LEOTP: An Information-centric Transport Layer Protocol for LEO Satellite Networks

Li Jiang*, Yihang Zhang*, Jinyu Yin*, Xinggong Zhang*, Bin Liu†
* Wangxuan Institute of Computer Technology, Peking University, Beijing, China
† Department of Computer Science and Technology, Tsinghua University, Beijing, China

Presentation for ICDCS 2023

Overview

The LEO satellite network is fast emerging. Unfortunately **CUBIC TCP** experiences **low bandwidth utilization (23.3%)** in it [1].

⇩

Segmented transmission control has advantages, but it face new challenges (connection, reliability, backlog)

⇩

Our idea: borrow the idea of **Information-Centric Networking (ICN)**, and re-design the **retransmission** mechanism and **congestion control** algorithm

[1] SaTCP: Link-Layer Informed TCP Adaptation for Highly Dynamic LEO Satellite Networks (INFOCOM 2023)

Background

LEO satellite networks are fast emerging

amazon | project kuiper

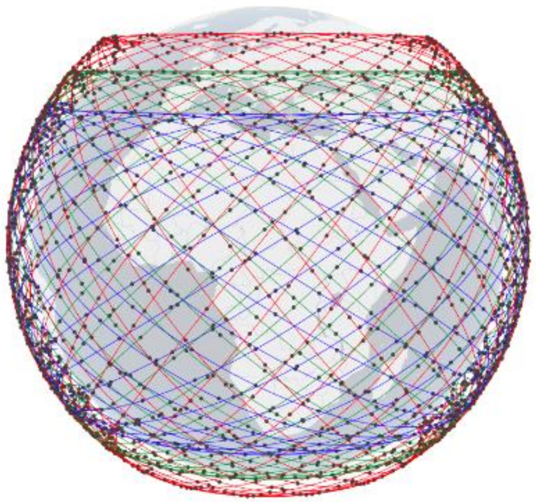

STARLINK

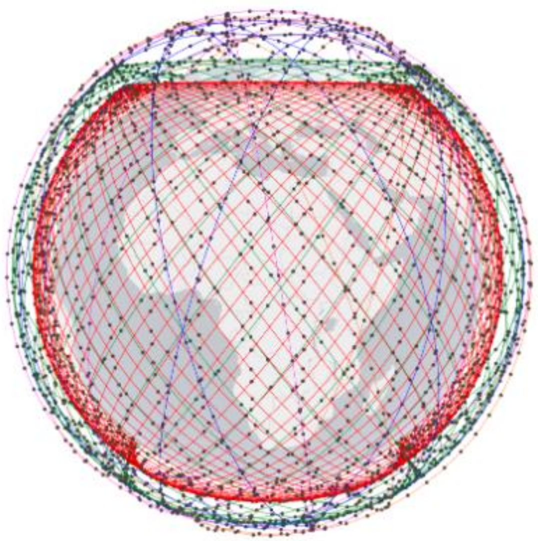

Telesat

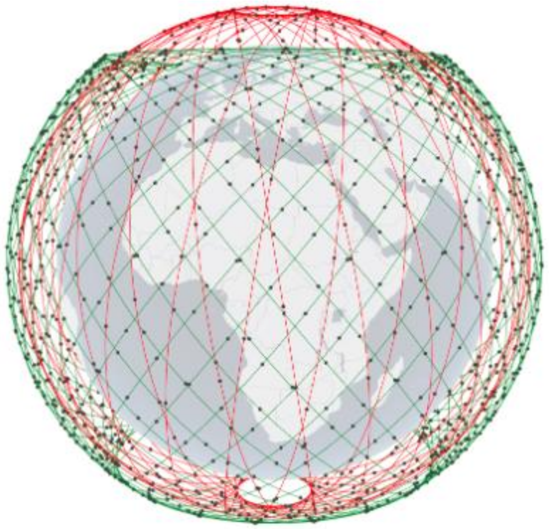

- Wide coverage
- High speed
- Low latency

Background

PEKING UNIVERSITY

Challenges for TCP in LEO satellite networks

- **High packet loss:** *1.56% for downloads and 1.96% for uploads[2]*
  - Degrade throughput of loss-based congestion control
  - High tail delay due to retransmission

[2] Francois Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. A first look at starlink performance. In Proceedings of the 22nd ACM Internet Measurement Conference, IMC ’22, page 130–136, New York, NY, USA, 2022. Association for Computing Machinery.

Background

Challenges for TCP in LEO satellite networks

➢ **High queuing delay:** *95th RTT is 175ms when propagation delay is only 20ms[2]*

- Hard to support latency-sensitive services

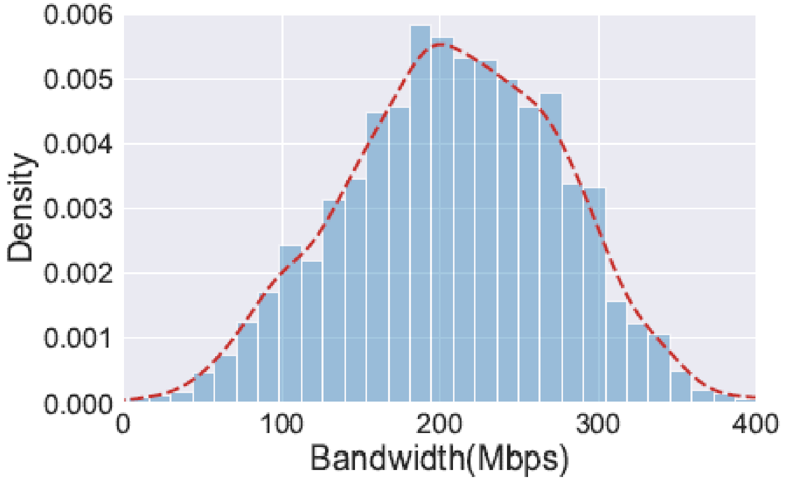

(a) Bandwidth distribution in Starlink

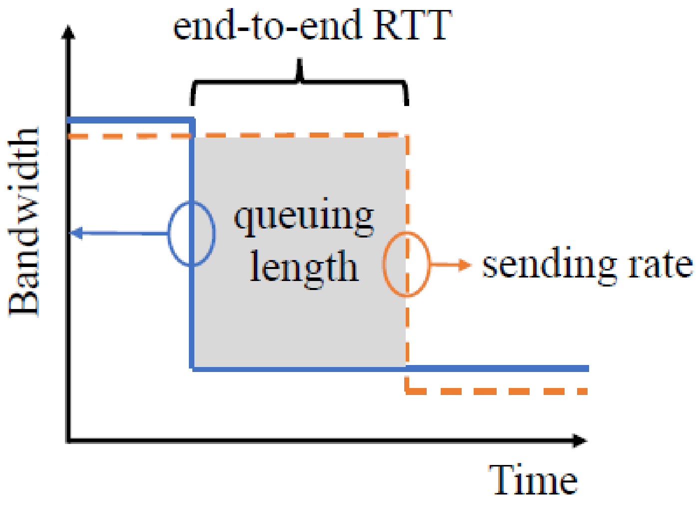

(b) Queuing due to bandwidth variations

Bandwidth variation with delayed feedback causes high queuing delay

Related work

End-to-end improvements

- TCP variants for satellite networks
  - Hybla
  - Peach

Aim at GEO networks

- Other transport layer protocols
  - QUIC

Limited by end-to-end transmission

Related work

In-network enhancements

Split TCP

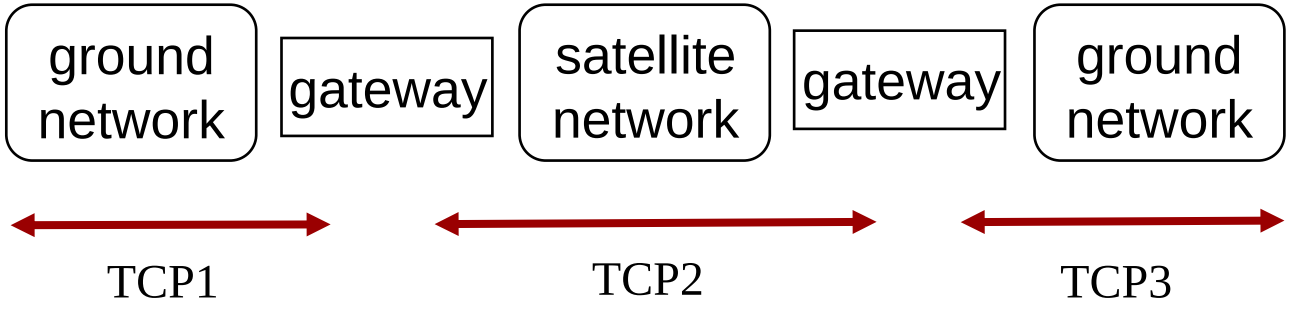

However, it does not work for LEO……

- It can not keep connection in dynamic topology
- It can not guarantee end-to-end reliability
- Packet backlog at intermediate nodes

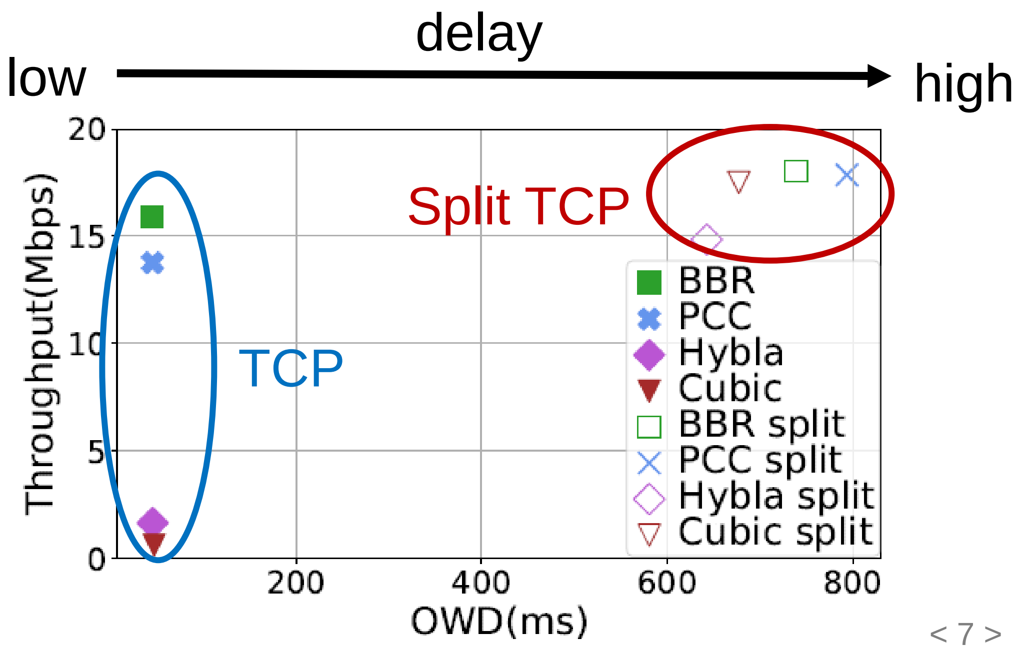

Idea

➢ Segmented transmission control

- Local loss detection -> Low cost for retransmission
- Congestion control -> Faster reaction to bandwidth variations

➢ Request-response model

- Location-independent naming
- In-network caching

↓

Support for mobility of intermediate nodes

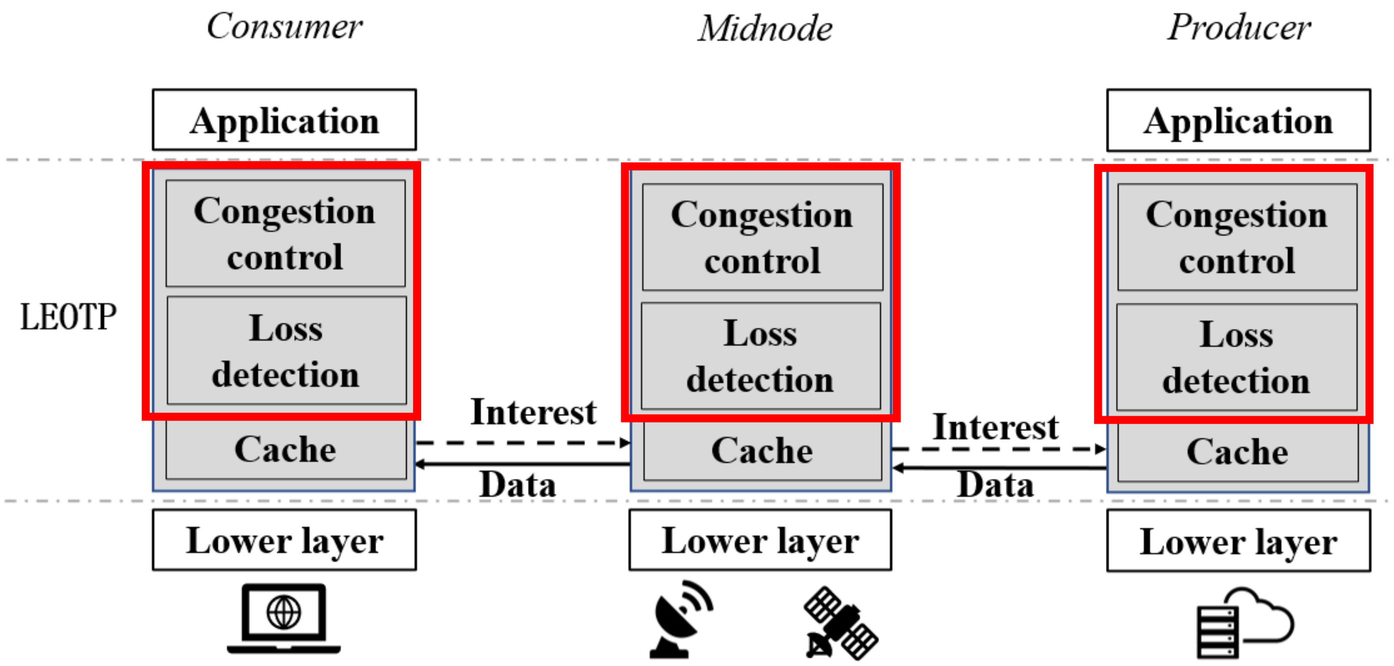

Idea

- Segmented transmission control
  - Local loss detection -> Low cost for retransmission
  - Congestion control -> Faster reaction to bandwidth variations

- Request-response model
  - Location-independent naming
  - In-network caching

↓

Support for mobility of intermediate nodes

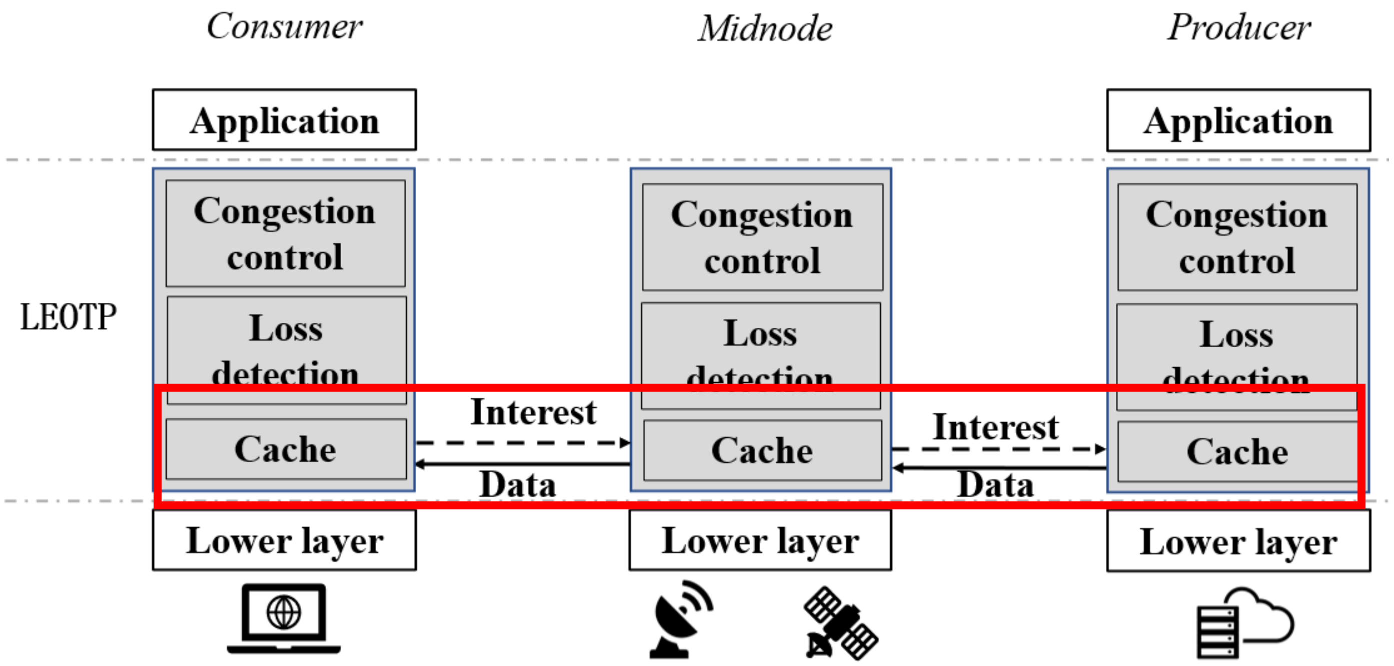

Design

	Endpoint	In-network
Hybrid retransmission	1. Consumer-driven retransmission	2. In-network retransmission
Hop-by-hop congestion control		3. Backpressure congestion control

Design-1: Consumer-driven retransmission

PEKING UNIVERSITY 北京大学

- Based on timeout mechanism

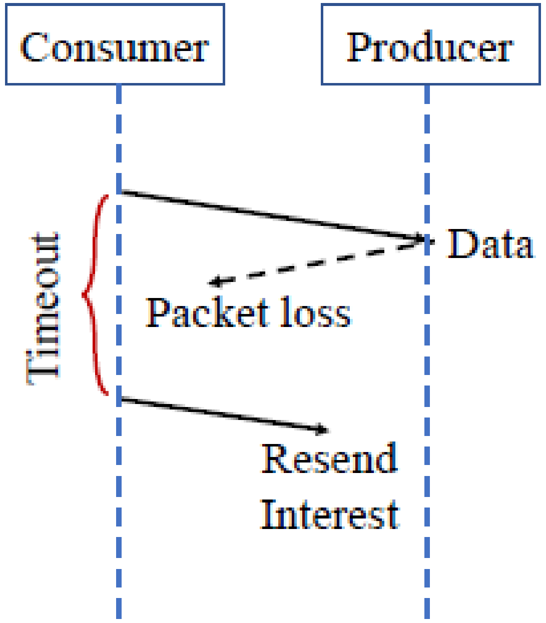

(a) TR mechanism

Guarantee end-to-end reliability at the last sort

Design-2: In-network retransmission

北京大学 PEKING UNIVERSITY

- Based on sequence number holes

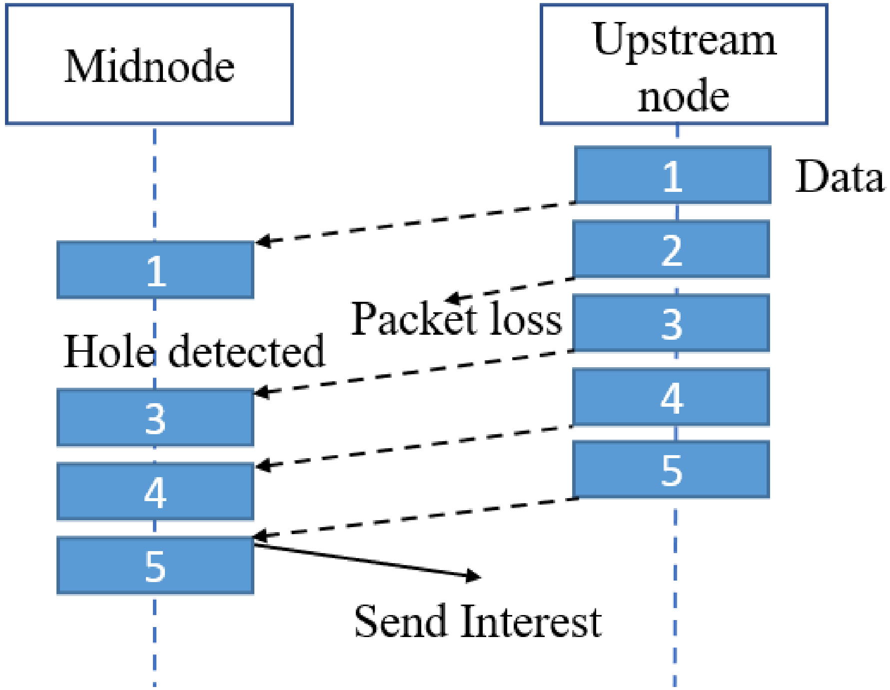

Recover most of the lost packets at a low cost of latency and bandwidth consumption

Design-2: In-network retransmission

How to avoid repeated retransmission?

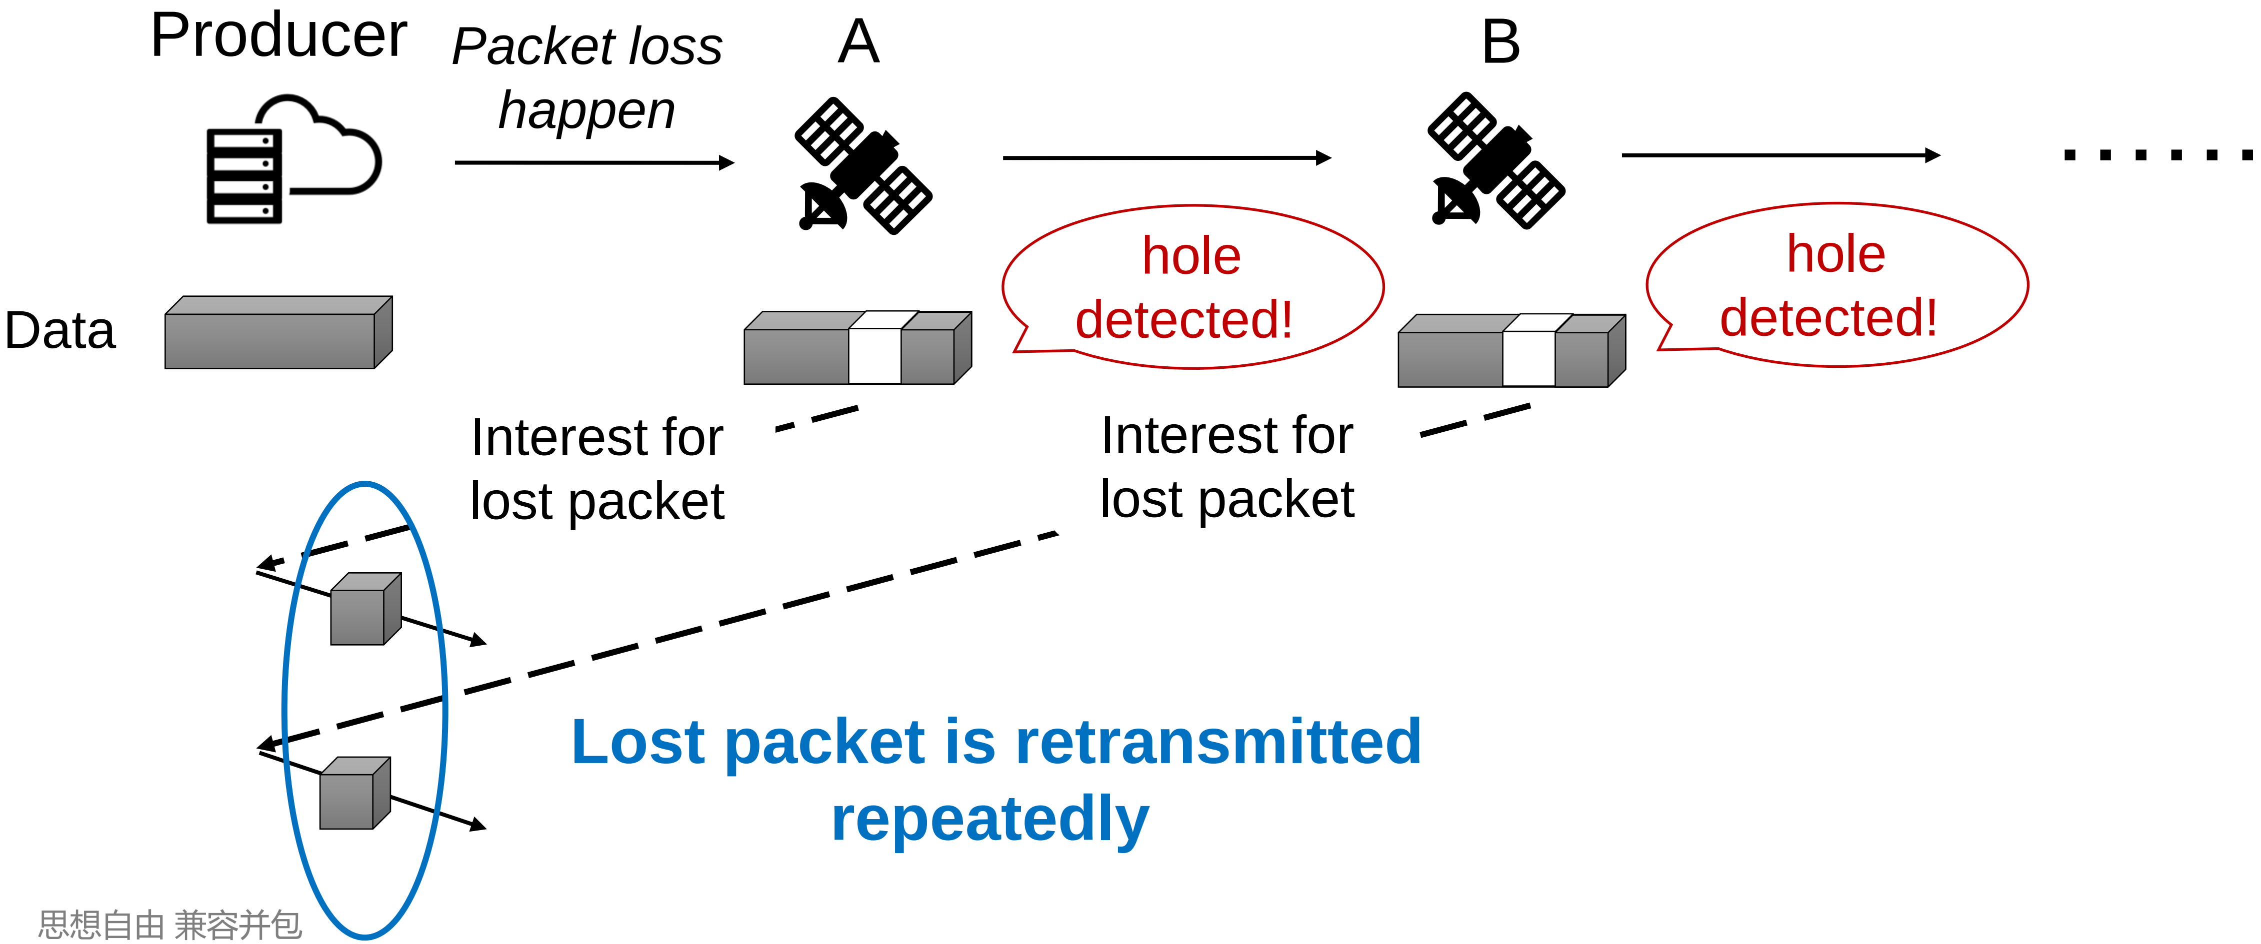

Lost packet is retransmitted repeatedly

Design-2: In-network retransmission

PEKING UNIVERSITY 北京大學

- **Void Packet Header (VPH) mechanism:** VPH are sent downstream as notification when detecting packet loss

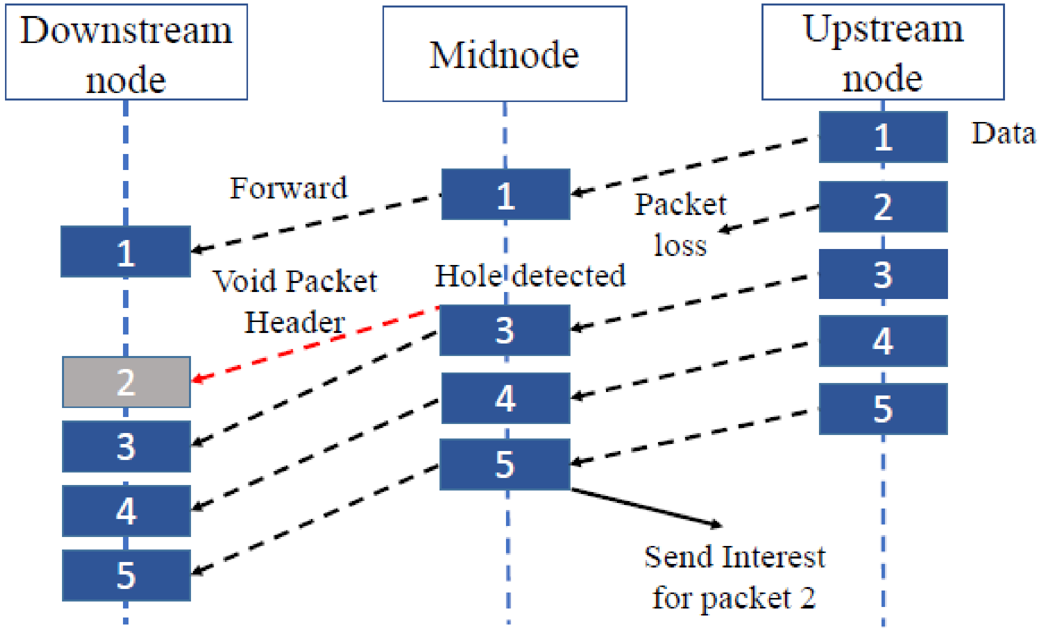

(b) SHR mechanism

Avoid duplicated retransmission

Design-3: Backpressure congestion control

北京大学 PEKING UNIVERSITY

Backpressure in hop-by-hop congestion control

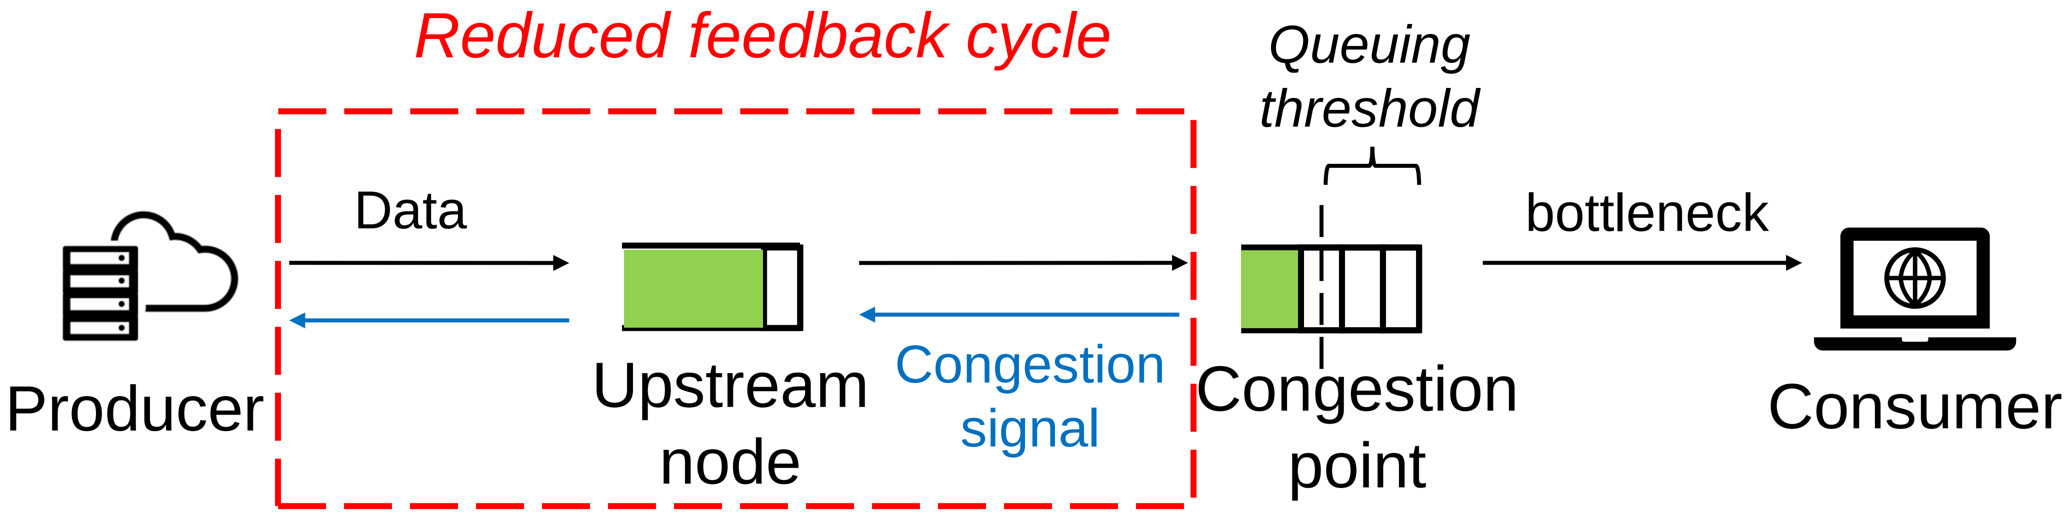

Avoid backlog at intermediate nodes

Design-3: Backpressure congestion control

- Requester-driven

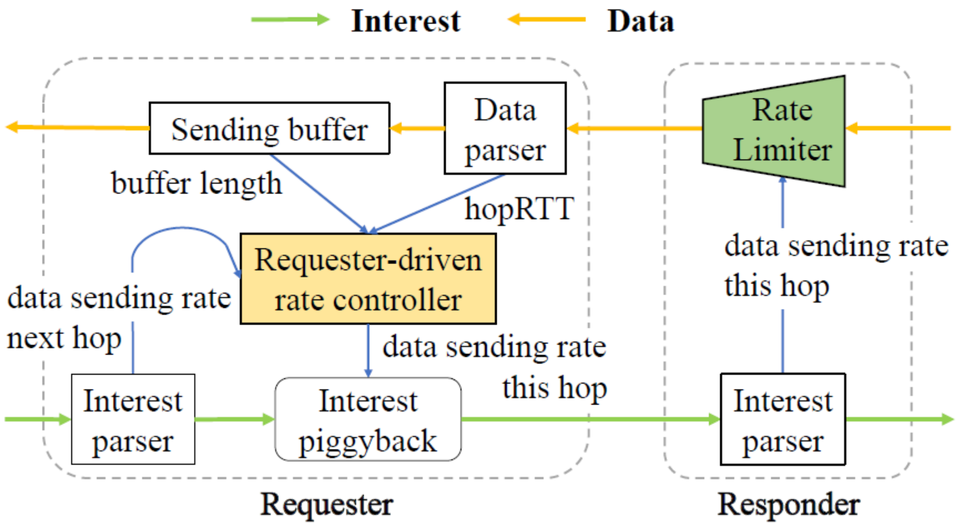

Design-3: Backpressure congestion control

北京大学 PEKING UNIVERSITY

- Requester-driven

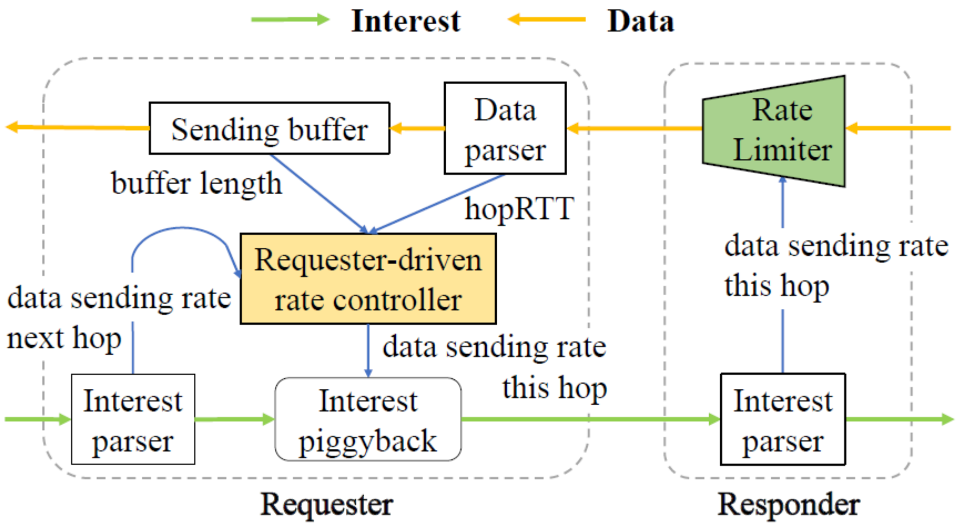

- RTT-based at individual hop

$$BDP = throughput * hopRTT_{min} \quad (6)$$

$$QueueLen = throughput * (hopRTT - hopRTT_{min}) \quad (7)$$

$$\boxed{cwnd} = \begin{cases} 2 * cwnd, & \text{if } state == SlowStart \\ cwnd + 1, & \text{else if } QueueLen \le M \\ k * BDP, & \text{otherwise} \end{cases} \quad (8)$$

Design-3: Backpressure congestion control

- Requester-driven

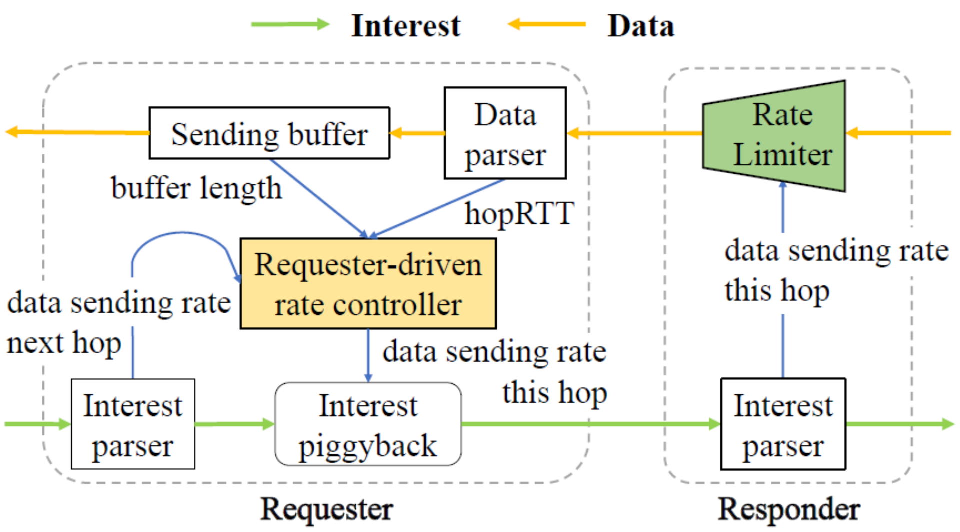

- RTT-based at individual hop

$$BDP = throughput * hopRTT_{min} \quad (6)$$

$$QueueLen = throughput * (hopRTT - hopRTT_{min}) \quad (7)$$

$$\boxed{cwnd} = \begin{cases} 2 * cwnd, & \text{if } state == SlowStart \\ cwnd + 1, & \text{else if } QueueLen \leq M \\ k * BDP, & \text{otherwise} \end{cases} \quad (8)$$

- Backpressure between hops

$$\boxed{Rate_{bp}} = Rate_{nextHop} + \frac{BL - BL_{tar}}{hopRTT} \quad (9)$$

=> $$\boxed{Rate} = min(\frac{\boxed{cwnd}}{hopRTT}, \boxed{Rate_{bp}}) \quad (10)$$

Evaluation

北京大學 PEKING UNIVERSITY

- Experiment setup
  - Network emulator Mininet [3]
  - Dynamic route is provided by Hypatia [4]
  - Packet loss and bandwidth variations are simulated close to real network
- Baseline methods
  - Cubic: default congestion control algorithm in Linux kernel(loss-based)
  - Hybla: a classic TCP variant designed for satellite networks (loss-based)
  - BBR & PCC: state-of-the-art TCP variants

[3] Bob Lantz, Brandon Heller, and Nick McKeown. A network in a laptop: rapid prototyping for software-defined networks. In Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks, pages 1–6, 2010.
[4] Simon Kassing, Debopam Bhattacherjee, Andre´ Baptista A´guas, Jens Eirik Saethre, and Ankit Singla. Exploring the ”internet from space” with hypatia. In Proceedings of the ACM Internet Measurement Conference, IMC ’20, page 214–229, New York, NY, USA, 2020. Association for Computing Machinery.

Evaluation

Controlled experiments

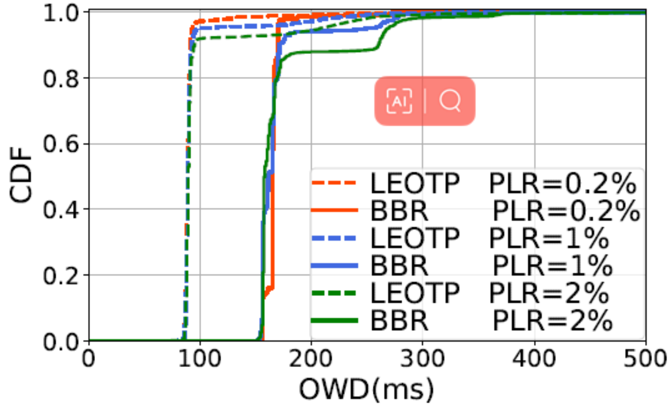

Fig. 10: The distribution of the re-transmitted packets' OWD in lossy link.

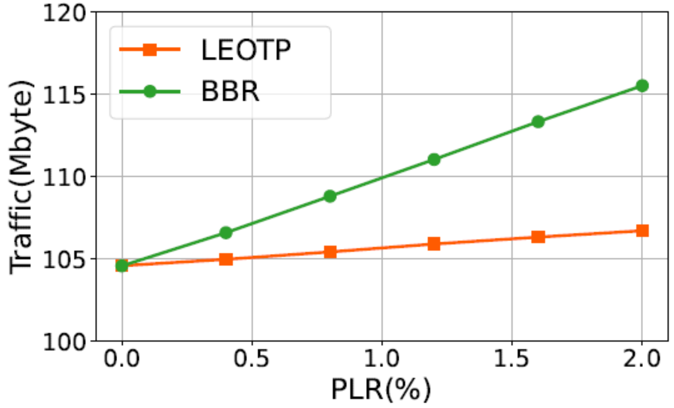

Fig. 11: The relation of loss rate and the traffic sent by sender for an 100MB file.

Low cost for retransmission

Evaluation

Controlled experiments

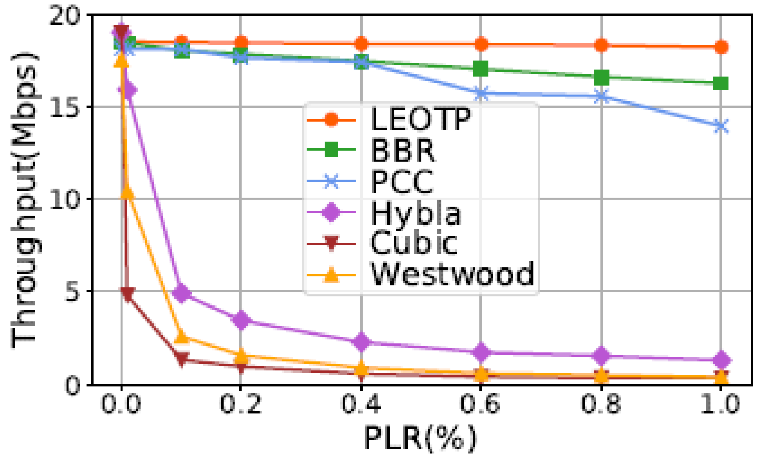

Fig. 12: The relation of loss rate and throughput.

High throughput against high PLR

Low latency under bandwidth variations

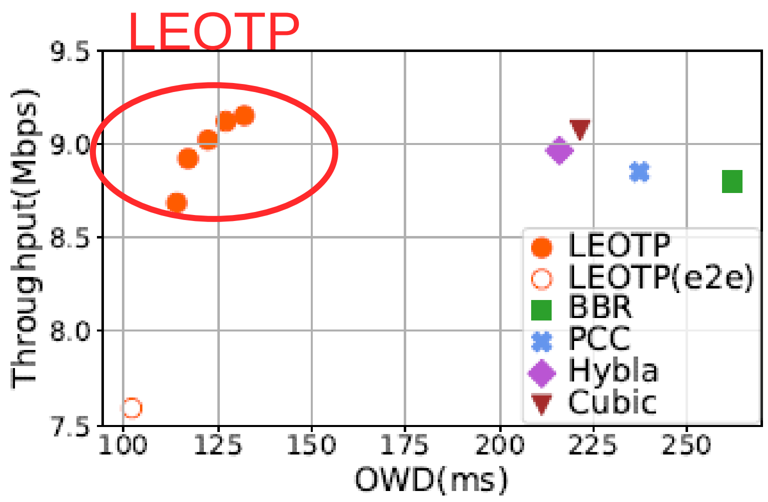

Fig. 14: Throughput-OWD trade-off under bandwidth fluctuations.

Evaluation

Emulation experiments

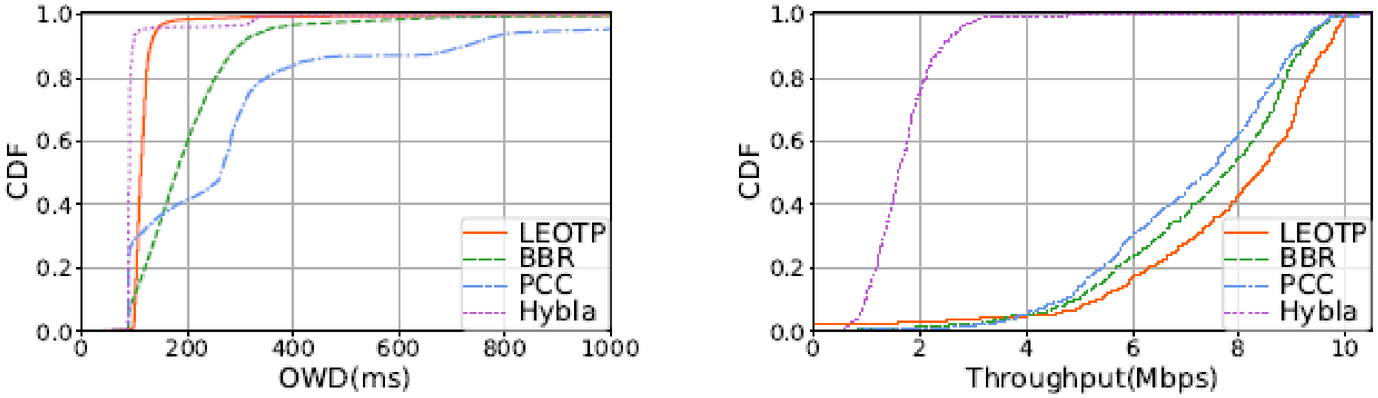

(a) OWD (b) throughput

Fig. 17: Cumulative distribution graph of OWD and throughput in Beijing-New York link with ISLs.

8%-12% higher throughput with 40%-60% delay reduction in transcontinental data transmission

Evaluation

Emulation experiments

TABLE II: The result of the ablation experiment

	BJ-HK		BJ-PR		BJ-NY	
	Throughput (Mbps)	OWD (ms)	Throughput (Mbps)	OWD (ms)	Throughput (Mbps)	OWD (ms)
A	7.82	49.17	7.70	76.57	7.91	118.64
B	7.78	51.39	7.67	80.74	7.73	126.10
C	7.38	40.15	7.23	66.40	6.80	103.63
D	7.24	42.05	7.03	70.38	6.52	112.20

	In-network retransmission	Hop-by-hop congestion control
A	√	√
B	×	√
C	√	×
D	×	×

Both **retransmission** module and **congestion control** module contribute to the better performance of LEOTP

Conclusion

北京大学 PEKING UNIVERSITY

- End-to-end transport protocols have limitations in LEO satellite networks
- We present LEOTP, an information-centric, cache-assisted transport protocol
- The in-network retransmission use VPH as notifications, reducing redundant retransmissions while providing fast loss recovery
- The backpressure-based congestion control provides quick reactions in long-distance networks
- Results: reliability, high throughput, and low latency
- More resources: *https://jl99888.github.io/LEOTP*

Thanks

Netvideo Group https://www.icst.pku.edu.cn/NetVideo/

Email: jl99888@pku.edu.cn