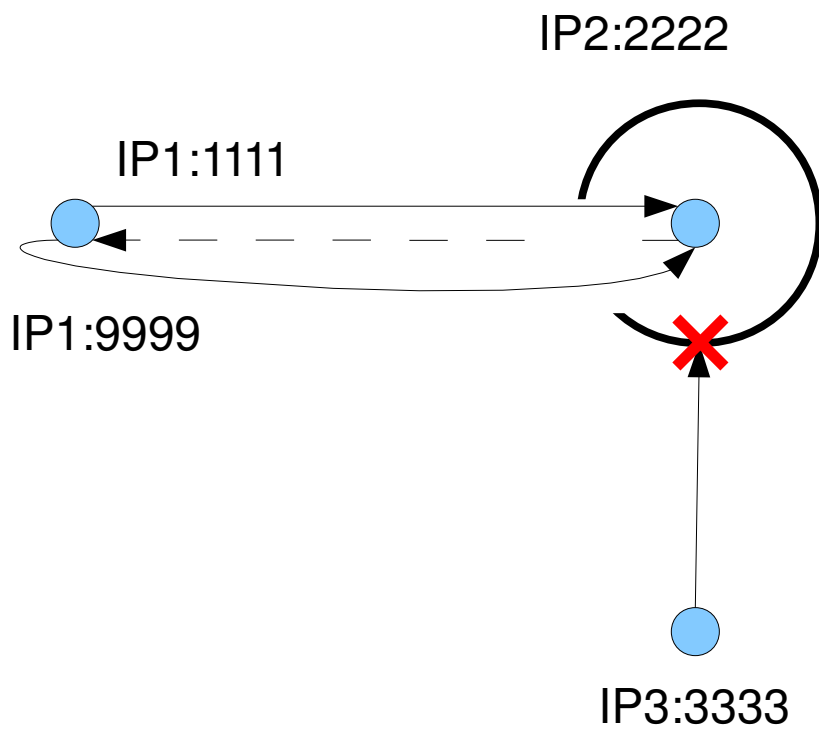
<b>Pattern</b>	<b>Nodes (%)</b>	<b>Possible cause(s)</b>
UUU-UUU	10.6	Firewall
RUU-UUU	<b>31.3</b>	Port restricted cone and symmetric NAT
RUU-RUU	2.8	
RRU-UUU	0.8	Restricted cone NAT
RRU-RRU	2.0	
RRR-UUU	2.7	Full cone NAT and real churn
RRR-RRR	<b>35.5</b>	
UUU-RRR	7.6	Behavior not matched
RRU-RRR	1.7	
Other	5	Rest of the cases

# NAT types




- Match IP and port (port restricted cone & symmetric)

# NAT types



- Match only IP  
(restricted cone)

# NAT types



- Match any connection (full cone)

Reference:

RFC3489 STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)

# What to do?

- Prevent my routing table from pollution
  - Check for **global reachability** before adding a node
- Don't pollute others' routing tables
  - When I'm aware of being partially reachable, tell others: “**don't add me to your routing table**”
    - Free riders?