


*Fast and Cautious:*  
**Leveraging Multi-path Diversity  
for Transport Loss Recovery  
in Data Centers**

**Guo Chen**

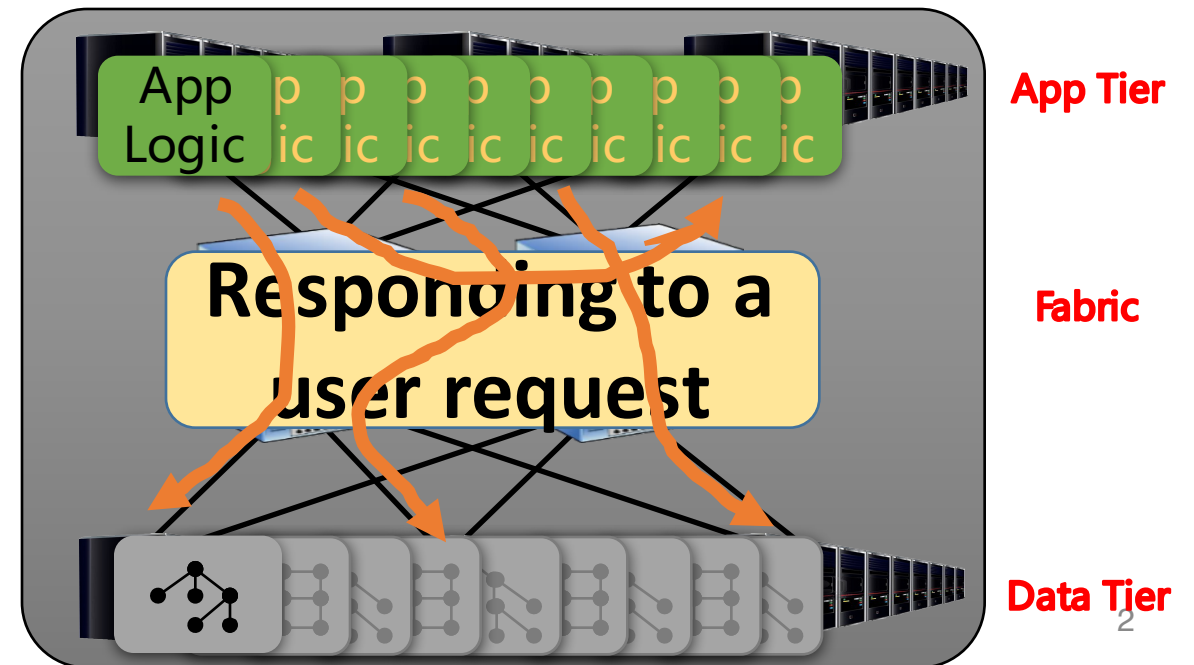
Yuanwei Lu, Yuan Meng, Bojie Li, Kun Tan, Dan Pei,  
Peng Cheng, Layong (Larry) Luo, Yongqiang Xiong,  
Xiaoliang Wang, and Youjian Zhao



# Motivation

- Services care about the **tail flow completion time** (tail FCT)
  - **Large number of flows** generated in each operation
  - Overall performance governed by the **last completed flows**

Large-scale web application  
hosted in  
**Data Center Network (DCN)**




# Motivation


- Services care about the **tail flow completion time** (tail FCT)
  - **Large number of flows** generated in each operation
  - Overall performance governed by the **last completed flows**
- But packet loss **hurts tail FCT**
  - **Real case** in a Microsoft Azure's DCN

DCN tail latency visualization

[Pingmesh (SIGCOMM'15)]



(a) Normal



(b) Spine failure

Spine switch **2%** random drop rate --> increase of 99<sup>th</sup> percentile latency of **all users**

# Outline

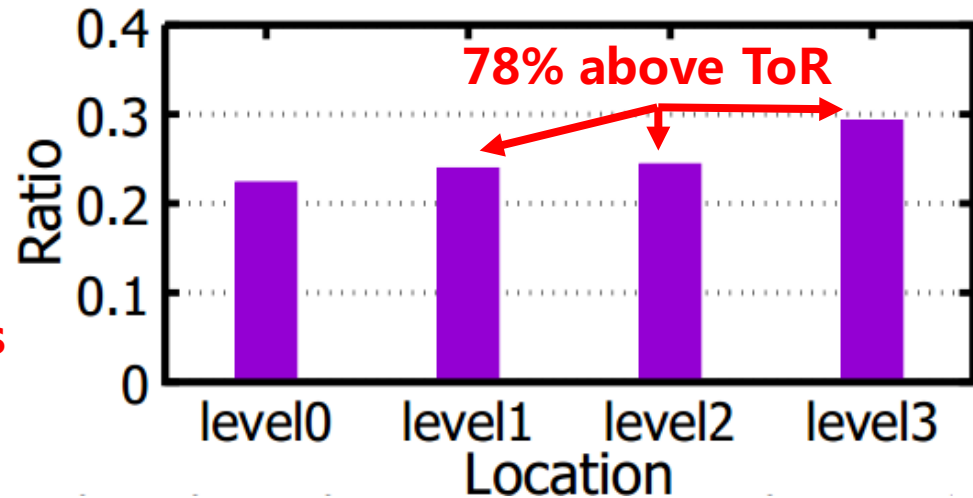
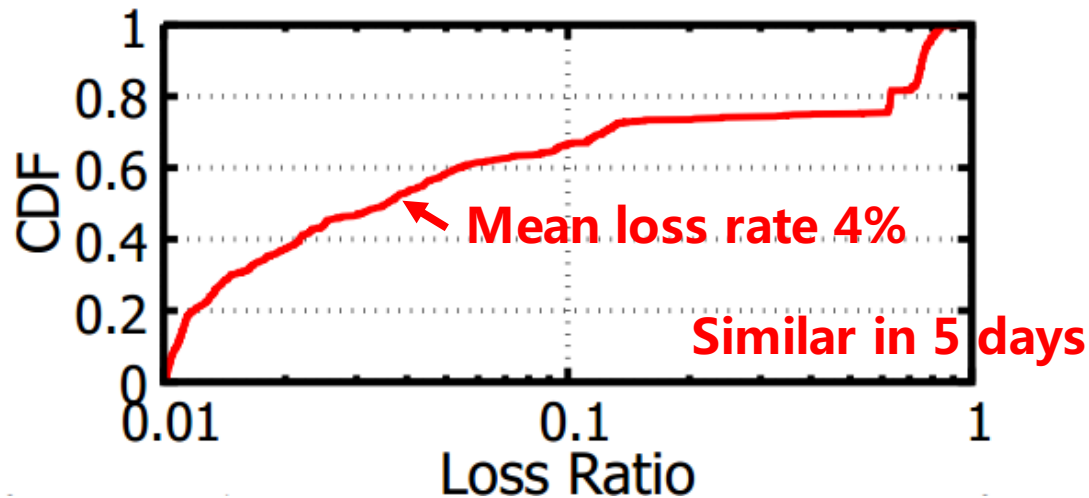
---

- Motivation
- **Packet Loss in DCN**
- Impact of Packet Loss
- Challenge for Loss Recovery
- FUSO Design
- Evaluation
- Summary

# Packet Loss in DCN

## ■ Loss characteristics

- Measured in a Microsoft production DCN during Dec. 1<sup>st</sup>-5<sup>th</sup>, 2015



*Loss rate and location distribution of lossy links (loss rate > 1%)*

- 1) Loss **frequently** happens (the overall loss rate is low)
- 2) Most losses happen **in the network** instead of the edge

# Packet Loss in DCN

## ■ Reasons causing loss

### □ Congestion loss

**Bursty; Transient**

- Uneven load-balance
- Incast

→ **Greatly mitigated  
(e.g., 1% → 0.01%)**

[Jupiter Rising SIGCOMM'15]

### □ Failure loss

**Complex; Hard to detect**

- Silent random drop
- Packet black-hole

→ **Common  
& Huge impact  
on performance**

[Pingmesh SIGCOMM'15]

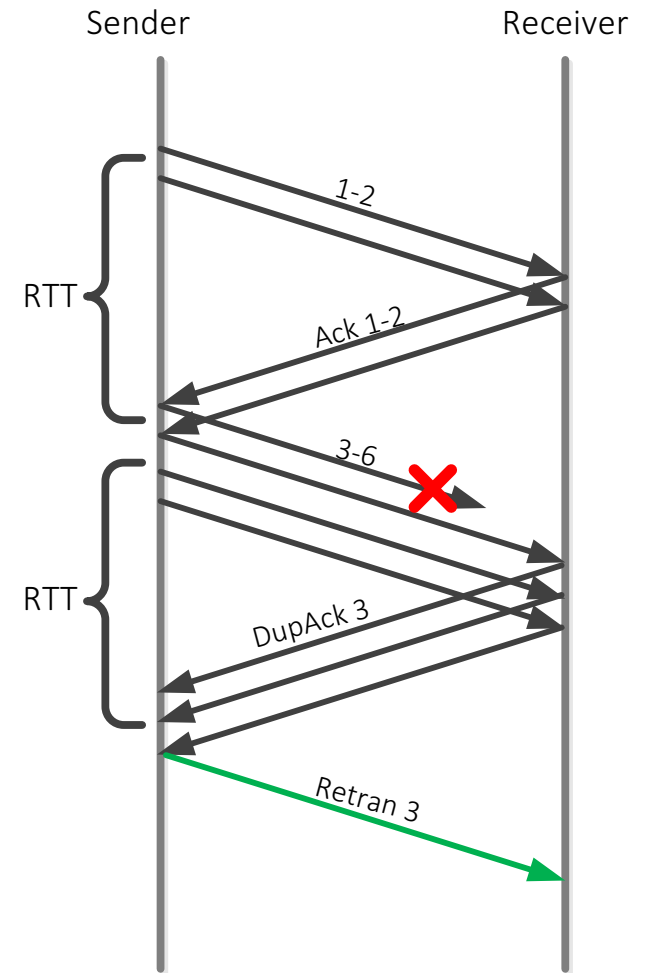
# Outline

---

- Motivation
- Packet Loss in DCN
- **Impact of Packet Loss**
  - **Why loss hurts the tail?**
  - **How hard loss hurts?**
- Challenge for Loss Recovery
- FUSO Design
- Evaluation
- Summary

# How TCP Handles Loss?

- Fast recovery
  - Wait for certain number of DACKs to detect the loss and retransmit






# How TCP Handles Loss?

- Fast recovery
  - Wait for certain number of DACKs to detect the loss and retransmit
- Timeout (RTO)
  - If not enough DACKs return, retransmit after a timeout


**RTO >> RTT** e.g. RTO=5ms, RTT<100us  
[Pingmesh (SIGCOMM'15), DCTCP (SIGCOMM'10)]



# How TCP Handles Loss?

- Fast recovery
  - Wait for certain number of DACKs to detect the loss and retransmit
- Timeout (RTO)
  - If not enough DACKs return, retransmit after a timeout


**RTO >> RTT** e.g. RTO=5ms, RTT<100us  
[Pingmesh (SIGCOMM'15), DCTCP (SIGCOMM'10)]



Encountering one **RTO** →  
dramatically increase the FCT

# Loss Incurs Timeout

- A little loss causes enough timeout to hurt the tail FCT




Timeout probability of flows with different sizes passing a path with different packet loss rate

1. 1% loss → more than 1% flows timeout
2. Larger flows (e.g. 100KB)
  - a. timeout ratio sharply grows when loss rate > 1%

# Loss Incurs Timeout

- A little loss causes enough timeout to hurt the tail FCT



Timeout probability of flows with different sizes passing a path with different packet loss rate

To avoid **RTO**

# Outline

---

- Motivation
- Packet Loss in DCN
- Impact of Packet Loss
- **Challenge for Loss Recovery**
- FUSO Design
- Evaluation
- Summary

# Challenge for TCP Loss Recovery

- Prior works **add aggressiveness** to congestion control to do loss recovery **before** timeout (RTO)
  - Tail Loss Probe (TLP) [SIGCOMM'13, RFC 5827]
    - transmit one prober **after 2RTT**
  - Instant Recovery (TCP-IR) [SIGCOMM'13, RFC 5827]
    - generate **an FEC packet** for every group of packets (up to 16)
    - FEC packets also act as probers, **delayed 1/4RTT** before sent
  - Proactive/RepFlow [SIGCOMM'13, INFOCOM'14]
    - **Duplicate** every packet/flow

# Challenge for TCP Loss Recovery

- **How long to wait** before sending recovery packets?
  - For congestion loss
    - Should **delay enough** in case of worsening congestion

**Bursty:  
Lead to multiple consecutive losses**

[Incast (WREN'09), DCTCP (SIGCOMM'10)]

# Challenge for TCP Loss Recovery

- **How long to wait** before sending recovery packets?
  - For congestion loss
    - Should **delay enough** in case of worsening congestion
  - For failure loss such as random drop
    - Should recover as **fast** as possible, otherwise already increase the FCT

- **Wait 2RTT is too costly** [TLP SIGCOMM'13, RFC 5827]
- **Accurate & high-precision RTT measurement is challenging**



# Brief Summary

- Loss easily incurs **timeout** to hurt the tail
- To prevent timeout, prior works **add fixed aggressiveness** to recover loss before timeout
- Hard to adapt to various loss conditions
  - Should be **fast** for failure loss
  - Should be **cautious** for congestion loss

How to **accelerate loss recovery** as soon as possible, under various loss conditions **without causing congestion?**

# Outline

---

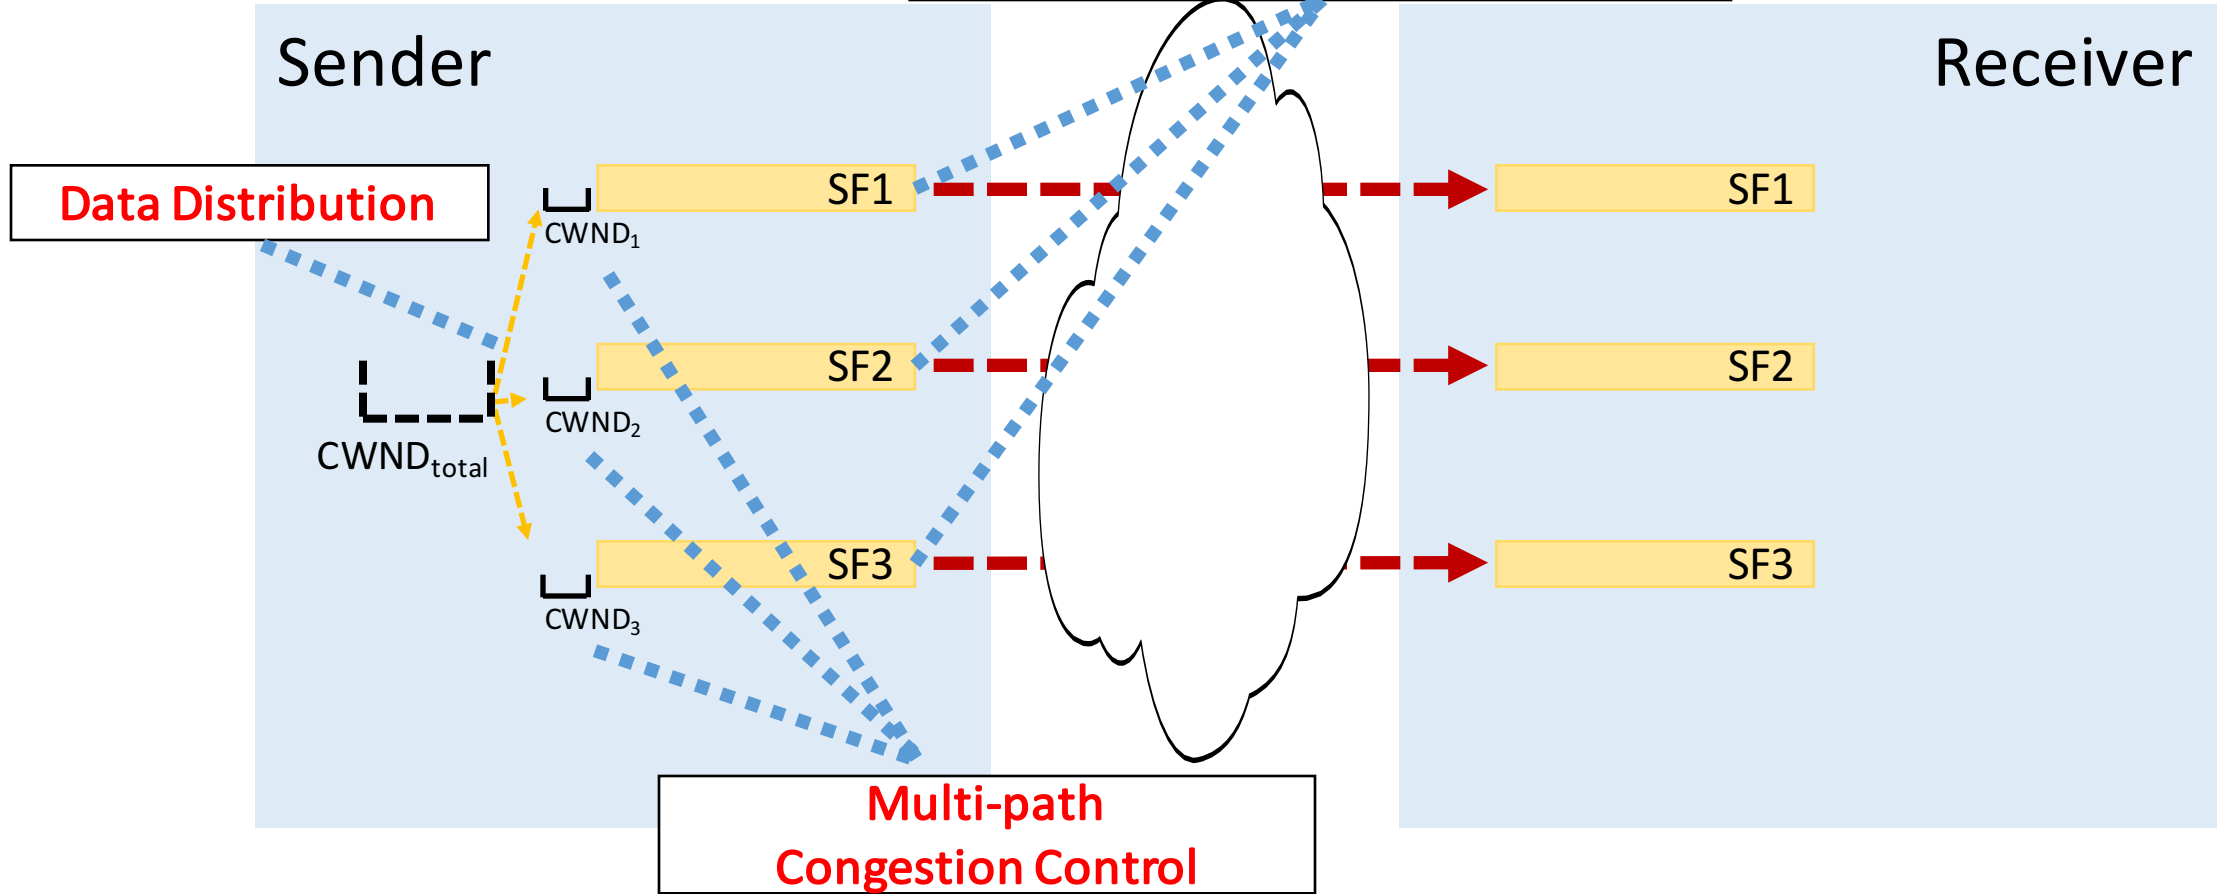
- Motivation
- Packet Loss in DCN
- Impact of Packet Loss
- Challenge for Loss Recovery
- **FUSO Design**
- Evaluation
- Summary

# FUSO: Fast Multi-path Loss Recovery

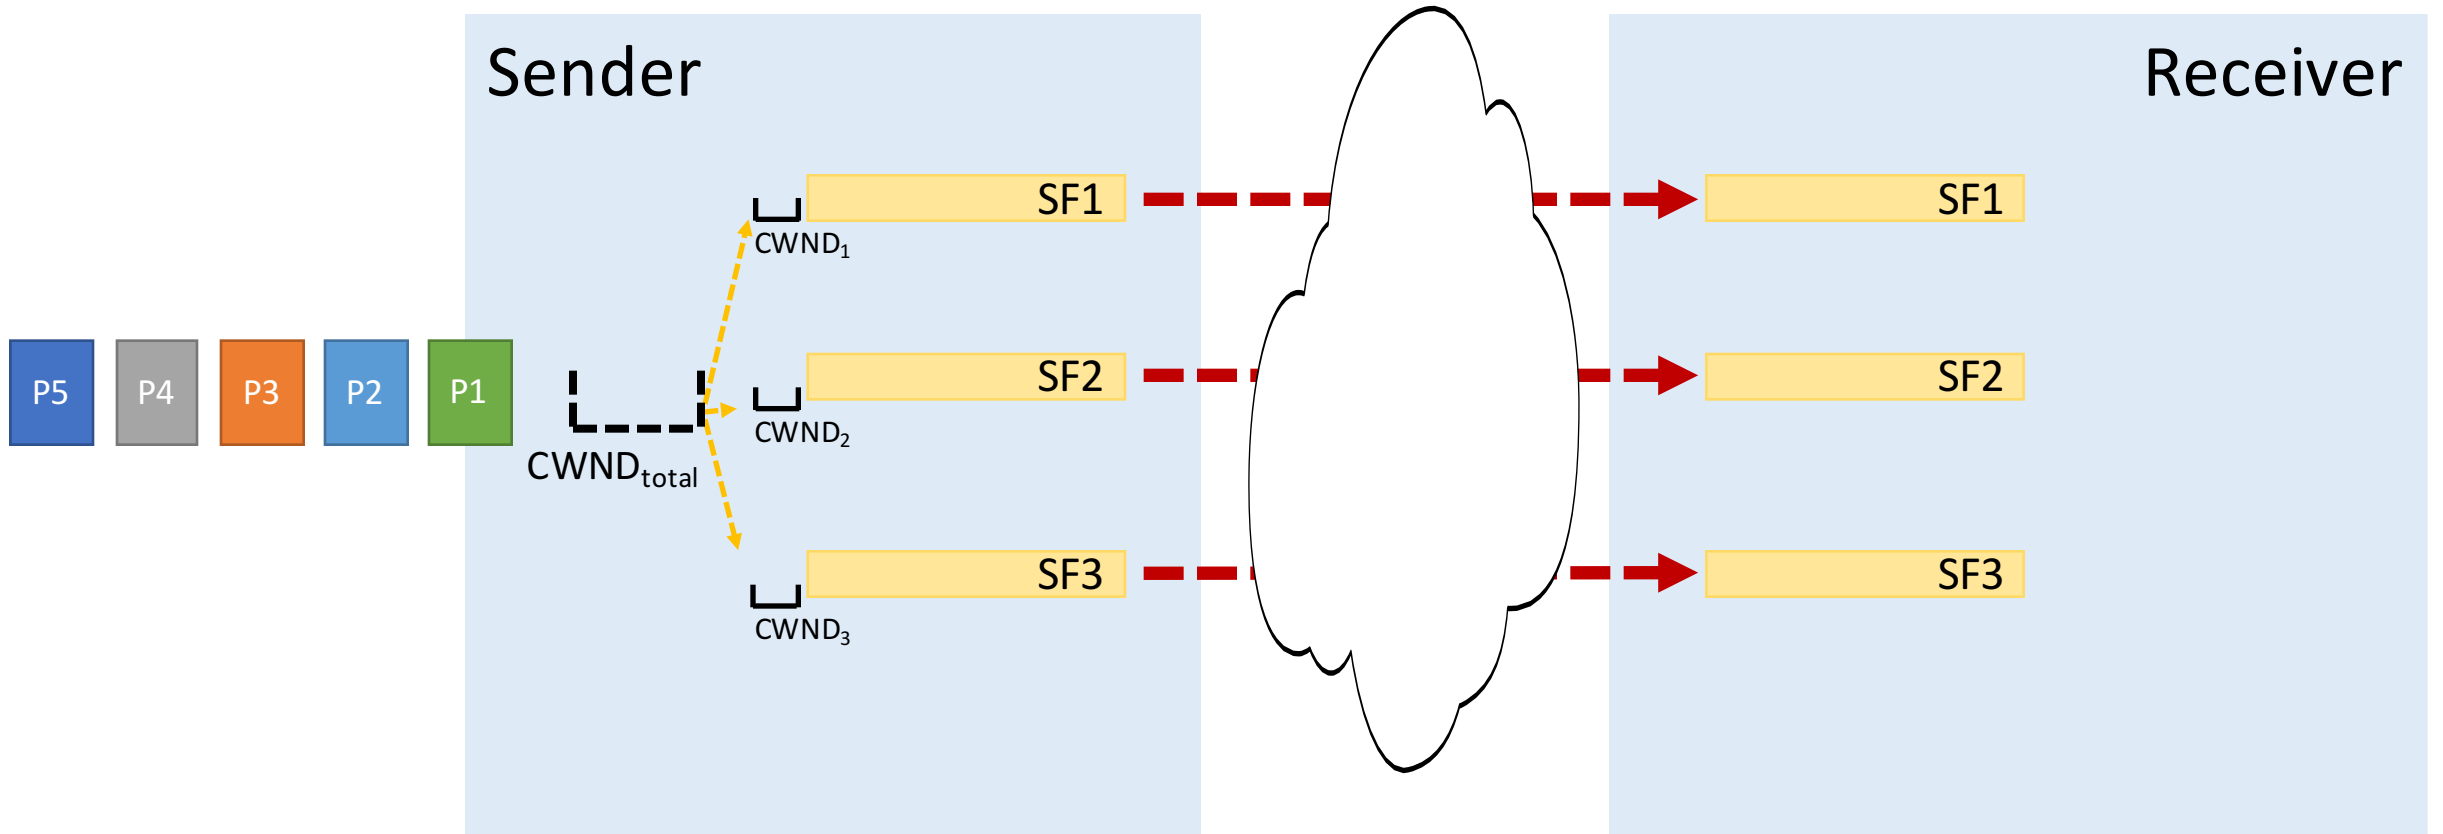
- Utilize the “good” paths to proactively conduct loss recovery for “bad” paths
  - Leveraging path diversity (multiple paths; a few encounter loss)
  
- Fast and Cautious
  - **Fast**
    - Proactive (immediate) recovery for potential packet loss utilizing spare transmission opportunity
  - **Cautious**
    - Strictly follow congestion control without adding aggressiveness

# Multi-path Transport Background

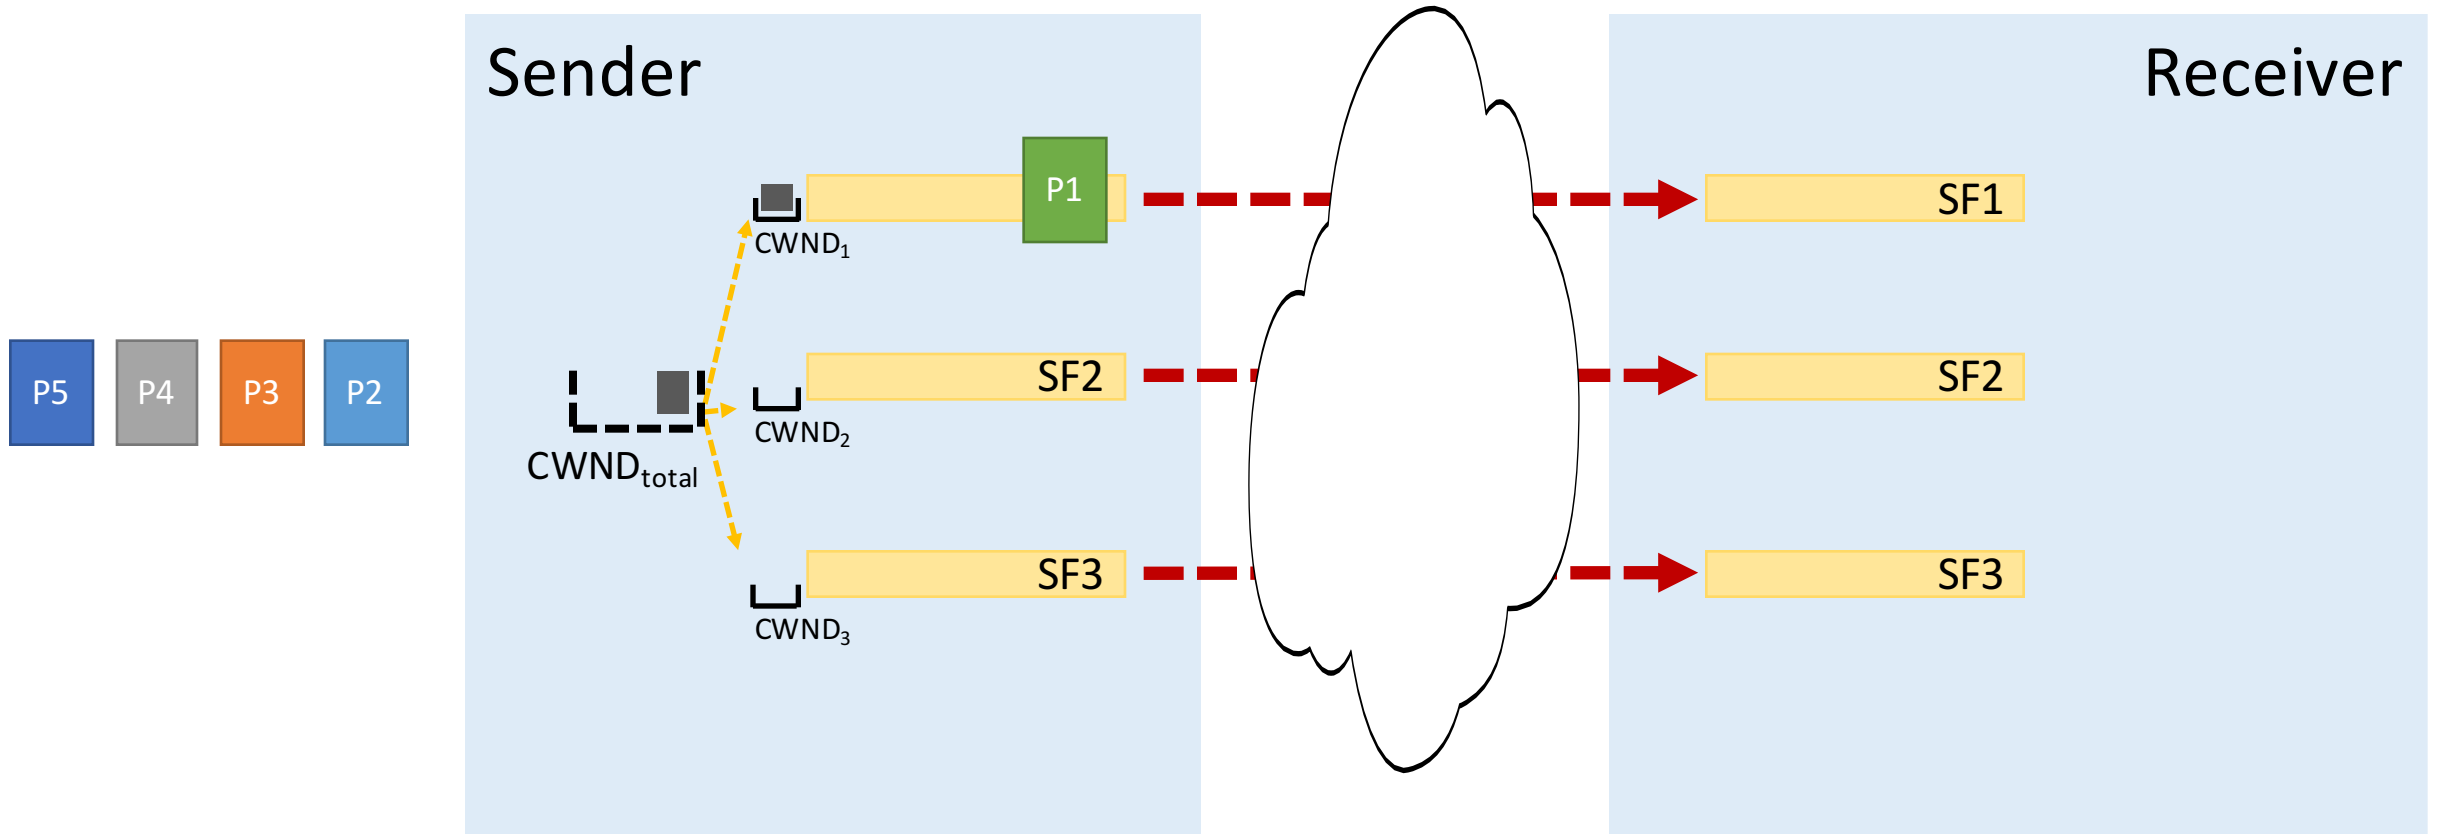
Sub-flows: Implicitly/Explicitly mapping to physical paths



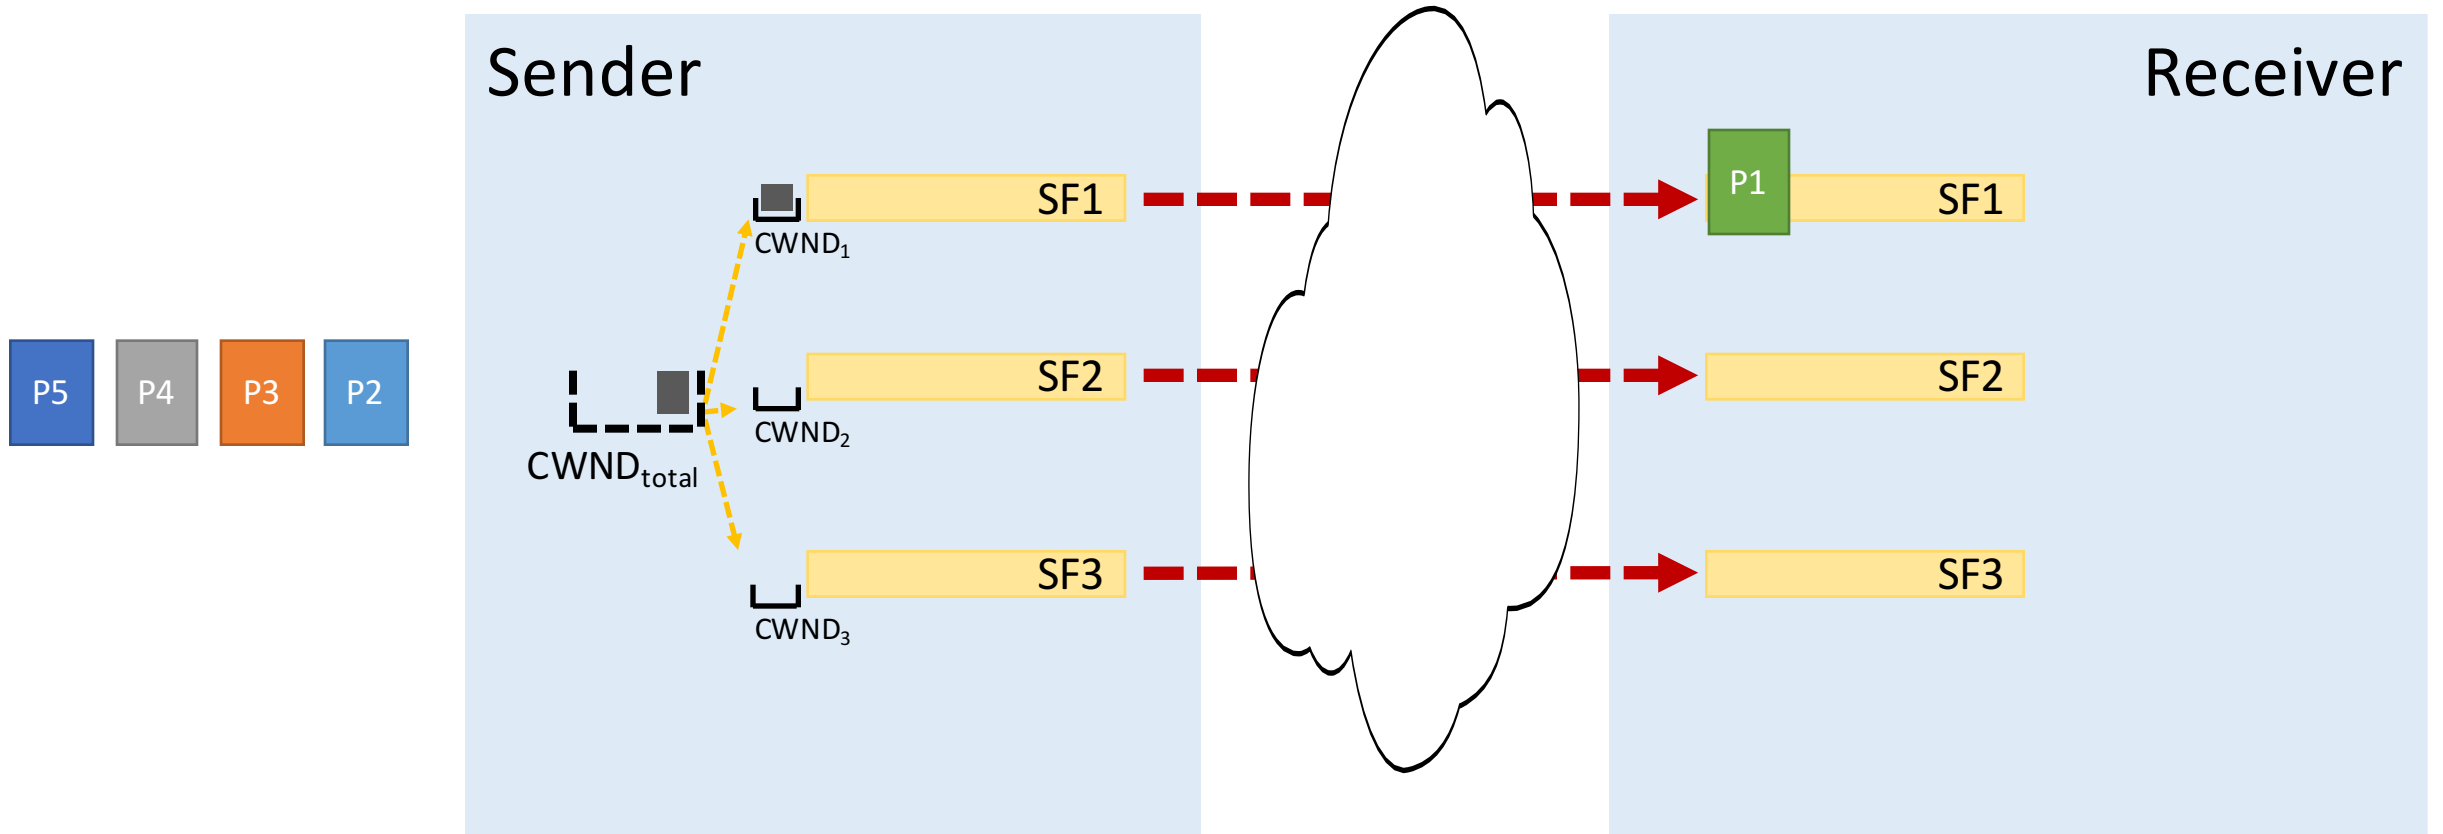
# FUSO



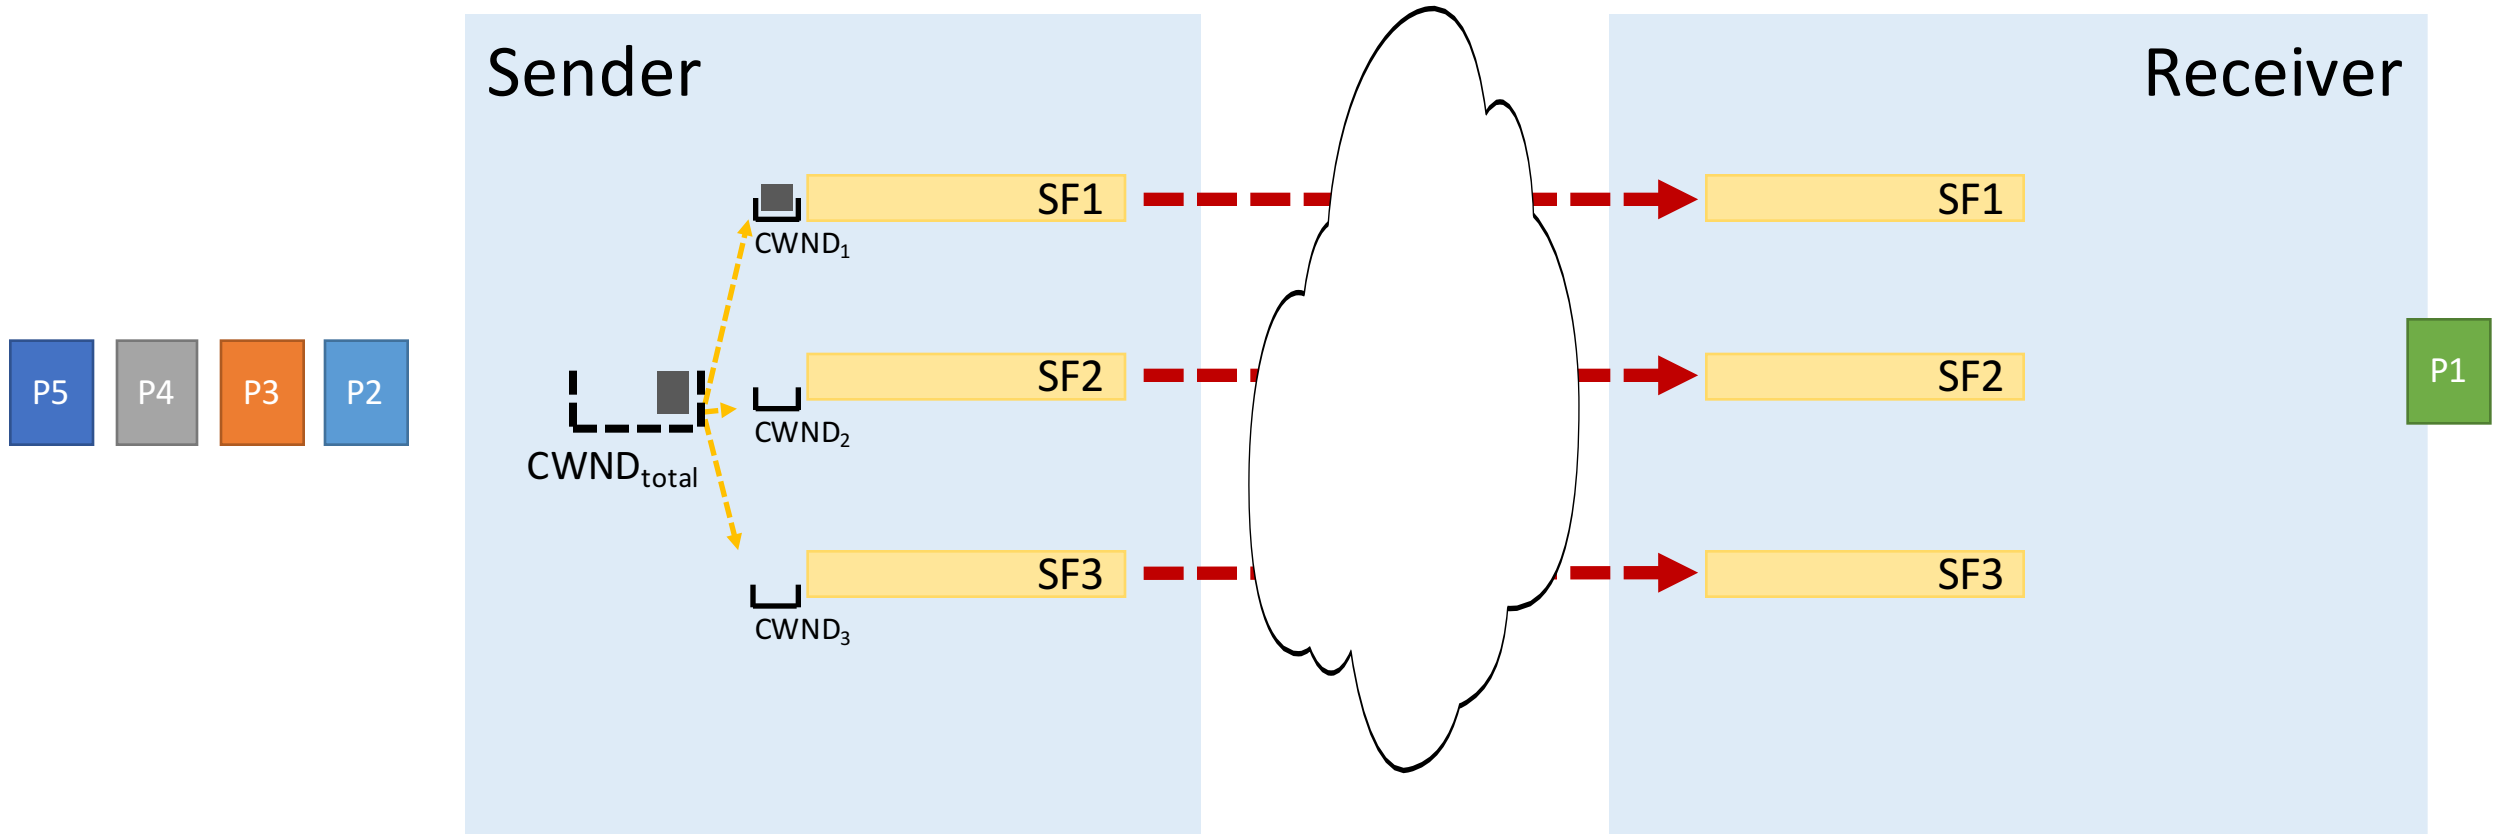
# FUSO



# FUSO

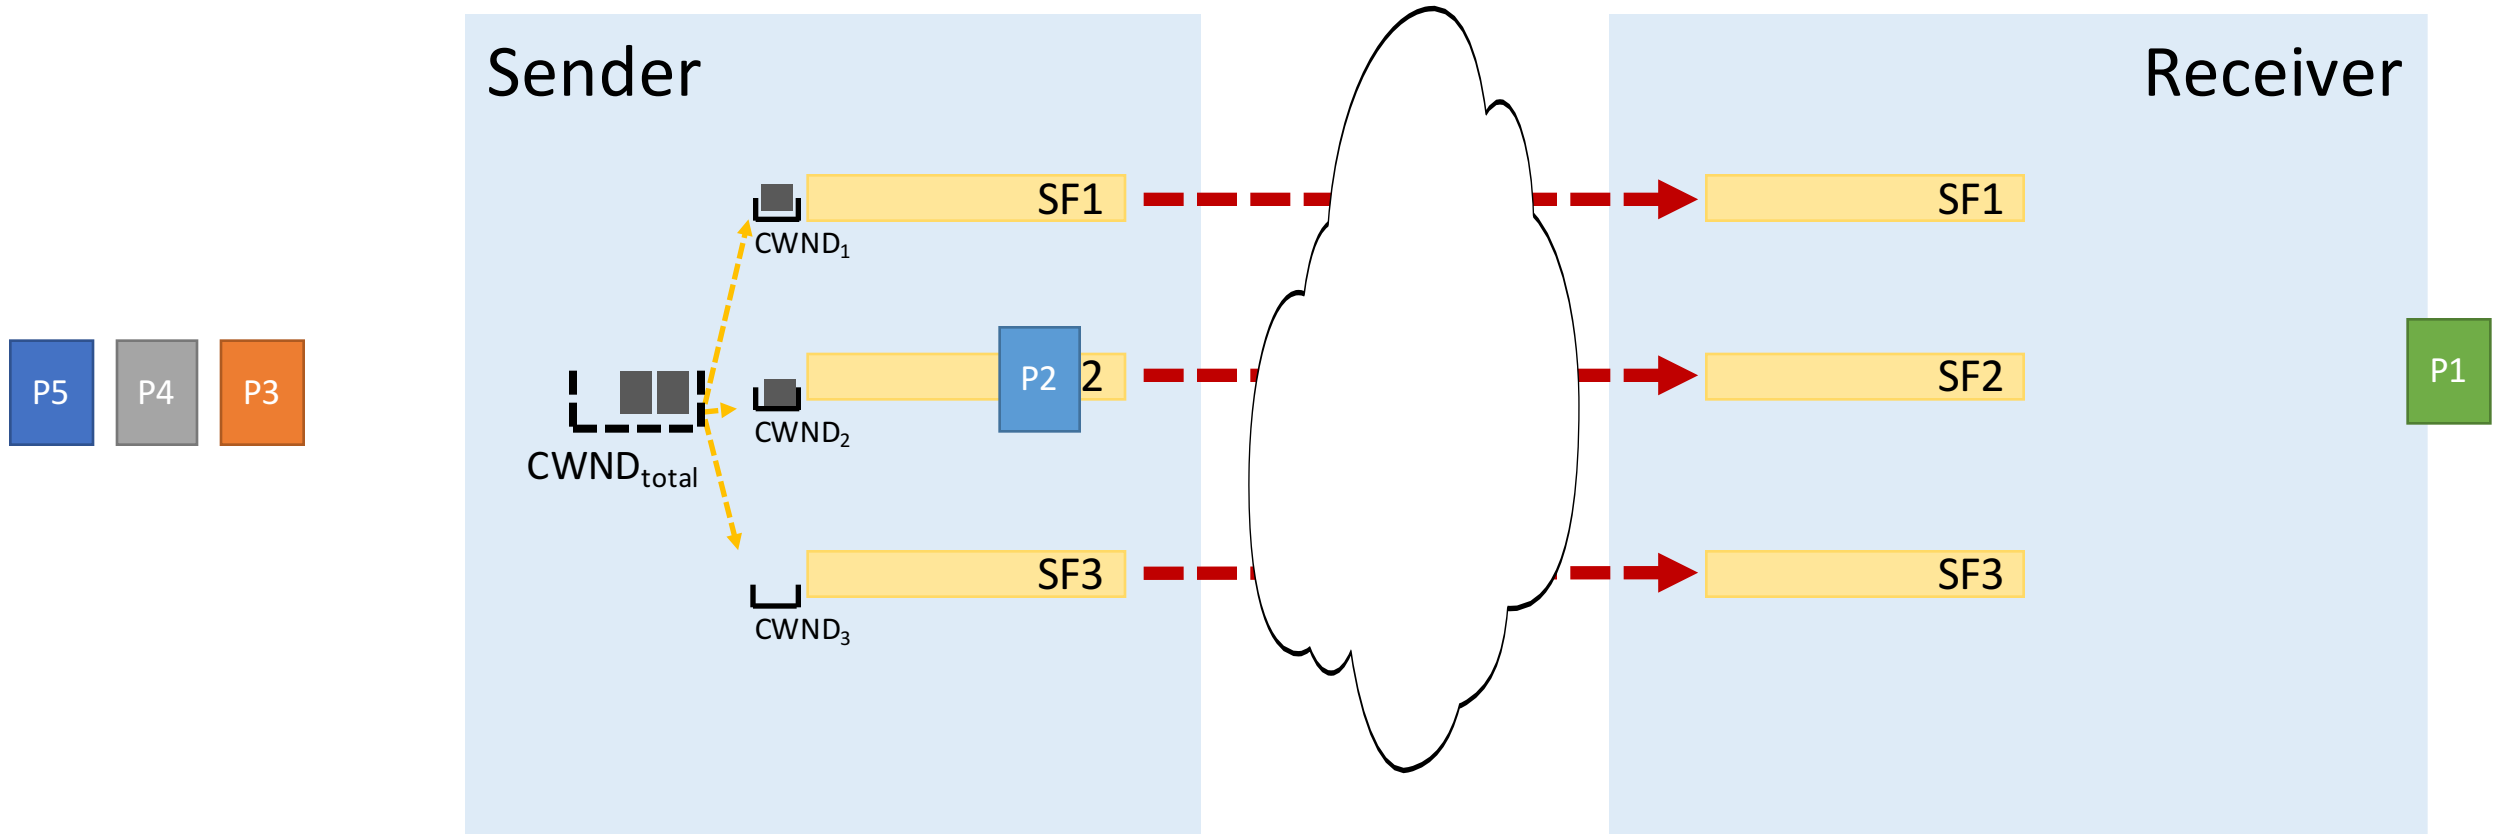


# FUSO






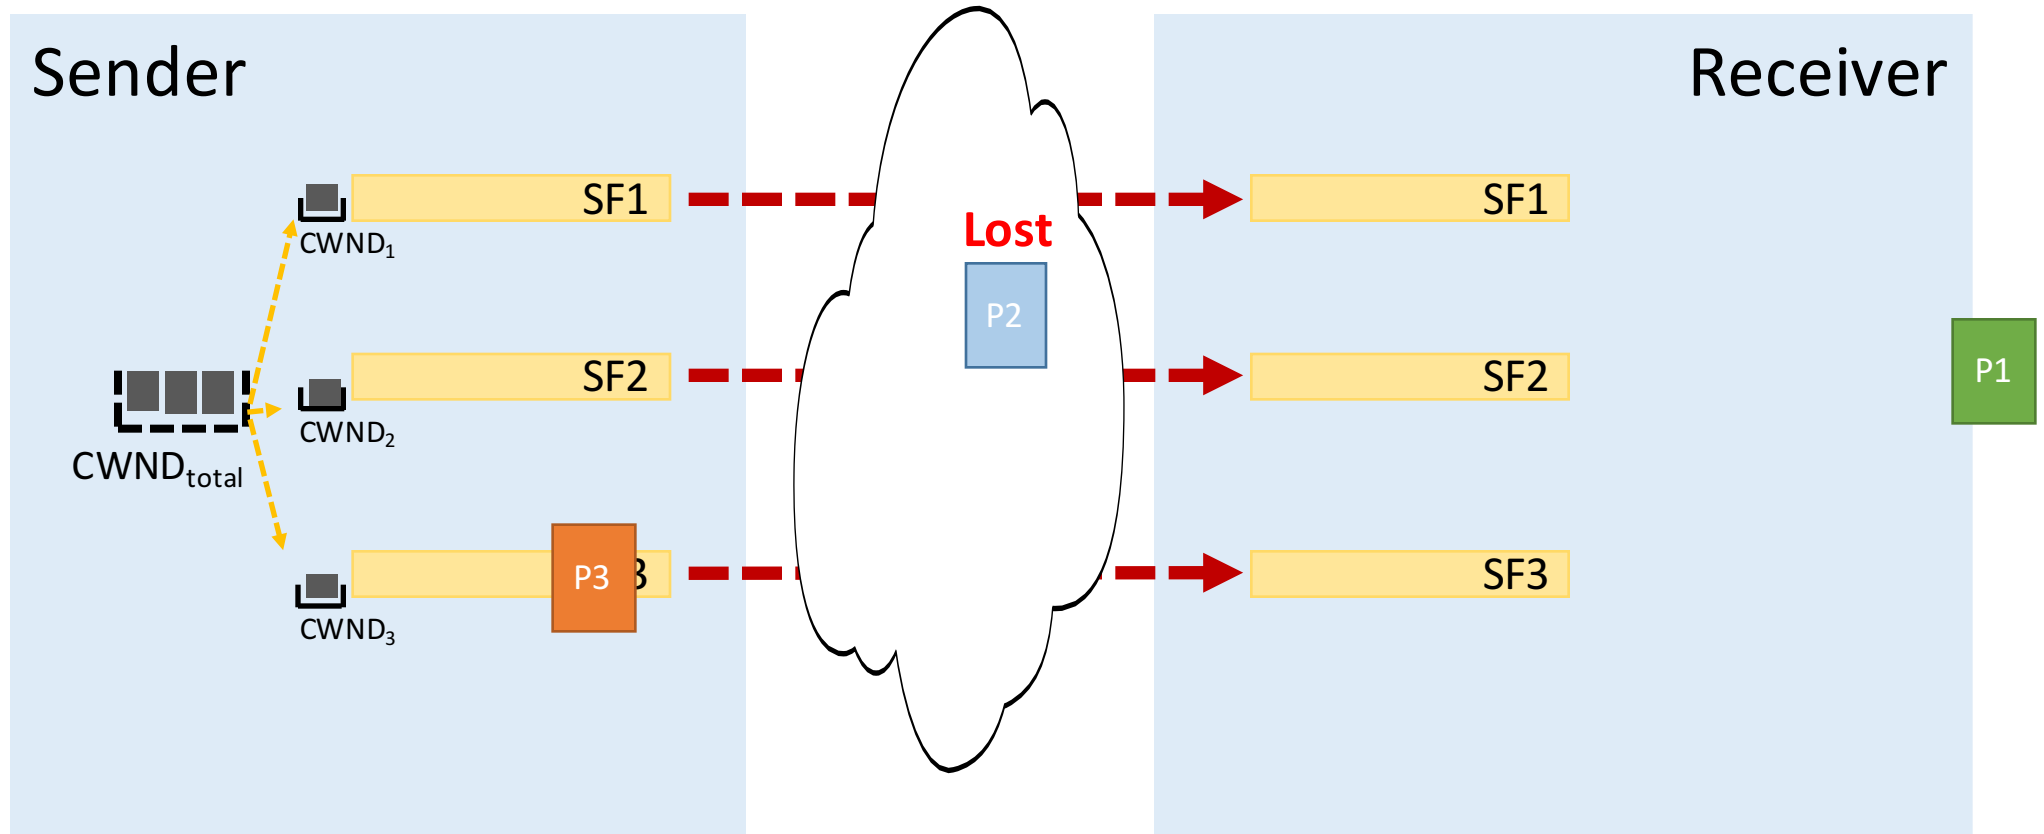
# FUSO



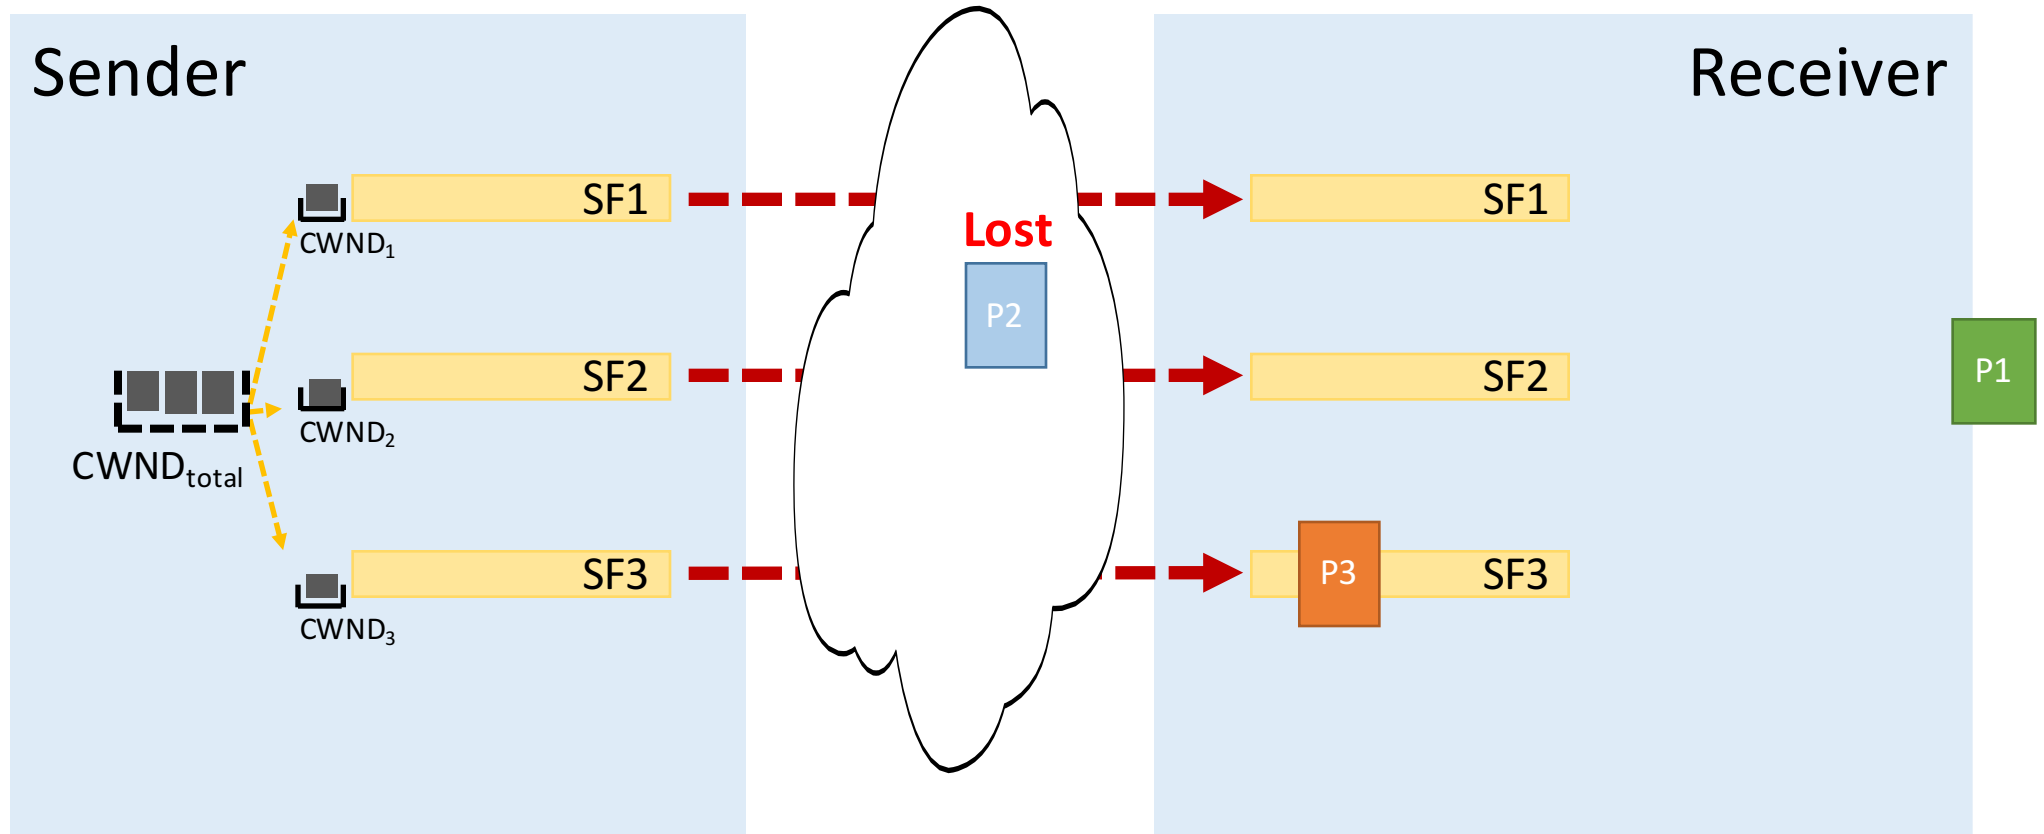
# FUSO




# FUSO



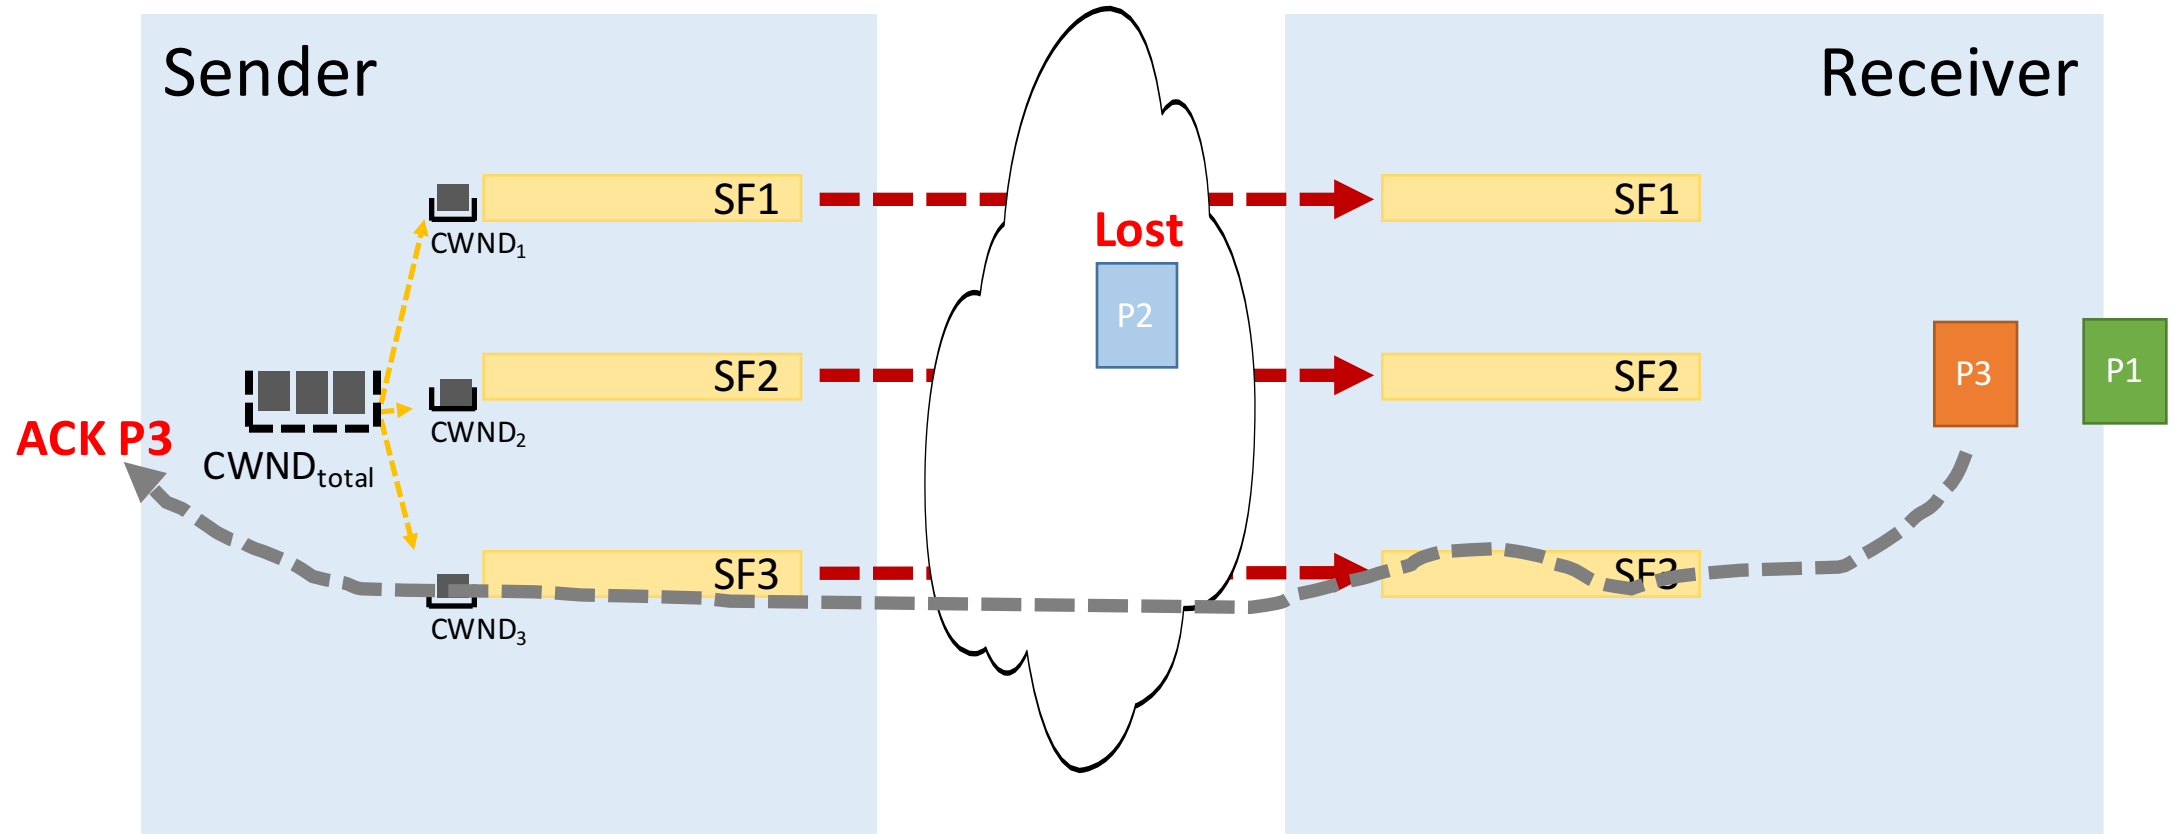
# FUSO




# FUSO



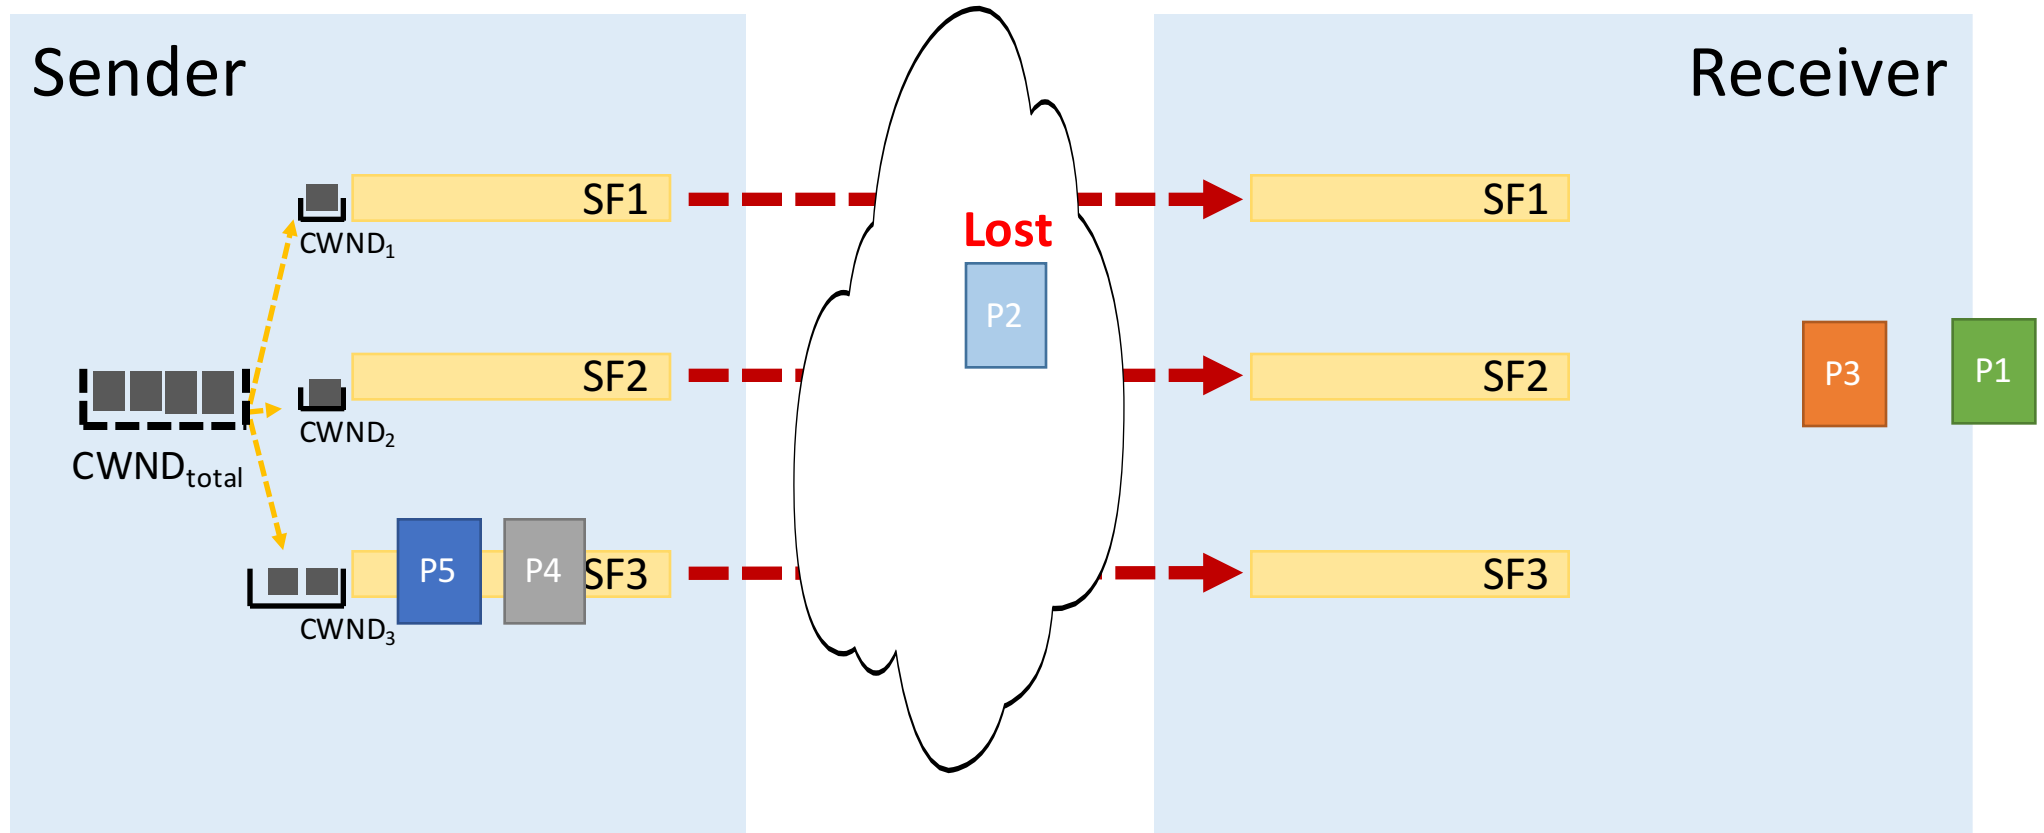
# FUSO



# FUSO

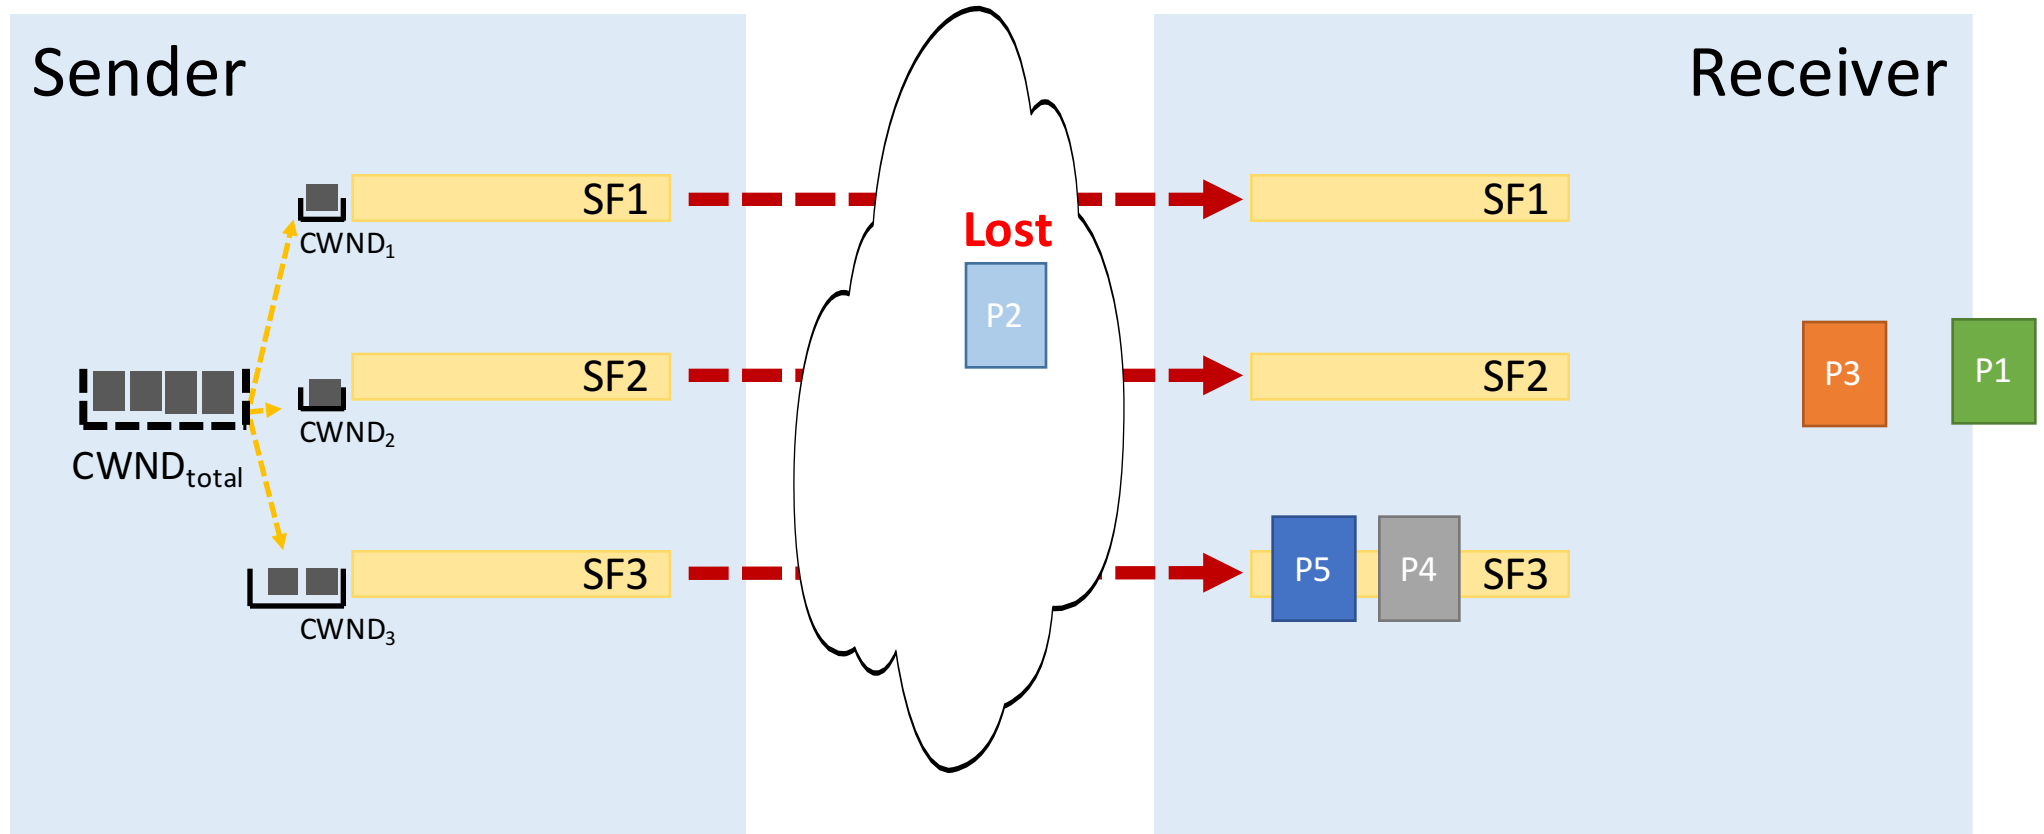


# FUSO






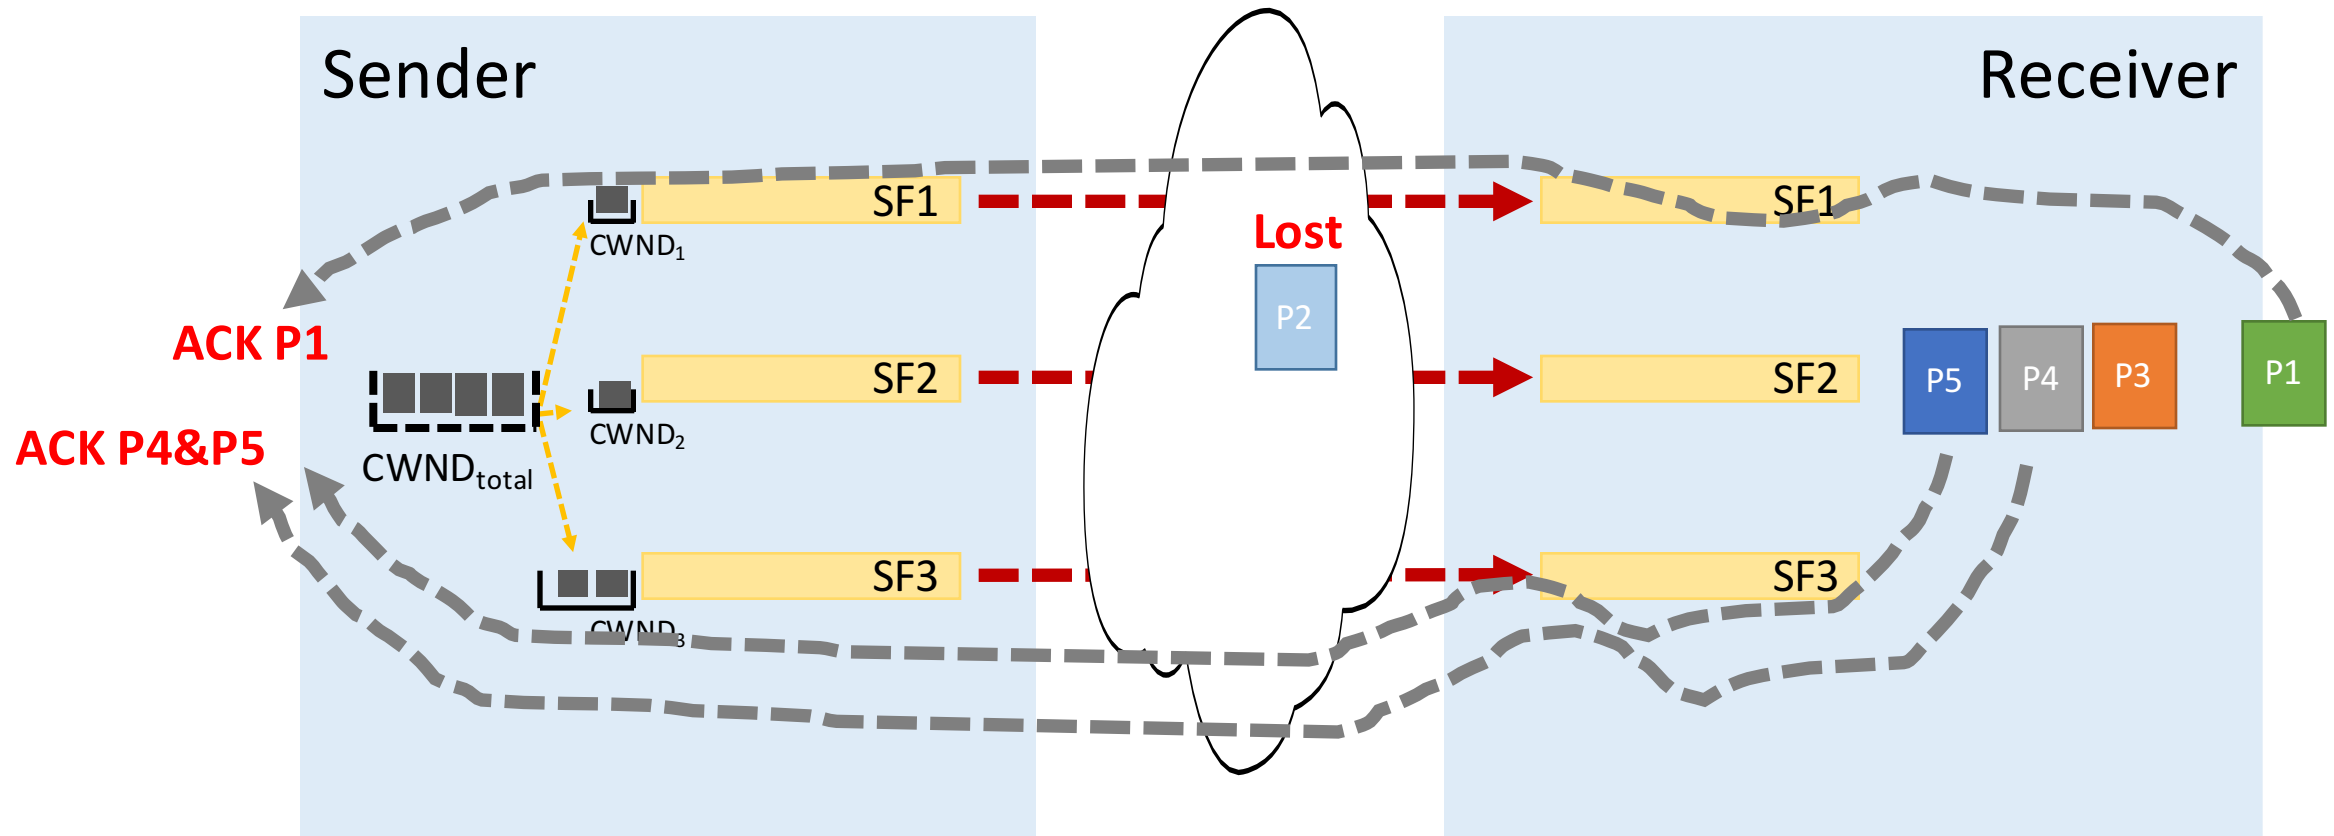
# FUSO



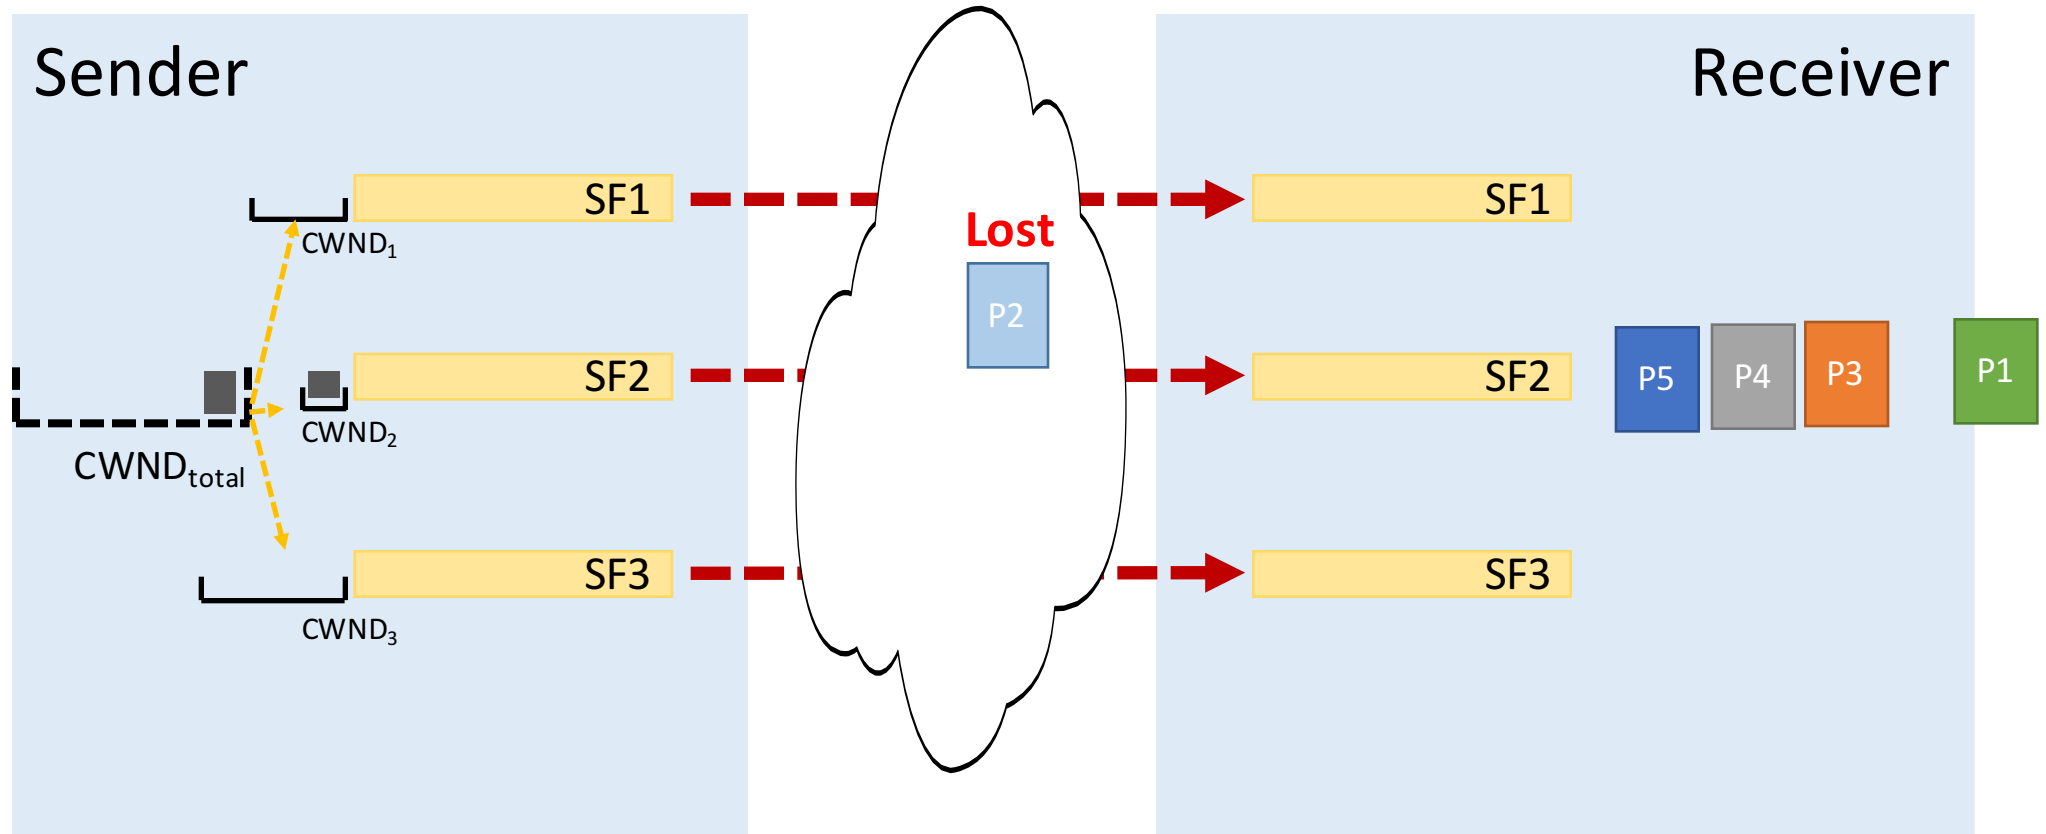
# FUSO



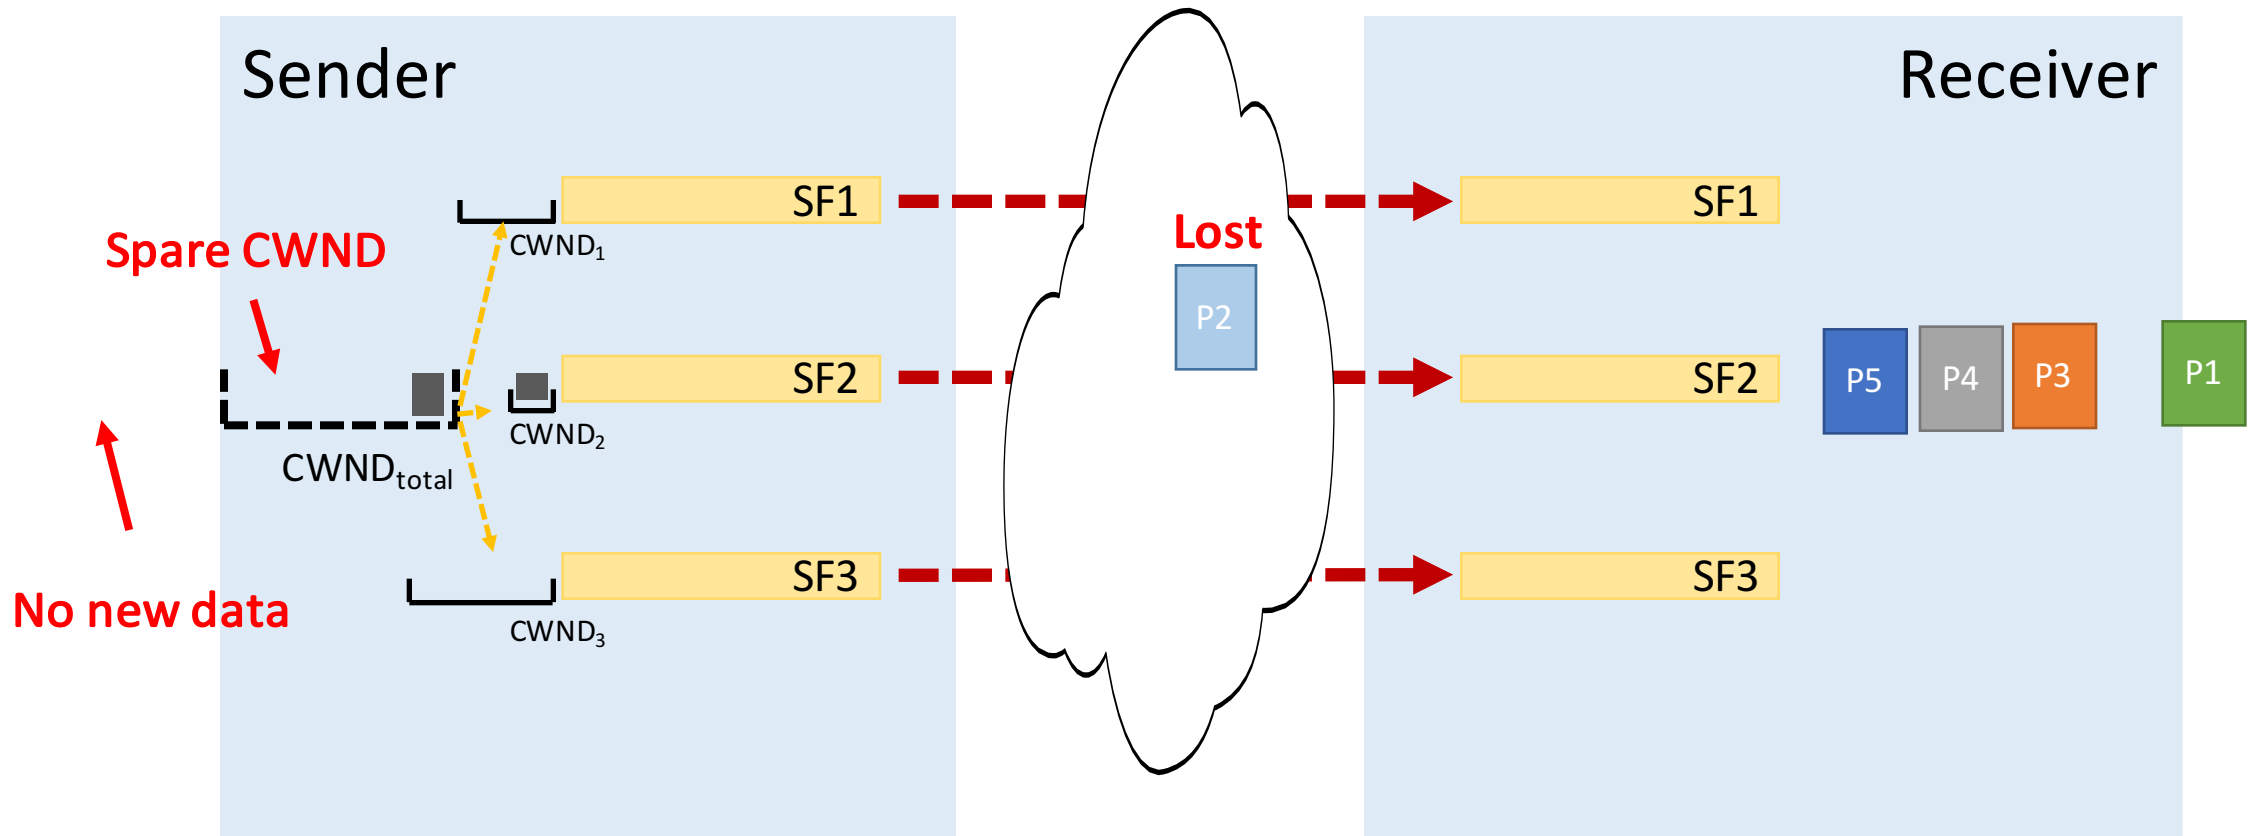
# FUSO



# FUSO

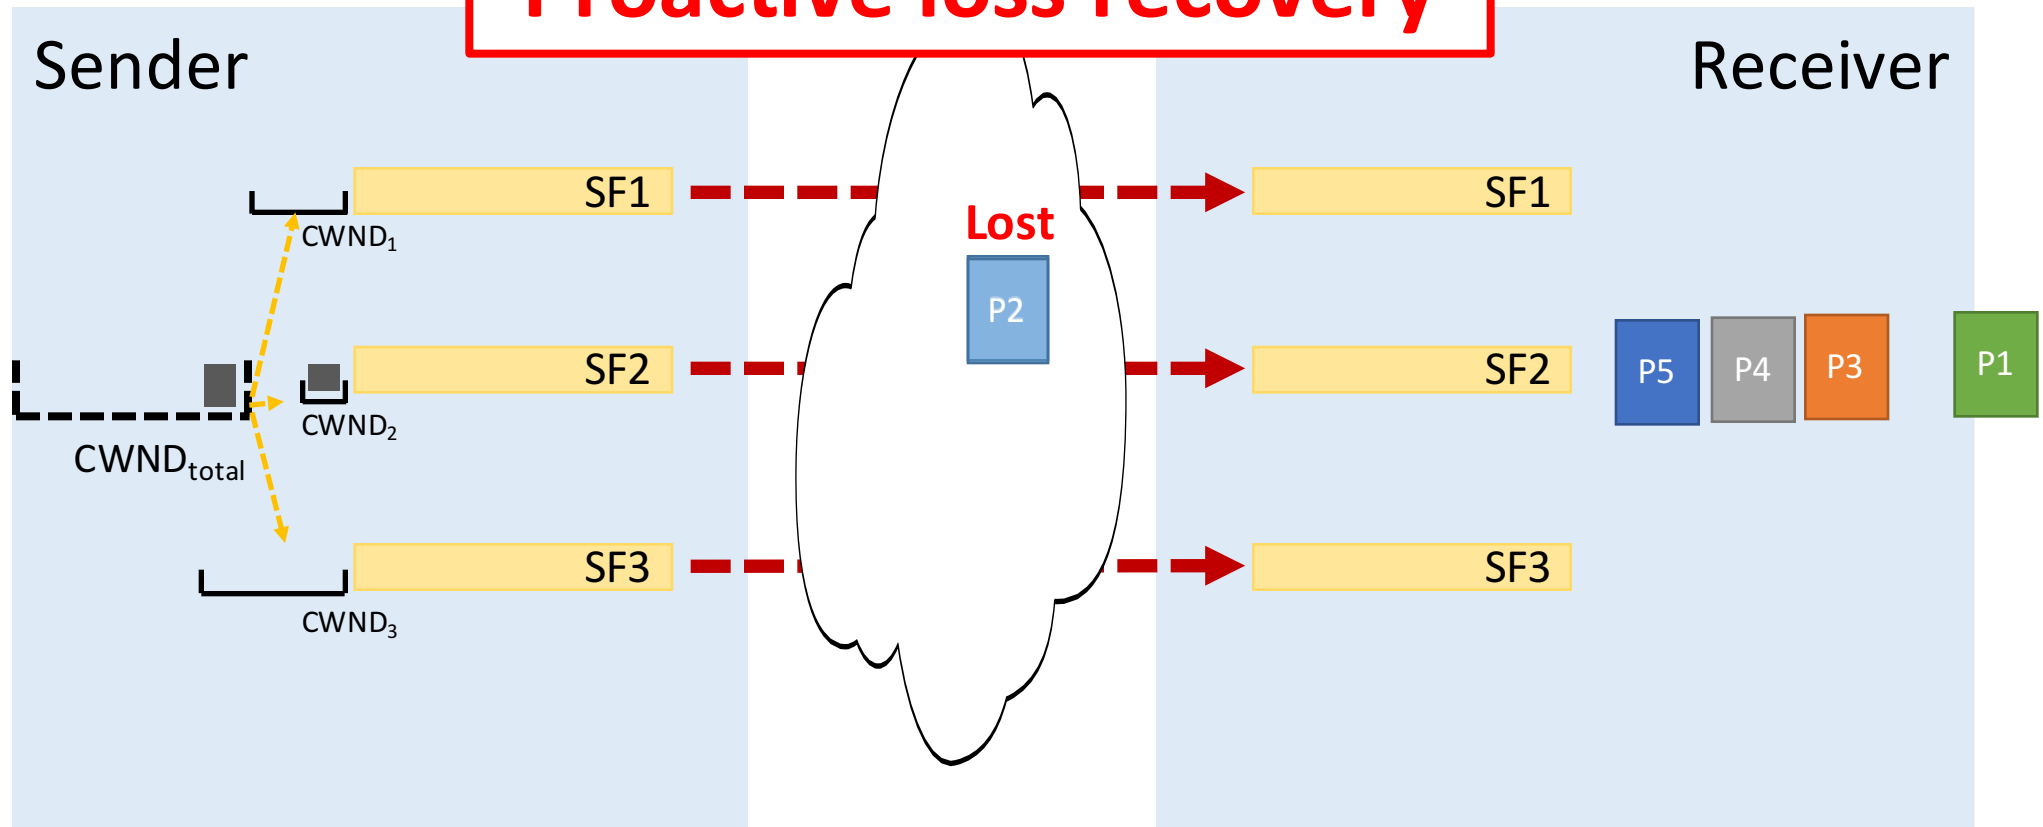


# FUSO



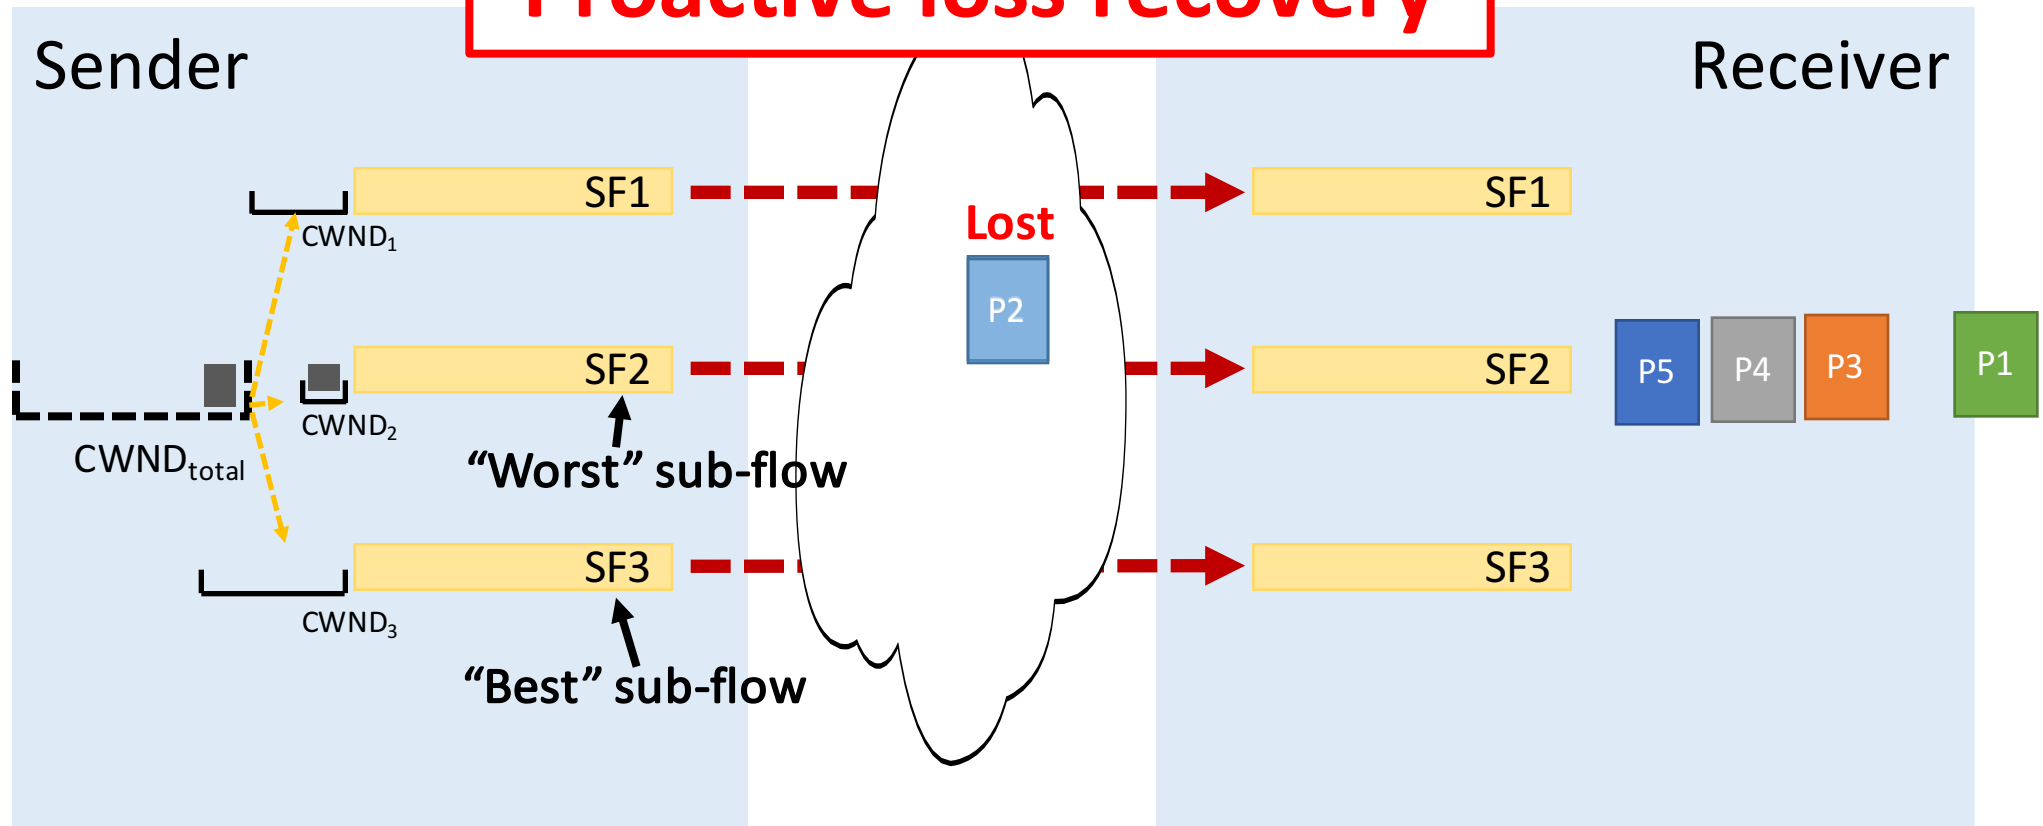
# FUSO

## Proactive loss recovery



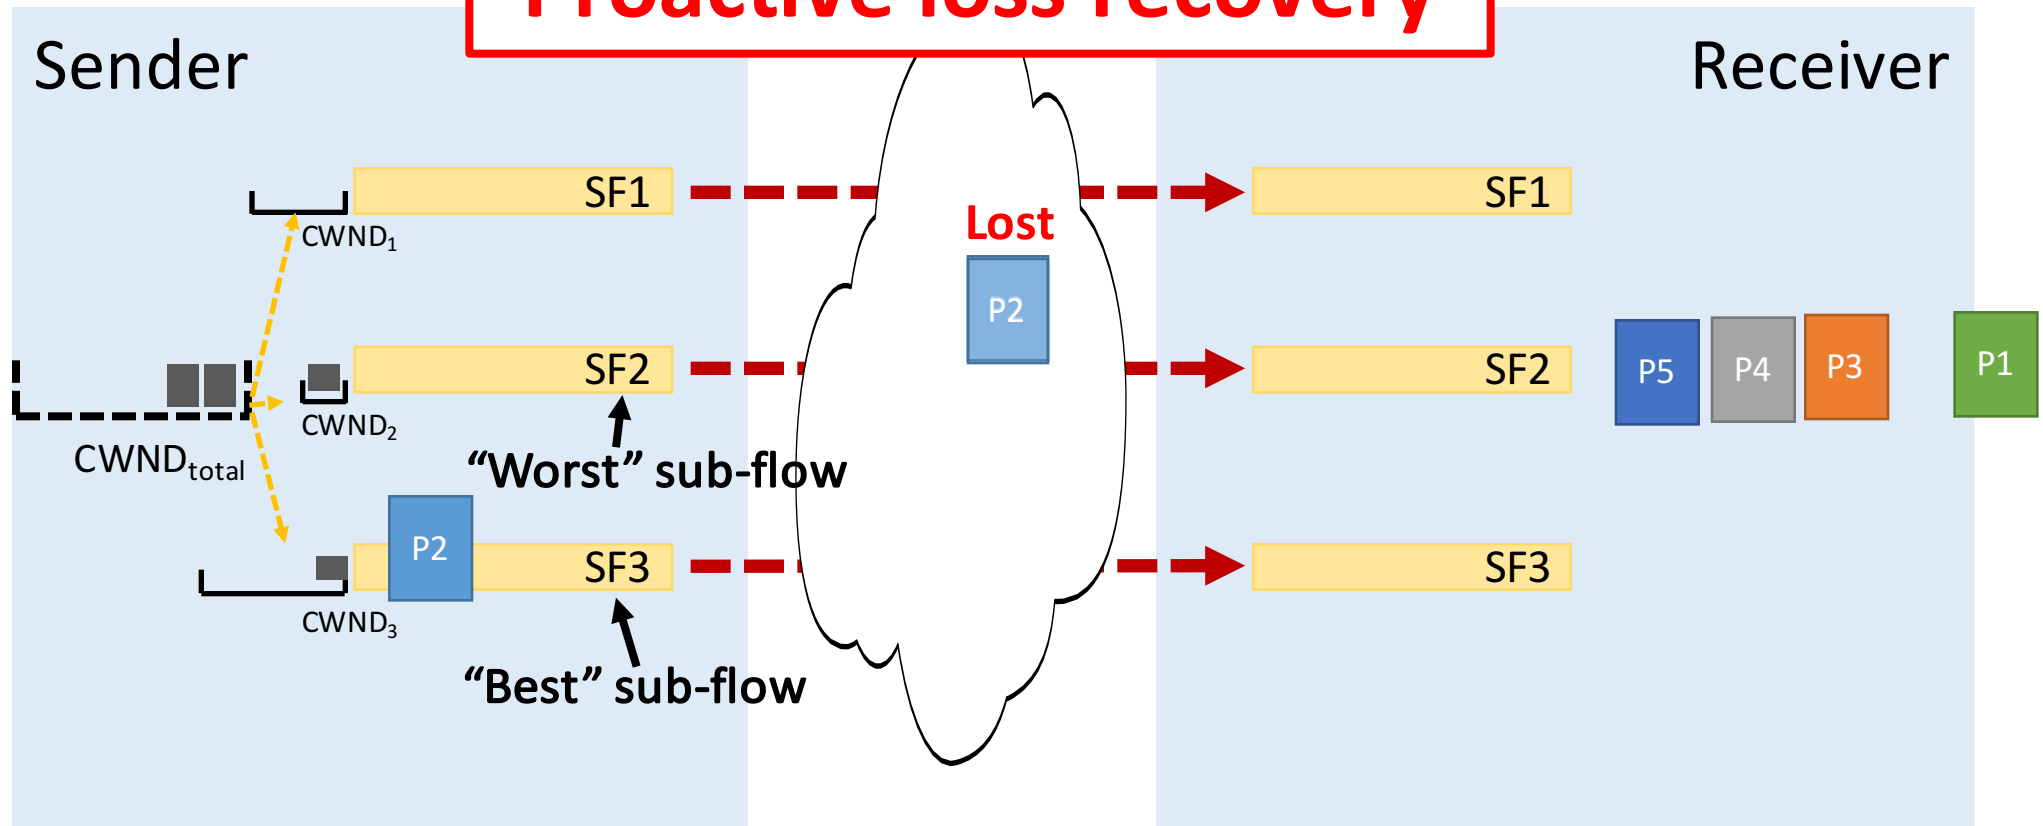
# FUSO

## Proactive loss recovery



# FUSO

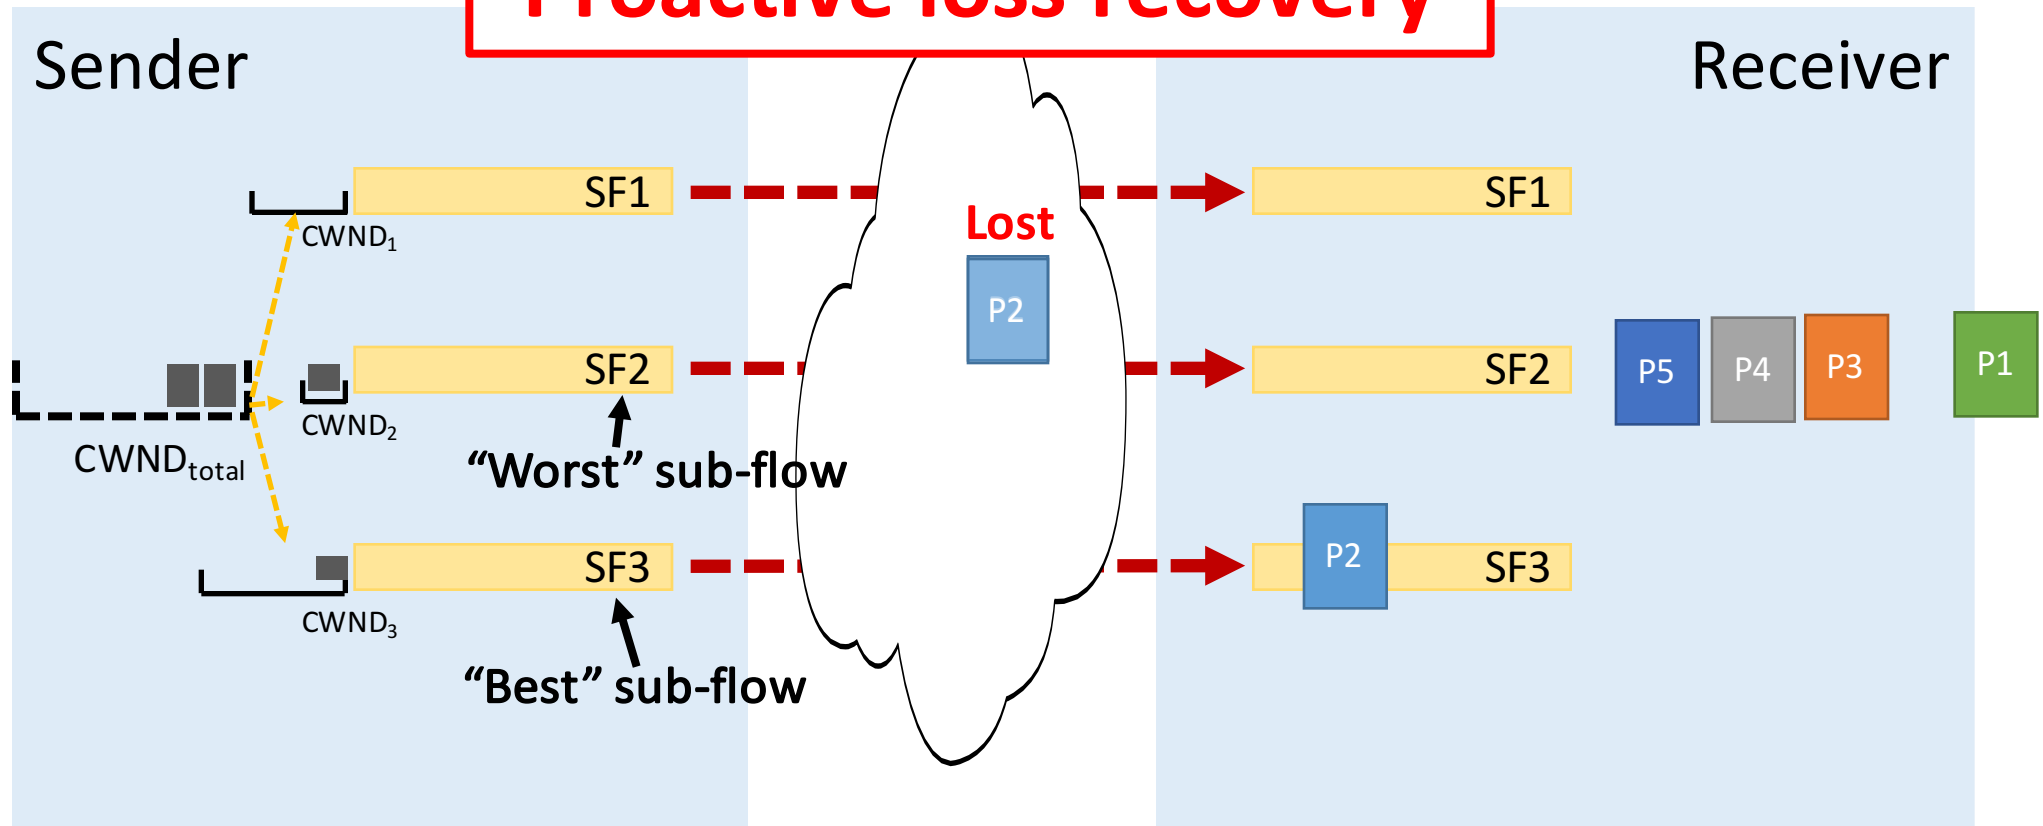
## Proactive loss recovery





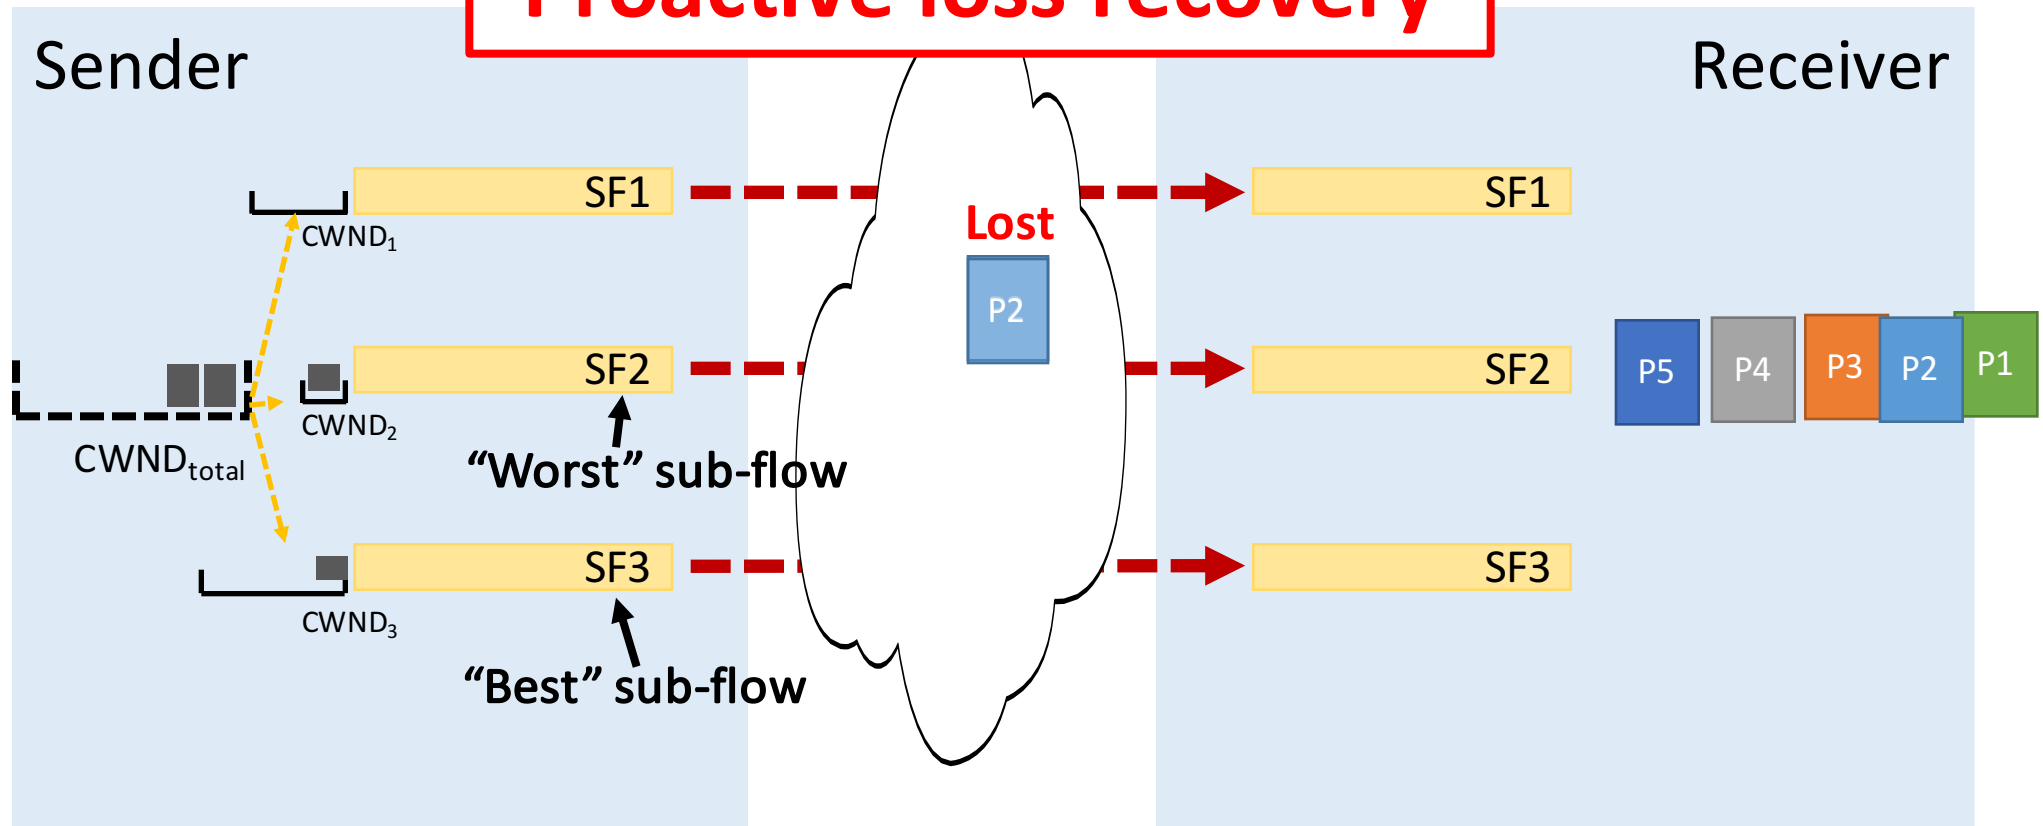
# FUSO

## Proactive loss recovery



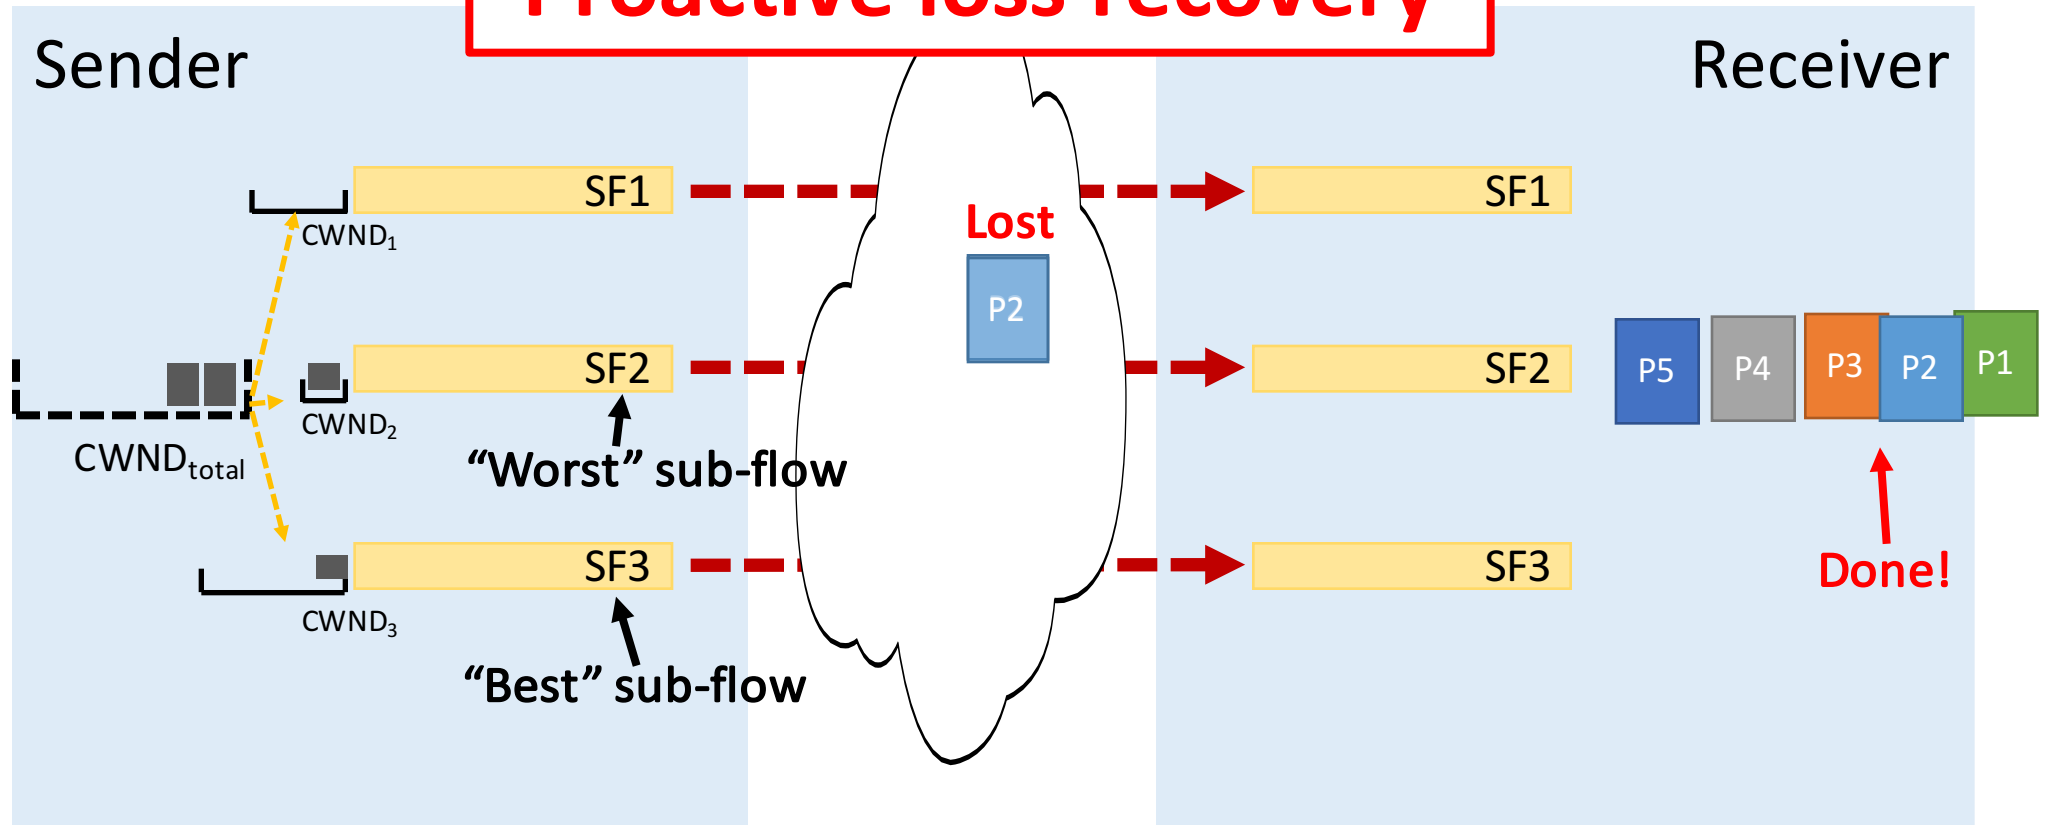
# FUSO

## Proactive loss recovery

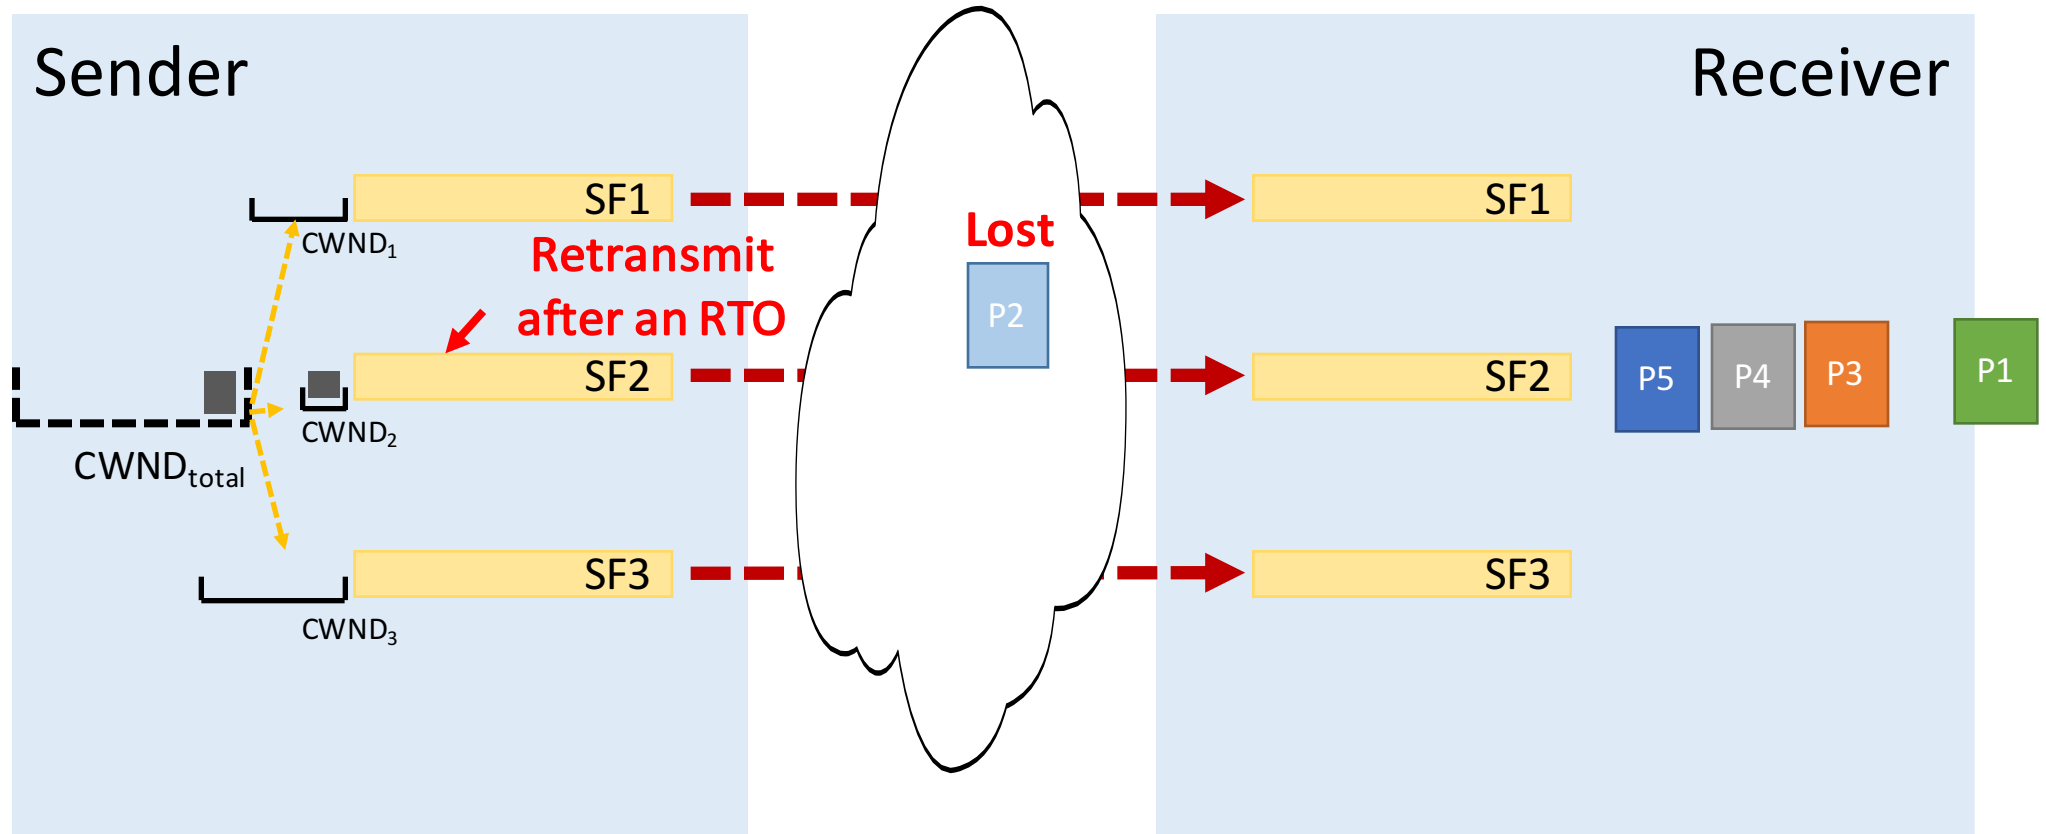


# FUSO

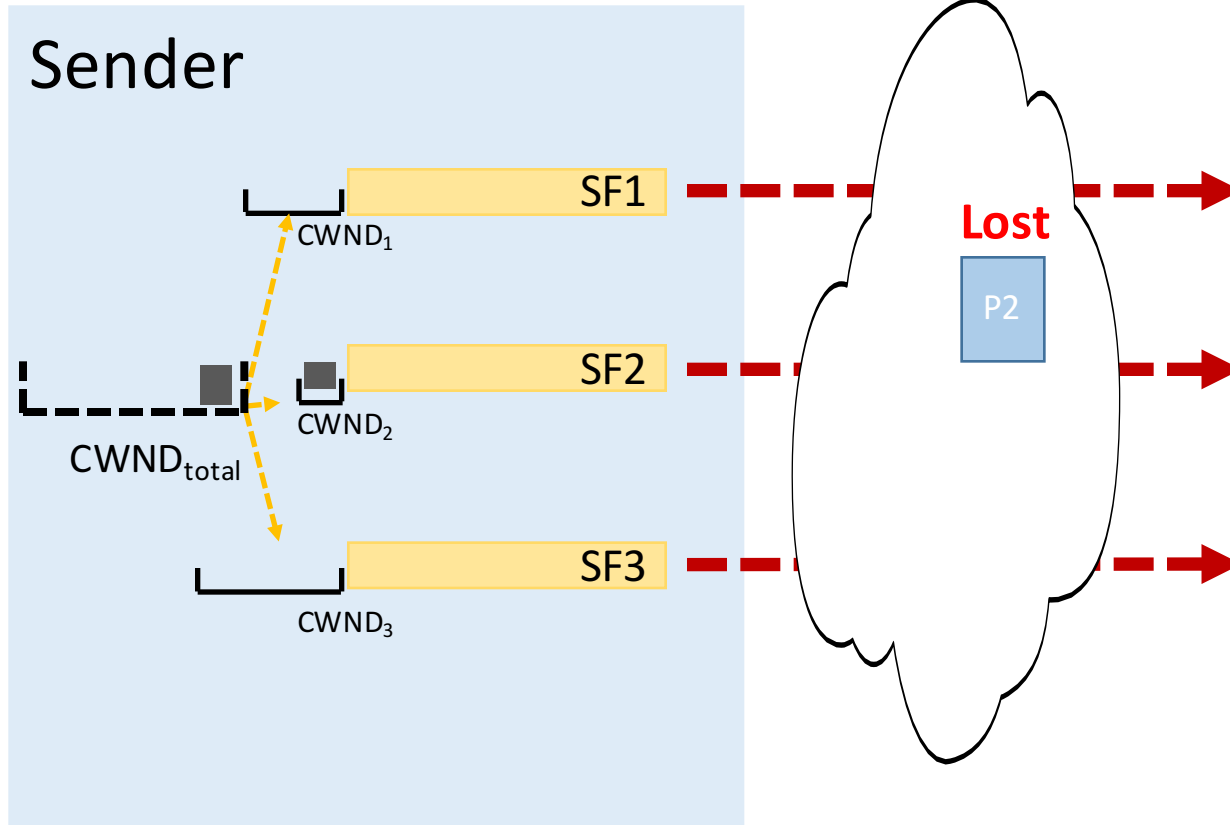
## Proactive loss recovery



# Standard MPTCP



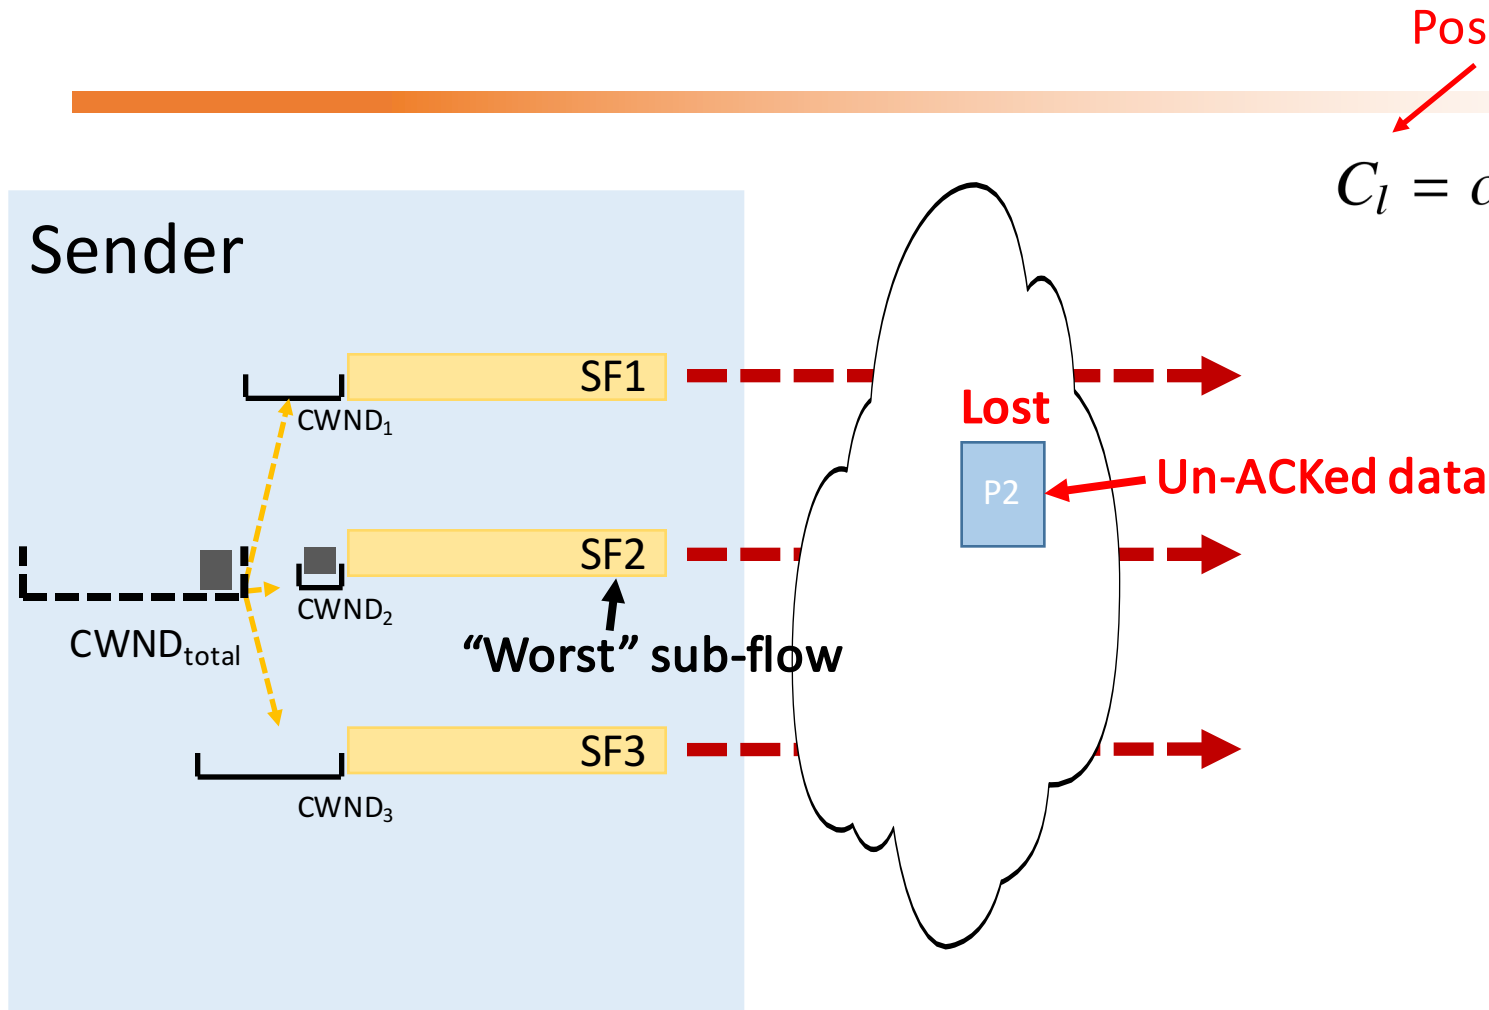
# FUSO: Path Selection



Possibility of encountering loss

$$C_l = \alpha \cdot \overline{\text{lossrate}} + \beta \cdot \text{lossrate}_{last}$$

# FUSO: Path Selection

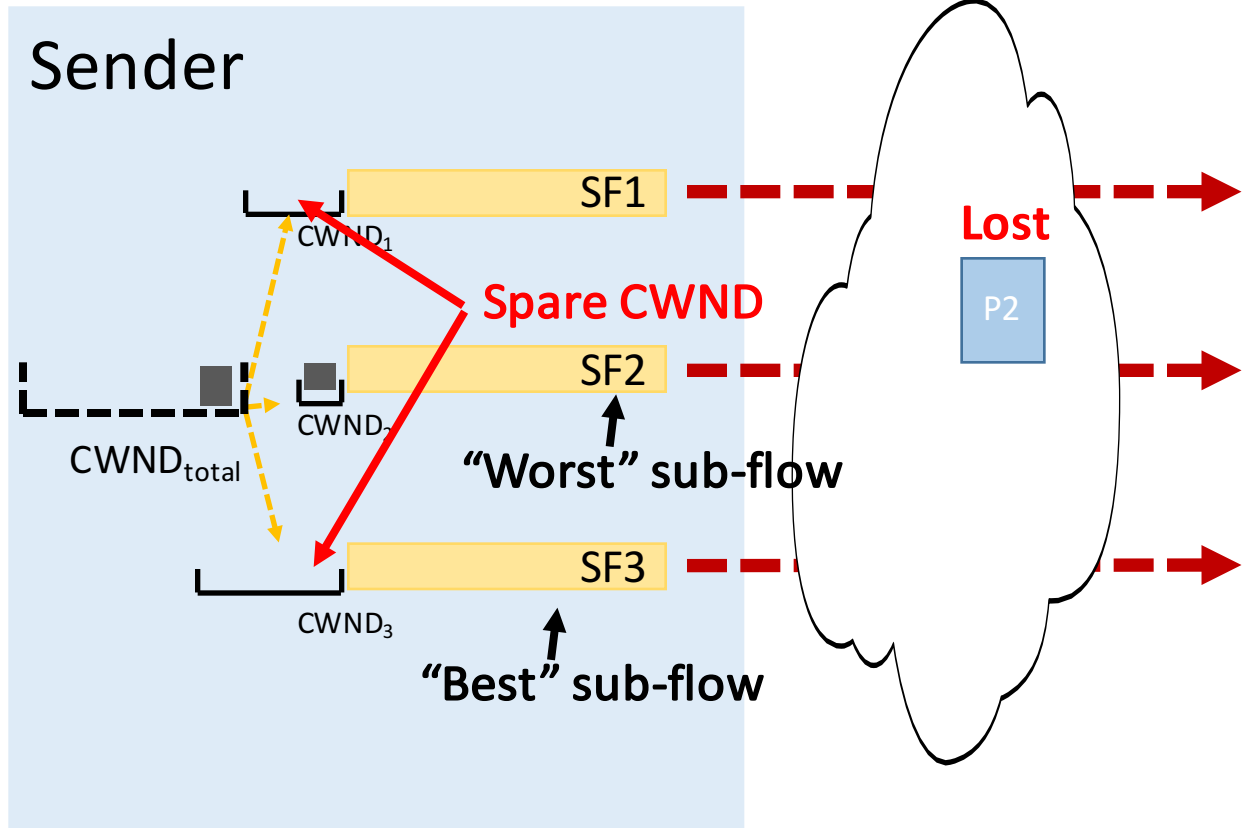


Possibility of encountering loss

$$C_l = \alpha \cdot \overline{\text{lossrate}} + \beta \cdot \text{lossrate}_{last}$$

- **“Worst” Sub-flow**
  - With un-ACKed data
  - Most likely having loss

# FUSO: Path Selection

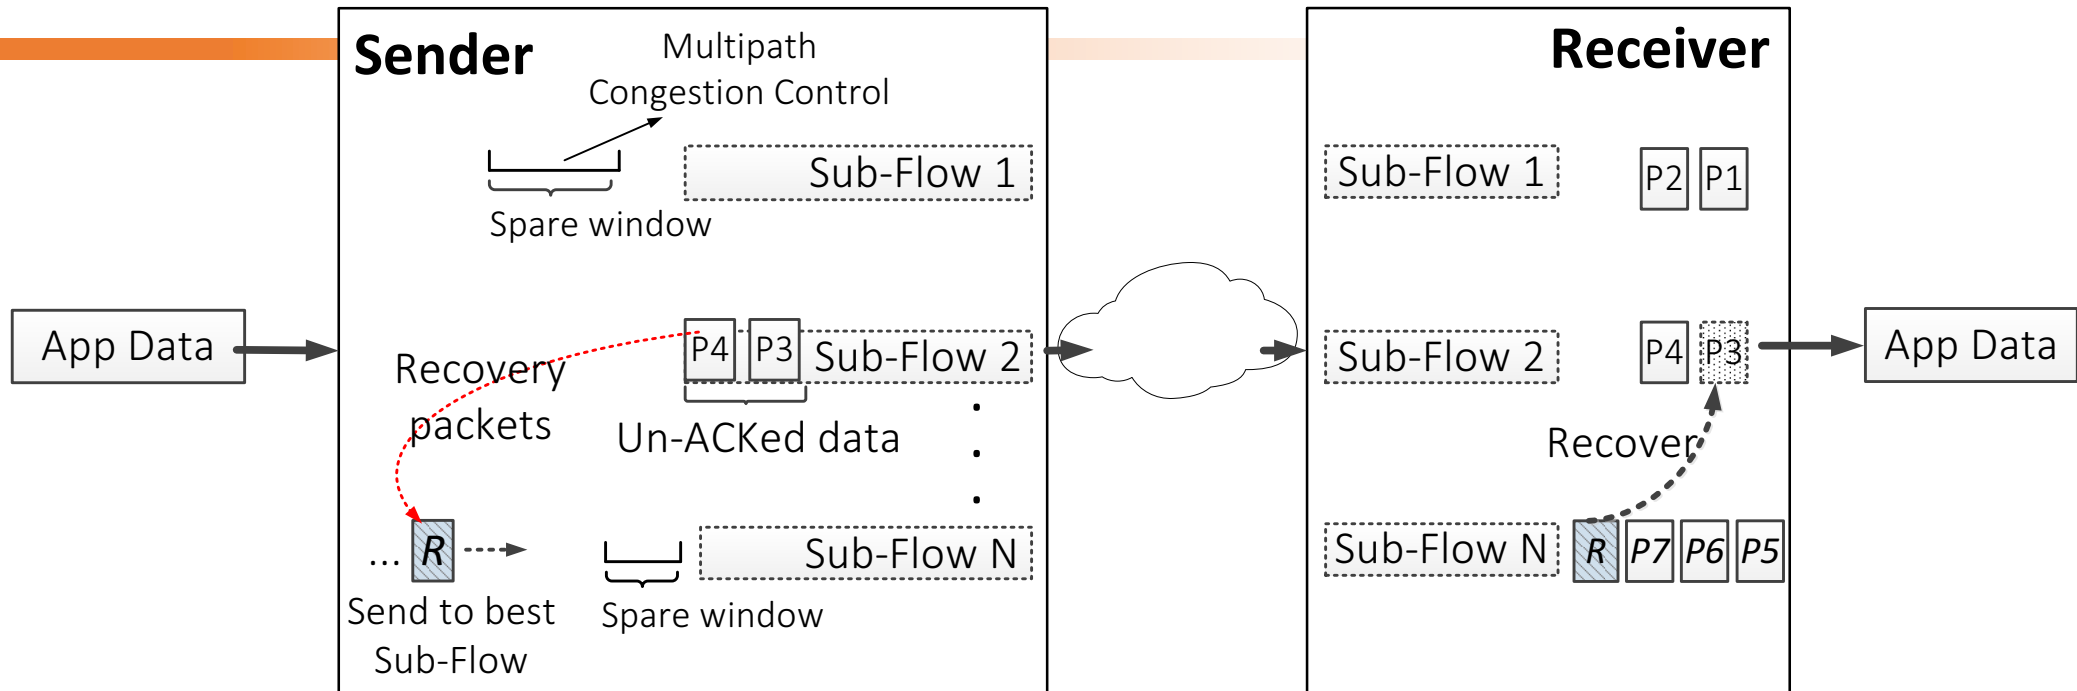


Possibility of encountering loss

$$C_l = \alpha \cdot \overline{lossrate} + \beta \cdot lossrate_{last}$$

- **“Worst” Sub-flow**
  - With un-ACKed data
  - Most likely having loss
- **“Best” Sub-flow**
  - With spare CWND
  - Least likely having loss

# FUSO in 1 Slide



- If (spare CWND) && (no new data)
  - Utilize the transmission opportunity to proactively recover
  - Use “good” paths to help “bad” paths
- Multi-path diversity offers many transmission opportunities
  - “Good” paths have spare window



# FUSO Implementation

<https://github.com/1989chenguo/FUSO>

- Implemented in Linux kernel; ~900 lines of code

```
1: function TRY_SEND_RECOVERIES()  
2:   while BytesInFlightTotal < CWNDTotal and no new data do  
3:     return ← SEND_A_RECOVERY()  
4:     if return == NOT_SEND then  
5:       break  
  
1: function SEND_A_RECOVERY()  
2:   FIND_WORST_SUB-FLOW()  
3:   FIND_BEST_SUB-FLOW()  
4:   if no worst found or no best sub-flow found then  
5:     return NOT_SEND  
6:   recovery_packet ← one un-ACKed packet of the worst sub-flow  
7:   Send the recovery_packet through the best sub-flow  
8:   BytesInFlightTotal += Sizerecovery_packet
```

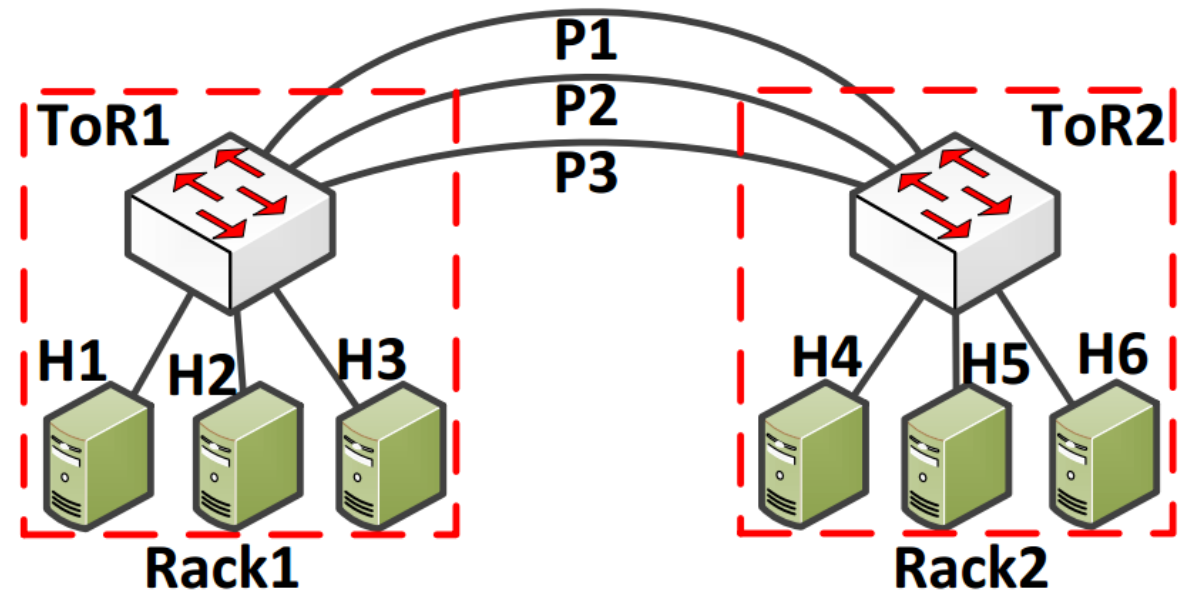
# Outline

---

- Motivation
- Packet Loss in DCN
- Impact of Packet Loss
- Challenge for Loss Recovery
- FUSO Design
- **Evaluation**
- Summary

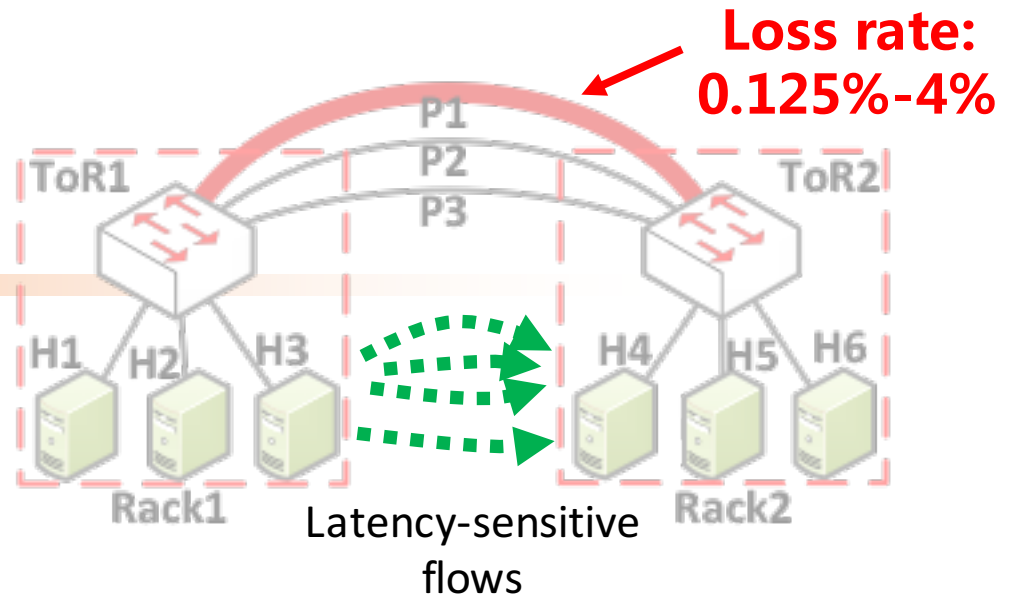
# Testbed Settings

- Network
  - 1Gbps fabric & 1Gbps hosts; ECMP routing; ECN enabled
- TCP
  - Init\_cwnd=16; min\_RTO=5ms

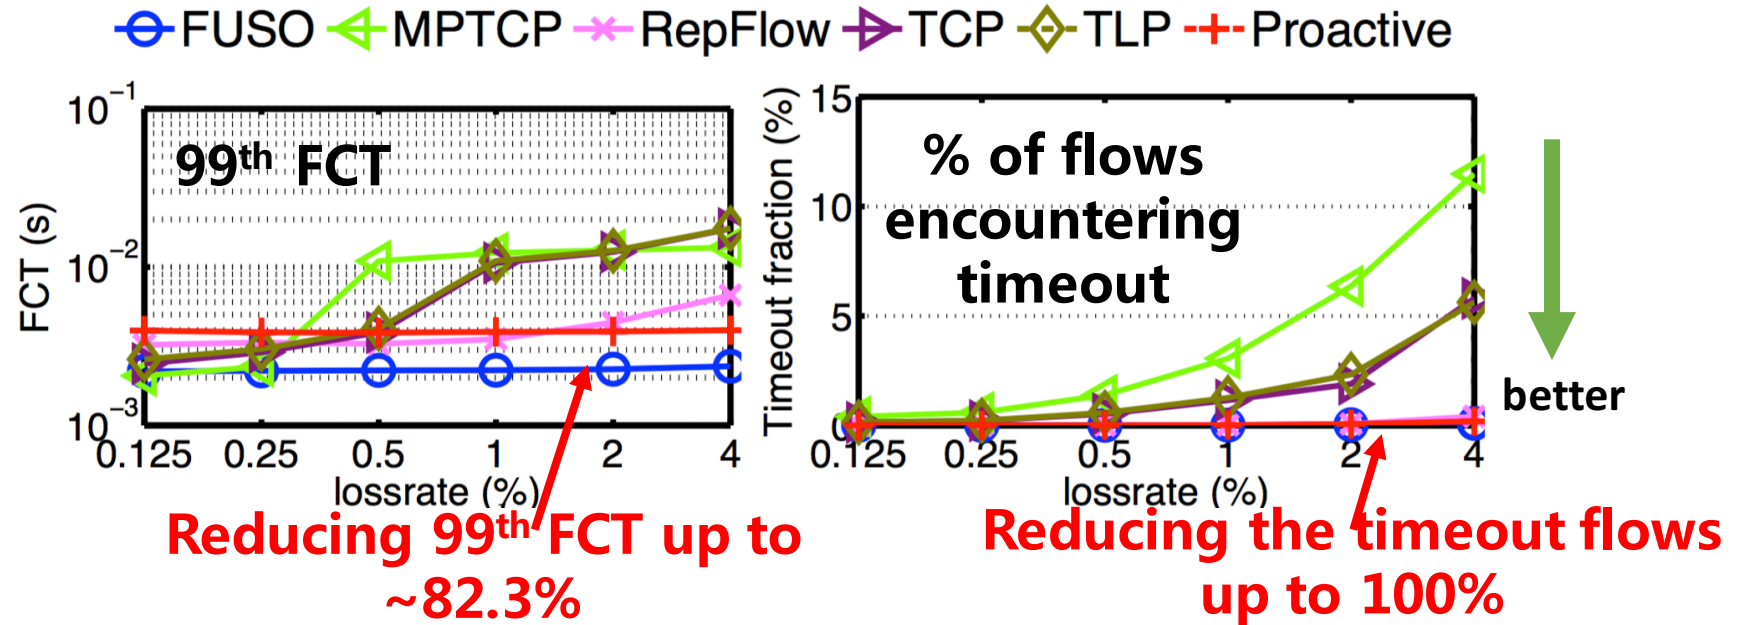


# Testbed Results

- Failure loss
  - Random-drop

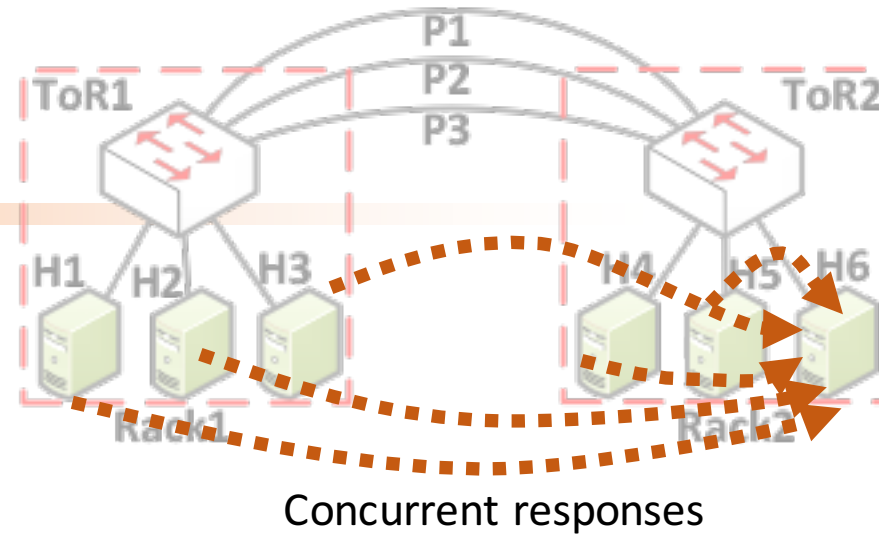


**Fast**




# Testbed Results

- Congestion loss
- Incast

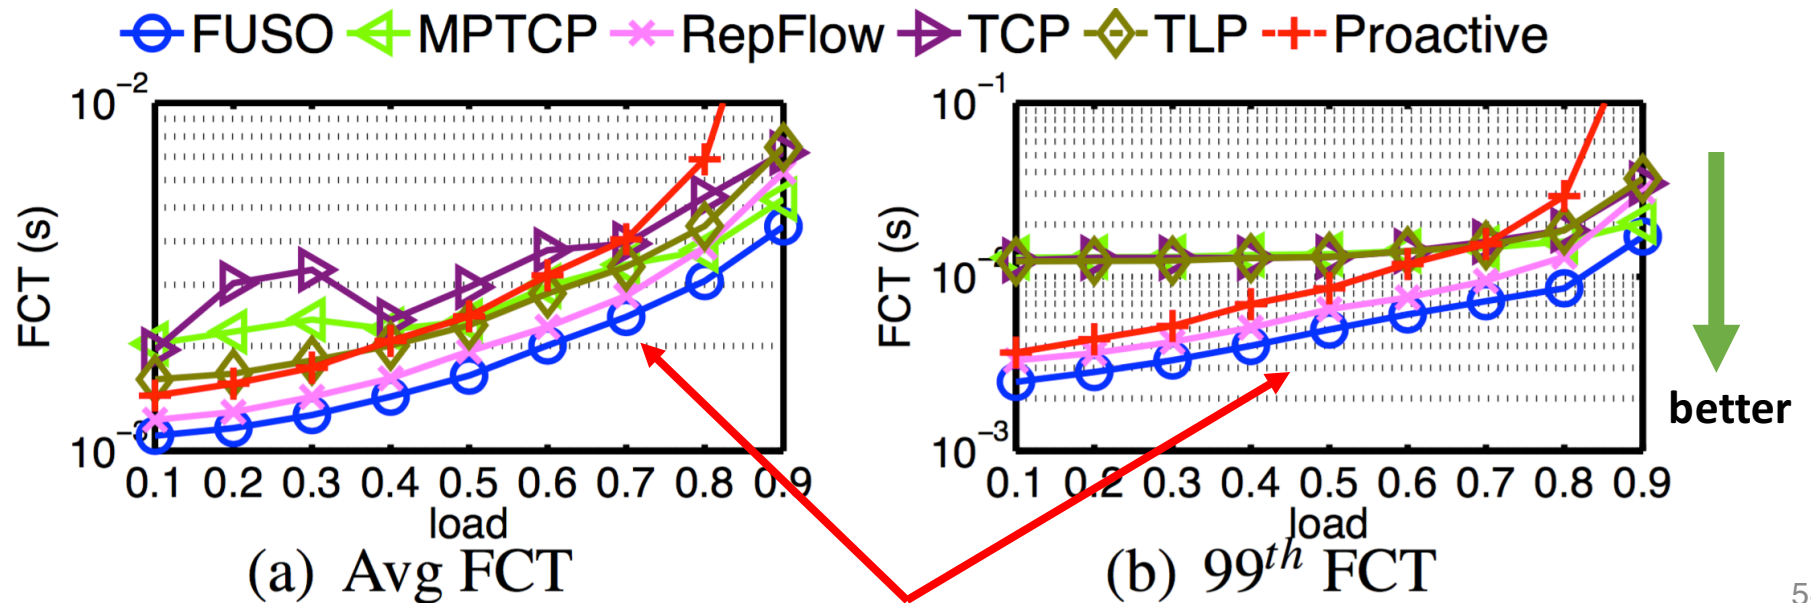



**Cautious**



# Testbed Results

- Failure loss & Congestion loss
  - From failure-loss-dominated to congestion-loss-dominated

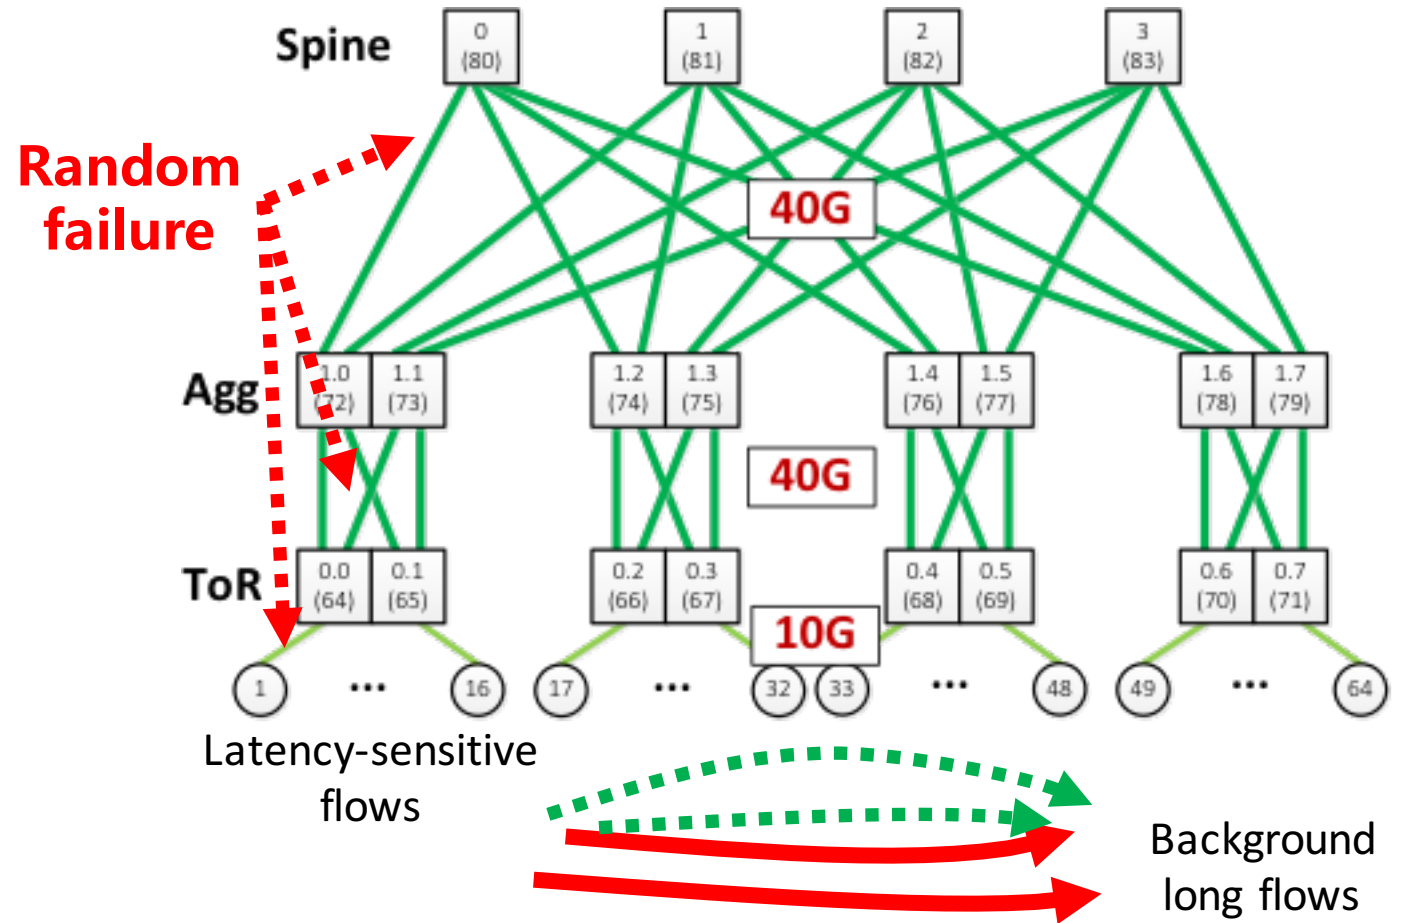


**Adapt to various loss condition**

# Larger-scale Simulations

## ■ Simulation settings

- NS2 simulator; 3-layer, 4-port FatTree
- 40Gbps fabric, 10Gbps host; 64 hosts, 20 switches
- Empirical failure generation

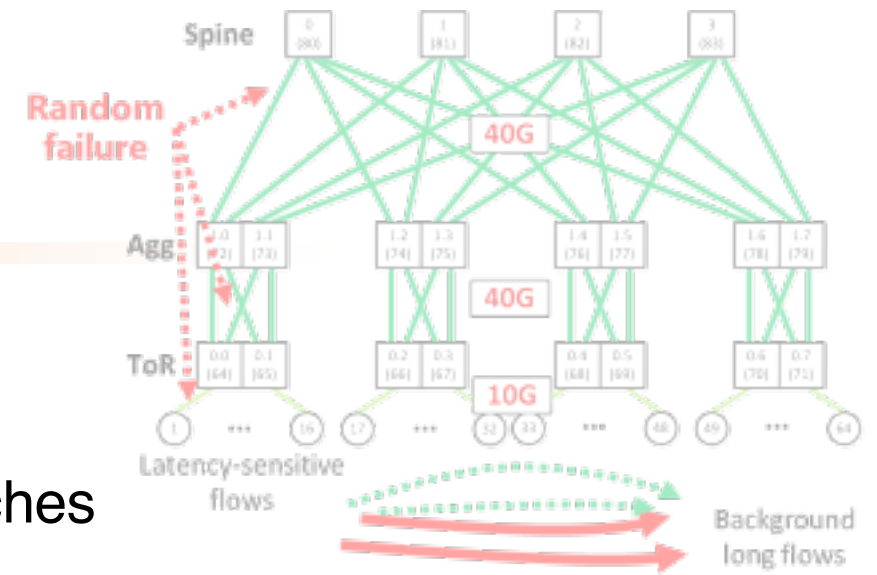




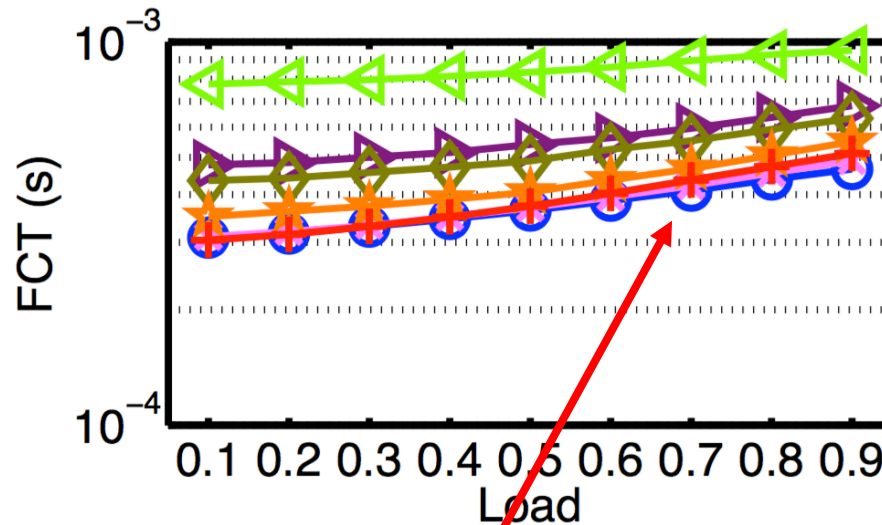
# Larger-scale Simulations

## Simulation settings


- NS2 simulator; 3-layer, 4-port FatTree fabric
- 40Gbps fabric, 10Gbps host; 64 hosts, 20 switches
- Empirical failure generation



- FUSO
- △ MPTCP
- × RepFlow
- ▽ TCP
- ◇ TLP
- ★ TCP-IR
- + Proactive



**Reducing the average FCT up to ~60.3%**



**Reducing the 99<sup>th</sup> FCT up to ~87.4%**

↓  
**better**



# Outline

---

- Motivation
- Packet Loss in DCN
- Impact of Packet Loss
- Challenge for Loss Recovery
- FUSO Design
- Evaluation
- **Summary**

# Summary

---

- Loss hurts tail latency
  - Loss is not uncommon
  - A little loss leads to enough timeout, hurting the tail
- Challenges for loss recovery
  - How to accelerate loss recovery under various loss conditions without causing congestion?
- Philosophy for FUSO
  - To be fast & cautious are equally important
  - **Fast:** Proactive loss recovery utilizing spare transmission opportunity, leveraging multipath diversity
  - **Cautious:** Strictly follows congestion control without adding aggressiveness

# Thanks

Q&A?