More than a Fair Share: Network Data Remanence Attacks against Secret Sharing-based Schemes

Leila Rashidi, Daniel Kostecki, Alexander James, Anthony Peterson, Majid Ghaderi, Samuel Jero, Cristina Nita-Rotaru, Hamed Okhravi, Reihaneh Safavi-Naini

Khoury College of Computer Science

Northeastern University

Appeared in NDSS 2021

Secure Communication

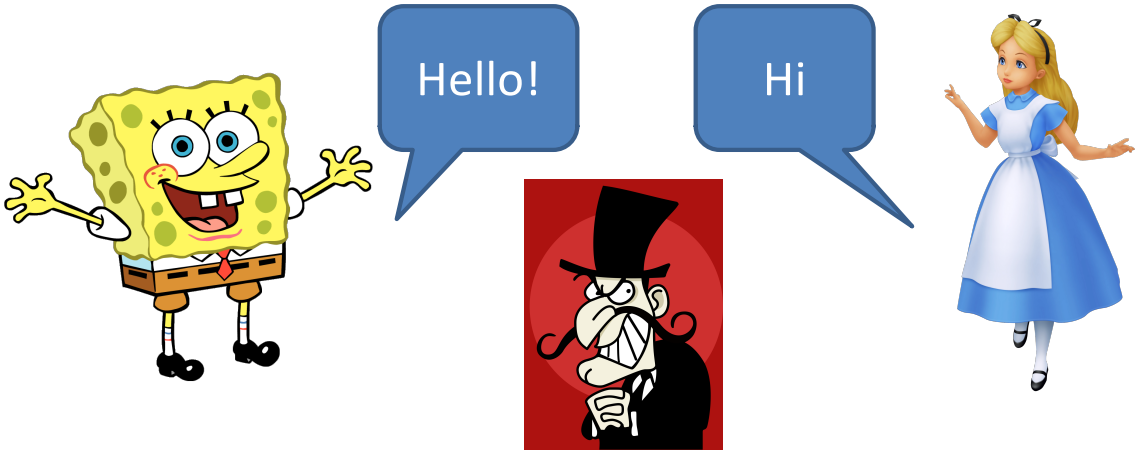

- Establish a secure and authenticated communication channel using standard protocols such as TLS or QUIC
- Security guaranteed by cryptographic primitives that assume computationally-bounded adversary

… disrupted by Quantum Computing

- Emergence of quantum computing breaks assumptions needed for the security of existing cryptographic primitives
- Design secure communication without relying on computational assumptions about the adversary
- Existing approaches
  - Information theory
    - Secret sharing
  - Computer networks
    - Multi-path routing
    - Path switching

Secret Sharing

A. Shamir. *How to Share a Secret. 1979*

- How to split and recreate a secret between participants n that do not trust each other
- A (k, n) scheme:
  - Divide a secret S into n pieces $s_1, \ldots s_n$
  - Any group of k or more users can jointly obtain the secret
  - Any group of k-1 or less users can not jointly obtain any information about the secret;
- Security: Secure as long as the adversary does not capture more than k-1 shares

Secret Sharing and Multi-path Routing

- The message remains perfectly secret as long as the adversary can access at most k – 1 paths
- Adversary bounded in terms of network access; does not know/observe ALL the paths

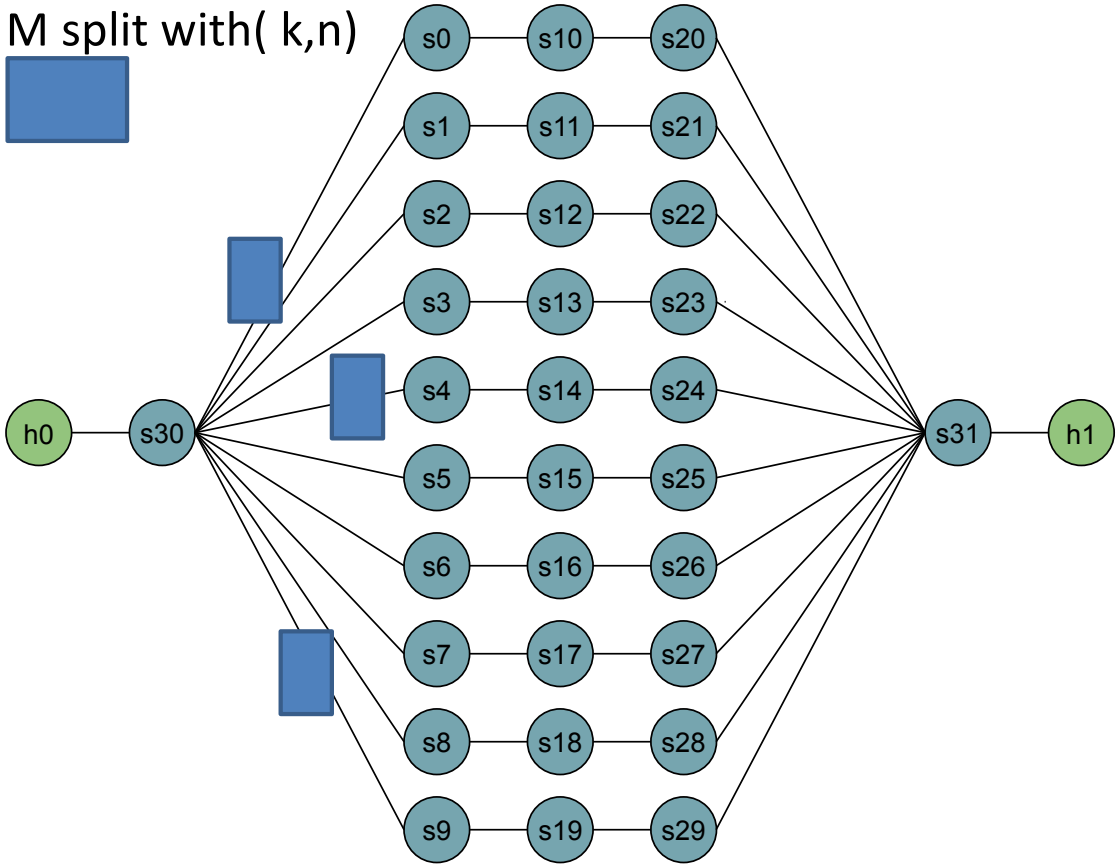

Multi-path Switching with Secret Sharing (MSSS)

- **Path-switching:** A random path is chosen for each message and used for transmission of the message
- **MSSS (k,n)**:
  - Sender splits the message in k shares
  - Sender sends the shares on k disjoint paths
  - Sender and receiver *switch* to a randomly selected set of paths out of the total set of n paths
- It provides information-theoretic security against an adversary with access to a quantum computer

R. Safavi-Naini, A. Poostindouz, and V. Lisy, "Path hopping: An MTD strategy for quantum-safe communication," in ACM Workshop on Moving Target Defense, 2017

This talk

Are practical implementations of multi-path switching with secret sharing schemes secure?

- Examine if assumptions made by the theoretical models to prove security are met in real networks
- Identify a side-channel (**Network Data Remanence**) and attacks exploiting it (**NDR Blind and NDR Planned**)
- Propose countermeasures and demonstrate their effectiveness

Multi-path Switching with Secret Sharing

System

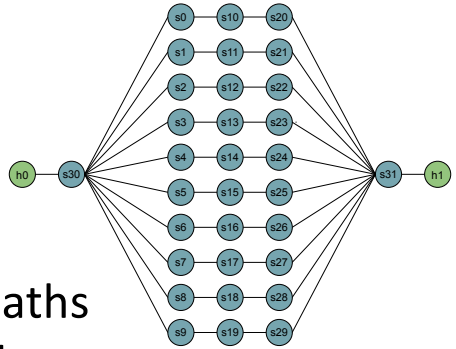

Network:
There are n disjoint paths known by sender and receiver and connecting them

Sender:
Each clock tick i:
Selects set K_i = {k paths our of n}
Splits M using (k,k) secret sharing
Sends them on the set of paths K_i

Receiver:
Listens to all paths; thus no need for secret key

Attacker

Can not observe and access all paths
Each clock tick j
Selects set K_j = {k paths out of n}
Accesses K_j to recover shares

Switch clock can be the same or not with the one of the sender

Security

It provides information-theoretic security and remains secure against an adversary with access to a quantum computer

Model Used for Security Analysis

- Model assumes that paths have same length and delay

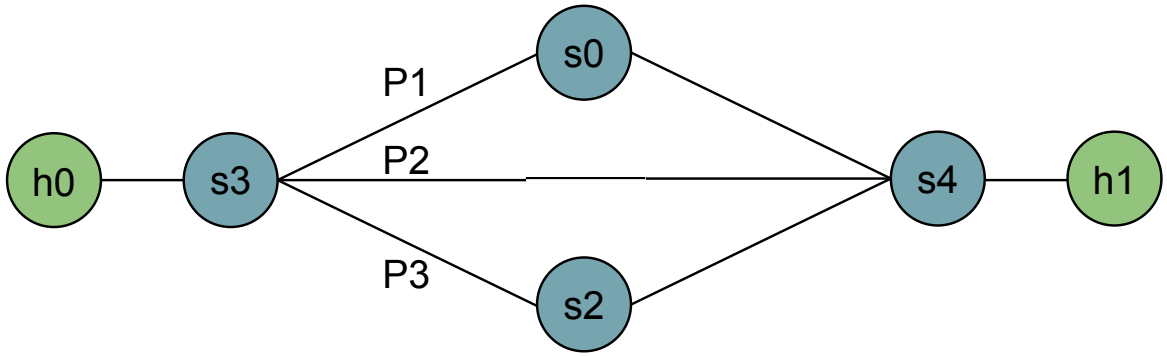

Real networks:

- Paths do not have the same number of hops
- Links (and paths) do not have the same delay

Attacker gets more chances at capturing a share on a path (than assumed by the model)

Network Data Remanence Side-Channel (NDR)

(5, 9) scheme, showing active paths – paths that have ongoing packets

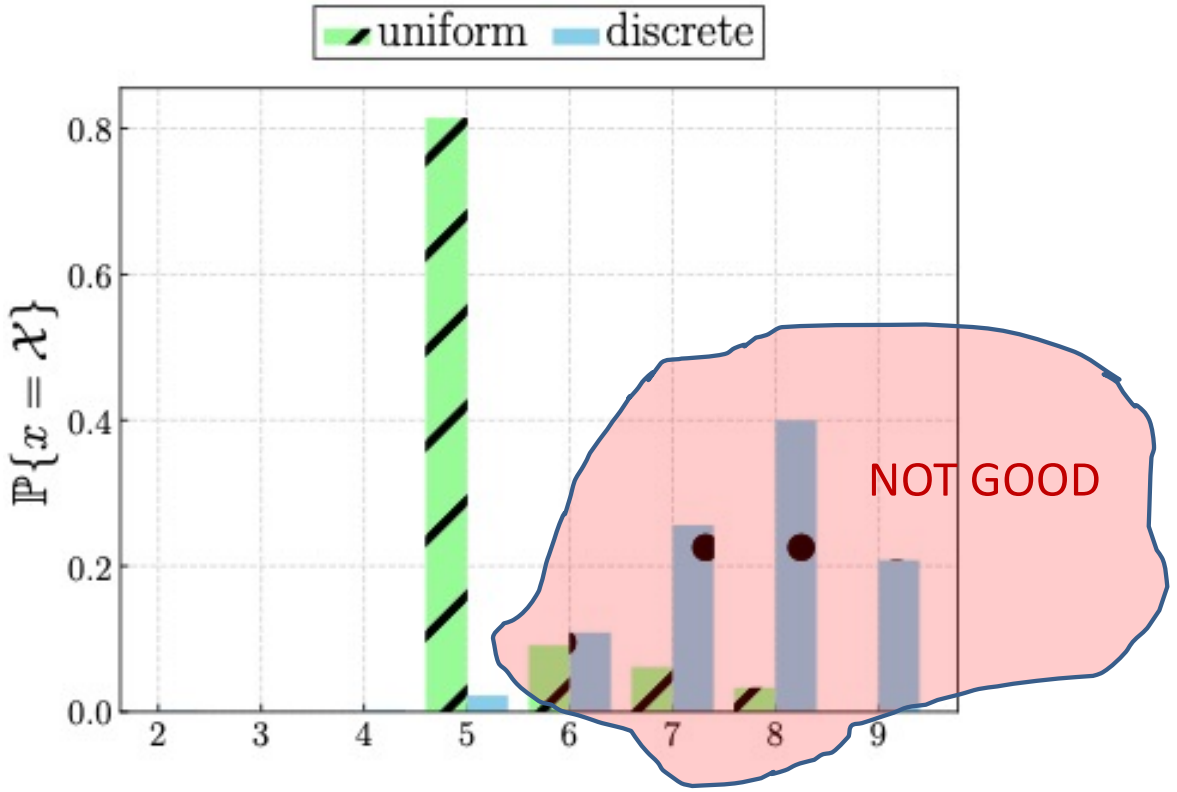

Packets linger longer in the network creating a side-channel

Attacker Capability

- Attacker captures packets at nodes
  - has access to all of the nodes, but they cannot possibly capture traffic from all of them at all times.
  - can only capture traffic at a fraction of nodes at each time.
- Attacker is able to listen to at most K nodes simultaneously (K is number of paths used by MSSS)
- Attacker can switch what paths they are listening to and at what intermediate nodes
- **Attacker chooses nodes, and can decide to stay on same path and select a node on the same path**

MSSS Attacks that Do Not Exploit NDR

- Attacker switches nodes or not, does it know or not the time the sender switching paths

- **Fixed** attacker: does not switch nodes
- **Independent** attacker: switches nodes but does not know switching time
- **Synchronized** attacker: switches nodes and knows switching time, i.e. it is synchronized with the sender

Network Data Remanence Attacks

- **NDR Blind**: selects K nodes from all nodes on all paths
- **NDR Planned**: follows shares as they travel along the paths in the network
  - listens to K random nodes of distance 1 from the sender
  - probes K random nodes of distance 2 from the sender during the second switching interval
  - and so on ….
- **NDR Planned Opt**: checks at each step to see if all shares needed to reconstruct a message are captured.
  - Starts at distance 1, instead of continuing with next hop

Attacks Summary

Name	Abrv.	Exploits NDR	Knows Switching Time	Switches Nodes	Knowledge of Path Composition
Fixed	FIX	No	Yes	No	Partial
Independent	IND	No	No	Yes	Partial
Synchronized	SYN	No	Yes	Yes	Partial
NDR Blind	BLD	Yes	Yes	Yes	Nothing
NDR Planned	PLN	Yes	Yes	Yes	Complete
NDR Planned Opt	OPT	Yes	Yes	Yes	Complete

NDR Planned Attack Analysis

$P_{pln}(m,t)$: probability that attacker has captured exactly m shares by tick t

$$P_{pln}(m,t) = \sum_{x=0}^{f(m)} P_{pln}(m-x, t-1) \times \frac{\binom{K-m+x}{x}\binom{N-K+m-x}{K-x}}{\binom{N}{K}}$$

$$P_{pln}(m,1) = \begin{cases} \frac{\binom{K}{m}\times\binom{N-K}{K-m}}{\binom{N}{K}}, & 0 \le m \le K \le \frac{N}{2} \text{ or} \\ & 0 < 2K - N \le m \le K \\ 0, & otherwise \end{cases}$$

$$f(m) = \begin{cases} \min(m-(2K-N), K), & 2K > N \\ \min(m, K), & 2K \le N \end{cases}$$

For a path of length L, $P_{pln}(K, L-1)$ is probability the attacker captures all the shares within the duration of transferring a message on that path

Probability of Data Recovery

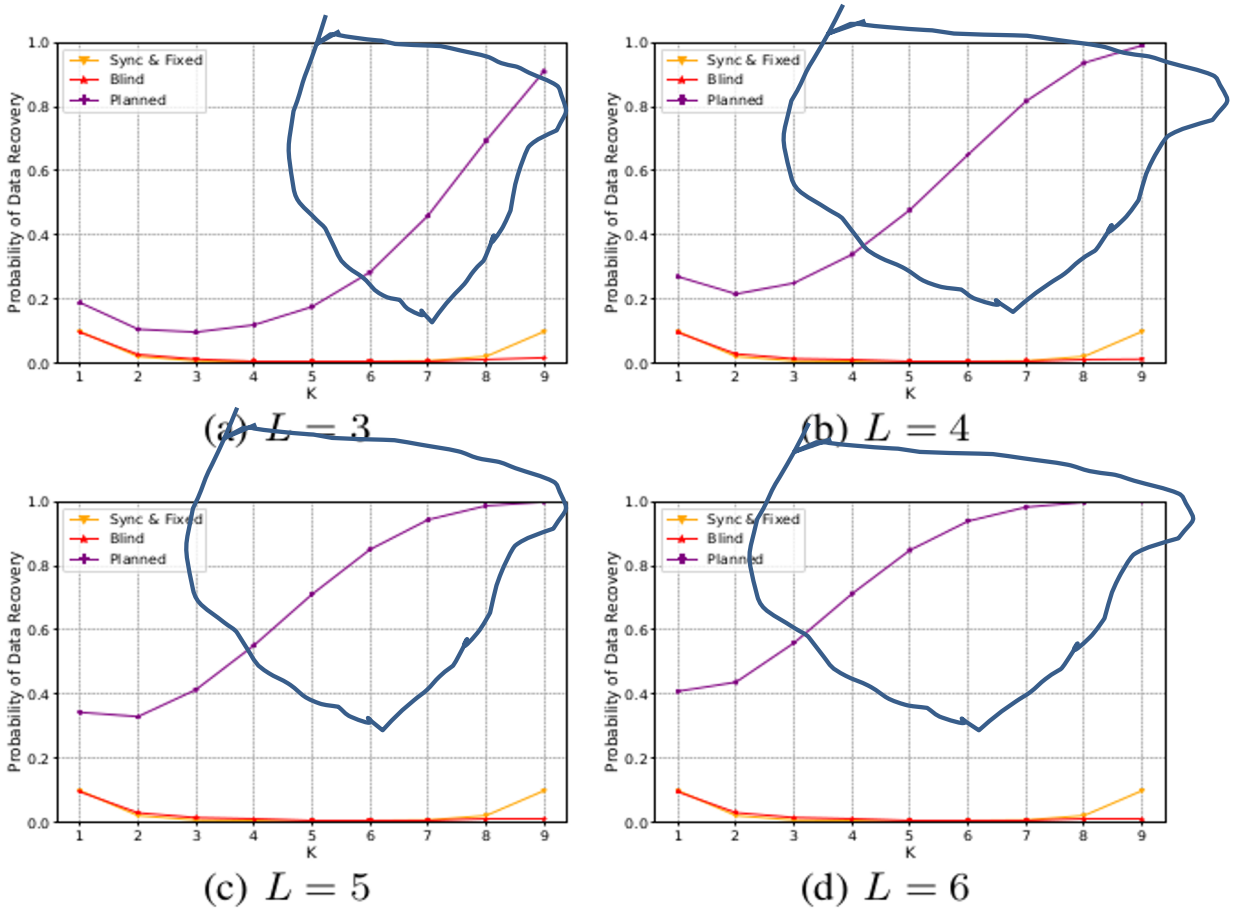

NDR Planned attacker is very effective

NDR Blind is not very effective

L: path length
K: # shares
N: # paths, 10

MSSS SDN-based Design

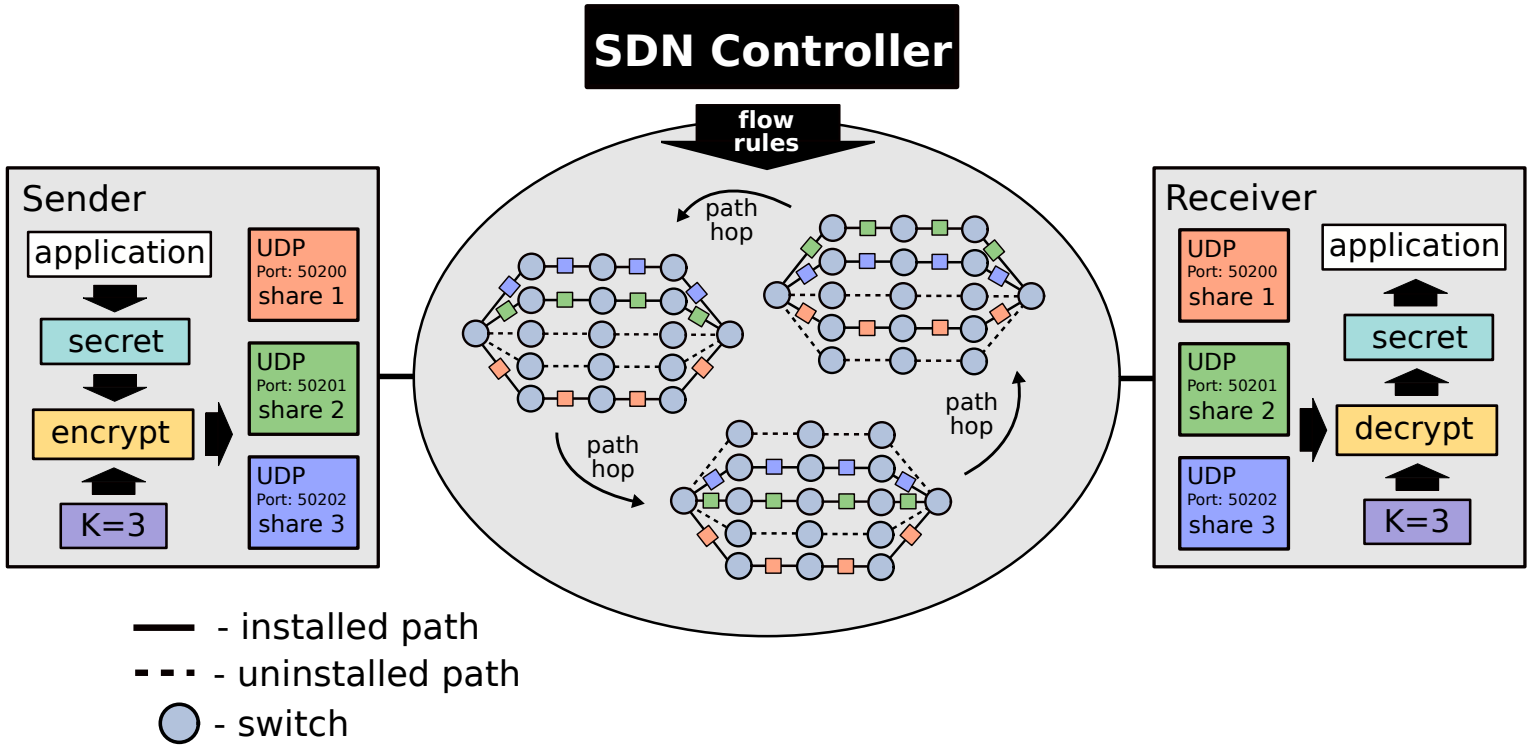

- UDP ports are used to distinguish between paths
- Receiver listens to all paths

Experimental Results

- MSSS implemented with ONOS and Mininet

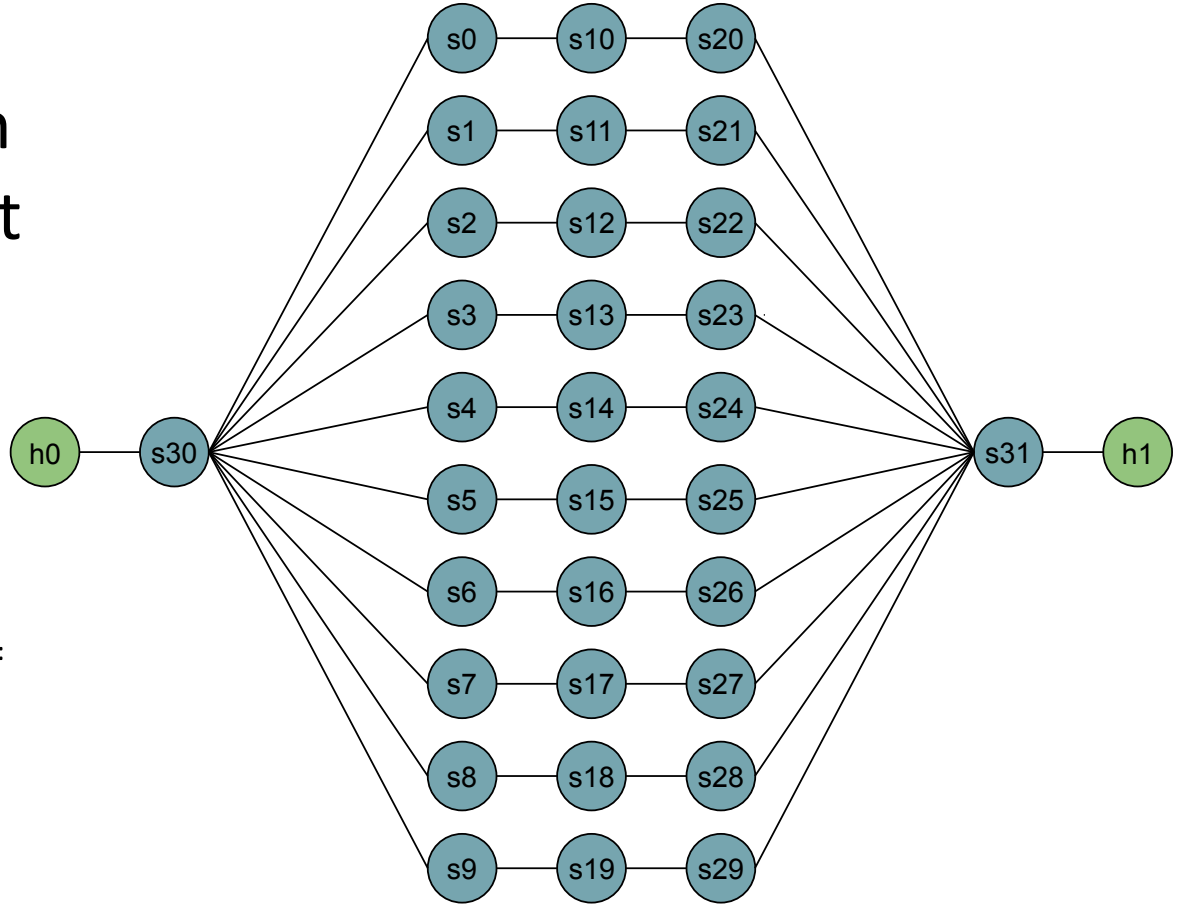

- N = 10
- L = 4
- K = 3
- Path switching interval δ = 100 ms
- File size = 10MB
- M = 512B

Impact of Path Length (each link has 50 ms delay)

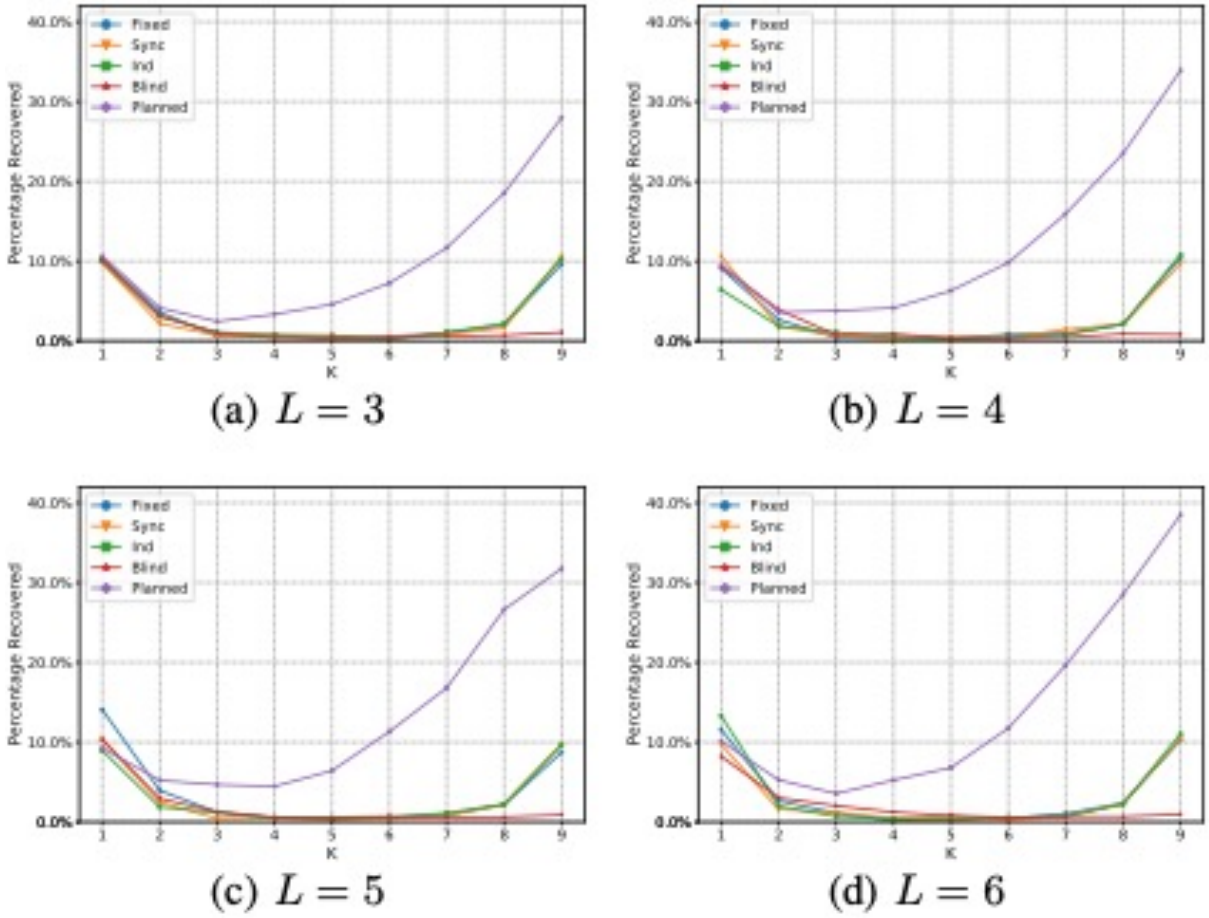

Impact of Path Delay

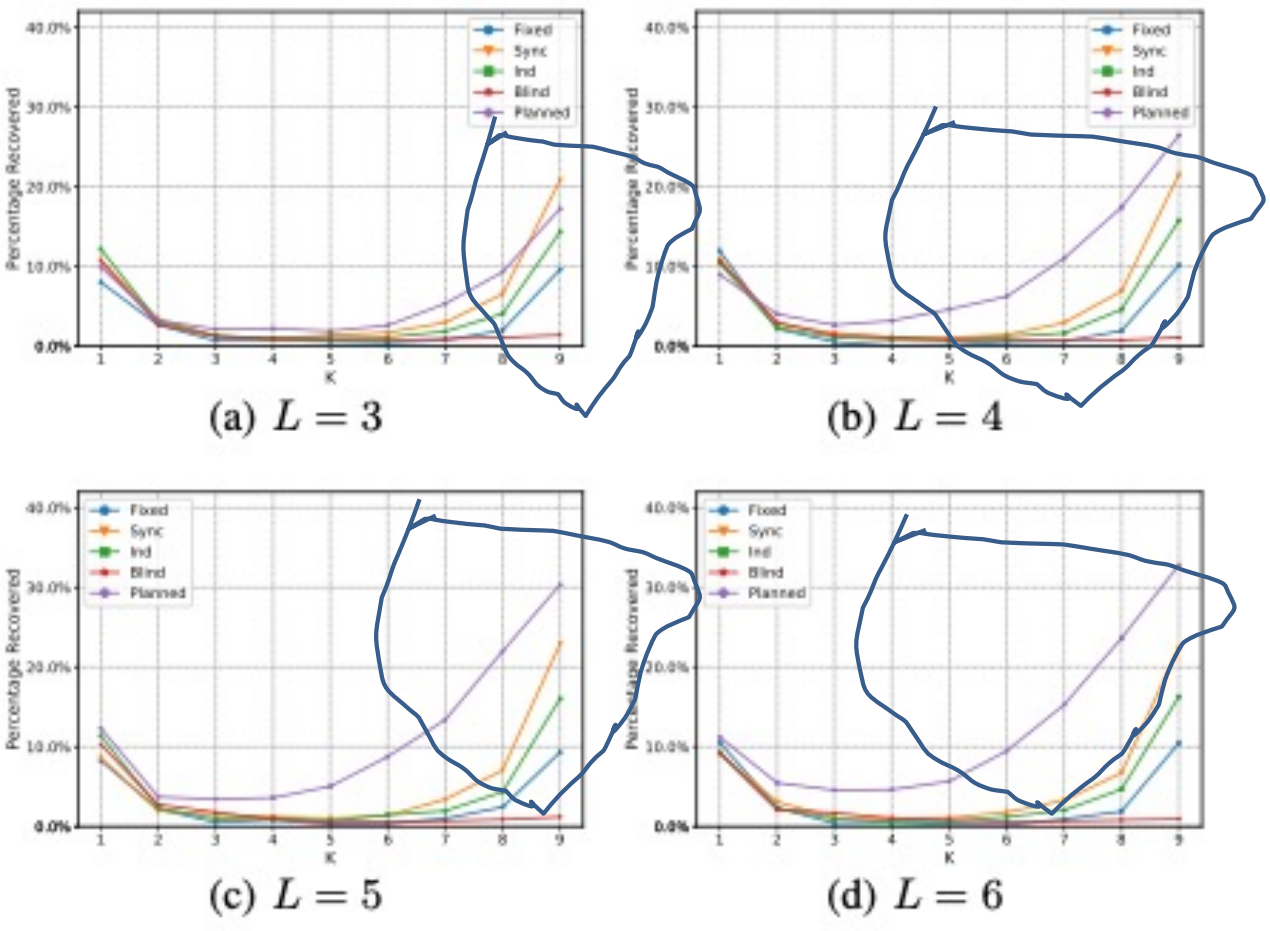

NDR Planned attacker is very effective in SDN –based implementation

How to Mitigate the Attacks?

We want to keep information theoretic security

Break the message into more shares

- How to send these shares:
  - Use more disjoint paths – need to also increase the attacker power to be fair
  - Use the same K paths repeatedly -- could result in reduced protection
- Our approach: distribute shares over both *time* and *space* instead of just space using a random set of paths to send a K-sized set of shares

Our Mitigation

- Generate more shares and spread them across both space and time
- Instead of (K, K), the sender uses (HK, HK) secret sharing
  - divide the shares into H sets of K shares
  - send these sets of shares, one at each consecutive clock tick
  - at t = 0, 1, . . . , H – 1, the sender chooses K paths uniformly at random, and then sends a share along each chosen path
- We call H *resilience factor*, a system parameter that can be configured by the sender

Analysis

$$P_{pln}(m,t) = \sum_{x=0}^{\min(K,m)} P_{pln}(m-x, t-1) \times D_{pln}(m,x)$$

$$D_{pln}(m,x) = \begin{cases} \frac{\binom{N-x}{K-x}}{\binom{N}{K}}, & m\%K = 0 \\ \frac{\binom{K-(m\%K-x)}{x} \times \binom{N-K+(m\%K-x)}{K-x}}{\binom{N}{K}}, & m\%K > 0 \end{cases}$$

Probability of data recovery for the NDR Planned Opt attacker is $P_{pln}(KH, L + H - 2)$

Effectiveness of Mitigation

NDR Blind

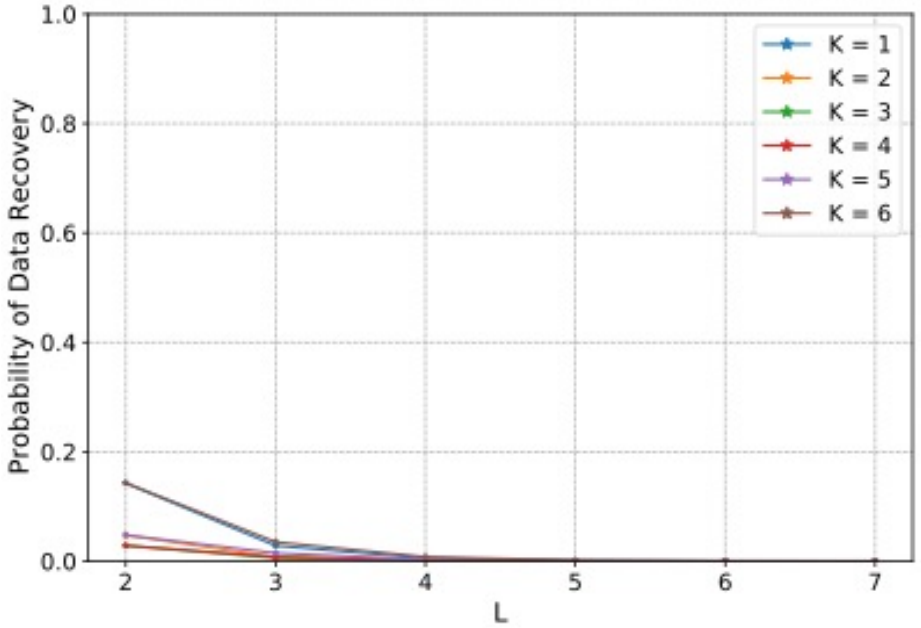

NDR Planned Opt

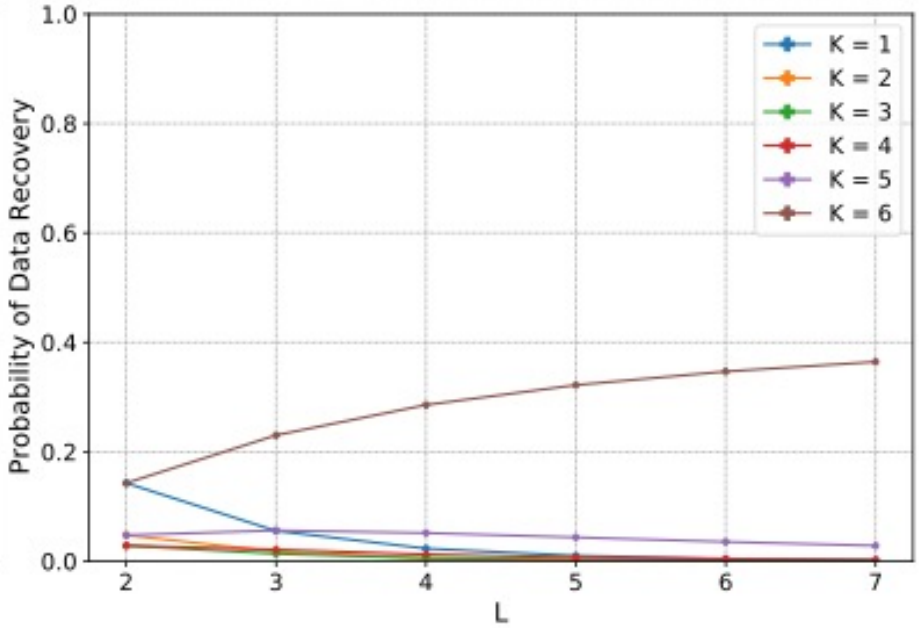

N = 7, H = L – 1

Overhead

End-to-end Latency

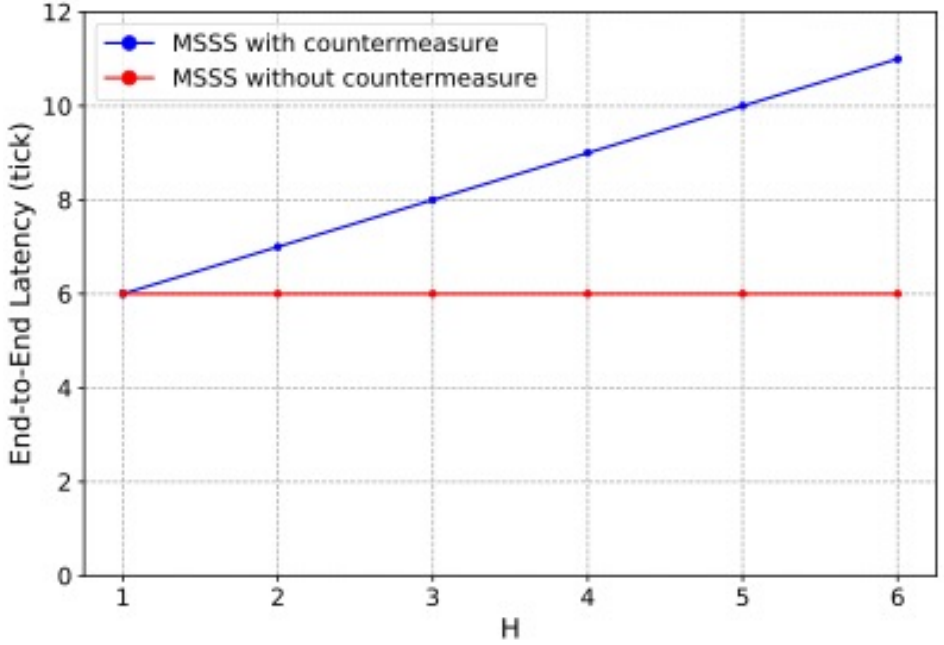

Coding Rate

(proportion of information over the total data generated by an encoder)

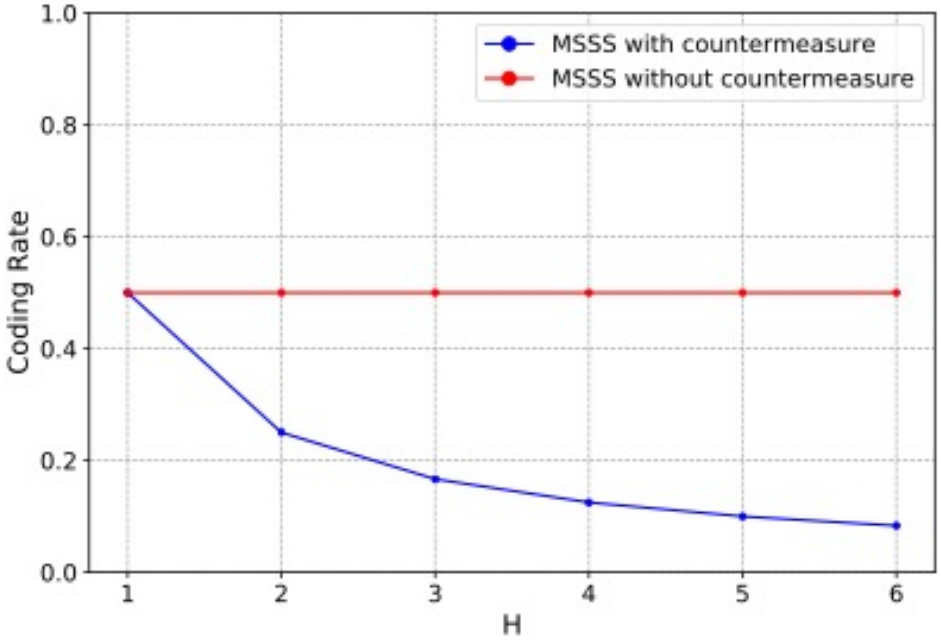

L = 6 and K = 2

Experimental Results: Probability Data Recovery

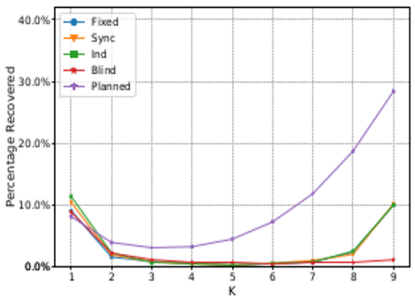

(a) $L = 3$, No Countermeasure

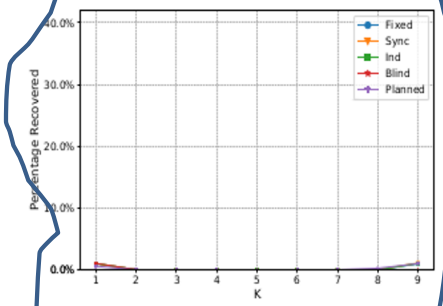

(b) $L = 3$, With Countermeasure

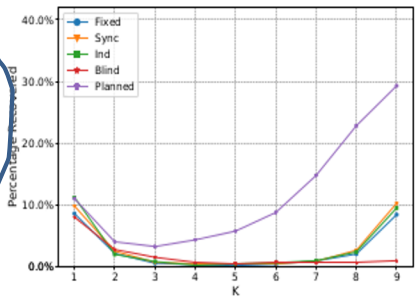

(c) $L = 4$, No Countermeasure

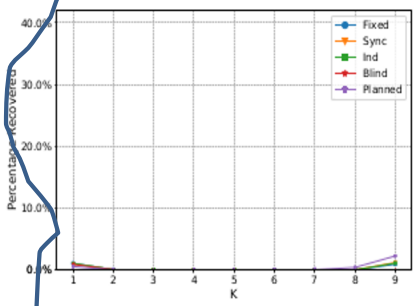

(d) $L = 4$, With Countermeasure

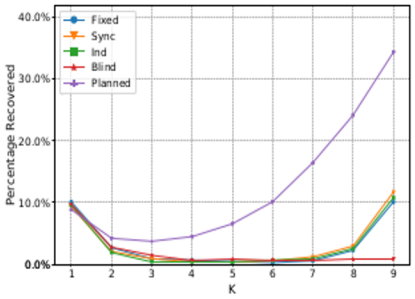

(e) $L = 5$, No Countermeasure

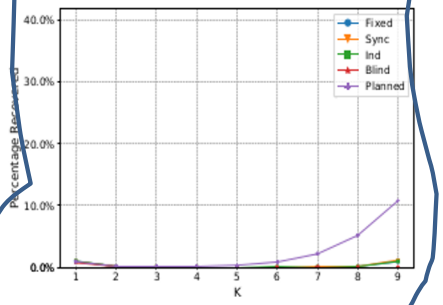

(f) $L = 5$, With Countermeasure

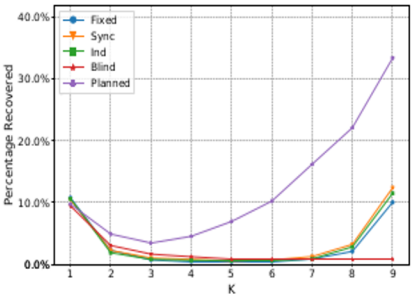

(g) $L = 6$, No Countermeasure

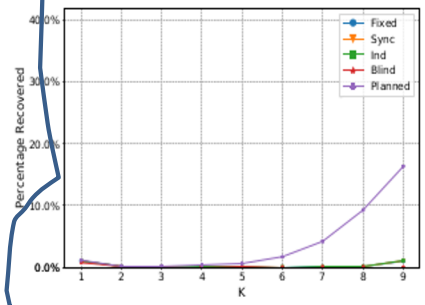

(h) $L = 6$, With Countermeasure

Effect of the countermeasure and number of shares on percentage of recovered data with varying path length. Fixed scenario with 2 *ms* delay between each node. The sender's $\delta = 4$ *ms* and the attacker's $\delta = 8$ *ms*. File size and the resilience factor, H, are set to 1 *MB* and 2, respectively.

Countermeasure mitigates the NDR Planned attack in SDN-based implementation

Experimental Results: Goodput

L = 3, H = 2, sender's δ = 4 *ms,* *a*ttacker's δ = 8 *ms*.

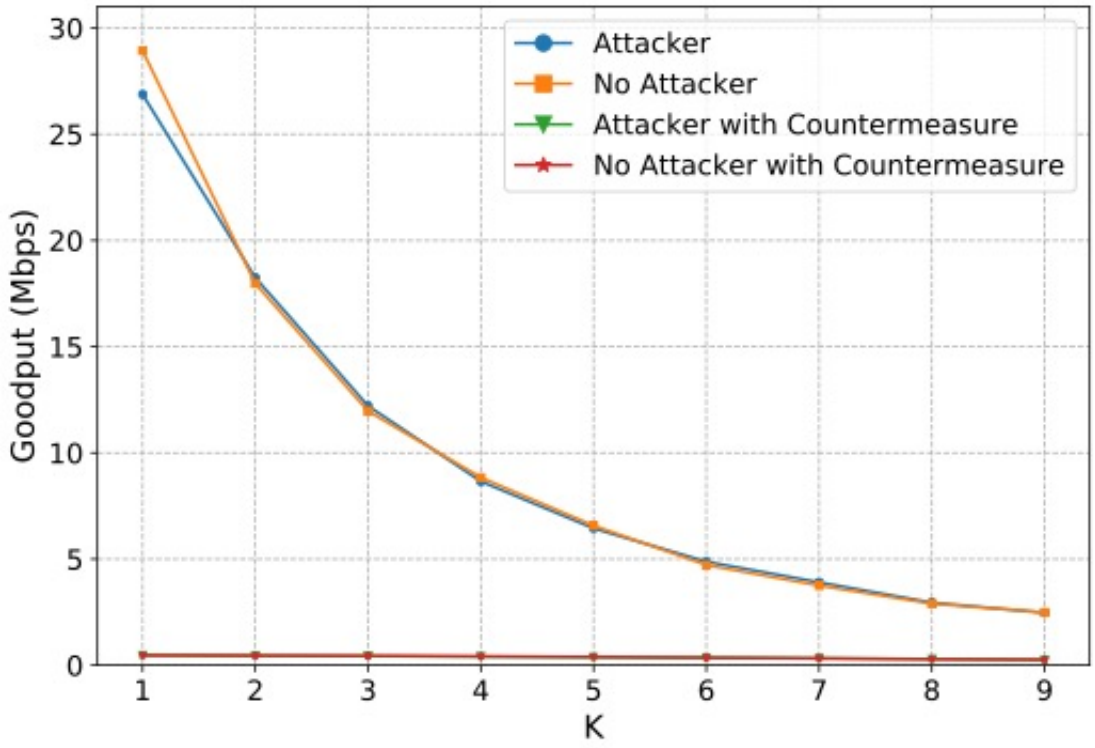

Increasing the number of shares, and spreading them through time, has a significant impact on performance

Summary

- Analyzed secure communication schemes that do not make computational assumptions about the attacker
- Identified a side-channel Network Data Remanence and analyzed and demonstrated attacks that exploit it in a SND-based implementation of MSSS
- Proposed a countermeasure, analyzed and demonstrated in the same SDN-based implementation