Northeastern University

Bandwidth-aware Multipath Secure Communication

By

Leila Rashidi

Postdoc Associate, Department of Computer Science, University of Calgary

In Collaboration with

Sogand Sadrhaghighi, Majid Ghaderi, Cristina Nita-Rotaru, Rei Safavi-Naini

December 2021

Outline

- Introduction
- Background
- Undetectable Attacks
- Trustworthiness of Network Infrastructure
- How to mitigate the leakage by untrusted switches?
- Our Solution
- Experimental Results
- Conclusion

Background: Secret Sharing

- A fundamental building block in
  - secure multiparty computation
  - distributed storage
  - side channel protection
- (t,n) threshold secret sharing scheme uses
  1. A **randomized share generation algorithm:**
     Takes a secret *S* and generates *n* shares
  2. A **deterministic reconstruction algorithm:**
     Takes any *t* shares and reconstructs the secret *S*

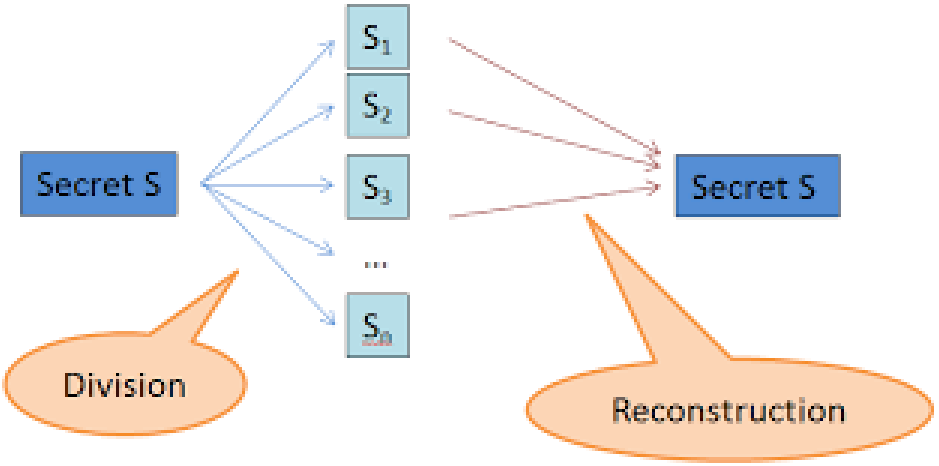

- Security property of (*t,n*)-Secret Sharing:
  The secret will be perfectly (information theoretically) secure if the adversary can have access to at most $t - 1$ shares.

Background: Multipath Routing

- A routing technique simultaneously using multiple alternative paths through a network.

- Benefits:
  - Fault tolerance
  - Increased bandwidth
  - Load Balancing
  - Improved Security

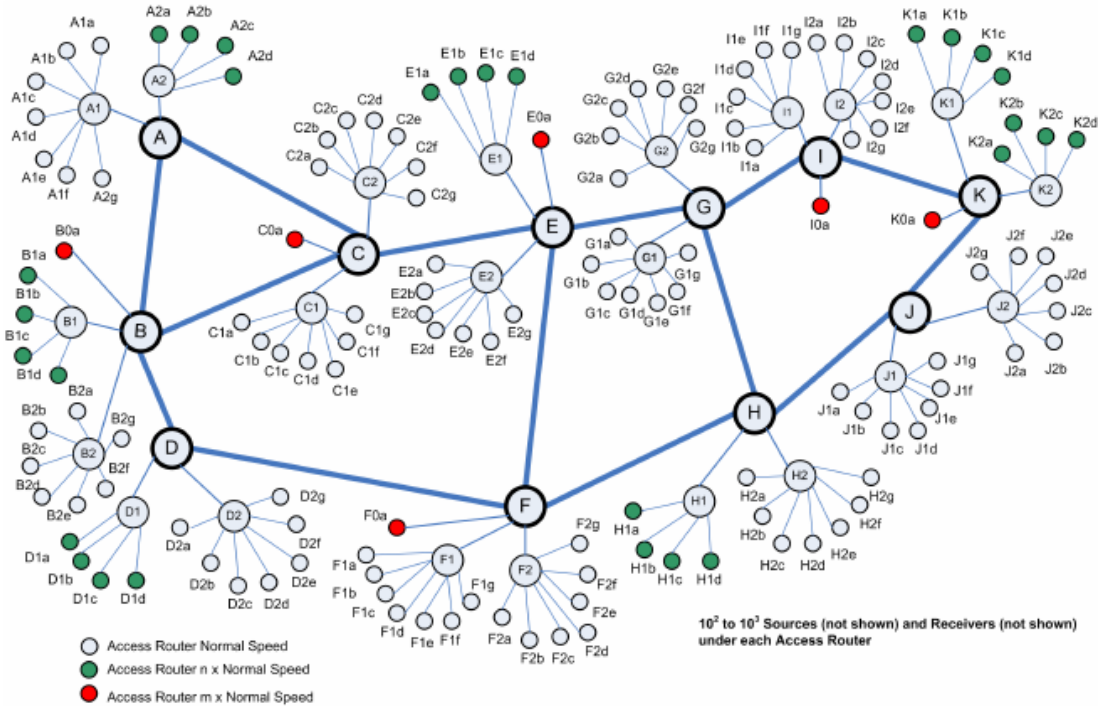

Undetectable Attacks

- Mechanisms to ensure steady, secure, and reliable communication:
  - Firewalls
  - Intrusion detection
  - Protection Systems
- These mechanisms can mitigate external attacks, but they may not detect silent attacks from within the network.
- Silent attack is very important.
  - **Example:** Solarwind attack
    - Initiated through a software backdoor
    - This backdoor remained undetected over a long period of time.

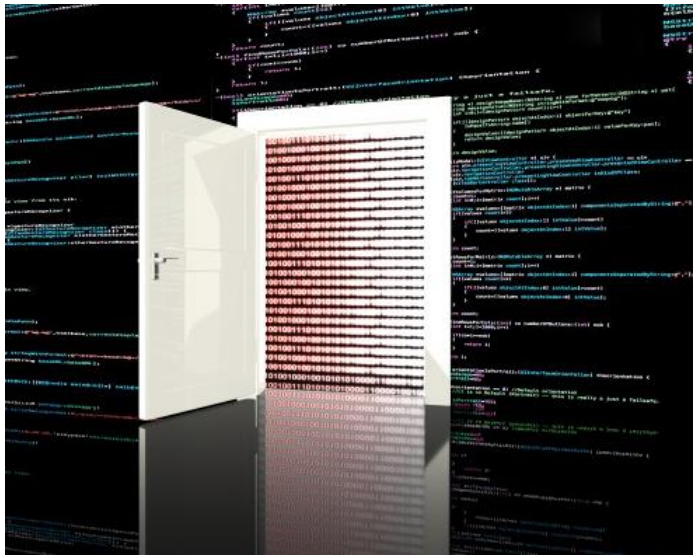

Undetectable Information leakage

- A back-doored switch could be instructed to slowly “leak” the network traffic that passes through it, to an external actor who will collect information about the network and/or users over time
- Detecting small leakages in particular is very difficult because the device does not deviate from its normal profile.

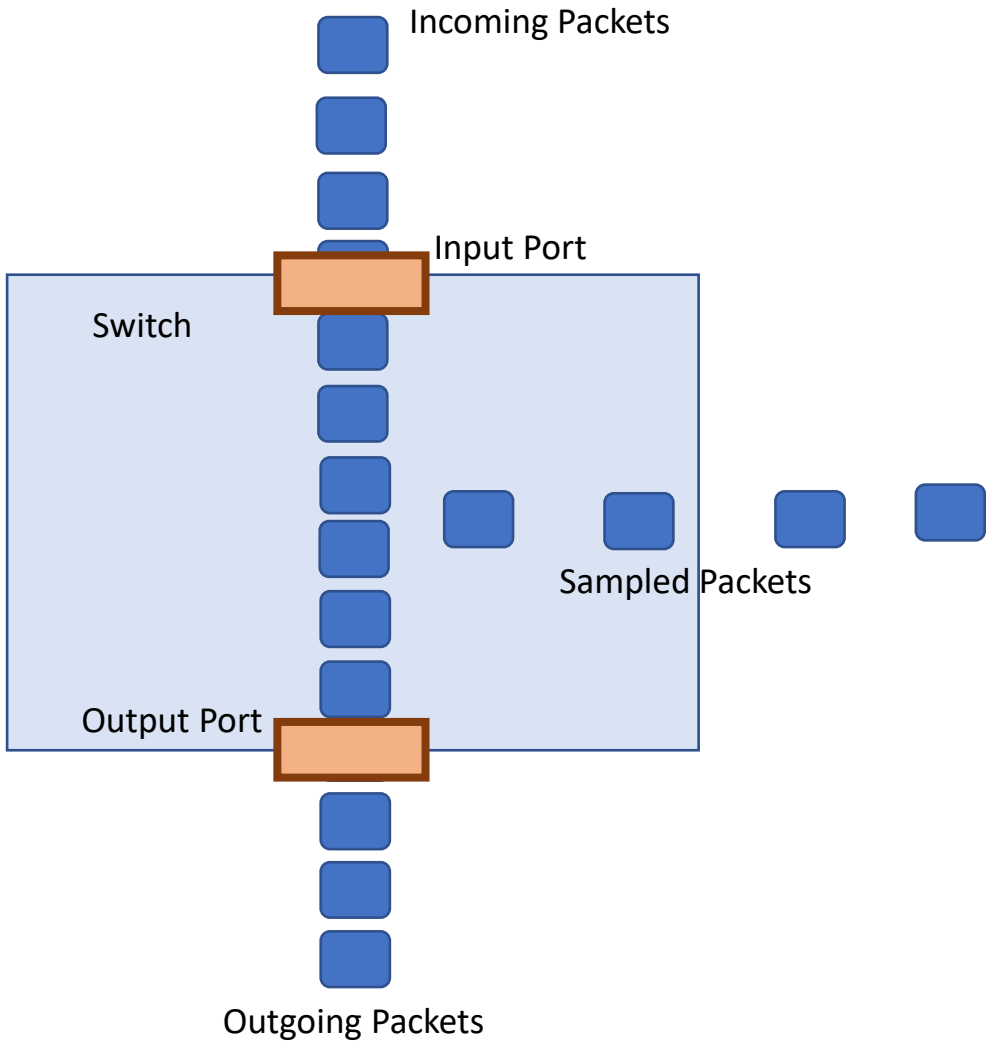

How to mitigate the information leakage by untrusted switches?

- **Solutions:**
  1. Using the end-to-end encryption at the edge router.
     - **Disadvantages:**
       - Significant complexity for key management at the edge of the network
       - Additional communication overheard of routing encrypted traffic
       - Computation cost of cryptographic protocols.
       - This solution will be even more costly in case of using quantum-secure algorithms.
  2. Excluding untrusted switches from communication and use only trusted switches if their locations are known
     - **Disadvantages:**
       - A drastic reduction of the network transmission capacity
       - In the worst case, communication is impossible.

A Promising Idea

- Some previous work has provided security against untrusted switches which capture all the traffic.
  - **Assumptions:**
    - There are k untrusted routers.
    - The sender knows the number of untrusted switches, but it does not know which switch is untrusted.
    - There are $k + 1$ node-disjoint paths between sender and receiver.
  - **Solution:**
    - The message is divided into $k + 1$ shares using $(k + 1, k + 1)$-secret sharing.
  - **Benefits:**
    - Post-quantum security
    - Efficient encoding of messages
    - Efficient reconstruction of messages

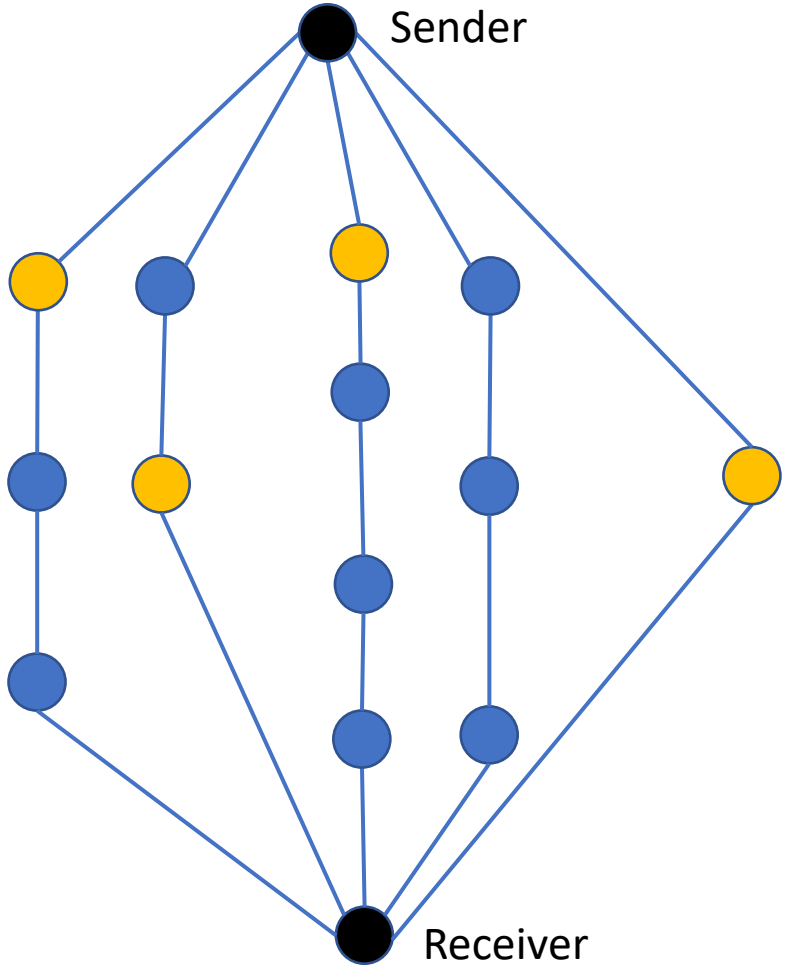

An example for $k = 4$

Network Model

- There are two edge switches, which are connected with multiple paths.
- Links have limited bandwidth (capacity).
- We assume communication is through packets of fixed size, each packet with fixed payload (and so fixed overhead).
- We assume packets are sent only from a sender edge switch to a receiver edge switch.
- The effect of other traffic in the network can be seen as reduced capacities of the links.
- We refer to payloads of packets sent to the sender edge switch as *messages*.

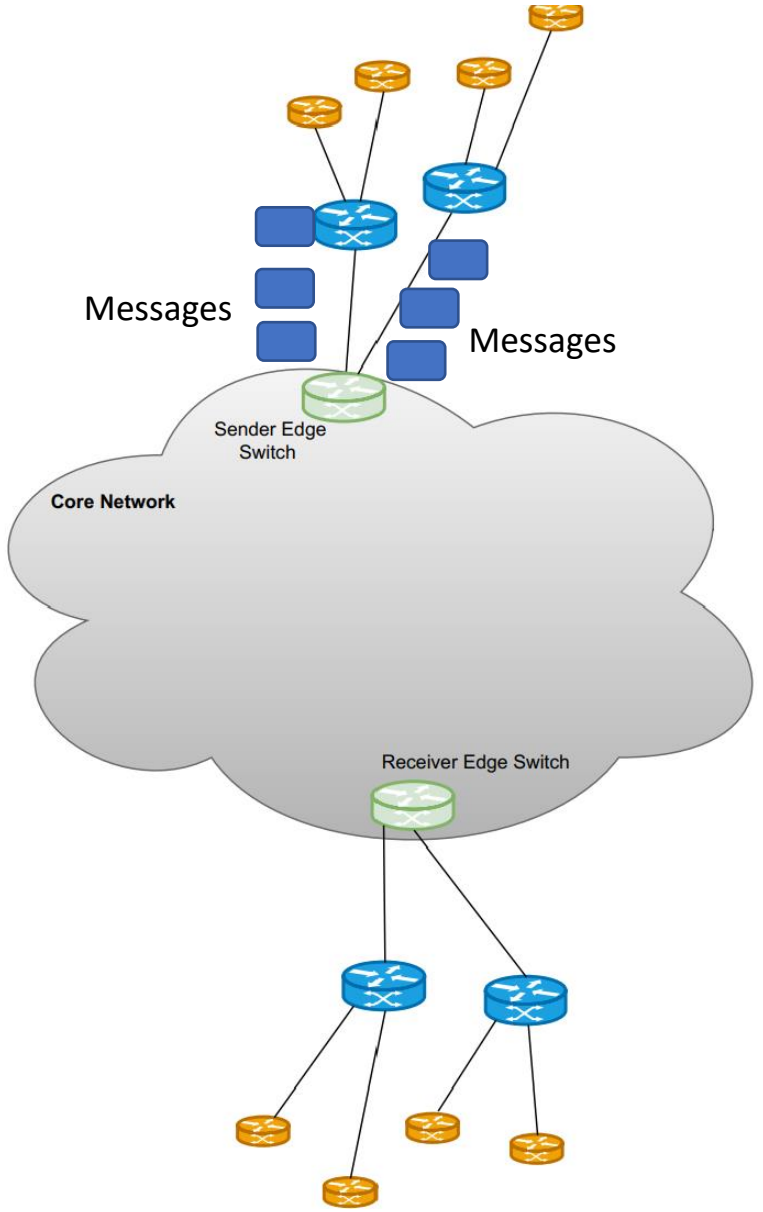

Network Model (Delays and Capacities)

- **Packet processing delay in switches:** negligible
- Link propagation delay and capacities are measured based on a short time interval, called *timeslot*.
- **Propagation Delay:**
  - It takes a number of timeslots for each packet to propagate along a link.
  - Propagation delay of a link is equal for all packets traversing it.
- **Link Capacity:** **T**he maximum number of packets which can be transmitted on the link within one timeslot.

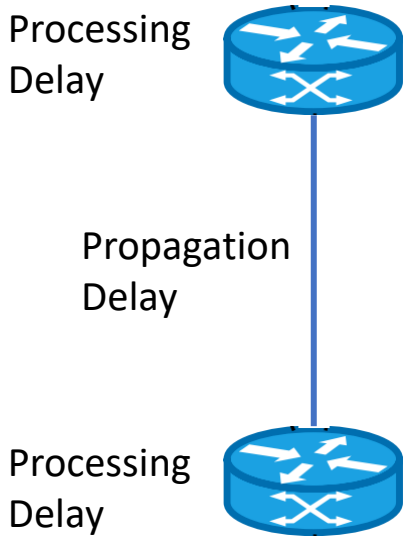

Network Model (Packet Drop)

- We do not model network buffers.
- Packets are not corrupted during transmission and propagation.
- Consider link l with capacity c is an outgoing link of switch s. If at a timeslot, this switch receives more than c packets for which their path include link l , then the excessive packets will be dropped.

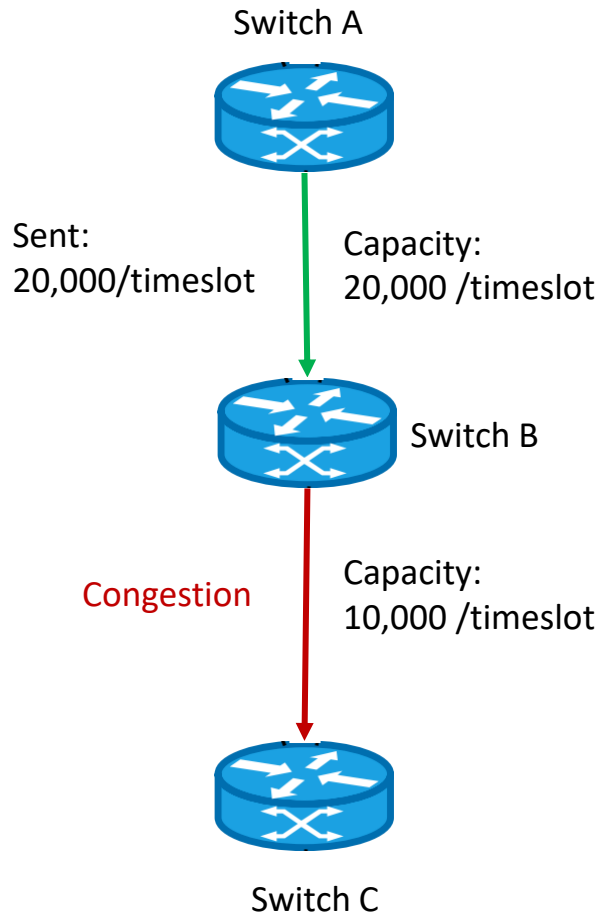

Threat Model

- A subset of switches in the network are leaky.
- A leaky switch samples each incoming packet with a fixed probability.
- All leaky switches leak information with the same probability.
- The network provider knows which switches are leaky and their sampling probability.
- All sampled packets are sent to a single adversary which is accessible from the network.

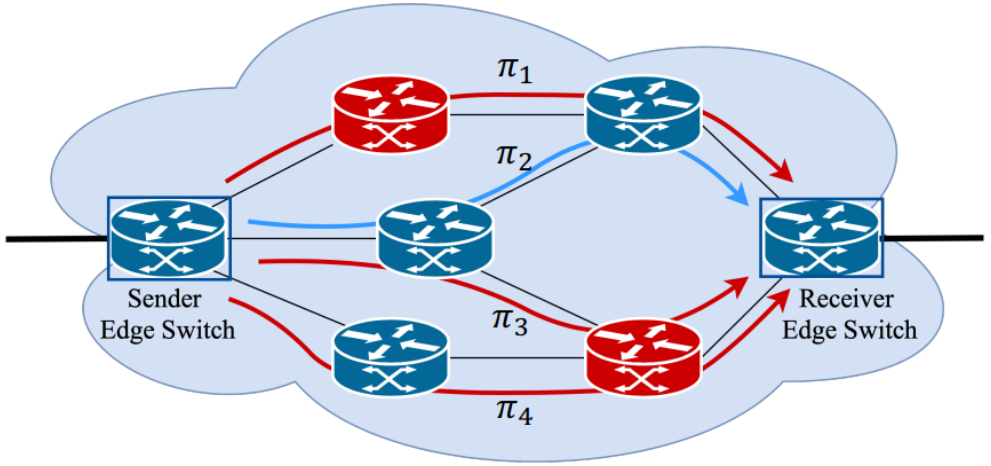

The red switches are untrusted.

Our Goals:

- Security Goal:
  - The probability of disclosing each message should not exceed a prespecified value, called leakage threshold (T).
  - Each message should be information-theoretically secure with probability $1 - T$.
- Reliability Goal:
  - Avoiding packet drop as much as possible.

Our Solution

- We propose a scheme, called Adaptive Multipath Secret Sharing (AMSS), to meet both security and reliability goals.
- This scheme divides each message into a number of shares and select path of each share such that the security guarantee is provided for the message.
- The number of shares is chosen on a per-message basis.
- Shares of a message can be sent over different paths or the same path.
- Generally, the number of shares is not independent of their paths.

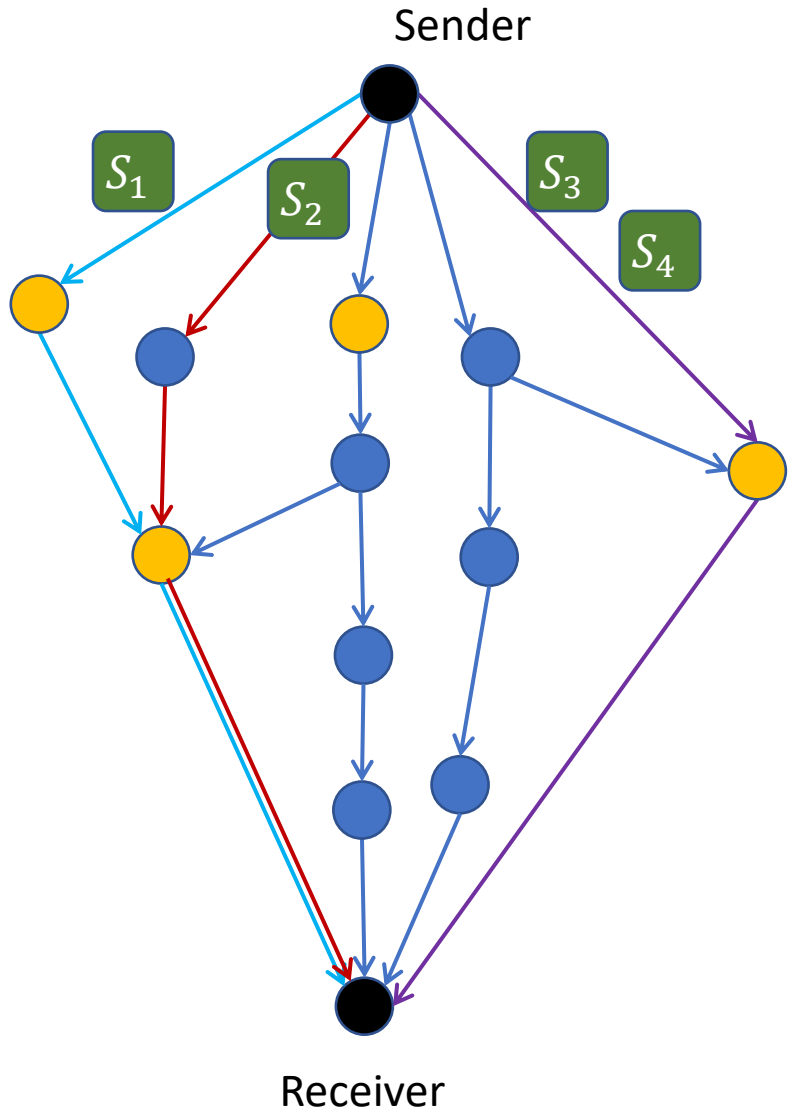

How does the AMSS scheme meet the security goal?

- Both secret sharing thresholds in AMSS are equal ((k, k)-secret sharing).
- Thus, in order to meet the security guarantee of a message, the probability of sampling all shares of a message should be less than or equal to the leakage threshold.
- This probability can be computed based on the number of shares sent on each path, the number of untrusted switches on each path, and the sampling probability:

$$\text{Leakage Probability of Path } i = 1 - (1-p)^{e_i}$$

$$\text{Message Leakage Probability} = \prod_{i=1,\, n_i>0}^{M} (1 - (1-p)^{e_i})^{n_i}$$

- The Security Condition:

$$\prod_{i=1,\, n_i>0}^{M} (1 - (1-p)^{e_i})^{n_i} \leq T$$

Notation	Definition
p	Sampling probability
T	Leakage threshold
M	Number of paths
e_i	Number of untrusted switches on path i
n_i	Number of shares sent on path i

How does the AMSS scheme avoid packet drops?

1. **Minimality Condition:** The share assignment should be minimal in the sense that removing even one share violates the security guarantee.
   - Thus, if $(n_1, n_2, \dots, n_M)$ is a share assignment vector, **none of** the following vectors provide the security condition.

$$(n_1 - 1, n_2, \dots, n_M)$$
$$(n_1, n_2 - 1, \dots, n_M)$$
$$\vdots$$
$$(n_1, n_2, \dots, n_M - 1)$$

2. Bandwidth-aware selection of the number of shares sent on each path

Trusted and Untrusted Paths

- Based on the leakage probabilities of paths and the leakage threshold, we label path as *trusted* and *untrusted*.
- **Trusted Path:** A path is trusted if its leakage probability is not greater than the leakage threshold.
  - Important Property: A message can be sent as a single share on a trusted path.
- **Untrusted Path:** A path is untrusted if its leakage probability is greater than the leakage threshold.

Share Assignment Vectors (SAVs)

- **Definition:** a vector of M non-negative integers which satisfy both the security and minimality conditions for a network with M paths and a specified leakage threshold.
- **Unit Share Assignment Vector:** Assigns only one share to only one trusted path
- **Non-Unit Share Assignment Vector:** Assigns shares to only untrusted paths. The total number of shares is more than one.
- The number of non-unit share assignment vectors grows super-exponentially with

$$\min\left(\left\lceil \log_{l_{min}} T \right\rceil, F\right)$$

- Thus, it is worthful to design an efficient algorithm to generate non-unit share assignment vectors.

Notation	Definition
F	Number of untrusted paths
T	Leakage threshold
l_{min}	Minimum path leakage probability among untrusted paths

Proposed Adaptive Multipath Secret Sharing (AMSS) Scheme

- AMSS has three phases.
- Given some theoretical results, we expect that phase 1 is NP-complete.
- **Secrecy Capacity:** The maximum number of messages which can be securely sent over the network per timeslot such that no link is overloaded.
- More precisely, in phase 2, we compute the maximum number of messages which can be sent using each SAV per timeslot such that no link is overloaded.

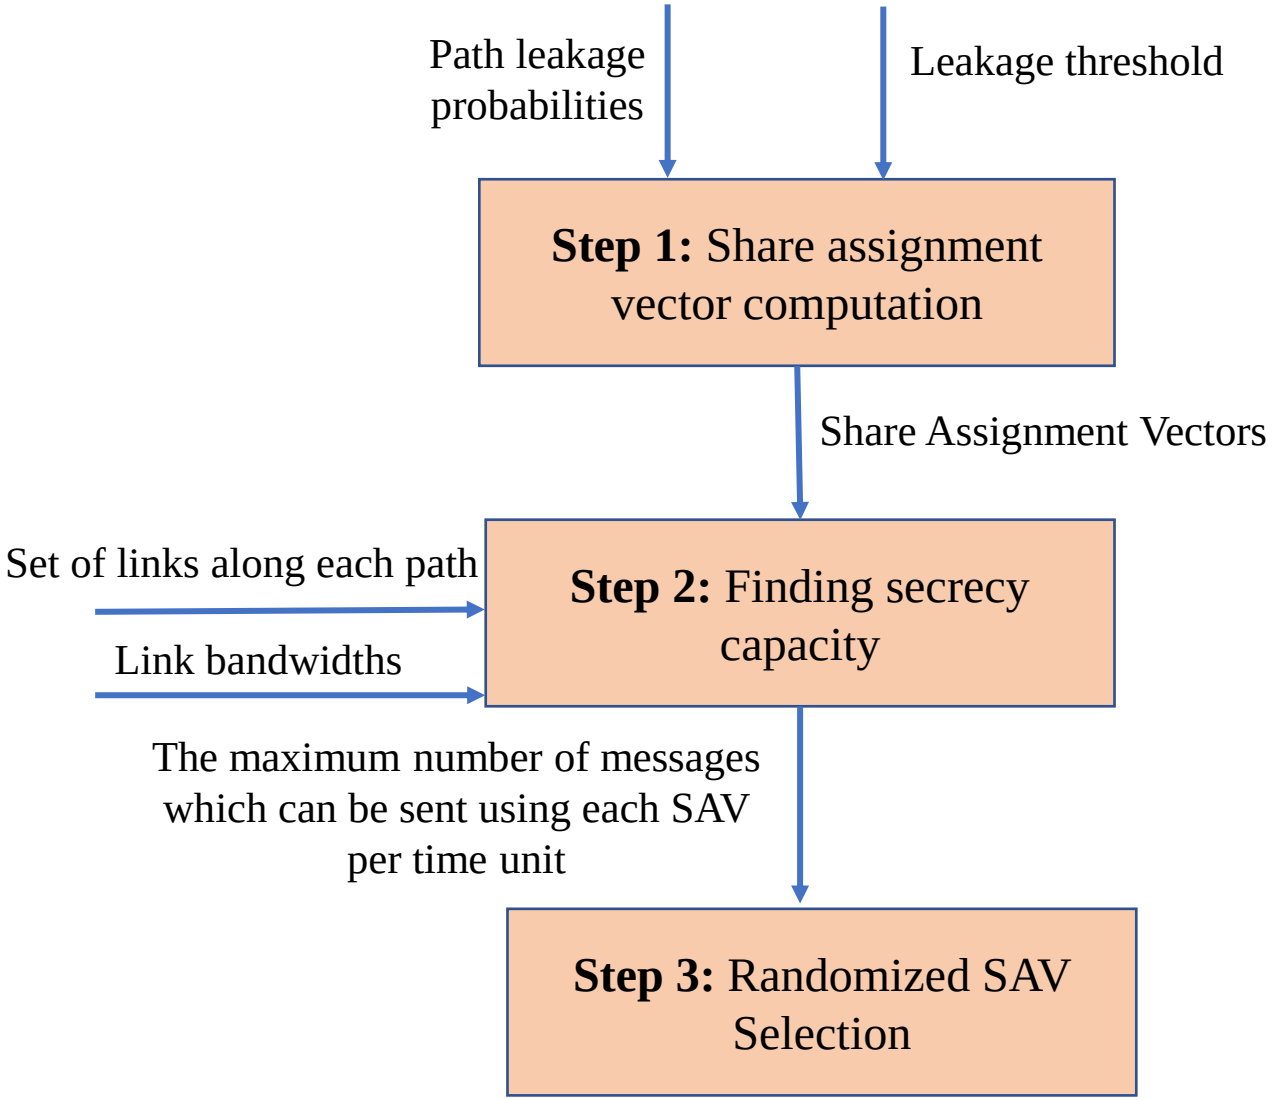

Phase 1: Computation of All Non-Unit SAVs

- For a network with M paths, we define a finite set of non-negative integer valued vectors with length M which includes all SAVs.
- Then, we propose an algorithm which iterates once over this set to find the SAVs.
- The list of vectors according to which this iteration is done has the following property:
  - A vector is a share assignment *iff* it satisfies the security condition and no previously found SAVs is strictly less than this vector according to a predefined partial order.
- **Example:**

Finite list of vectors: $v_1, v_2, v_3, \ldots, \boldsymbol{v_i}, v_{i+1}, \ldots, v_n$

SAVs found so far: $v_2, v_{10}, \ldots, v_j$ $(j < i)$

If v_i satisfies the security condition and none of $v_2, v_{10}, \ldots, v_j$ is strictly less than v_i, then v_i is added to the list of SAVs as follows.

Finite list of vectors: $v_1, v_2, v_3, \ldots, v_i, \boldsymbol{v_{i+1}}, \ldots, v_n$

SAVs found so far: $v_2, v_{10}, \ldots, v_j, v_i$

Complexity of Phase 1

Notation	Definition
F	Number of untrusted paths
T	Leakage threshold
l_{max}	Maximum leakage probabilities among untrusted paths

- $\log_{l_{max}} T$ is the minimum number of shares required to satisfy the security condition if share are sent on the most untrusted path.
- **Complexity:**

$$O\left(F\left(\log_{l_{max}} T\right)^2 \times \min\left(F^F, \left(\log_{l_{max}} T\right)^{2F-2}\right)\right)$$

Phase 2: Finding Secrecy Capacity

Problem 1: Computation of the Secrecy Capacity

Given a network with M paths and the links along each one with their capacities, the computation of the secrecy capacity can be formulated as follows.

Input: Set of links which are on path i represented as Γ_i,
Link capacities $(c_1, c_2, \ldots, c_z)$,
SAVs represented as $(n_1^i, n_2^i, \ldots, n_M^i)$, $1 \leq i \leq I$.
Output: $N_{max}, x_1^*, x_2^*, \ldots, x_I^*$

$$\max_{x_1, x_2, \ldots, x_I} N \tag{9a}$$

$$\text{s.t. } N = \sum_{i=1}^{I} x_i \tag{9b}$$

$$\sum_{j=1, k \in \Gamma_j}^{M} \sum_{i=1}^{I} (x_i \cdot n_j^i) \leq c_k, \ \forall k : 1 \leq k \leq z \tag{9c}$$

$$x_1, x_2, \ldots, x_I \in \mathbb{Z}^* \tag{9d}$$

Zero Packet Loss Guarantee

- **Theorem:** If the edge switch sends at most x_i^* messages using share assignment vector i per timeslot, $1 \leq i \leq I$, then no packet is dropped.

Output: $N_{max}, x_1^*, x_2^*, \ldots, x_I^*$

$$\max_{x_1, x_2, \ldots, x_I} N$$

$$\text{s.t. } N = \sum_{i=1}^{I} x_i$$

$$\sum_{j=1, k \in \Gamma_j}^{M} \sum_{i=1}^{I} (x_i \cdot n_j^i) \leq c_k, \ \forall k : 1 \leq k \leq z$$

$$x_1, x_2, \ldots, x_I \in \mathbb{Z}^*$$

- **Key assumptions of the network model used in proof:**
  - Processing delay in switches are negligible.
  - It takes a number of timeslots for each packet to propagate along a link.
  - Propagation delay of a link is equal for all packets traversing it.

How to reduce packet Loss when assumptions do not hold?

- If one of the assumptions do not hold, some packets may be dropped even if the total number of messages sent per timeslot does not exceed the secrecy capacity and buffers are considered.
- In order to reduce the packet loss, we should avoid sending a burst of packets on the same path.
- Thus, the order according to which the SAVs are selected for sending messages matters.
- Sending a large number of messages according to the same SAV may result in buffer overflow.
- Thus, at each timeslot, we consider a pool of SAVs in which the initial number of each SAV equals the number obtained from optimization (x_1^*, x_2^*, …, x_I^*).
- As a message arrives, we select one SAV from the pool uniformly at random and take it out of pool.

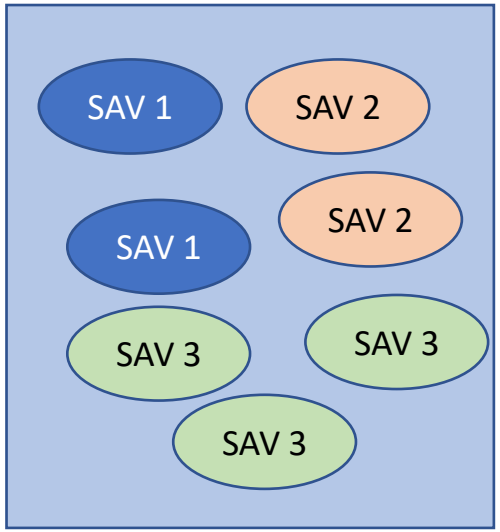

The initial items in pool when the number of SAVs is three, $x_1^* = x_2^* = 2$, and $x_3^* = 3$.

Baselines

- Only Trusted Switches (OTS) Scheme:
  - This scheme use only trusted switches for communication.
  - It behaves similar to AMSS for zero leakage threshold.
  - No message is leaked under OTS scheme.
- Multipath Secret Sharing Scheme:
  - This scheme choose a maximal set of node-disjoint paths hoping that at least one path consists of only trusted switches.
  - It divides each message into as many shares as the number of paths.
  - One share is sent over each path.

Mininet Experiments

- **Mininet** is a network emulator that creates a realistic virtual network composed of virtual hosts, switches, controllers, and links.
- Thus, experimental results obtained in Mininet are expected to closely resemble those obtained
  in a physical network.
- Network Setting:
  - Three switches are leaky and sample each packet with probability 0.01.
  - Link bandwidth are in the order are *Mbps*.
  - Packet size and message size are 500 *B* and 454 *B*, respectively.
  - Default leakage threshold is 0.0001.
  - Propagation delay of each link equals 1 *ms*.
  - Capacity of each port buffer is 50 *KB*.

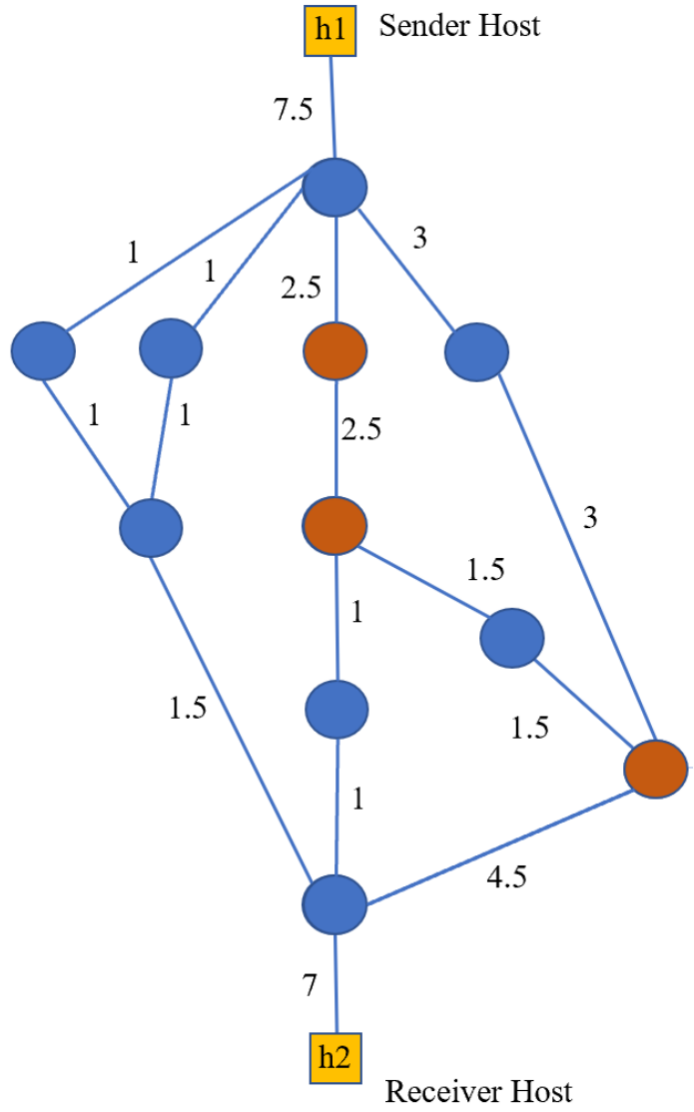

Comparison with Baselines

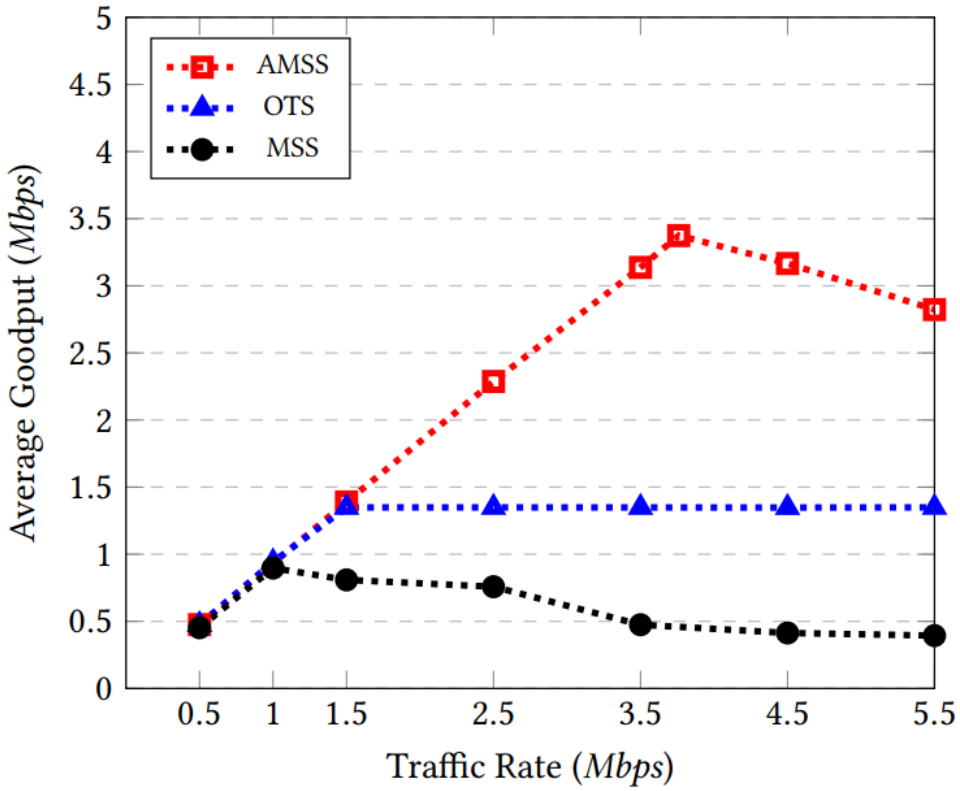

Scheme	Secrecy Capacity (*Mbps*)	Traffic Rate (*Mbps*)
AMSS	3.446	3.795
MSS	0.908	1.000
OTS	1.362	1.500

Evidence for Providing the Required Security Guarantee by AMSS

- We sent a sequence of 2 million messages according to AMSS with a speed which was less than the secrecy capacity.
- Then, we measured the leakage rate as follows.

$$\text{Leakage Rate} = \frac{\text{Number of disclosed messages}}{\text{Number of messages sent according to nonunit SAVs}}$$

- **Expectation:**

$$\text{Leakage Rate} \leq \text{Leakage Threshold}$$

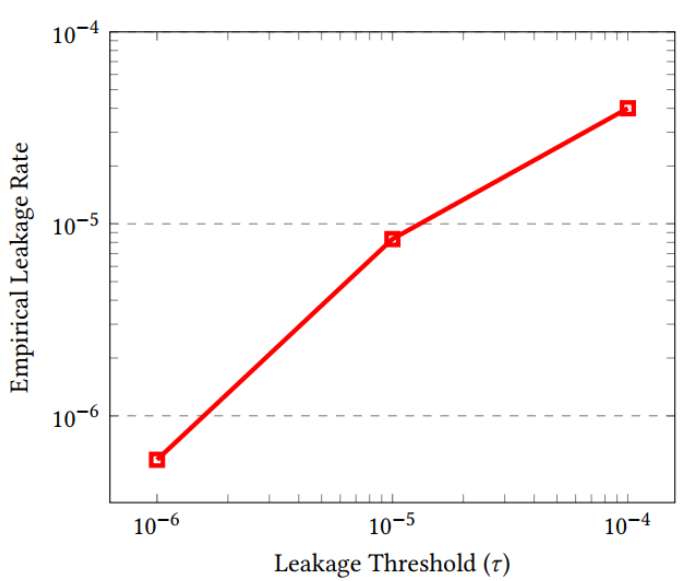

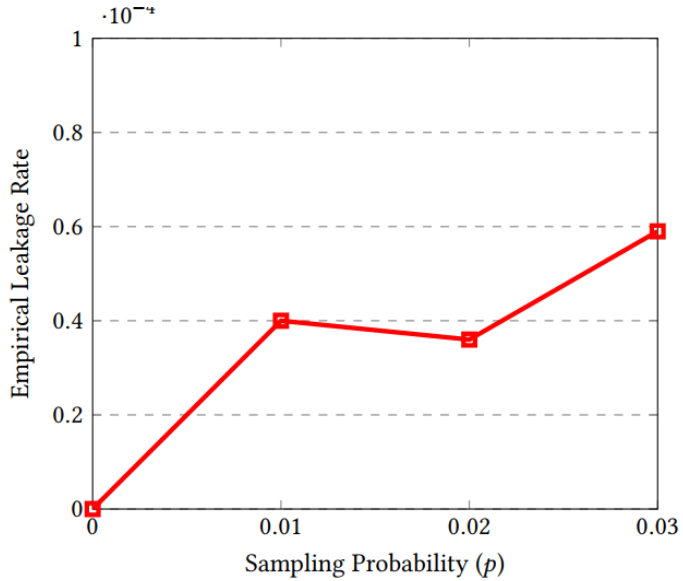

Scalability Analysis Using Discrete-Event Simulation

- We chose two switches in a real-world ISP topology as the sender and receiver switches.
- Then, we computed the 15 shortest paths between the sender and receiver switches.
- 30% of switches which are along the 15 shortest paths were untrusted (with sampling probability 0.0001).
- Bandwidth of all links except the two links shown by dashed style were 10 Gbps as reported in the dataset.
- Bandwidth of the other two links were 2.4 Gbps.
- Two hosts are connected to the sender and receiver switches to send and receive the traffic using the network.

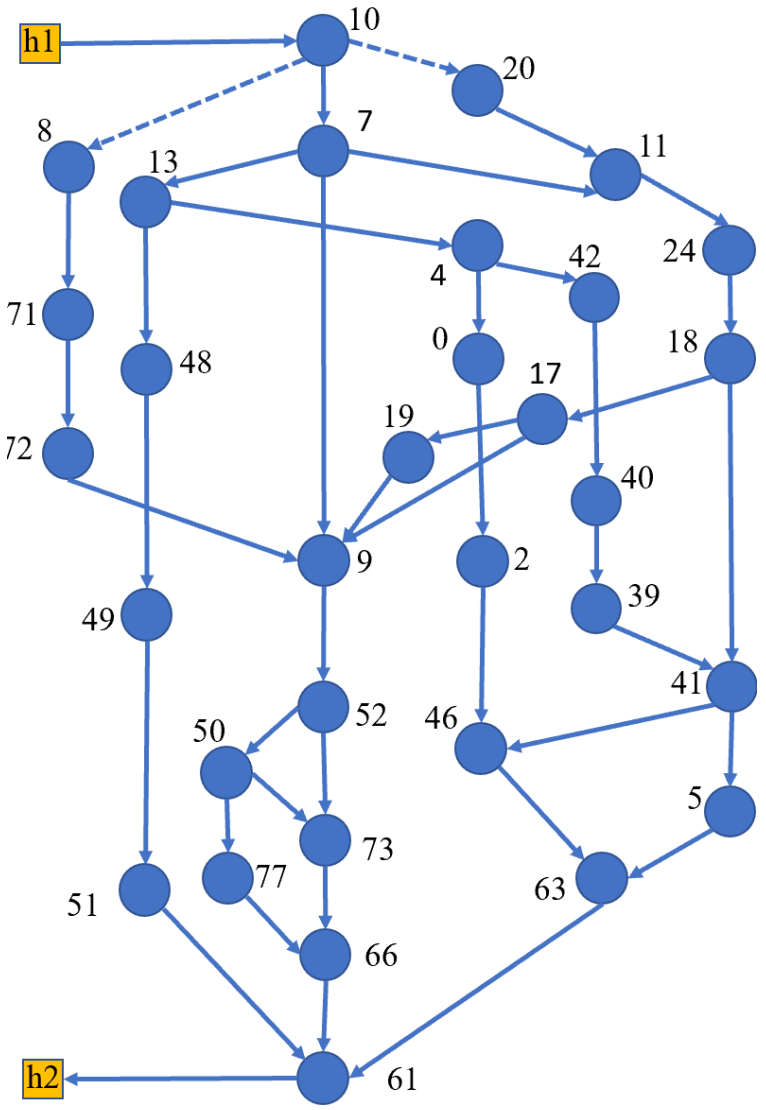

AMSS achieves the highest secrecy capacity!

- The distribution of the untrusted switch over paths impacts on the secrecy capacity.
- The following figure shows the secrecy capacity for different positioning of the leaky switches.
- In each experiment, 30% of switches were untrusted.

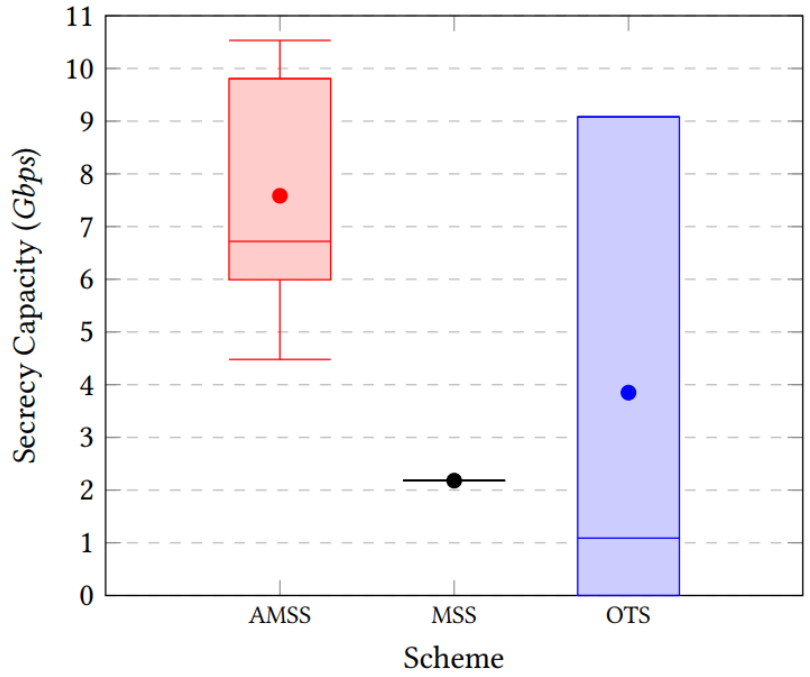

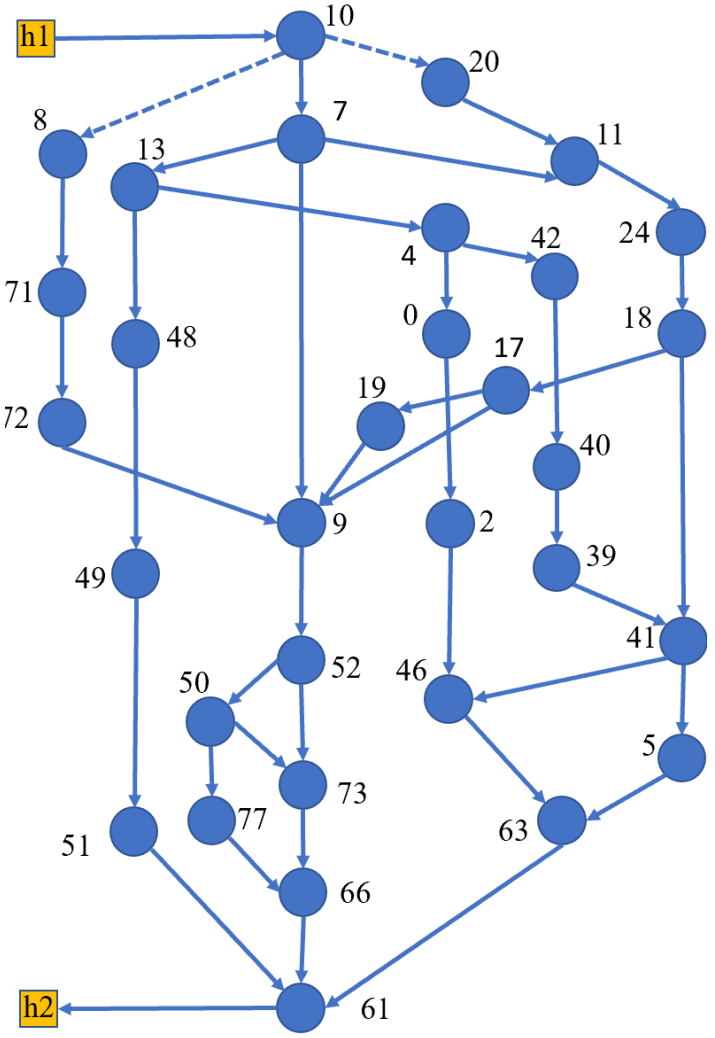

Effect of the Number of Untrusted Switches on the Secrecy Capacity

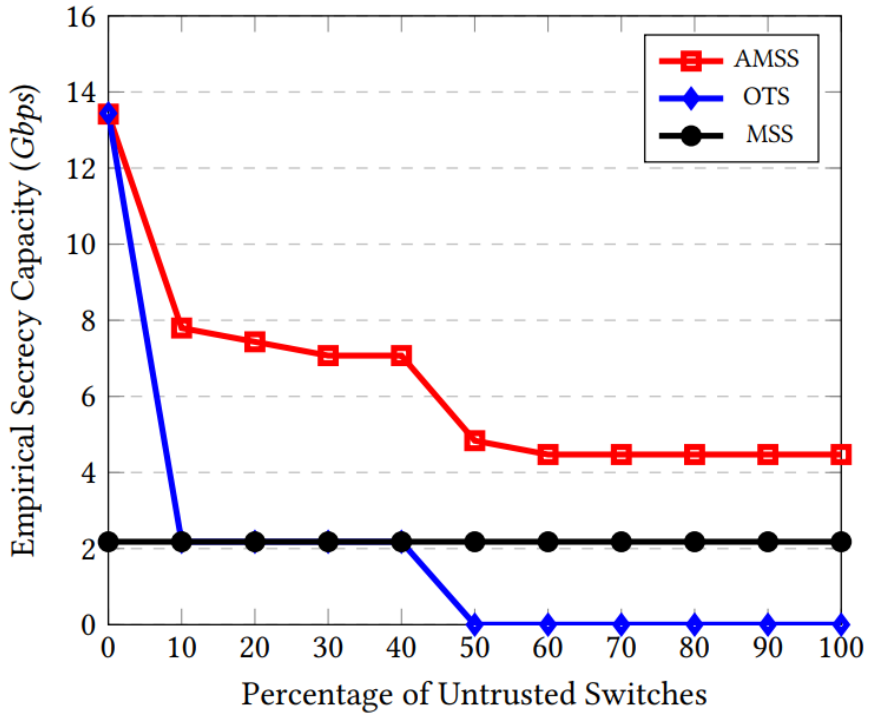

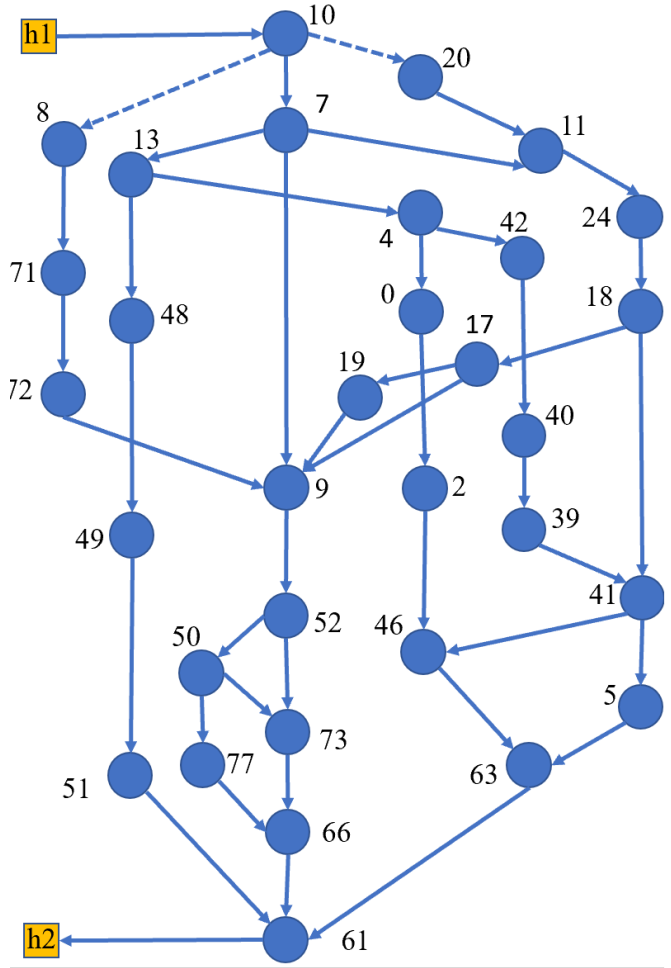

AMSS reduces the packet loss rate.

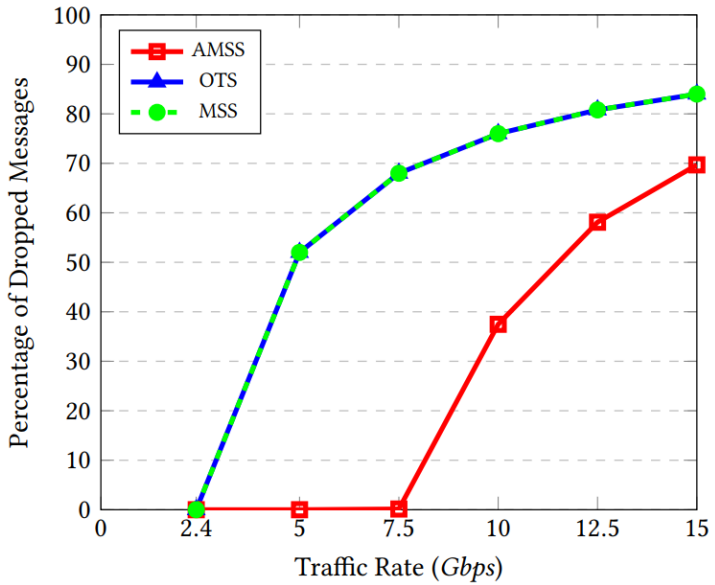

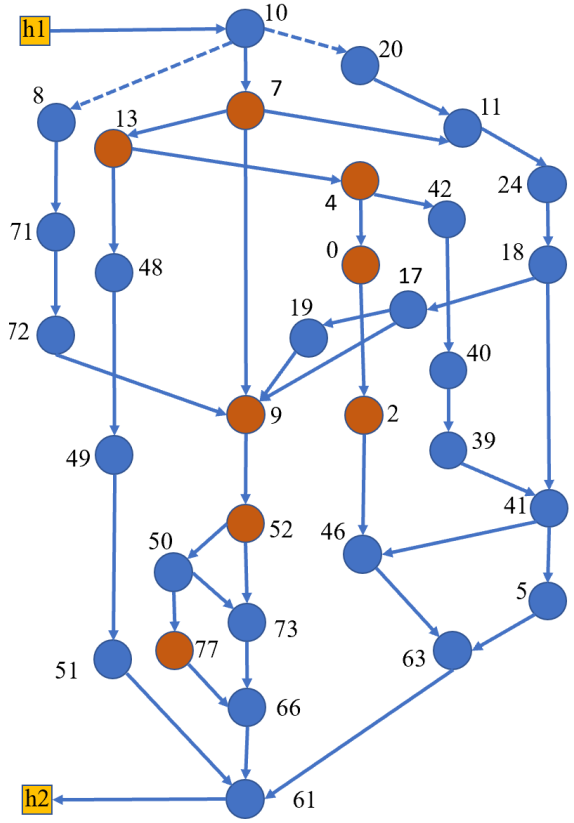

Table: The secrecy capacity of different schemes when packet length is 500 B out of which 46 B is header. The third column shows the corresponding traffic rate including headers

Scheme	Secrecy Capacity (*Gbps*)	Traffic Rate (*Gbps*)
AMSS	7.082	7.500
OTS	2.179	2.400
MSS	2.179	2.400

Conclusion

- In this talk, we proposed an approach, called AMSS, to mitigate the untrusted switches which leak information in an undetectable way.
- The proposed approach is based on multipath routing and secret sharing and limits the probability of leaking each message to a threshold.
- Messages may be divided into different number of shares under this approach.
- AMSS maximizes the number of messages which can be sent into the network over a timeslot securely without overloading links.
- Experimental results shows the advantage of AMSS over the baseline with respect to reducing packet loss rate and increasing the secrecy capacity.

Thanks for your attention

Any question?