LeoSVCSch: SVC-Aware Multipath Scheduling for Wireless Real-Time Video over Dynamic LEO Networks

Xiankang Peng, Wang Yang*, Lihuan Hui
School of Computer Science and Engineering
Central South University
{xiankangpeng, yangwang, lh_hui}@csu.edu.cn

***Abstract*—Low Earth Orbit (LEO) satellite networks can shorten wide-area paths for long-distance real-time video, but satellite handovers and weather-induced attenuation can cause RTT inflation, bandwidth drops, burst losses, and short outages. Existing multipath schedulers react only after impairments appear and do not explicitly use handover information or video-layer importance. This paper presents LeoSVCSch, a Scalable Video Coding (SVC)-aware multipath scheduler for real-time video over candidate LEO paths. LeoSVCSch prioritizes base-layer packets on safer paths, drops enhancement-layer packets when they compete with protected traffic, and uses bounded duplicate protection during persistent degradation. We implement LeoSVCSch in a testbed carrying WebRTC video over multipath QUIC and replaying LEO traces in Mininet. Compared with representative transport-centric and redundancy-aware multipath schedulers, LeoSVCSch improves average display frames per second (FPS) by about 32% and reduces freeze ratio by about 76%, with lower QUIC goodput because it drops less critical enhancement-layer packets and occasionally uses bounded duplicate protection.**

***Index Terms*—LEO Satellite Networks, Wireless Multimedia, Real-Time Video, Scalable Video Coding, Multipath QUIC, Multipath Scheduling**

I. Introduction

Real-time video applications, including conferencing, remote monitoring, and interactive collaboration, depend on timely frame delivery. Unlike elastic transfers, delay variation and delivery interruptions directly appear as frame-rate instability or freezes [1]. Long-distance sessions are especially vulnerable because client-to-cloud and inter-domain traffic can follow congested or indirect terrestrial routes [2], [3].

Low Earth Orbit (LEO) satellite networks provide broad coverage and relatively low propagation delay, offering an alternative for long-distance real-time video [2], [4]. However, satellite mobility, beam switching, and routing updates create handover windows, whereas cloud blockage, rain attenuation, and atmospheric disturbance cause longer-lived bandwidth reduction, loss, and jitter [5], [6]. Measurements report latency variation, packet and frame loss, and audio interruptions around satellite switching and adverse weather [7]–[9].

Multipath transport can reduce dependence on any single LEO path by steering traffic to other candidates [10], [11]. Existing schedulers, however, react mainly to round-trip time (RTT), loss, bandwidth, or completion-time estimates [12], [13]. A path may still appear favorable immediately before a predictable handover, and transport-centric policies neither distinguish short handover disruptions from persistent attenuation nor account for video-layer importance [6], [14]. For example, a MinRTT-style scheduler may still select a path close to a handover because its current RTT remains low. If decoding-critical video packets enter the following burst-loss interval, the receiver may fail to decode subsequent frames, turning a short path event into a visible freeze.

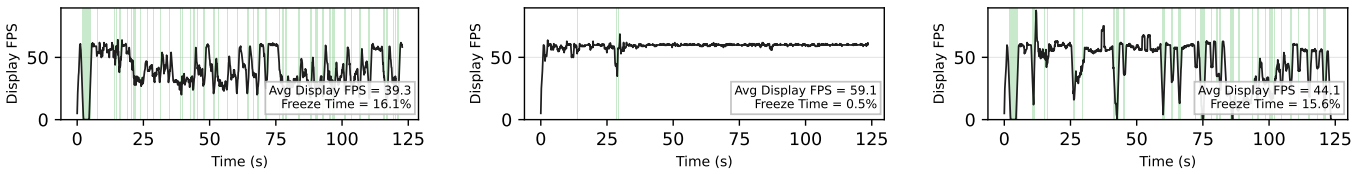
To make scheduling decisions under such LEO dynamics, the scheduler needs a way to decide which packets should be protected first. Scalable Video Coding (SVC) provides an application-level priority signal by separating a stream into a base layer required for decoding and enhancement layers that improve quality or temporal smoothness [15], [16]. Under impaired conditions, a scheduler can therefore protect base-layer delivery and send enhancement-layer packets only when safe capacity remains.

We propose LeoSVCSch, an SVC-aware multipath scheduler for real-time video over multiple candidate LEO satellite paths. LeoSVCSch addresses the limitations of transport-centric multipath scheduling in dynamic LEO links by separating predictable handover disruption from persistent weather-induced attenuation. For handovers, it steers base-layer packets away from paths entering unstable windows. For persistent degradation, it drops less critical enhancement-layer packets when safe capacity is insufficient and applies bounded duplicate protection to selected protected packets. Under favorable conditions, it uses residual multipath capacity for enhancement-layer quality; when degradation persists, it keeps capacity and duplicate budget for base-layer packets.

Our main contributions are summarized as follows:

- **Layer-aware multipath scheduling.** We design a multipath QUIC scheduler that uses SVC labels to prioritize base-layer delivery and allocate residual capacity to enhancement traffic.
- **Separate actions for handovers and persistent degradation.** LeoSVCSch combines proactive handover avoidance with enhancement-layer control and bounded duplicate protection during persistent degradation.
- **Prototype and evaluation.** We implement LeoSVCSch in a testbed carrying WebRTC video over multipath QUIC

979-8-3195-0167-7/26/$31.00 ©2026 IEEE



(a) Terrestrial baseline (b) Ideal LEO (c) Impaired LEO

Fig. 1. Playback performance under the three network conditions. Each subplot reports average FPS and freeze ratio.

and replay LEO traces in Mininet, evaluating display FPS, freeze ratio, FPS jitter, one-way delay, delay jitter, and QUIC goodput.

II. Background and Motivation

A. Overview of SVC

SVC encodes a video stream into multiple dependent layers. The base layer provides the minimum decodable stream, while enhancement layers refine resolution, frame rate, or visual quality [15]. A receiver can still display video when only the base layer arrives; additional enhancement layers improve quality when link conditions permit [17]. We use the temporal layering mode of VP9-SVC [18]. For clarity, we refer to the lower temporal layer as the base layer and the higher temporal layer as the enhancement layer. The base layer provides the essential frame stream, whereas the enhancement layer inserts additional frames to improve motion smoothness. This distinction gives the scheduler an explicit packet-importance signal and motivates layer-aware multipath scheduling in LeoSVCSch [19].

B. Motivation for LeoSVCSch

To better understand the potential and limitations of LEO links for long-distance real-time video, we conduct a single-path transmission experiment under three conditions using the same source, encoding configuration, and transmission pipeline: a cross-regional terrestrial baseline, an ideal LEO link without handovers or weather-induced attenuation, and an impaired LEO link with both effects. The ideal LEO case captures the low-latency potential of stable satellite operation, while the impaired case shows how these impairments affect playback [20].

As shown in Fig. 1, the ideal LEO link increases average FPS from 39.3 to 59.1 by 50.4% and reduces the freeze ratio from 16.1% to 0.5% by 96.9% relative to the terrestrial baseline. With handovers and weather-induced attenuation, FPS decreases from 59.1 to 44.1 by 25.4%, while the freeze ratio increases from 0.5% to 15.6%. These results show that LEO can improve playback when the link is stable, but handovers and weather-induced attenuation can remove much of this benefit through outages, delay jitter, and burst loss, consistent with prior measurements of Starlink handover behavior [14], [21].

LEO satellite connectivity can provide shorter wide-area paths, but the gain is limited when a handover window, bandwidth drop, or weather-related loss burst occurs. Multiple candidate LEO paths allow a scheduler to move base-layer packets away from a risky path and use residual capacity for enhancement-layer packets. However, conventional multipath schedulers mainly react to RTT, bandwidth, and loss after degradation occurs. They neither anticipate predictable handover windows nor distinguish short handover disruptions from longer weather-induced attenuation.

This distinction matters because the preferred action depends on event duration: short handover windows favor proactive path migration, whereas longer degradation periods may require enhancement-layer suppression and bounded duplicate protection.

SVC-aware scheduling gives these path decisions a video-level priority order. Rather than maximizing aggregate throughput alone, LeoSVCSch prioritizes timely base-layer delivery during handover disruption and persistent weather-induced attenuation, while limiting enhancement-layer traffic when capacity is insufficient. It implements these layer-aware decisions at the packet scheduler, where SVC-layer labels and LEO link-state evidence are available together.

III. Design of LeoSVCSch

LeoSVCSch is an SVC-aware multipath scheduler for real-time video delivery over candidate LEO paths. Its design is motivated by two observations. First, not all video packets have the same importance: base-layer packets require stronger delivery protection, whereas enhancement-layer packets can tolerate more aggressive dropping when resources are limited. Second, LEO-path impairments arise from both predictable events, such as satellite handovers, and unpredictable events, such as weather-induced attenuation. Treating all losses uniformly can therefore cause unnecessary path changes, excessive duplication, or inefficient resource use.

To address these challenges, LeoSVCSch combines packet-layer awareness with path-level network evidence.

A. System Overview and Decision Flow

LeoSVCSch makes per-packet forwarding decisions at the Selective Forwarding Unit (SFU)/forwarder using four inputs: SVC layer labels, QUIC path metrics, QUIC DATAGRAM delivery evidence, and path-specific handover windows obtained from satellite ephemeris or switch plans. As shown in Fig. 2, these inputs support path ranking, handover avoidance, Persistent Degradation State Detection, enhancement-layer control, and bounded duplicate protection. The signals remain separate: Candidate Ranking represents ordinary path

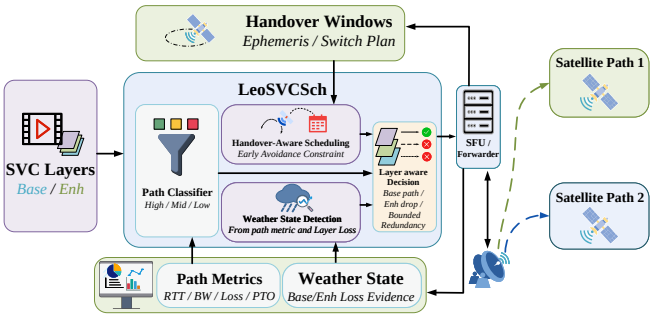


Fig. 2. Overview of LeoSVCSch.

quality, handover windows impose a path-specific handover-risk constraint, and DATAGRAM evidence drives connection-level degradation detection.

SVC labels distinguish base-layer packets needed for decoding from enhancement-layer packets that improve temporal smoothness. For each active and sendable path, the Path Classifier uses availability, RTT, bandwidth condition, recent loss, and probe timeout (PTO) or abnormal state. Its ranking key is

$$\Phi_i(t) = \big(C_i(t), r_i(t), \pi_i(t), \ell_i(t), id_i\big), \tag{1}$$

where $C_i(t)$ encodes availability and a High/Mid/Low quality class, $r_i(t)$ is smoothed RTT, $\pi_i(t)$ is the PTO or abnormal state, $\ell_i(t)$ is recent loss rate, and id_i provides deterministic ordering. Bandwidth contributes to $C_i(t)$ rather than appearing separately in the tuple.

The scheduler updates Candidate Ranking and the handover mask $m(t)$ before using unmasked DATAGRAM evidence to update the Persistent Degradation State. At send time, it selects the primary base-layer path, sends enhancement-layer packets only when safe residual capacity remains, and considers duplication only after scheduling the primary copy.

B. Handover-Aware Base-Layer Scheduling

LeoSVCSch treats LEO handovers as predictable, short-term, and path-specific risks. Algorithm 1 applies handover-aware base-layer path migration (Early Avoidance): base-layer packets avoid the affected path when possible, while enhancement-layer packets use only capacity remaining after protected traffic.

Let $\mathcal{P} = \{p_1, p_2, \ldots, p_N\}$ be the candidate path set, and let $a_i(t) \in \{0,1\}$ indicate whether p_i is active and sendable. For each known handover interval $[T_i^s, T_i^e]$, LeoSVCSch adds an adaptive lead margin and a fixed tail margin:

$$\begin{aligned} \Delta^{\text{lead}}(t) &= \min\!\big\{\Delta^{\text{lead}}_{\max}, \max\!\big\{\Delta^{\text{lead}}_{\min}, r_{\text{ref}}(t)\big\}\big\}, \\ T_i^-(t) &= T_i^s - \Delta^{\text{lead}}(t), \\ T_i^+ &= T_i^e + \Delta^{\text{tail}}, \\ h_i(t) &= \begin{cases} 1, & T_i^-(t) \le t \le T_i^+, \\ 0, & \text{otherwise.} \end{cases} \end{aligned} \tag{2}$$

The lead margin adapts to smoothed RTT $r_{\text{ref}}(t)$ within configured bounds, while the tail margin prevents premature

Algorithm 1 Handover-Aware Base-Layer Scheduling

Require: packet x, layer $L(x)$, paths $\mathcal{P}$, path metrics, handover windows
Ensure: selected path or enhancement-layer dropping decision

```
Obtain Candidate Ranking Φ_i(t)
Compute handover-risk constraint h_i(t) for each path
if x belongs to the base layer then
    Select the best active path with h_i(t) = 0 using Φ_i(t)
    Fall back to the best active path if all paths are risky
    Hold the current path while it remains safe and stable
    Schedule x on the selected path
else
    if safe capacity remains after base-layer protection then
        Follow the residual-capacity policy
    else
        Drop the enhancement-layer packet
    end if
end if
```

path reuse after the switch. The external-event-driven state $h_i(t)$ constrains path selection without redefining the ordinary quality class $C_i(t)$.

A base-layer packet first uses the highest-ranked active path outside the risk interval:

$$p_B^*(t) = \arg \min_{\substack{p_i \in \mathcal{P} \\ a_i(t)=1,\, h_i(t)=0}} \Phi_i(t). \tag{3}$$

If no safe path exists, the scheduler falls back to the best active path under $\Phi_i(t)$. Path holding retains the current base-layer path while it remains safe and stable. Enhancement-layer packets may use residual capacity on another path or be dropped when they would compete with protected traffic.

C. Persistent Degradation State and Bounded Duplicate Protection

Weather-induced attenuation is persistent, uncertain, and may affect multiple paths, so avoiding one known path is insufficient. LeoSVCSch instead derives a connection-level Persistent Degradation State from QUIC DATAGRAM delivery evidence and reserves capacity for protected traffic while the state is active. Algorithm 2 summarizes this detection and protection process.

Let $\mathcal{Q}$ be the set of packet classes observed by the forwarder: base-layer, enhancement-layer, and timing-critical control packets. For each $q \in \mathcal{Q}$, let $S_q(t)$ and $L_q(t)$ be the transmitted and lost DATAGRAM counts in the recent detection window. The aggregate loss ratio is

$$\ell_{\text{sys}}(t) = \frac{\sum_{q\in\mathcal{Q}} L_q(t)}{\sum_{q\in\mathcal{Q}} S_q(t) + \epsilon}, \tag{4}$$

where ϵ prevents division by zero. Recent-window and smoothed loss estimates, together with consecutive bad and recovery observations, form persistent degradation and recovery evidence.

Algorithm 2 Persistent-Degradation Protection

Require: packet x, layer $L(x)$, Candidate Ranking $\Phi_i(t)$, DATAGRAM counters, mask $m(t)$, state $w(t)$
Ensure: selected path set or enhancement-layer dropping decision

if $m(t) = 1$ **then**
  Exclude the current sample and keep $w(t) = w(t-1)$
else
  Update recent-window and smoothed loss estimates
  Update bad/recovery counters
  Update the Persistent Degradation State $w(t)$
end if
if $w(t) = 0$ **then**
  Return the normal layer-aware decision
else if x belongs to the enhancement layer **then**
  if safe residual capacity is available **then**
    Follow the residual-capacity policy
  else
    Drop the enhancement-layer packet
  end if
else
  Select the primary path using $\Phi_i(t)$ and the handover-risk constraint
  Schedule x on the primary path
  Evaluate $d_{\mathrm{dup}}(x,t)$
  if $d_{\mathrm{dup}}(x,t) = 1$ **then**
    Schedule at most one duplicate on another available path
  end if
end if

The handover mask is tied to the expanded intervals above: $m(t) = 1$ when the current detection sample lies within $[T_i^-(t), T_i^+]$ for at least one active path p_i, and $m(t) = 0$ otherwise. Let $E_{\mathrm{deg}}(t)$ and $E_{\mathrm{rec}}(t)$ denote persistent degradation and recovery evidence. The state evolves with hysteresis:

$$w(t) = \begin{cases} 1, & m(t) = 0 \text{ and } E_{\mathrm{deg}}(t), \\ 0, & m(t) = 0 \text{ and } E_{\mathrm{rec}}(t), \\ w(t-1), & \text{otherwise.} \end{cases} \tag{5}$$

The mask prevents a predictable handover burst from becoming degradation evidence, while hysteresis filters isolated losses and prevents recovery after a single good sample.

When $w(t) = 1$, LeoSVCSch selects the primary path using Candidate Ranking and the current handover-risk constraint. The duplicate decision is

$$d_{\mathrm{dup}}(x,t) = \begin{cases} 1, & L(x) \in \mathcal{L}_{\mathrm{prot}},\ w(t) = 1, \\ & \text{protection policy permits,} \\ 0, & \text{otherwise,} \end{cases} \tag{6}$$

where $\mathcal{L}_{\mathrm{prot}}$ contains protected packet classes, including base-layer and timing-critical control packets. When $d_{\mathrm{dup}}(x,t) = 1$, at most one extra copy is sent, preferably on a different available path. When $w(t)$ returns to zero, duplication is disabled and enhancement-layer packets resume the residual-capacity policy.

IV. Evaluation

A. Experimental Setup

Fig. 3 shows a WebRTC testbed in which a sender browser sends a 720p, 60-fps VP9-SVC stream to the local SFU. The local SFU forwards Real-time Transport Protocol packets to the remote SFU over multipath QUIC, and the remote SFU delivers them to the receiver. SVC labels distinguish base-layer and enhancement-layer packets. All schedulers use Bottleneck Bandwidth and Round-trip propagation time (BBR) with identical video, paths, and traces.

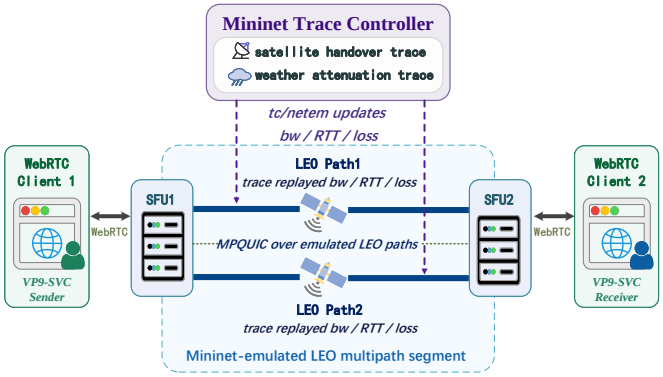
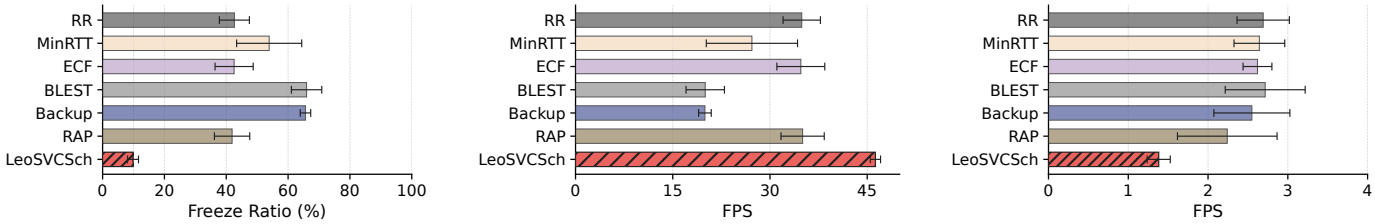


Fig. 3. Experimental testbed. A WebRTC video stream is forwarded between two SFUs over multipath QUIC, while Mininet replays LEO link traces on the candidate paths.

Network and traces. Mininet emulates two independently updated LEO paths between the SFUs, with bandwidth, RTT, and packet loss replayed through `tc/netem`. The base traces use measured LEO link dynamics [22]; persistent weather-induced attenuation is injected as sustained bandwidth reduction, loss, and additional delay. Handover disruption is short and path specific, whereas attenuation is longer-lived and may affect multiple paths.

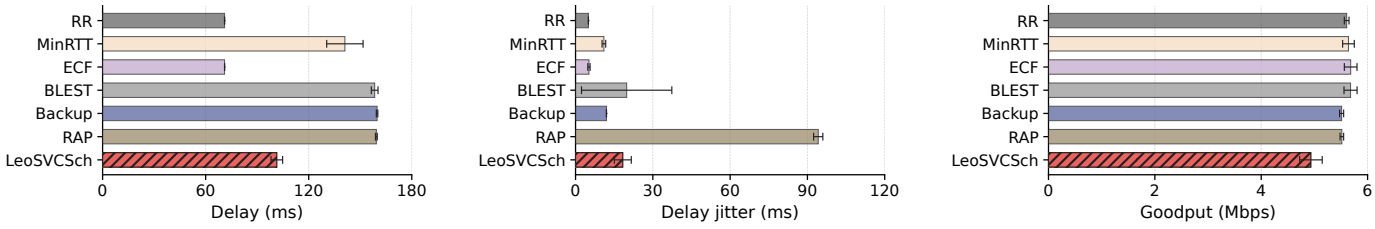
Baselines. We compare Round Robin (RR), MinRTT, Earliest Completion First (ECF), Blocking Estimation-based Multipath TCP scheduler (BLEST), Backup, Reinjection on Any Path (RAP), and LeoSVCSch. They respectively represent load splitting, lowest-RTT selection, completion-time scheduling [12], blocking avoidance [13], primary-backup failover, redundancy-aware scheduling, and the proposed layer-aware scheduler.

Metrics. We report receiver display FPS, freeze ratio, FPS jitter, end-to-end (E2E) one-way delay (OWD) p95, E2E delay jitter p95, and QUIC goodput. Display FPS is the rendered frame rate, and freeze ratio is the fraction of playback time with no newly rendered frame. FPS jitter measures frame-rate variation across adjacent windows, p95 denotes the 95th percentile, and goodput is the video data delivered by QUIC per second.



(a) Freeze ratio (b) Average FPS (c) FPS jitter

Fig. 4. Video playback quality under LEO multipath traces: (a) freeze ratio, (b) average FPS, and (c) FPS jitter.



(a) E2E one-way delay (b) E2E one-way delay jitter (c) QUIC goodput

Fig. 5. Interactive latency and transport performance under LEO multipath traces: (a) E2E OWD p95, (b) E2E delay jitter p95, and (c) QUIC goodput.

B. Overall Performance

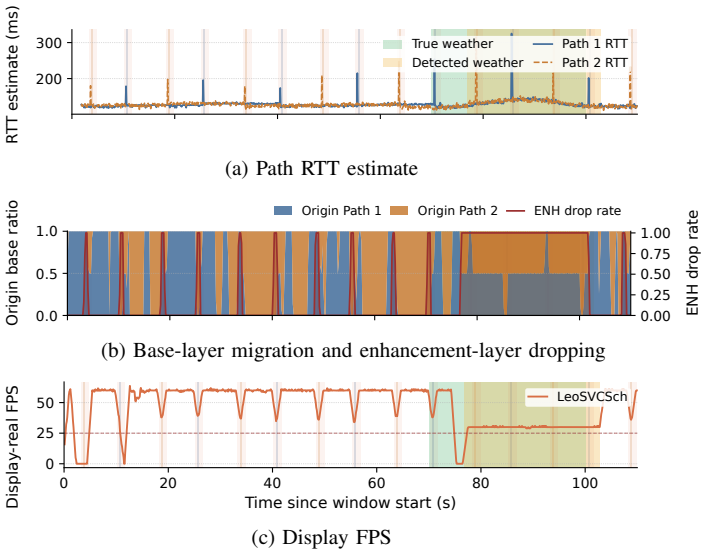
Fig. 4 shows that LeoSVCSch reduces freeze ratio from 41.9% under RAP to 9.8%, a reduction of about 76%, while increasing average display FPS from 35.0 to 46.3, an improvement of about 32%. FPS jitter decreases from 2.24 to 1.38, or about 38%. These gains follow from moving base-layer packets away from handover-risk paths, dropping enhancement-layer packets when protected traffic needs capacity, and applying bounded duplicate protection during persistent degradation. The baselines optimize path use, delay, failover, or redundancy but do not jointly use video-layer importance and predictable handover information.

Fig. 5 shows the latency and goodput tradeoff. RR and ECF obtain lower E2E OWD p95 and E2E delay jitter p95, while LeoSVCSch keeps lower OWD than MinRTT, BLEST, Backup, and RAP. LeoSVCSch may avoid a currently low-RTT path before a handover or reserve capacity for protected packets, so it does not minimize every packet-level delay sample. Its QUIC goodput is 4.93 Mbps, below ECF's 5.68 Mbps, because it drops some enhancement-layer packets and uses capacity for selected duplicates. The lower goodput therefore accompanies higher display FPS and fewer freezes.

C. Temporal Behavior Analysis

Fig. 6 connects network events, scheduling decisions, and display FPS in one representative trace. Short RTT spikes correspond to handover-related path changes, while the persistent-degradation interval lasts across multiple updates. Early Avoidance moves base-layer packets away from the affected path before a handover. During persistent degradation, enhancement-layer packet dropping and bounded duplicate protection reserve capacity for protected traffic.

Display FPS follows base-layer delivery and recovers as link conditions improve. Extreme impairment can still reduce



(a) Path RTT estimate

(b) Base-layer migration and enhancement-layer dropping

(c) Display FPS

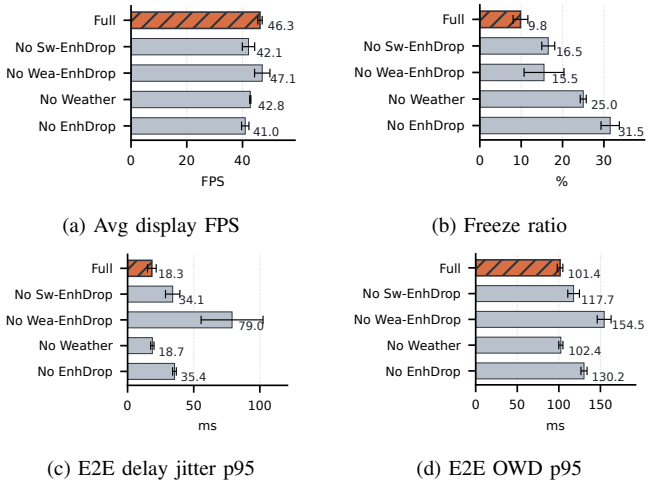
Fig. 6. Temporal behavior under a representative LEO multipath trace. The trace includes both handover-induced short-term degradation and a longer persistent degradation period.

FPS because LeoSVCSch protects selected packets and sends at most one duplicate; the remaining drops occur when the available path capacity and redundancy budget are insufficient.

D. Ablation Study

The ablation variants disable handover-period enhancement-layer dropping (**No Sw-EnhDrop**), persistent-degradation-period enhancement-layer dropping (**No Wea-EnhDrop**), Persistent Degradation State Detection and bounded duplicate protection (**No Weather**), or enhancement-layer dropping in both periods (**No EnhDrop**). **Full** denotes the complete design.

Fig. 7 shows that Full achieves 46.3 average display FPS, 9.8% freeze ratio, 18.3 ms E2E delay jitter p95, and 101.4 ms



(a) Avg display FPS

(b) Freeze ratio

(c) E2E delay jitter p95

(d) E2E OWD p95

Fig. 7. Ablation study of LeoSVCSch in terms of average display FPS, freeze ratio, latency stability, and one-way delay.

E2E OWD p95. No Sw-EnhDrop lowers FPS to 42.1 and raises freeze ratio to 16.5%. No Wea-EnhDrop slightly raises average FPS to 47.1 but also raises freeze ratio to 15.5% and degrades latency. No Weather lowers FPS to 42.8 and raises freeze ratio to 25.0%, while No EnhDrop performs worst at 41.0 FPS and 31.5% freeze ratio. These results show that path migration, bounded duplicate protection, and enhancement-layer control address different effects of LEO link impairment.

V. Conclusion

This paper presents LeoSVCSch, an SVC-aware multipath scheduler for long-distance real-time video over candidate LEO paths. LeoSVCSch treats two impairment patterns separately: predictable, short-term, and path-specific handover disruption, and persistent weather-induced attenuation that may affect multiple paths. The scheduler prioritizes base-layer packets by handover-aware base-layer path migration, enhancement-layer packet dropping when safe capacity is insufficient, and bounded duplicate protection during persistent degradation. We implement LeoSVCSch in a testbed that carries WebRTC video over multipath QUIC and replays LEO link traces in Mininet. Compared with representative transport-centric and redundancy-aware multipath schedulers, LeoSVCSch improves average display FPS by about 32% and reduces freeze ratio by about 76%, with lower QUIC goodput caused by enhancement-layer dropping and bounded duplicate protection. These results suggest that, for long-distance LEO real-time video, a multipath scheduler should protect base-layer packets separately from enhancement-layer packets and should treat handover windows differently from longer degradation periods.

References

[1] ITU-T, "G.114: One-way transmission time," https://www.itu.int/rec/T-REC-G.114-200305-I/en, 2003.

[2] Z. Lai, W. Liu, Q. Wu, H. Li, J. Xu, Y. Wang, Y. Li, and J. Liu, "SPACE-RTC: Unleashing the Low-Latency Potential of Mega-Constellations for Wide-Area Real-Time Communications," *IEEE Trans. Mob. Comput*, vol. 24, no. 2, pp. 642–657, Feb. 2025.

[3] Q. Zhang, Q. Wu, Z. Lai, J. Li, H. Li, Y. Liu, Y. Li, and J. Liu, "Spache: Accelerating Ubiquitous Web Browsing via Schedule-Driven Space Caching," in *Proceedings of the ACM Web Conference 2025*, ser. WWW '25. Sydney, NSW, Australia: Association for Computing Machinery, 2025.

[4] Z. Lai, H. Li, and J. Li, "StarPerf: Characterizing Network Performance for Emerging Mega-Constellations," in *Proceedings of the 2020 IEEE 28th International Conference on Network Protocols*, ser. ICNP. IEEE, 2020, pp. 1–11.

[5] H. Fang, H. Zhao, J. Shi, M. Zhang, G. Wu, Y. C. Chou, F. Wang, and J. Liu, "Robust Live Streaming over LEO Satellite Constellations: Measurement, Analysis, and Handover-Aware Adaptation," in *Proc. ACM MM*, ser. MM '24, New York, NY, USA, 2024, pp. 5958–5966.

[6] M. A. Ullah, A. Heikkinen, M. Uitto *et al.*, "Impact of Weather on Satellite Communication: Evaluating Starlink Resilience," in *Proc. IEEE VTC-Spring*, ser. VTC2025-Spring, 2025, pp. 1–7.

[7] H. Zhao, H. Fang, F. Wang, and J. Liu, "Realtime Multimedia Services over Starlink: A Reality Check," in *Proc. ACM NOSSDAV*, ser. NOSSDAV '23, New York, NY, USA, 2023, pp. 43–49.

[8] M. M. Kassem, A. Raman, D. Perino, and N. Sastry, "A Browser-side View of Starlink Connectivity," in *Proc. ACM IMC*, ser. IMC '22, New York, NY, USA, 2022, pp. 151–158.

[9] F. Michel, M. Trevisan, D. Giordano, and O. Bonaventure, "A First Look at Starlink Performance," in *Proc. ACM IMC*, ser. IMC, 2022, pp. 130–136.

[10] P. Du, X. Li, Y. Lu *et al.*, "Multipath TCP over LEO Satellite Networks," in *Proc. IEEE Int. Conf. Wireless Commun. Mobile Comput. (IWCMC)*, ser. IWCMC 2015. IEEE, 2015, pp. 1–6.

[11] X. Zhang, S. Sun, M. Tao, Q. Huang, and X. Tang, "Multi-Satellite Cooperative Networks: Joint Hybrid Beamforming and User Scheduling Design," *IEEE Trans. Wireless Commun.*, vol. 23, no. 7, pp. 7938–7952, 2024.

[12] Y.-s. Lim, E. M. Nahum, D. Towsley, and R. J. Gibbens, "ECF: An MPTCP Path Scheduler to Manage Heterogeneous Paths," in *Proc. ACM CoNEXT*, ser. CoNEXT '17, 2017, pp. 147–159.

[13] S. Ferlin, Ö. Alay, O. Mehani, and R. Boreli, "BLEST: Blocking Estimation-Based MPTCP Scheduler for Heterogeneous Networks," in *Proc. IFIP Networking*. IEEE, 2016, pp. 431–439.

[14] L. Jiang, Y. Zhang, Y. Hu *et al.*, "StarTCP: Handover-Aware Transport Protocol for Starlink," in *Proc. ACM APNet*, ser. APNet '24, 2024, pp. 169–170.

[15] T. Schierl, T. Stockhammer, and T. Wiegand, "Mobile Video Transmission Using Scalable Video Coding," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 17, no. 9, pp. 1204–1217, 2007.

[16] B. Fu, D. Staehle, G. Kunzmann *et al.*, "QoE-Based SVC Layer Dropping in LTE Networks Using Content-Aware Layer Priorities," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 12, no. 1, pp. 1–23, 2015.

[17] G. Bakar, R. A. Kirmizioglu, and A. M. Tekalp, "Motion-Based Rate Adaptation in WebRTC Videoconferencing Using Scalable Video Coding," *IEEE Trans. Multimedia*, vol. 21, no. 2, pp. 429–441, 2019.

[18] ——, "Motion-Based Adaptive Streaming in WebRTC Using Spatio-Temporal Scalable VP9 Video Coding," in *Proc. IEEE GLOBECOM*, 2017, pp. 1–6.

[19] Y. Wei, Q. Ye, K. Qu, W. Zhuang, and X. Shen, "E2E Performance Modeling for Slice-Based Video Streaming With Layered Encoding," *IEEE Trans. Netw*, vol. 33, no. 6, pp. 3363–3377, 2025.

[20] M. Hui, S. Zhai, D. Wang *et al.*, "A Review of LEO Satellite Communication Payloads for Integrated Communication, Navigation, and Remote Sensing: Opportunities, Challenges, Future Directions," *IEEE Internet of Things J.*, 2025.

[21] T. Wang, T. Liu, Y. Li *et al.*, "Packet Loss Modeling and Forward Erasure Correction for LEO Satellite Networks," *IEEE Trans. Commun.*, 2026.

[22] Z. Lai, Z. Li, Q. Wu, H. Li, J. Li, X. Xie, Y. Li, J. Liu, and J. Wu, "LeoCC: Making Internet Congestion Control Robust to LEO Satellite Dynamics," in *Proc. ACM SIGCOMM*, ser. SIGCOMM '25, New York, NY, USA, 2025, pp. 129–146.