SWAP: Protecting Pull-Based P2P Video Streaming Systems From Inference Attacks

Giang Nguyen
Communication Networks, TU Dresden, Germany
&
Stefanie Roos, Benjamin Schiller and Thorsten Strufe
Privacy and Data Security, TU Dresden Germany

Coimbra, Portugal
22.06.2016

Motivation

ComNets

- Appearance of video streaming services
- Demand for videos: consistent growth
- Bandwidth supply at Internet core: increases slowly
- Consequences:
  - High cost for service provider

    YouTube: $1million/day
  - Lower quality of experience

    ~400 kbps
- Service unavailability:
  - Under sudden demand
  - Due to sabotage

HUFF POST LIVE amazon instant video YouTube Live

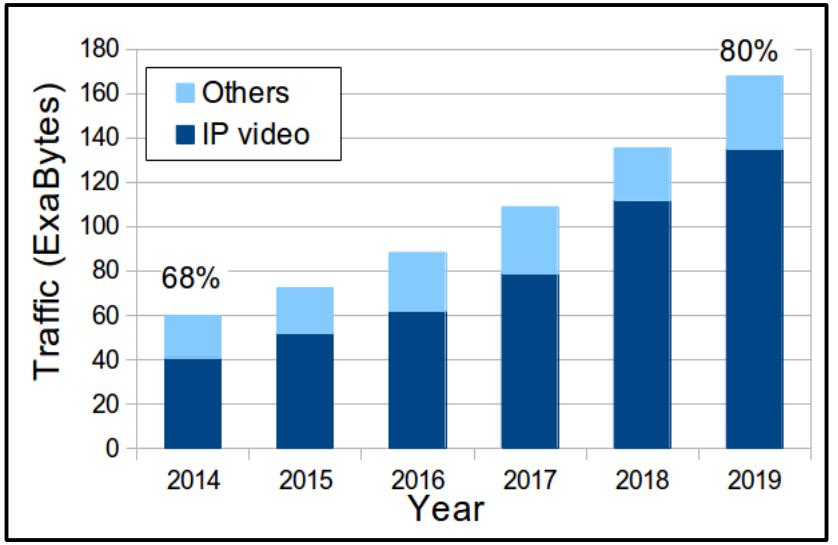

(Source: Cisco)

How to deliver reliable video streams to a large number of users in a cost-effective and resource-efficient way?

Content

Peer-to-Peer

- Incorporates peers bandwidth
- Advantages:
  - Lowered cost
  - Scalable
- Disadvantages:
  - Affected by churn (including failure)
  - Vulnerable to attacks
- Fundamental building blocks:
  1) Membership management
  2) Overlay construction
  3) Dissemination of video chunks

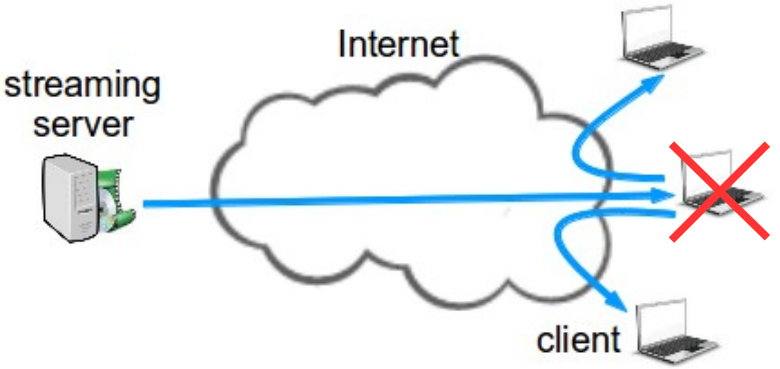

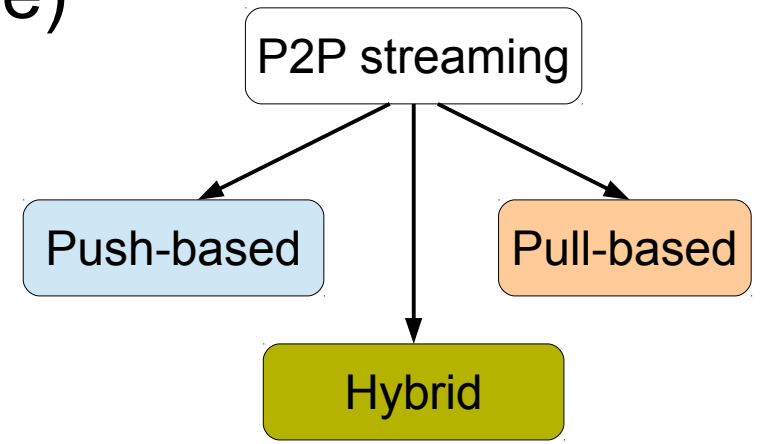

Related work

Push-based

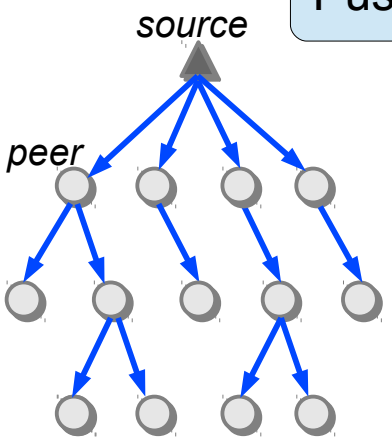

- Low delay
- Low overhead
- Low robustness

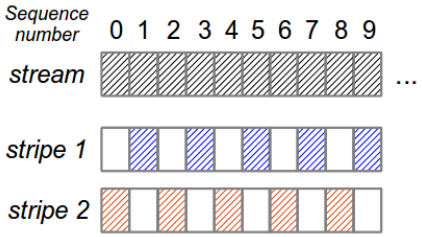

Pull-based

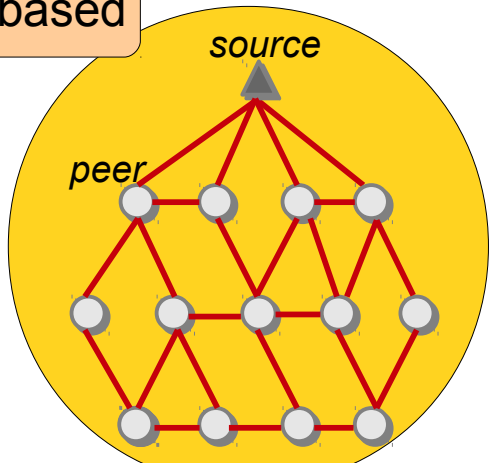

- Higher delay
- Higher overhead
- Higher robustness

(Example: DONet)

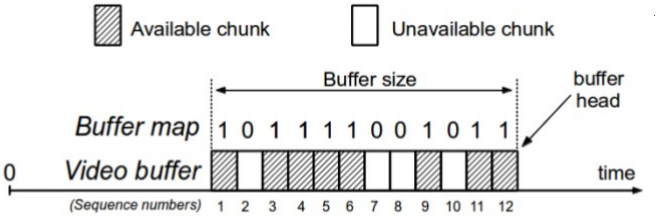

Hybrid

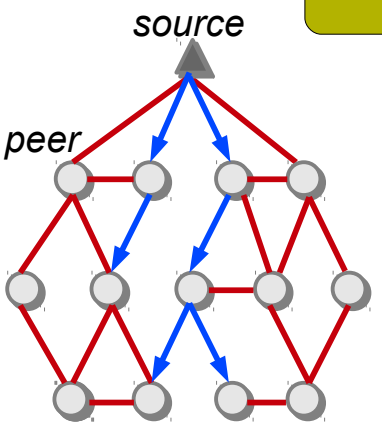

- Low delay & overhead
- Increased robustness
- Fragile backbone (Example: mTreebone)

How to improve the resilience of P2P streaming systems against both churn and attacks?

Tiering effects

- Insights from PPLive
  - Tiering effects
  - Stable over time

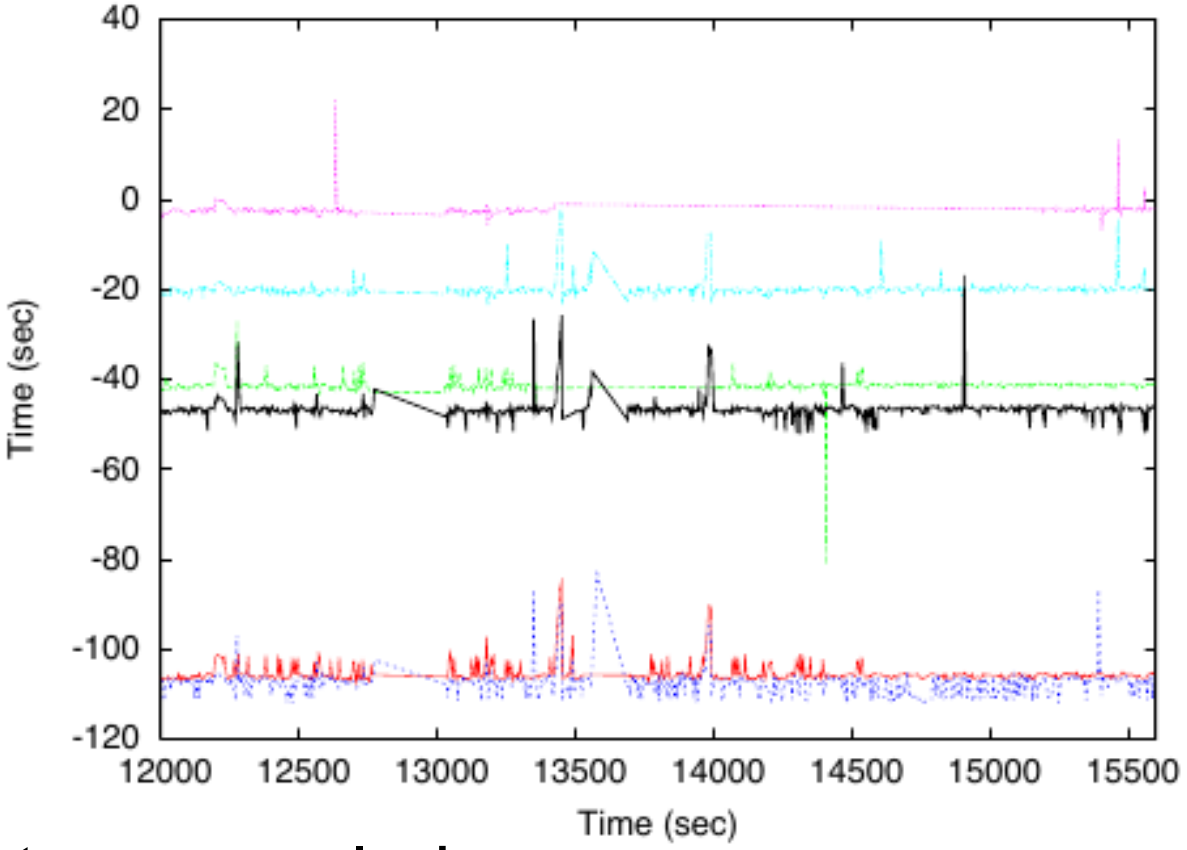

(Hei et al., 2007)

- Consequences
  - Overlay network structure revealed
  - Disrupting the flow of the video distribution
  - Especially when targeting head nodes

Attacker model

- Inferring overlay structure
  - Collect buffer maps
  - Offset

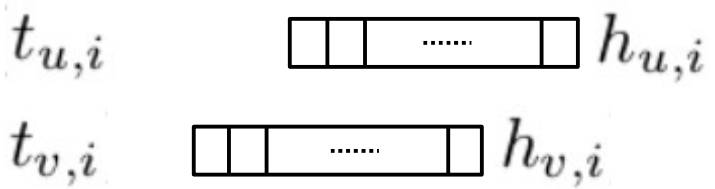

$$\delta_{u,i} := (h_{u,i} - h_{v,i}) - (t_{u,i} - t_{v,i}) \cdot r/L$$

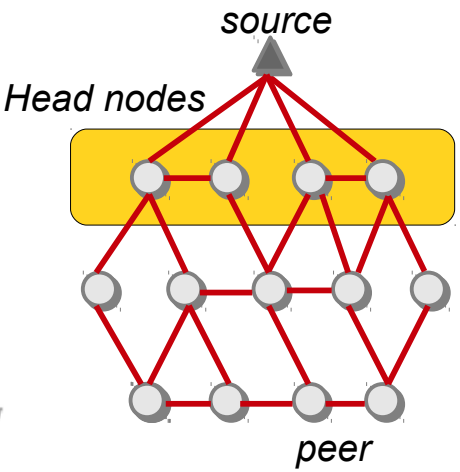

- Inference attacker model
  - Probing buffer maps
    - probability ***q*** to get a buffer map
    - ***m*** rounds
  - Identifying head nodes
  - Shutting down head nodes

SWAP scheme

- General idea:

  Increasing dynamics to mitigate attacks

- SWAP's operations:
  1) Partner nomination
  2) Nomination forwarding
  3) Periodic swapping

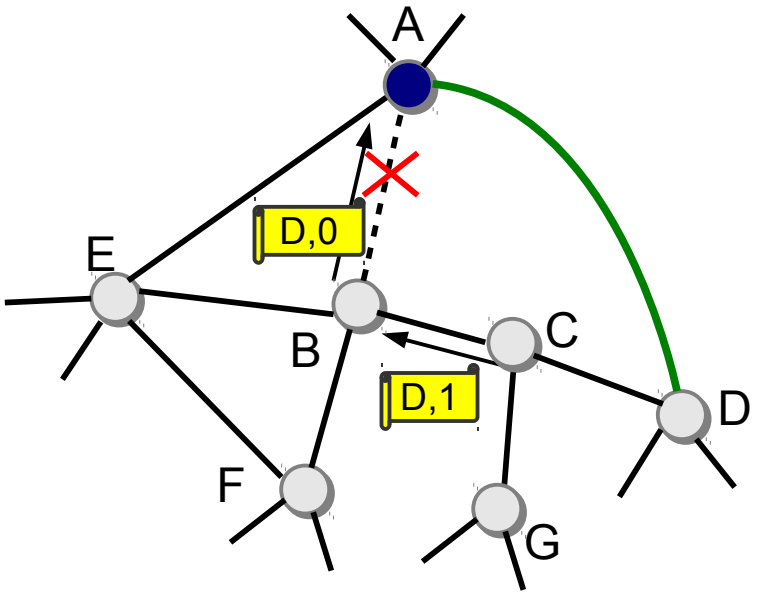

Evaluation (1/4)

ComNets

- Questions:
  1) How accurately does the inference attacker identify head nodes?
  2) How efficiently does SWAP mitigate the inference attacker?
  3) To which extent does SWAP increase the resilience of pull-based systems against the inference attacker?
- Metrics:
  - Attack's accuracy
  - Average and maximum miss ratios
- Settings:
  - 2000 peers, 1200 seconds simulation time
  - Upload bandwidth: Source 8 Mbps, peer 1 Mbps,
  - Coordinated attack at 800 seconds after start of stream

Evaluation (2/4)

How accurately does the inference attacker identify head nodes?

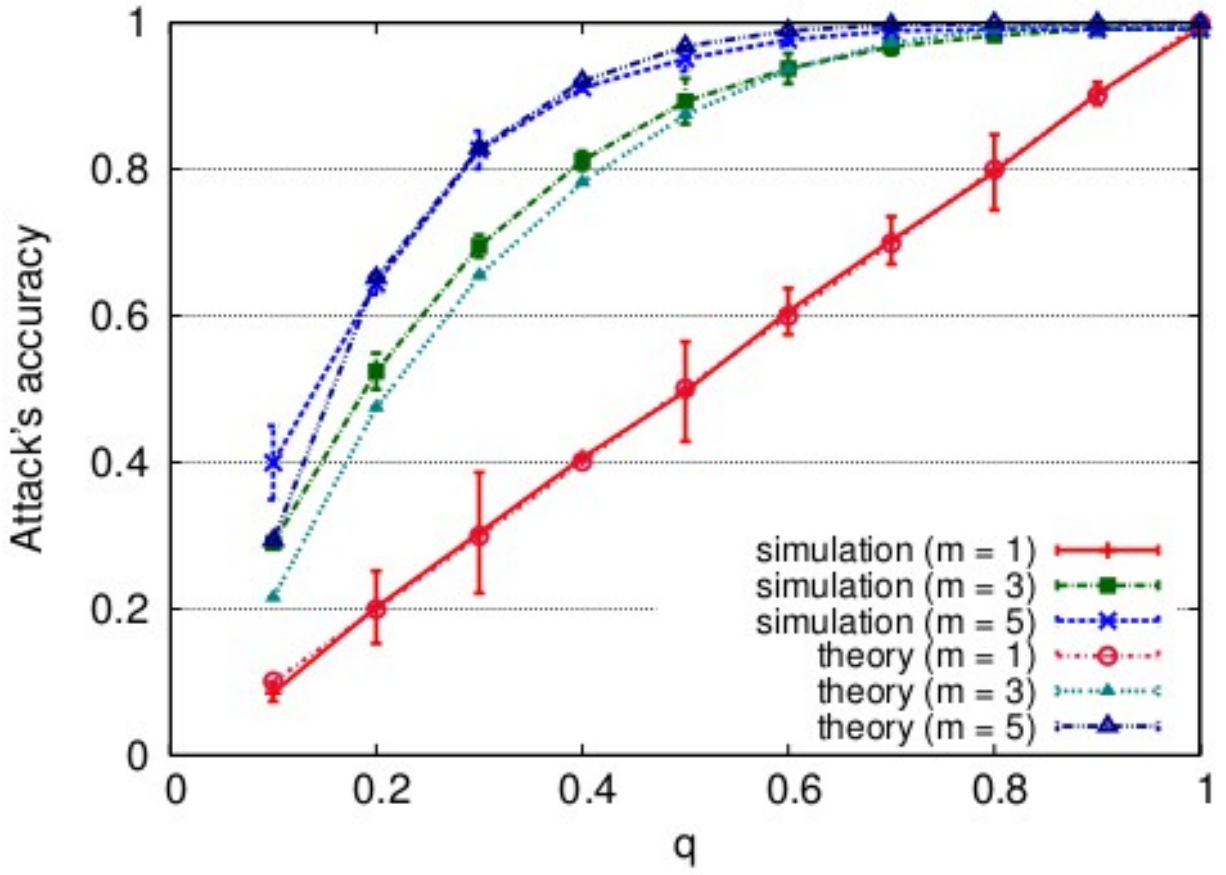

Evaluation (3/4)

How efficiently does SWAP mitigate the inference attacker?

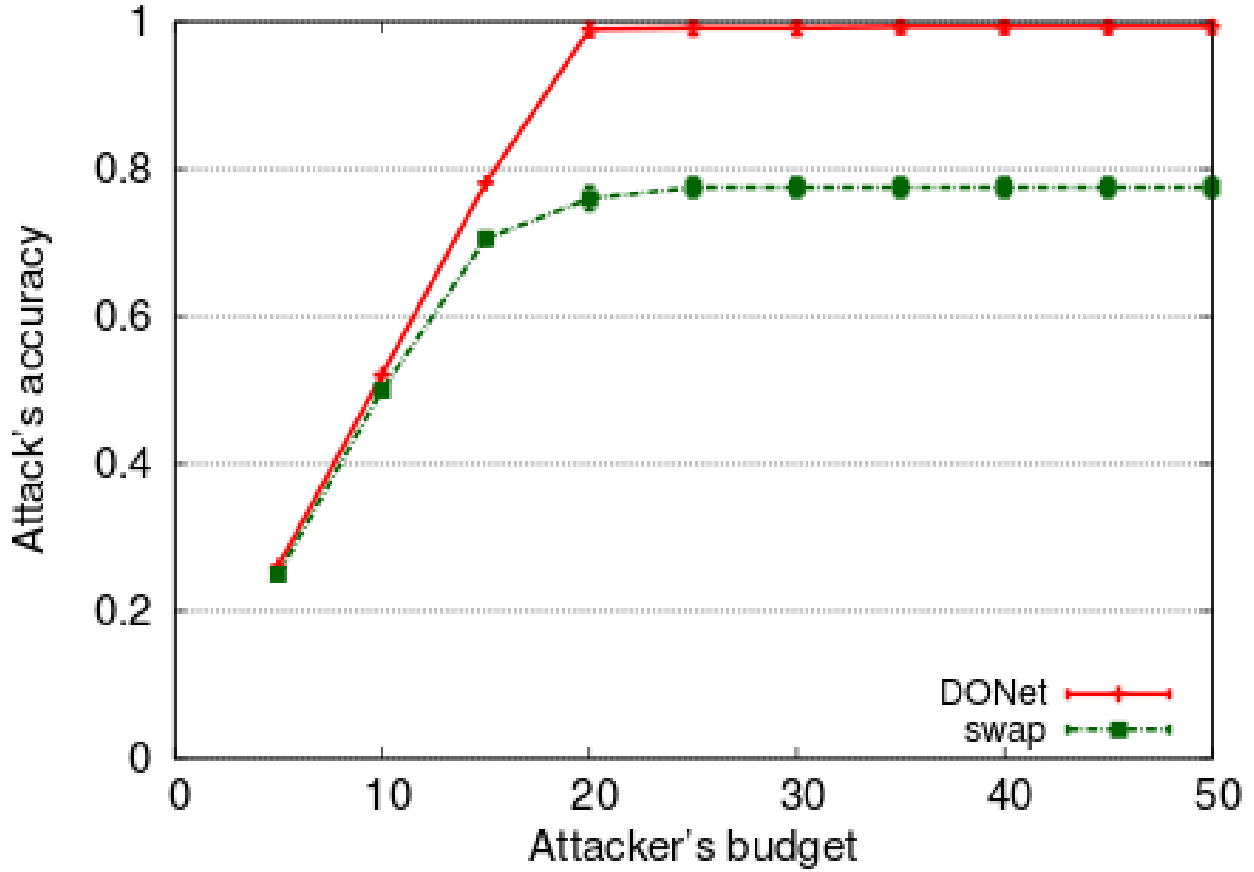

Evaluation (4/4)

To which extent does SWAP increase the resilience of pull-based systems against the inference attacker?

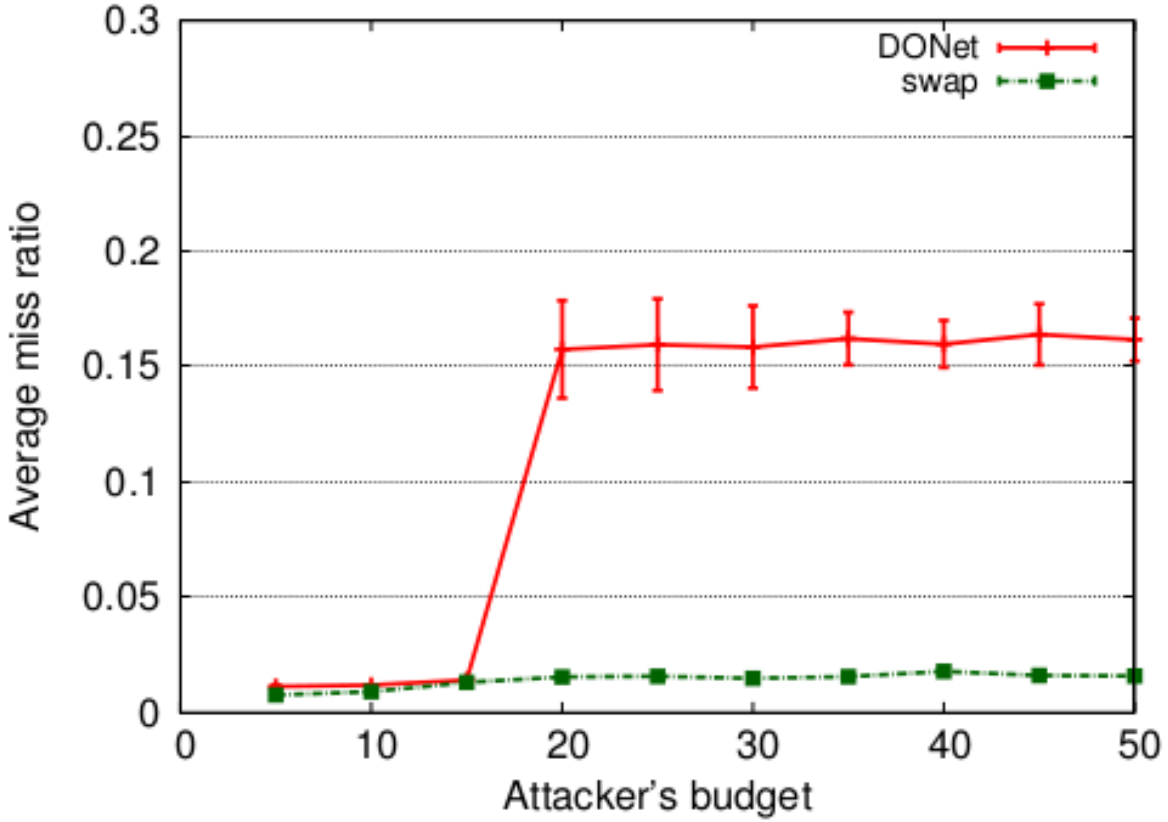

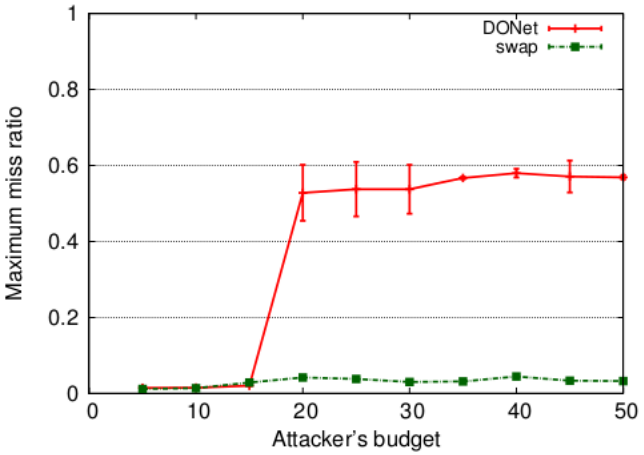

Conclusion

- Challenges: Inference attacker
  1) Collects buffer maps to identify head nodes
  2) Shuts down head nodes

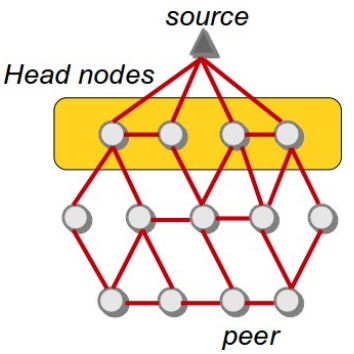

- Countermeasure: SWAP scheme
  1) Nominates a partner as replacement
  2) Forwards nomination message
  3) Swaps partners proactively

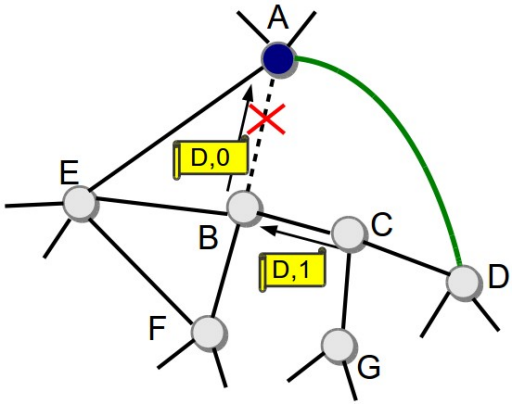

- SWAP drastically reduces chunk miss ratios