

AutoRec: Accelerating Loss Recovery for Live Streaming in a Multi-Supplier Market

Tong Li¹, Member, IEEE, Xu Yan¹, Bo Wu¹, Cheng Luo, Fuyu Wang, Jiuxiang Zhu¹, Haoyi Fang, Xinle Du, and Ke Xu¹, Fellow, IEEE

Abstract—Due to the limited permissions for upgrading dual-side (i.e., server-side and client-side) loss tolerance schemes from the perspective of CDN vendors in a multi-supplier market, modern large-scale live streaming services are still using the automatic-repeat-request (ARQ) based paradigm for loss recovery, which only requires server-side modifications. In this paper, we first conduct a large-scale measurement study with up to 50 million live streams. We find that loss shows *dynamics* and live streaming contains frequent *on-off mode switching* in the wild. We further find that the recovery latency, enlarged by the ubiquitous retransmission loss, is a critical factor affecting live streaming’s client-side QoE (e.g., video freezing). We then propose an enhanced recovery mechanism called AutoRec, which can transform the disadvantages of on-off mode switching into an advantage for reducing loss recovery latency without any modifications on the client side. AutoRec allows users to customize overhead tolerance and recovery latency tolerance and adaptively adjusts strategies as the network environment changes to ensure that recovery latency meets user demands whenever possible while keeping overhead under control. We implement AutoRec upon QUIC and evaluate it via testbed and real-world commercial services deployments. The experimental results demonstrate the practicability and profitability of AutoRec.

Index Terms—QUIC, loss recovery, on-off traffic pattern, multi-supplier market.

Received 11 January 2025; revised 3 July 2025 and 30 September 2025; accepted 23 November 2025; approved by IEEE TRANSACTIONS ON NETWORKING Editor A. Balasubramanian. This work was supported in part by the National Natural Science Foundation of China under Grant 62572473, Grant 62202473, Grant 62441230, and Grant 62272466; in part by the Science Fund for Creative Research Groups of the National Natural Science Foundation of China under Grant 62221003; in part by the Key Program of the National Natural Science Foundation of China under Grant 61932016 and Grant 62132011; in part by the National Science Foundation for Distinguished Young Scholars of China under Grant 62425201; and in part by the Tencent Basic Platform Technology Rhino-Bird Focused Research Program. (Corresponding author: Bo Wu.)

Tong Li, Xu Yan, Jiuxiang Zhu, and Haoyi Fang are with the Key Laboratory of Data Engineering and Knowledge Engineering, School of Information, Renmin University of China, Beijing 100872, China (e-mail: tong.li@ruc.edu.cn; yanxu@ruc.edu.cn; jiuxiangzhu@ruc.edu.cn; 2021201680@ruc.edu.cn).

Bo Wu, Cheng Luo, and Fuyu Wang are with the Department of Cloud Architecture and Platform, Tencent Technologies, Shenzhen 518057, China (e-mail: wub14@tsinghua.org.cn; lancelotluo@tencent.com; ivanfywang@tencent.com).

Xinle Du is with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China (e-mail: xinledu18.thucsnet@gmail.com).

Ke Xu is with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China, and also with Zhongguancun Laboratory, Beijing 100094, China (e-mail: xuke@tsinghua.edu.cn).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TON.2025.3642008>, provided by the authors.

Digital Object Identifier 10.1109/TON.2025.3642008

I. INTRODUCTION

INTERNET live services such as Youtube Live, TikTok Live, and Twitch have gradually become a fundamental element for enriching daily life and work [1], [2]. This underscores the critical need to enhance live streaming transmission performance. The ubiquitous packet loss is an essential factor affecting client-side quality-of-experience (QoE) [3], [4], which will introduce head-of-line (HOL) blocking and even incur long-time video freezing if the available frames in the player buffer are all consumed. Therefore, loss tolerance control matters in live-streaming services.

The existing loss tolerance schemes mainly focus on designing dual-side (i.e., service-side CDN and client-side application) control policies, including Forward Error Correction (FEC) [5], [6], [7], [8], multi-path retransmissions [9], [10], semi-reliable transmissions [11], [12], and application-level controls [13], [14], [15]. However, as shown in Fig. 1, live-streaming application operators (e.g., TikTok Live) usually apply the Multi-Supplier Strategy [16] in the CDN market (see § II-A). As a result, it is the CDN vendor’s duty that optimize the transmission performance (e.g., loss tolerance), according to which the application operators will choose the better-performed CDN vendors to carry more traffic (i.e., larger market share). In this case, only server-side sending policies can be adjusted by the selected CDN vendors, which lack the proper authority to synchronize client-side control rules. Thus, the aforementioned loss tolerance schemes face significant deployment challenges in a multi-supplier CDN market.

In this case, most modern CDN vendors only apply the automatic-repeat-request (ARQ) paradigm [17], [18] to control loss tolerance as the commercial solution, which retransmits only one replica of the lost packet when a loss is detected. However, we find that the legacy ARQ-based loss recovery is far from satisfactory according to our performed large-scale measurements. For example, the proportion of connections with maximum retransmission times of two or more exceeds 43%. Among them, a considerable portion of the connections has certain packets that are retransmitted even more than 10 times (§ III-A). The retransmission loss enlarges the loss recovery latency by 123.2 ms on average and 279.3 ms in the worst case. This enlarged recovery latency further increases the probability of empty buffer space on the client side, thereby increasing the risk of video freezing (§ III-C). Thus from a philosophical standpoint, it is worth asking: Can we accelerate loss recovery solely through CDN servers without modifying clients? How can it be addressed with controllable overhead?

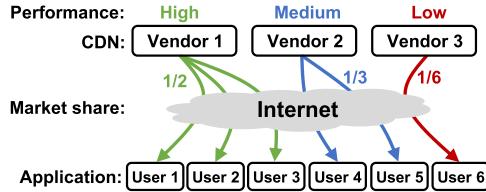


Fig. 1. An example of a multi-supplier market for CDN vendors in large-scale live streaming services.

Our key insight is that the on-off mode switching ubiquitously occurs in current live streams (§ III-B), where the bandwidth during the “off” periods is not fully utilized. Our measurements show that each live stream spends 464 ms in off-mode per second and enters off-mode 20 times per second on average. It is well-studied that the on-off traffic pattern is not conducive to transmission control [19], [20], [21], [22], [23], [24]. However, we argue that it can act as an advantage for the loss tolerance control of live streaming. In this paper, we present AutoRec. This enhanced loss recovery mechanism can transform the disadvantages of frequent on-off mode switching into an advantage of loss tolerance control in live streaming.

A naive implementation of AutoRec is to directly reinject a fixed number of replicas of loss packets upon entering the off-mode. However, this approach faces two challenges. **First**, the fixed settings of redundant replicas cannot adapt well to the dynamics of packet loss, while excessive reinjection of packets results in non-trivial recovery overhead. **Second**, according to our testbed experiments, the off-mode can be unevenly distributed throughout an entire live session, and the sending of replicas might be delayed without timely off-mode entry. To tackle these issues, we propose a two-step solution: Redundancy Adaption and Reinjection Control.

Redundancy Adaption intelligently determines how many replicas of the lost packets should be reinjected. It sets the number of replicas based on the user-customizable overhead tolerance, recovery latency tolerance, and the network environment. This assures loss recovery is accelerated to meet user demands whenever possible while keeping the overhead controllable.

Reinjection Control determines when to retransmit the replicas. Generally, it enables the replica reinjection during the off-modes. AutoRec further adopts the *opportunistic reinjection* to trigger loss reinjection even lacking the desired opportunity of off-mode. This assures that each replica can be reinjected in time even under uneven distribution of off-modes.

We implement the AutoRec prototype in the user-space QUIC protocol and deploy it on both testbed and real-network CDN proxy for 6 months. The experimental results demonstrate the practicability and profitability of AutoRec, in which the average times and duration of client-side video freezing can be lowered by 10.10% and 4.74%, respectively.

The rest of the paper is organized as follows: § background introduces the multi-supplier CDN market and its impact on CDN vendors’ accelerating loss recovery. § III motivates our work with a large-scale measurement study. Then, the high-level architecture and design details of AutoRec are depicted in § IV and § V, respectively. § VI describes the detailed

implementation of AutoRec. § VII gives the experimental evaluations of AutoRec. § VIII presents a discussion of AutoRec’s overhead and its interaction with QUIC’s native mechanisms. § IX overviews the related work and § X concludes the paper.

II. BACKGROUND

In this section, we first introduce the Multi-Supplier Strategy in the modern CDN market (§ II-A). We then discuss the appeal of accelerating loss recovery (§ II-B), along with the restriction on the loss recovery mechanism (§ II-C) from the perspective of CDN vendors in the multi-supplier CDN market.

A. Multi-Supplier CDN Market

Live-streaming application operators typically subscribe to multiple CDN vendors simultaneously and periodically evaluate the performance of these vendors’ services based on weighted QoE metrics, then reallocate subscription shares accordingly. For example, as shown in Fig. 1, a live-streaming application operator subscribes to the services of three CDN vendors: Vendor 1, Vendor 2, and Vendor 3. In the previous performance evaluation period conducted by the live-streaming application operator, Vendor 1 had the highest performance, Vendor 2 had medium performance, and Vendor 3 had the lowest performance. Therefore, the live-streaming application operator will allocate 1/2 of all service requests to Vendor 1, 1/3 to Vendor 2, and 1/6 to Vendor 3. For live-streaming application operators, the Multi-Supplier Strategy not only allows for a comparative analysis of different CDN vendors’ service performance, facilitating the selection of higher-performing CDN vendors but also continually drives CDN vendors to enhance their service performance, leading to better CDN services for the operators themselves. Thus, live-streaming application operators are motivated to adopt the Multi-Supplier Strategy, leading to the formation of a multi-supplier CDN market where multiple CDN vendors serve the same video live-streaming application operator.

B. Appeal of Accelerating Loss Recovery

Loss recovery influences video freezing. Live streaming is highly sensitive to data timeliness, as traffic that fails to reach receivers on time results in video freezing, thereby affecting client-side QoE. Video freezing occurs when the player buffer becomes empty, that is, when no video frame can be uploaded before all cached frames have been consumed. One important reason for the player buffer becoming empty is that the lost data cannot be recovered in time, causing longer intra-stream HOL blocking time, especially in the common designs of TCP and QUIC protocol. This paper focuses on enhancing the timeliness of loss recovery to reduce video freezing.

Video freezing influences profit. In the multi-supplier CDN market, for CDN vendors, improving their QoE metrics can lead to a larger share of subscriptions, resulting in higher profits. The video freezing metrics (i.e., video freezing times per 100 seconds and video freezing duration per 100 seconds) are important QoE indicators for evaluating the service performance of CDN providers. Thus, the CDN vendors are motivated to optimize video freezing metrics.

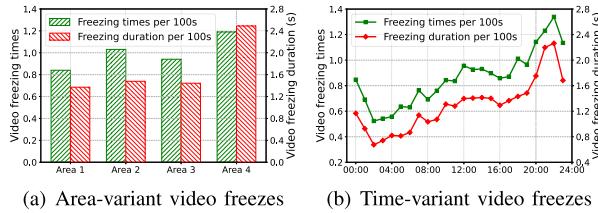


Fig. 2. The times and duration of L7-player freezes.

The current video freezing is suboptimal. We conduct large-scale measurements and collect one-year values of video freezing metrics from a famous live platform in four areas (i.e., Southeast Asia, Latin America, the Middle East, and a selected country). Fig. 2(a) presents the spatial variations of video freezing metrics across different regions, while Fig. 2(b) depicts their temporal variations in Area 1. As shown in Fig. 2(a), over each 100 s interval,¹ clients experience on average more than 1.3 s of video freezing (up to \sim 2.5 s in Area 4), where even a single freeze event can significantly degrade user experience. Furthermore, Fig. 2(b) indicates that video freezing events exceeding 2 s or occurring more than twice are frequently observed after 20:00.

C. Restriction on the Loss Recovery Mechanism

In the multi-supplier CDN market, CDN vendors face difficulties in deploying loss recovery mechanisms that require modifications on both ends, which poses challenges for the universal deployment of mechanisms like FEC in production networks. The live-streaming client software is designed and developed by the live-streaming application operators themselves, which includes proprietary information, so its source code is prohibited from being accessed or modified by CDN vendors. Therefore, loss recovery mechanisms that require modifications on both ends need the cooperation of live-streaming application operators. In a multi-supplier market, developing adaptations for a specific CDN vendor may compromise the generality of their software protocols, pose interaction risks with other vendors, and increase operational costs. Consequently, live-streaming application operators have no incentive to develop adaptations for the mechanism upgrades of a specific CDN vendor.

III. MEASUREMENT STUDY

In this section, we conduct large-scale measurements² to motivate our work, in which the characteristics of loss (§ III-A) and live streaming (§ III-B) in the wild are further explored. We then analyze the performance of live streaming in the presence of packet loss, particularly retransmission loss (§ III-C).

A. Characteristics of Loss

We analyzed the characteristics of loss using real-world logs spanning up to two weeks. These logs documented specific details related to packet loss, including the frequency of loss events and the magnitude of each loss. Our measurements cover over 200,000 connections across diverse application scenarios and geographic regions.

¹The average QUIC connection lifetime of each live stream is 99.6 s.

²The large-scale measurements are performed on the server side that runs NGINX architecture with QUIC BBR.

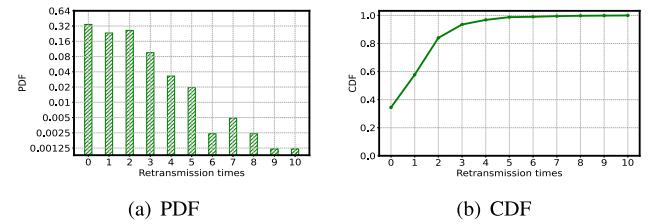


Fig. 3. The maximum retransmission times in the wild.

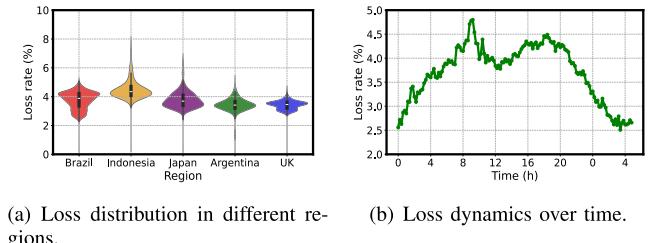


Fig. 4. Examples of loss dynamics in the wild.

Observation #1: Retransmission loss is ubiquitous.

A packet may undergo multiple retransmissions before the receiver correctly receives it. The number of retransmissions per packet before successful reception can be further analyzed. For any given connection, the maximum retransmission times can be calculated as the highest count of retransmissions across all packets within that connection. Fig. 3 illustrates the distribution, i.e., PDF (Probability Density Function) and CDF (Cumulative Distribution Function) of the maximum retransmission times in the production network. The results show that the proportion of connections with maximum retransmission times of two or more exceeds 43%. Among them, a considerable portion of the connections have certain packets that are retransmitted even more than 10 times. Traditional ARQ mechanism focuses solely on promptly retransmitting lost packets after each loss event, but it overlooks the total number of retransmissions and the total time required for the retransmitted data to be successfully received by the receiver. Therefore, the traditional ARQ mechanism hardly meets the timeliness requirements for data in certain scenarios (e.g., live streaming).

Observation #2: Loss shows dynamics. We then investigate the distribution of the loss rate (every 5 minutes) for each connection across various global regions. The results are depicted in Fig. 4(a). While some regions exhibit similar average loss rates, their packet loss rate deviations differ, as indicated by the shapes of the violin plots. For instance, both Brazil and Japan have an average loss rate of 3.78%, but Japan has the highest packet loss rate of 7.1% while Argentina has 5.7%. Fig. 4(b) further illustrates how the packet loss rate evolves. Specifically, the loss rate is dynamic, varying between 2% and 5% over 24 hours. This confirms that packet loss exhibits dynamic behavior in real-world scenarios. This also reveals that an adaptive loss tolerance scheme should adapt to differentiated conditions and the dynamic nature of networks.

B. Characteristics of Live Streaming

Unlike video-on-demand or file traffic, the current live streams frequently and extensively reveal off-mode, in which

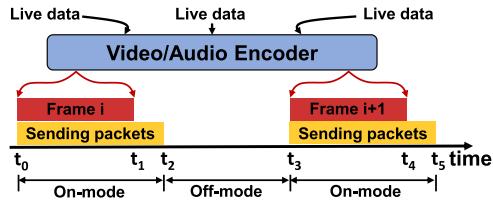


Fig. 5. The sketch of on-off mode in live streams.

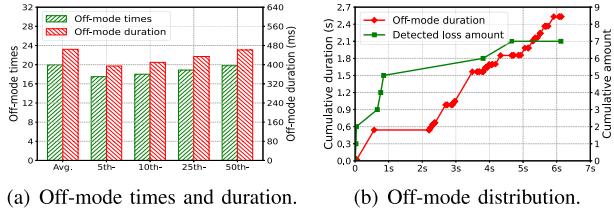


Fig. 6. On-off mode measurements and experiments.

a sender³ temporarily has no data (i.e., becomes application-limited [21], [25]) for continuous transmissions after sending one or more frames. On the one hand, the new generation rate of live data might be slower (e.g., than the sending rate), requiring senders to wait until their sending queues are replenished. For CDN vendors, the data generation rate reflects the traffic transmission rate from anchors to servers, which can be easily affected by real-time network status. On the other hand, the live data will be encoded into video or audio frames based on selected frame rate(s) and bitrate before its transmission, which can also introduce time intervals between adjacent frames. If one frame has been sent out while the follow-up frame has not yet been encoded, the CDN sender enters off-mode.

As Fig. 5 shows, the data of live streaming is encoded (by video/audio encoder) into frame i which is delivered to the sending queue (from t_0 to t_1) and is sent out before t_2 . In this case, the sender must wait (from t_2 to t_3) for frame $i+1$ which will be encoded based on the follow-up live data. In this paper, the on-off mode that appears on the sender side can be recognized with the following conditions.

- **On-mode:** A mode that occurs when enough data exists in the sending queue (i.e., $t_0 \sim t_2$ and $t_3 \sim t_5$ in Fig. 5), which can be sent by a sender at the subsequent time.
- **Off-mode:** A mode that occurs when no data can be obtained for traffic transmissions (i.e., $t_2 \sim t_3$ in Fig. 5), making senders have to wait for the follow-up frame.

We make large-scale measurements and gain the following observations to further explore the on-off mode of live streaming.

Observation #1: The on-off mode switching commonly exists in live streaming. Fig. 6(a) presents our measurements of the duration each stream spends in off-mode every second, as well as the frequency of each stream entering off-mode on a per-second basis, indicating that more than 95% of the live streams can be in off-mode for 394 ms per second and enter off-mode 17 times per second. On average, each live stream

³Since we collect and analyze data as well as conduct experimental evaluations all from the CDN edge server, without special instructions, the term “sender” refers to the CDN edge server instead of the original streaming source.

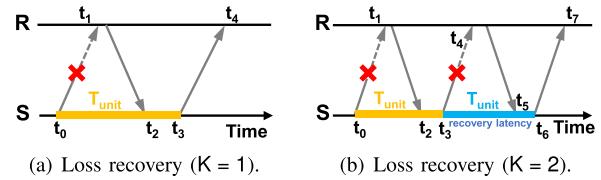


Fig. 7. Loss recovery between sender (S) and receiver (R).

spends 464 ms in off-mode per second and enters off-mode 20 times per second.

Observation #2: The off-modes can be unevenly distributed throughout the live-streaming lifetime. Fig. 6(b) depicts the cumulative values of both off-mode duration and detected loss amount in a connection of our testbed experiments. We can find (i) off-mode mainly occurs after 2 s while only 2 out of 7 packets are detected lost during this period; (ii) most packet losses (5 of 7) occur with minimal off-mode exposure within the first second of the measured stream, during which the stream remains in on-mode when the 4th to 6th losses are detected.

C. Live Streaming Performance Under Loss

Live streaming requires high data timeliness which can affect client-side QoE. Due to the characteristics of packet loss and the limitations of traditional ARQ (§ III-A), the timeliness issues caused by packet loss in live streaming are critical. We introduce *recovery latency* to measure the timeliness of packet loss recovery and conduct large-scale measurements on live streams in real production networks.

Recovery latency is defined as the duration from when any data is detected lost to when resending a recovery packet *that will be successfully received*. *Loss detection time*, denoted as T_{unit} , is defined as the time elapsed for a data packet from being sent to being retransmitted. Recovery latency consists of zero or more T_{unit} , reflecting the additional time for loss recovery besides the first T_{unit} . In Fig. 7(a), recovery latency is 0. Fig. 7(b) shows the resent data is detected lost again and another recovery packet (*that will be successfully received* at t_7) is sent at t_6 , where recovery latency = $t_6 - t_3$. Besides, the *maximum recovery latency* is employed to evaluate the largest recovery latency for loss recoveries.

We make large-scale measurements and collect the transmission logs of 50 million live streams. We then classify the average values of measured recovery latency and maximum recovery latency based on the ranges of loss rate and Smooth Round-Trip Time (SRTT), as Table I shows.

Observation #1: The current recovery latency of live streaming is far from satisfactory. For the lost data that requires more than one retransmission, traffic senders waste 123.2 ms (i.e., recovery latency = 123.2 ms), on average, before sending out the recovery packets that their receivers will indeed acknowledge. Furthermore, the average maximum recovery latency is up to 279.3 ms, which can easily cause video freezing.

Observation #2: Higher loss rate and larger SRTT can easily introduce more deteriorated recovery latency. As Table I shows, with the increase in loss rate, the recovery latency increases accordingly. This is because the recovery data may experience another packet loss in the networks with

TABLE I
LIVE-STREAMING RECOVERY LATENCY MEASUREMENT RESULTS

metrics	value	loss rate				SRTT(ms)			
		0% ~ 3%	3% ~ 10%	10% ~ 30%	30% ~ 50%	0 ~ 50	50 ~ 200	200 ~ 500	500 ~ 2000
recovery latency (ms) [*]	123.2	26.8	94.5	186.4	530.2	68.2	170.2	273.6	418.0
maximum recovery latency (ms)	279.3	51.8	204.9	451.3	1288.4	149.2	391.3	698.8	960.9

^{*} The displayed recovery latency (maximum recovery latency) only records the each-stream's average (maximum) recovery latency of the lost data, whose recovery requires two or more retransmissions.

a higher loss rate. By contrast, SRTT has serious impacts on recovery latency. For example, maximum recovery latency will be improved to ~ 1 s, on average, under the SRTT of $500 \sim 2000$ ms, which is unacceptable for the timeliness of loss recovery.

Therefore, an enhanced recovery scheme is highly required to further optimize the current unsatisfied recovery latency and promote the recovery timeliness of live streams.

IV. THE AUTOREC OVERVIEW

In this section, we first discuss the design principles of AutoRec, whose architecture overview will then be given.

A. Design Principles

AutoRec regards the control-unfriendly on-off mode as an essential opportunity for enhancing the recovery latency of live streaming and mitigating the negative effects on player freezing caused by potential HOL blocking. Our goal is to accelerate loss recovery to meet user demands whenever possible while keeping the overhead controllable. AutoRec should follow two design principles for optimizing the timeliness of loss recovery in live streams.

Principle #1: The overhead introduced by AutoRec must be controllable. If the overhead is not carefully managed, AutoRec could result in negative returns. From the perspective of a CDN vendor, AutoRec primarily introduces two types of overhead. The first type is the traffic cost incurred by replicas under traffic billing conditions, which increases costs. The second type is the potential goodput reduction caused by redundant replicas, which can lead to a decrease in QoE and thus reduce revenue. Therefore, AutoRec should have clear and simple control logic to manage the overhead introduced by injecting replicas.

Principle #2: AutoRec should ensure that recovery latency is reduced to meet user demands whenever possible. Different application scenarios have different requirements for recovery latency. AutoRec should allow for the setting of recovery latency requirement and, while strictly adhering to overhead constraints, inject few but enough packets to reduce the recovery latency to meet the recovery latency requirement whenever possible.

B. The Architecture of AutoRec

AutoRec is a sender-side modification to the protocol stack whose key modules are illustrated in Fig. 8. Particularly, AutoRec adopts redundancy adaption to compute the number of replicas of a lost packet that should be retransmitted next periodically. Given the number of replicas, AutoRec then adopts reinjection control to determine the specific order and time for sending out each replica from the sender.

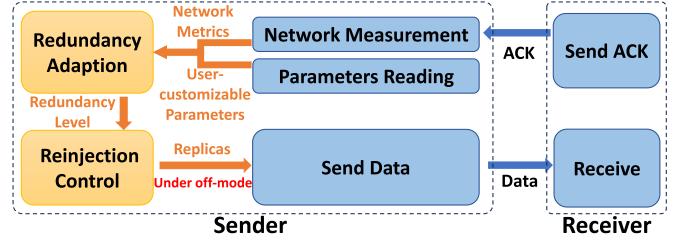


Fig. 8. The architecture of AutoRec.

Redundancy adaption. This module answers the question of how many replicas should be sent to accelerate loss recovery. The *redundancy level* (denoted by K , $K \in \mathbb{N}$) is defined as the number of replicas that should be resent for a specific lost packet. To adapt to the dynamics of loss, we incorporate redundancy adaptation to allow the redundancy level to vary dynamically. To achieve this, we introduce the *Redundancy Adapter*, which carefully selects the most appropriate redundancy level for the lost packets.

Reinjection control. This module answers the question of how to send the given number of replicas determined by the Redundancy Adapter. For each lost packet, more than one replica might be injected into the network. To avoid bandwidth contention for unlost data transmission, we introduce the *Reinjection Controller* to enable sending replicas of the lost packet if the stream is in off-mode. To further reduce the loss recovery latency, the Reinjection Controller also opportunistically captures the ideal chance to reinject replicas even when the stream is not in the off-mode. This ensures that AutoRec can accelerate packet loss recovery even when off-mode is absent or unevenly distributed.

V. DETAILED DESIGN

In this section, we will depict the detailed designs of AutoRec for optimizing loss recovery.

A. Redundancy Adapter

We first define *redundancy cost* as the ratio of the total number of injected replicas to the total number of sent packets (excluding retransmitted packets and injected replicas). *Goodput reduction* is defined as the reduction percentage in goodput when injected replicas compared to the goodput without injected replicas during transmission. The Redundancy Adapter provides three user-customizable parameters: the recovery latency tolerance (denoted by α), the redundancy cost tolerance (denoted by β), and the goodput reduction tolerance (denoted by γ).

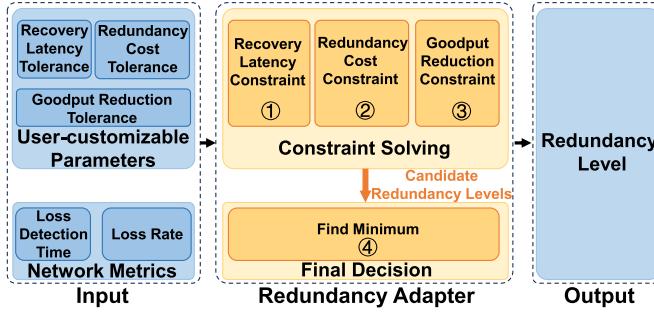


Fig. 9. The redundancy adapter of AutoRec.

The redundancy level is determined periodically.⁴ AutoRec will measure the network metrics within each decision interval. The network metrics include the average loss detection time (§ III-C) and the average loss rate. When there is no packet loss (i.e., loss rate = 0) in the last decision interval, the Redundancy Adapter directly sets the redundancy level to 0. Otherwise (i.e., $0 < \text{loss rate} < 1$), taking the user-customizable parameters and network metrics as input, the Redundancy Adapter will execute the following four steps to calculate the output, which is the redundancy level, as Fig. 9 shows. **Step #1:** Calculate the minimum redundancy level satisfying the recovery latency constraint (denoted by K_α). **Step #2:** Calculate the maximum redundancy level satisfying the redundancy cost constraint (denoted by K_β). **Step #3:** Calculate the maximum redundancy level satisfying the goodput reduction constraint (denoted by K_γ). **Step #4:** Determine the final redundancy level (denoted by K_θ) according to the above K_α , K_β , K_γ . Next, we will elaborate step by step.

Step #1: Calculate the minimum redundancy level satisfying the recovery latency constraint. Let K denote the redundancy level, $F(K)$ denote the average recovery latency for lost packets⁵ under network conditions with the maximum average recovery latency (i.e., the entire transmission process remains in on-mode), and α denote the recovery latency tolerance. The formal description of the problem to calculate the minimum redundancy level that satisfies the recovery latency constraint (denoted by K_α) is as follows:

$$K_\alpha = \arg \min_K K \quad \text{s.t. } F(K) \leq \alpha \text{ and } K \in \mathbb{N}$$

Given the loss rate (denoted by R), the loss detection time T_{unit} and the redundancy level K , when $0 < R < 1$, $T_{unit} \geq 0$ and $K \in \mathbb{N}$, $F(K)$ can be computed as follows (please refer to Appendix B for the detailed derivation):

$$F(K) = \frac{(1 + KR^K)}{(1 - R)(1 + K)} T_{unit} \quad \text{where } K \in \mathbb{N}, 0 < R < 1 \text{ and } T_{unit} \geq 0 \quad (1)$$

We further find that $F(K)$ is *monotonically decreasing* for $K \in \mathbb{N}$ (please refer to Appendix C for the detailed proof). In this case, we can determine K_α using the binary search method.

Step #2: Calculate the maximum redundancy level satisfying the redundancy cost constraint. Let K denote

⁴The decision interval (denoted by D) is usually set several Round-Trip Times (RTTs), e.g., $D = 5$ RTTs.

⁵In this paper, we only focus on the lost packets that are retransmitted over once.

the redundancy level, $G(K)$ denote the average redundancy cost, and β denote the redundancy cost tolerance. The formal description of the problem to calculate the maximum redundancy level that satisfies the redundancy cost constraint (denoted by K_β) is as follows:

$$K_\beta = \arg \max_K K \quad \text{s.t. } G(K) \leq \beta \text{ and } K \in \mathbb{N}$$

Given the loss rate (denoted by R) and the redundancy level K , when $0 < R < 1$ and $K \in \mathbb{N}$, $G(K)$ can be computed as follows (please refer to Appendix D for the detailed derivation):

$$G(K) = KR \quad \text{where } K \in \mathbb{N} \text{ and } 0 < R < 1 \quad (2)$$

We further find that $G(K)$ is *monotonically increasing* for $K \in \mathbb{N}$ (please refer to Appendix E for the detailed proof). In this case, we can determine K_β using the binary search method.

Step #3: Calculate the maximum redundancy level satisfying the goodput reduction constraint. Let K denote the redundancy level, $H(K)$ denote the average goodput reduction under network conditions with the maximum average goodput reduction (i.e., the entire transmission process remains in on-mode), and γ denote the goodput reduction tolerance. The formal description of the problem to calculate the maximum redundancy level that satisfies the goodput reduction constraint is (denoted by K_γ) as follows:

$$K_\gamma = \arg \max_K K \quad \text{s.t. } H(K) \leq \gamma \text{ and } K \in \mathbb{N}$$

Given the loss rate (denoted by R) and the redundancy level K , when $0 < R < 1$ and $K \in \mathbb{N}$, $H(K)$ can be computed as follows (please refer to Appendix F for the detailed derivation):

$$H(K) = 1 - \frac{1}{1 + KR - (1 + K)R^2 + R^{K+2}} \quad \text{where } K \in \mathbb{N} \text{ and } 0 < R < 1 \quad (3)$$

We further find that $H(K)$ is *monotonically increasing* for $K \in \mathbb{N}$ (please refer to Appendix G for the detailed proof). In this case, we can determine K_γ using the binary search method.

Step #4: Determine the redundancy level for this decision interval. Since AutoRec aims to satisfy the constraint of recovery latency whenever possible while strictly satisfying the constraints of the redundancy cost and the goodput reduction, the Redundancy Adapter determines the K_θ as follows.

$$K_\theta = \min\{K_\alpha, K_\beta, K_\gamma\} \quad (4)$$

B. Reinjection Controller

We use D_{ul} to denote the lost data that has been resent but unacknowledged by its receiver. The data D_{ul} is added and stored in a reinjection queue that is managed at the sender. To decide when to retransmit the replicas, the Reinjection Controller is introduced, which reinjects replicas when the current stream enters off-mode. The Reinjection Controller also adopts opportunistic reinjection to further reduce loss recovery

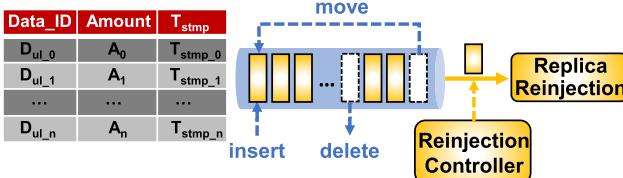


Fig. 10. Sender-side reinjection queue.

latency while keeping the bandwidth contention for unlost data transmission within acceptable bounds.

Queue management. In AutoRec, a reinjection queue will be created by a traffic sender when a new connection of live streaming is established, which will be removed once this connection is closed. AutoRec enables senders to update the reinjection queue (using the following operations) for each live stream whenever lost data is retransmitted or acknowledged, as Fig. 10 shows. The detected lost data will be inserted into the end of the reinjection queue after it has been retransmitted. The D_{ul} will be deleted from the reinjection queue when (i) the resent D_{ul} is acknowledged by its receiver, or (ii) the performed reinjection times exceeds the decided K_θ . The D_{ul} will be moved from the head to the end of the reinjection queue if it has been resent again, which is sorted by the timestamp of its last reinjection or retransmission. In addition, the AutoRec sender maintains a status table for each reinjection queue, which records the reinjection times (A_i) that has been performed, timestamp (T_{stmp}) of D_{ul} 's last reinjection or retransmission and D_{ul} identification ($Data_ID$),⁶ as Fig. 10 shows.

Opportunistic reinjection. To address the issue of unevenly distributed on-off mode (§ III-B) and conduct replica reinjection more timely, AutoRec enables traffic senders to inject packets even during on-mode. As illustrated in Fig. 10, in the status table of the reinjection queue, T_{stmp} can be used to compute D_{ul} 's “silence” duration since its last retransmission or reinjection. Once it exceeds the threshold T_{thres} , D_{ul} will be fetched from the reinjection queue and then be resent regardless of whether the off-mode is entered or not. To achieve well-distributed loss reinjections, at each decision interval, the Reinjection Controller will use the average loss detection time (denoted by T_{unit}) measured by AutoRec and the K_θ calculated by the Redundancy Adapter (§ V-A) as input, and calculate T_{thres} through the following formula:

$$T_{thres} = \frac{T_{unit}}{K_\theta + 1} \quad (5)$$

Fig. 11 illustrates an example of how the Reinjection Controller performs lost injections with several critical events. Assume the redundancy level determined by the Redundancy Adapter is 2. In this example, a lost packet is detected and resent at t_1 . At time t_2 , the time elapsed since the last transmission of this packet has exceeded threshold T_{thres} . Therefore, the Reinjection Controller will reinject a replica of this packet even if the sender is currently not in off-mode. From t_3 to t_5 , the AutoRec sender enters off-mode, which enables another loss replica to be reinjected at t_3 . After this point, no further replicas of this packet are injected, because for a single packet,

⁶In this paper, $Data_ID$ can be recognized as the packet number (pkt_num) in TCP or the stream offset ($stream_offset$) in QUIC.

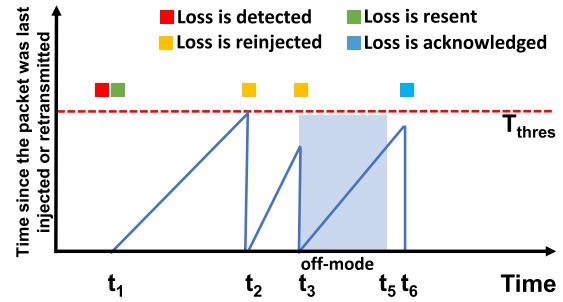


Fig. 11. An example of how the Reinjection Controller handles lost injections.

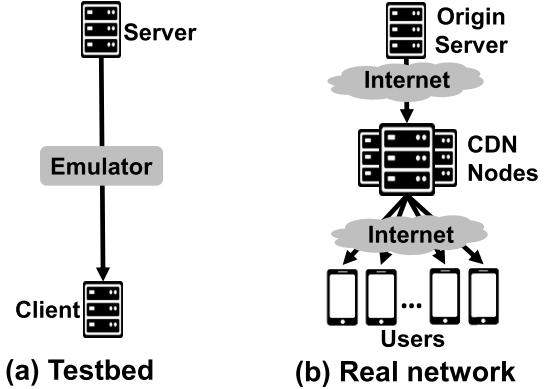


Fig. 12. Testbed and real-world deployments.

we cannot inject more replicas than the redundancy level determined by the Redundancy Adapter. At t_6 , one of the reinjected replicas is acknowledged, which means that the loss recovery process of the packet is completed. Note that only the last injected replica is checked for loss. Even if all replicas are lost, the existing ARQ mechanism will retransmit the last injected replica when it detects the loss of the last injected replica. Thus, the packet will not fail to be retransmitted by ARQ, nor will it be retransmitted excessively.

VI. IMPLEMENTATION

We integrate all the components described in Section V into only sender-side upgrades and implement AutoRec prototype based on the user-space QUIC protocol (with the version of LSQUIC Q043) [26] and NGINX architecture (with the version of ias-nginx 1.17.3) [27]. Our implementation consists of 900+ lines of code without any client-side modification. Please refer to Appendix H for details on the AutoRec implementation.

VII. EXPERIMENTAL EVALUATION

We perform the experimental evaluations on our established testbed and the real networks, respectively. In the testbed experiments, servers deploying AutoRec generate live-streaming data and transmit it to clients through a network environment simulated by the simulator, as Fig. 12(a) shows. In real network experiments, CDN nodes deploying AutoRec deliver live-streaming data obtained from the origin servers to real user clients, as Fig. 12(b) shows.

The testbed server and client are running on CentOS Linux release 7.9 with Intel(R) Xeon(R) CPU E5-2670 v3 @ 2.30GHz (E5-2620 v3 @ 2.40GHz), 48 (24) processors, 62 GB memory, and 1000 Mbps NIC. The employed network emulator is HoloWAN ultimate 2600u which supports 0~1000 Mbps

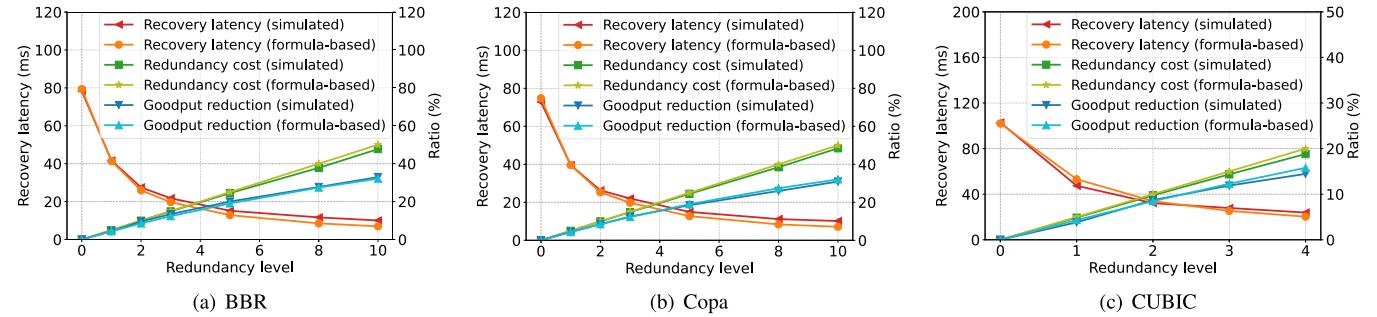


Fig. 13. Formula accuracy validation.

bandwidth, 0~10 s delay, 0~100% loss rate and 0~1000 GB buffer length. In this section, unless otherwise declared, the BBR (with version 1) [28] scheme is leveraged for congestion controls. Additionally, to enhance the robustness of AutoRec, we limit the value range of the redundancy level to [0, 10]. The baseline scheme is the typical ARQ paradigm that will retransmit one recovery packet once a packet is detected lost. In this section, each stream's recovery latency is displayed only for the lost data, whose recovery requires at least two retransmissions.

To better evaluate AutoRec, we define a new metric called *recovery deterioration rate*, which refers to the ratio between the amount of lost data (D_k) that takes two or more T_{unit} (i.e., recovery latency $\geq T_{unit}$) to be recovered, to the amount of all lost data. For example, there are 2 packets (one is in Fig. 7(a) and the other is in Fig. 7(b)) detected lost, which require one T_{unit} and two T_{unit} for their successful recoveries (at t_4 and t_7), respectively. Then, we can learn $D_k = 1$ and recovery deterioration rate = 50%.

A. Testbed Evaluation

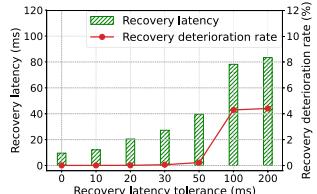
The testbed experiments are performed to validate the accuracy of the formulas used by AutoRec, the parameter sensitivity of AutoRec, the effectiveness of AutoRec's opportunistic reinjection, the advantages compared to the ART [29], the performance of AutoRec in Real-Time Communication (RTC) scenarios, the robustness of AutoRec across diverse packet loss models and loss rate magnitudes, and the performance comparison of AutoRec against multiple FEC strategies [7], [8], [30], [31]. The basic environment is configured as follows: 5% loss rate, 60 ms Round-Trip Time (RTT), 12 Mbps bandwidth. Unless otherwise declared, we use recovery latency tolerance $\alpha = 30$ ms, redundancy cost tolerance $\beta = 30\%$, and goodput reduction tolerance $\gamma = 20\%$, without loss of generality. The QUIC's loss recovery mechanism employs Tail Loss Probe (TLP) [32], Retransmission Timeout (RTO) [33], and Forward Acknowledgement (FACK) [34] for packet loss detection. Upon detecting lost packets, it immediately initiates retransmissions. Crucially, QUIC imposes no maximum retransmission limit per packet to fulfill its reliable data delivery guarantee. Each obtained metric is the average value of over 100 sets of experiments, in which each live stream will last for 60 seconds.

Formulas accuracy validation. In this experiment, we validate the accuracy of three formulas used by AutoRec for calculating the benefits and overhead varying with the

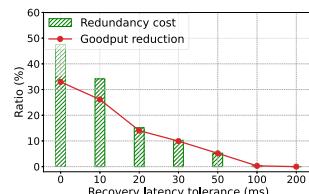
redundancy level (i.e., Formula 1, Formula 2, and Formula 3 in § V-A). We use a live stream with a bitrate of 12 Mbps to ensure that the entire transmission process of the stream remains in on-mode, thereby maximizing the negative impact of redundancy injection on goodput. Additionally, we slightly modify the Redundancy Adapter so that the redundancy level remains a fixed preset value throughout the transmission process, allowing us to measure the benefits and overhead of AutoRec at the specified redundancy level. We use the formulas employed by AutoRec, which describe the benefits and overhead varying with the redundancy level, to calculate the benefits and overhead of AutoRec at each fixed redundancy level. Then, we use the testbed to simulate and measure the benefits and overhead of AutoRec at each fixed redundancy level. By comparing the values calculated using the formulas with those measured through simulation, we can validate the accuracy of the formulas used by AutoRec for calculating the benefits and overhead varying with the redundancy level. Due to the significant impact of congestion control algorithms on transmission control, we conduct formulas accuracy validation based on different congestion control algorithms separately.

As shown in Fig. 13, whether using BBR, Copa, or Cubic, the calculated values from the formulas and the testbed simulation values for recovery latency, redundancy cost, and goodput reduction for a given redundancy level are in close agreement. When using BBR and Copa, the absolute error of recovery latency does not exceed 3.2 ms, and the absolute errors of redundancy cost and goodput reduction do not exceed 2.4%. When using CUBIC, due to the larger values of recovery latency, the absolute error of recovery latency is slightly higher but does not exceed 5.8 ms, and the absolute errors of redundancy cost and goodput reduction do not exceed 1.4%. As the redundancy level increases, differences between the calculated values from the formulas and the testbed simulation values increase. This is because the actual amount of injected replicas is limited by the sending rate determined by the congestion control algorithm, meaning that the amount of replicas injected within a T_{unit} is capped. The larger the redundancy level, the greater the difference between the actual average amount of replicas injected per packet and the redundancy level, resulting in a larger difference between the simulated measurement values and the calculated values from the formulas for AutoRec's benefits and overhead.

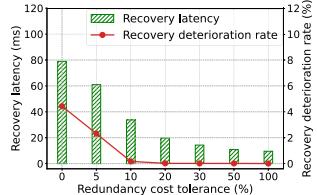
In summary, the formulas used by AutoRec can accurately describe the benefits and overhead of AutoRec at each fixed redundancy level.



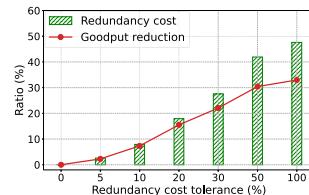
(a) Benefits with different recovery latency tolerance.



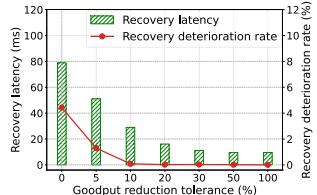
(b) Overhead with different recovery latency tolerance.

Fig. 14. Recovery latency tolerance (α) sensitivity validation.

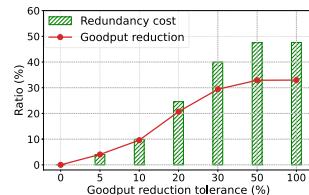
(a) Benefits with different redundancy cost tolerance.



(b) Overhead with different redundancy cost tolerance.

Fig. 15. Redundancy cost tolerance (β) sensitivity validation.

(a) Benefits with different goodput reduction tolerance.



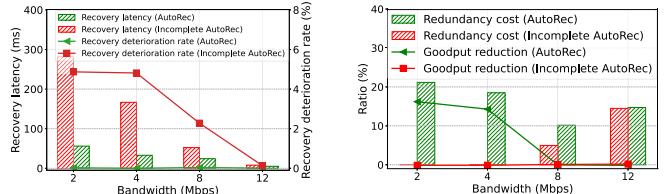
(b) Overhead with different goodput reduction tolerance.

Fig. 16. Goodput reduction tolerance (γ) sensitivity validation.

The sensitivity of user-customizable parameters. In this experiment, we test the performance of AutoRec under different user-customizable parameters (i.e., recovery latency tolerance, redundancy cost tolerance, and goodput reduction tolerance in § V-A). We use a live stream with a bitrate of 12 Mbps to ensure that the entire transmission process remains in on-mode, thereby maximizing the impact of injection on goodput. We first set both the redundancy cost tolerance and the goodput reduction tolerance to 100% to observe how AutoRec’s benefits and overhead change with the recovery latency tolerance. Next, we set recovery latency tolerance to 0 ms and observe how AutoRec’s benefits and overhead change with the redundancy cost tolerance or goodput reduction tolerance by setting the goodput reduction tolerance to 100% or the redundancy cost tolerance to 100%, respectively.

As shown in Fig. 14, the benefits and overhead of AutoRec gradually decrease as the recovery latency tolerance increases. As shown in Fig. 15 and Fig. 16, as the redundancy cost tolerance or goodput reduction tolerance increases, AutoRec’s benefits and overhead also increase. When the recovery latency tolerance is set to 0 ms and both the redundancy cost tolerance and goodput reduction tolerance are set to 100%, AutoRec’s benefits and overhead are at their maximum, with the values of recovery latency, recovery deterioration rate, redundancy cost, and goodput reduction being 9.5 ms, ~0%, 47.6%, and 33.0%, respectively.

As shown in Fig. 14, due to factors such as the limit on the maximum redundancy level, AutoRec cannot achieve



(a) Benefits with different bandwidth.

(b) Overhead with different bandwidth.

Fig. 17. Opportunistic reinjection effectiveness validation.

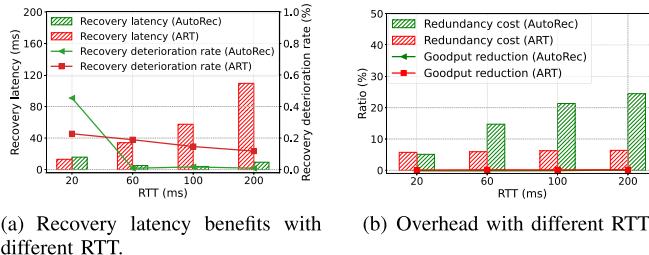
the theoretical minimum recovery latency. When the recovery latency tolerance is set between 10 ms and 200 ms, AutoRec can inject appropriate replicas to control the recovery latency within no more than 3 ms of the recovery latency tolerance. When the recovery latency tolerance exceeds 100 ms, the recovery latency can already meet the recovery latency tolerance requirements without injecting redundant replicas by AutoRec. At this point, AutoRec injects negligible replicas, and the redundancy cost and goodput reduction approach 0%.

As shown in Fig. 15 and Fig. 16, regardless of the values set for the redundancy cost tolerance or goodput reduction tolerance, AutoRec can maintain the measured redundancy cost or goodput reduction no more than 1% of the redundancy cost tolerance or goodput reduction tolerance. When the redundancy cost tolerance or goodput reduction tolerance gradually increases from 0% to 30%, AutoRec can quickly reduce the recovery latency to below 20 ms. Further increasing the redundancy cost or goodput reduction tolerance beyond 30% provides diminishing returns in terms of recovery latency optimization, while incurring higher overhead.

In summary, AutoRec can accelerate recovery latency to meet user-customizable recovery latency tolerance whenever possible while satisfying the constraints of the user-customizable redundancy cost tolerance and goodput reduction tolerance.

Opportunistic reinjection effectiveness validation. In this experiment, we verify the effectiveness of using opportunistic reinjection in the Reinjection Controller (§ V-B). We implement two versions of AutoRec: one with the opportunistic reinjection mechanism and one without. We set the bitrate of the live stream to 4 Mbps and observe how the benefits and overhead of the two versions of AutoRec change as the bandwidth gradually increases from 2 Mbps to 12 Mbps.

When the bandwidth is between 2 Mbps and 4 Mbps, the entire transmission process is in on-mode. As shown in Fig. 17, in this situation, the version of AutoRec without the opportunistic reinjection mechanism injects negligible replicas, resulting in higher recovery latency and recovery deterioration rate, which are 281.4 ms ~ 166.7 ms and 4.9% ~ 4.8%, respectively, while its redundancy cost and goodput reduction are almost 0. In contrast, AutoRec with the opportunistic reinjection mechanism injects an appropriate amount of redundant replicas, optimizing the recovery latency to 56.1 ms ~ 32.9 ms and the recovery deterioration rate to nearly 0%. This optimization effect is significantly better compared to AutoRec without the opportunistic reinjection mechanism. Furthermore, its redundancy cost and goodput reduction reach 21.2% ~ 18.5% and 16.2% ~ 14.3%, respectively, which are higher than those of AutoRec without the opportunistic reinjection



(a) Recovery latency benefits with different RTT.

(b) Overhead with different RTT.

Fig. 18. AutoRec vs. ART.

mechanism. Nevertheless, it still meets the constraints of the redundancy cost tolerance and goodput reduction tolerance.

When the bandwidth is between 8 Mbps and 12 Mbps, part of the transmission process will be in off-mode. As shown in Fig. 17, the version of AutoRec without the opportunistic reinjection mechanism can inject more redundant replicas as the bandwidth increases, optimizing the recovery latency and recovery deterioration rate to 52.3 ms \sim 7.9 ms and 2.3% \sim 0.1%, respectively, with its redundancy cost and goodput reduction being 5.0% \sim 14.5% and nearly 0%. Similarly, the version of AutoRec with the opportunistic reinjection mechanism can also inject an appropriate amount of redundant replicas, optimizing the recovery latency and recovery deterioration rate to 24.3 ms \sim 5.1 ms and nearly 0%, respectively, which is better than that of the version without the opportunistic reinjection mechanism. Furthermore, its redundancy cost and goodput reduction are 10.2% \sim 14.7% and nearly 0%, respectively. Although the overhead is higher compared to the version without the opportunistic reinjection mechanism, they still meet the constraints of redundancy cost tolerance and goodput reduction tolerance.

In summary, compared to the version of AutoRec without the opportunistic reinjection mechanism, the version with the opportunistic reinjection mechanism can inject an appropriate amount of replicas to accelerate loss recovery when there is no off-mode or when the off-mode distribution is uneven.

AutoRec vs. ART. This experiment compares two loss recovery acceleration schemes: AutoRec and ART [29]. We set the bitrate of the live stream to 4 Mbps and observe how the benefits and overhead of AutoRec and ART change as the RTT gradually increases from 20 ms to 200 ms.

As Fig. 18 shows, when the RTT is 20 ms, ART injects more replicas than AutoRec, optimizing the recovery latency and recovery deterioration rate to 13.02 ms and 0.23%, respectively. Its redundancy cost and goodput reduction are 5.73% and nearly 0%, respectively. In this case, AutoRec optimizes the recovery latency and recovery deterioration rate to 15.76 ms and 0.45%, respectively. Although the optimization effect is slightly worse than ART, it meets the recovery latency tolerance. Its redundancy cost and goodput reduction are 5.10% and nearly 0%, respectively, which is slightly lower compared to ART.

As Fig. 18 shows, when the RTT is between 60 ms and 200 ms, ART's recovery latency increases from 34.03 ms to 109.44 ms as the RTT increases, with the recovery deterioration rate between 0.12% and 0.19%. Its redundancy cost and goodput reduction are 5.93% \sim 6.37% and nearly 0%, respectively. On the other hand, AutoRec injects more redundant replicas than ART, maintaining the recovery latency

below 30 ms and the recovery deterioration rate nearly 0%. Its redundancy cost increases with the RTT, ranging from 14.72% to 24.44%, while its goodput reduction remains nearly 0%.

In summary, compared to ART, AutoRec adjusts the amount of redundant replica injection based on different RTTs to optimize the loss recovery latency.

AutoRec vs. FEC. In this experiment, we implement AutoRec on two RTC simulation platforms to: (1) validate its performance in RTC scenarios, (2) assess its robustness across diverse packet loss models and loss rate magnitudes, and (3) benchmark its effectiveness against multiple FEC strategies.

This experimental evaluation consists of the following mechanisms:

- **AutoRec:** Our proposed packet loss recovery mechanism. In RTC scenarios, we use parameter settings recovery latency tolerance $\alpha = 0$ ms (minimizing recovery latency), redundancy cost tolerance $\beta = 20\%$, and goodput reduction tolerance $\gamma = 20\%$ as the standard configuration.
- **AutoRec_{LC}:** The AutoRec mechanism with $\alpha = 0$ ms, $\beta = 5\%$, and $\gamma = 5\%$. Compared to standard AutoRec ($\beta = 20\%$, $\gamma = 20\%$), AutoRec under this configuration aims to reduce bandwidth cost by accepting higher susceptibility to video freezing.
- **AutoRec_{HC}:** The AutoRec mechanism with $\alpha = 0$ ms, $\beta = 50\%$, and $\gamma = 50\%$. Compared to standard AutoRec ($\beta = 20\%$, $\gamma = 20\%$), AutoRec under this configuration aims to reduce video freezing by tolerating higher bandwidth costs.
- **RTX:** A fundamental packet loss recovery mechanism that detects lost packets via ACK sequencing and retransmits them.
- **WebRTC₁₄:** The original Web Real Time Communication (WebRTC) [35] FEC mechanism proposed in Google's 2014 paper [30], which dynamically adjusts redundancy based on RTT measurements.
- **WebRTC_{NOW}:** The current industry-standard WebRTC FEC implementation (used in Stadia [36], Meet [37], etc.) that employs more aggressive redundancy than WebRTC₁₄.
- **Hairpin:** A novel packet loss recovery mechanism for edge-based interactive streaming that adaptively applies FEC encoding to lost packets [8]. The standard coefficient λ of Hairpin is set to 10^{-4} .
- **Hairpin_{LC}:** The Hairpin mechanism configured with $\lambda = 10^{-1}$. Compared to standard Hairpin ($\lambda = 10^{-4}$), Hairpin under this configuration aims to reduce bandwidth cost by accepting higher susceptibility to video freezing.
- **Hairpin_{HC}:** The Hairpin mechanism configured with $\lambda = 10^{-7}$. Compared to standard Hairpin ($\lambda = 10^{-4}$), Hairpin under this configuration aims to reduce video freezing by tolerating higher bandwidth costs.
- **Tambur:** An efficient loss recovery mechanism for videoconferencing that leverages streaming codes—a specialized FEC variant optimized for burst losses [7]. We set the maximum FEC redundancy ratio to 25% by default.

- **Tambur_{LC}:** The Tambur mechanism was configured with its maximum redundancy parameter set to the minimum value, resulting in the generation of exactly one FEC redundancy packet per video frame. Compared to standard Tambur (maximum redundancy ratio = 25%), Tambur under this configuration aims to reduce bandwidth cost by accepting higher susceptibility to video freezing.
- **Tambur_{HC}:** The Tambur mechanism configured with maximum redundancy ratio = 50%. Compared to standard Tambur (maximum redundancy ratio = 25%), Tambur under this configuration aims to reduce video freezing by tolerating higher bandwidth costs.

We evaluate the following performance metrics:

- **Deadline miss rate (DMR):** Proportion of video frames failing to render before their deadline due to packet loss.
- **Bandwidth cost (BWC):** The ratio of non-video-frame bytes (e.g., retransmission redundancy bytes, FEC redundancy bytes) to total transmitted bytes.
- **Percent of video spent frozen:** The ratio of the cumulative duration of video freezing events to the total video playback time, expressed as a percentage.

We configure the environment with 12 Mbps bandwidth, 4 Mbps bitrate, 60 fps frame rate, 30 ms RTT, and 100 ms playback deadline.

First, using an ns3-based WebRTC simulator [38], we evaluate the deadline miss rate and bandwidth cost of AutoRec, AutoRec_{LC}, AutoRec_{HC}, WebRTC_{NOW}, WebRTC₁₄, Hairpin, Hairpin_{LC}, and Hairpin_{HC} under a variety of conditions, including four fixed packet loss rates (0.5%, 1%, 5%, 10%) and two real-world network traces. Specifically, Trace #1 represents a high-loss, stable-latency environment, characterized by an average bandwidth of 11.3 Mbps, an average RTT of 13.7 ms, an RTT variance of 13, an average packet loss rate of 10.1%, and a 95th percentile packet loss rate of 75.0%. In contrast, Trace #2 exemplifies a low-loss, high-jitter RTT scenario, with an average bandwidth of 4.6 Mbps, a longer average RTT of 38.1 ms, a substantially larger RTT variance of 178, an average packet loss rate of 2.8%, and a 95th percentile packet loss rate of 25.0%.

Fig. 19 illustrates the experimental results. As shown in Fig. 19(a) and Fig. 19(b), at low packet loss rates (0.5% and 1%), AutoRec_{HC}, AutoRec, and AutoRec_{LC} all achieve 0% deadline miss rate with modest bandwidth cost, where the highest cost are 9.92%, 9.92%, and 4.39% respectively, satisfying redundancy cost constraint. While WebRTC_{NOW} and WebRTC₁₄ also achieve 0% deadline miss rate, they incur substantially higher bandwidth costs reaching 32.81% and 16.16%. Hairpin_{HC} achieves 0% deadline miss rate at both loss rates, whereas Hairpin fails at 0.5% loss and Hairpin_{LC} fails at both rates. Notably, Hairpin in all three configurations maintains minimal bandwidth costs under 3%.

As shown in Fig. 19(c) and Fig. 19(d), at high loss rates (5% and 10%), AutoRec_{HC} maintains a 0% deadline miss rate with bandwidth cost up to 46.28%. However, AutoRec fails to achieve 0% deadline miss rate at 10% loss rate while incurring up to 17.23% bandwidth cost. AutoRec_{LC} completely fails to achieve 0% deadline miss rate and performs comparably to RTX at 10% loss rate. AutoRec in all three configurations

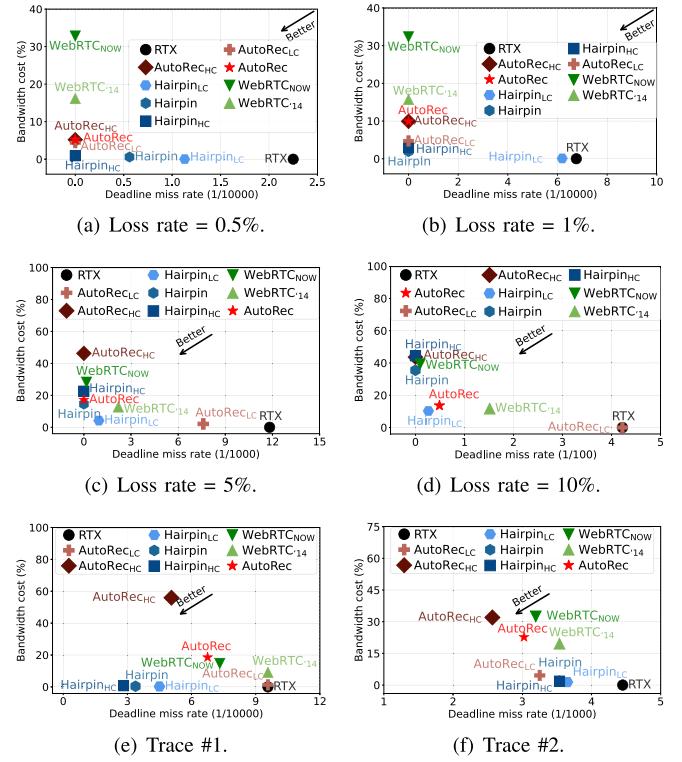
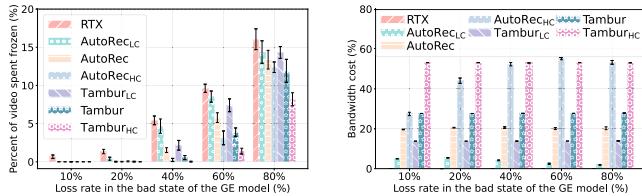


Fig. 19. The performance of the AutoRec and FEC mechanisms under both controlled fixed packet loss rates and real-world network traces.

satisfies its respective redundancy cost constraints. Neither WebRTC_{NOW} nor WebRTC₁₄ achieves a 0% deadline miss rate under either loss rate, with bandwidth costs below 39.38% and 12.77% respectively. Hairpin_{HC} and standard Hairpin achieve 0% deadline miss rate with bandwidth cost up to 44.48% and 35.48% respectively, while Hairpin_{LC} fails with bandwidth cost up to 10.26%. Notably, at 10% loss rate, Hairpin_{LC} reduces deadline miss rate more significantly than AutoRec while consuming less bandwidth.

As shown in Fig. 19(e), under Trace #1, AutoRec performs considerably worse than Hairpin but outperforms WebRTC. Specifically, in this high-loss environment, the AutoRec_{LC} instance injects negligible redundant packets to avoid exceeding the overhead limit, resulting in minimal improvement to the deadline miss rate. In contrast, the AutoRec and AutoRec_{HC} instances reduce the deadline miss rate to 0.067% and 0.051%, with associated bandwidth costs of 18.5% and 55.9%, respectively. However, owing to the superior efficiency of FEC over simple retransmission, even the best-performing AutoRec instance (AutoRec_{HC}) is surpassed by Hairpin_{LC}. The latter achieves a lower deadline miss rate of 0.045% at a significantly lower bandwidth cost of only 0.27%. Among the WebRTC versions, WebRTC_{NOW} achieves the best result, yet its deadline miss rate (0.073%) remains higher than that achieved by the AutoRec instance, with a bandwidth cost of 14.7%.

As shown in Fig. 19(f), under Trace #2, AutoRec demonstrates stronger performance than both Hairpin and WebRTC. Owing to the low loss rate and high RTT variability, Hairpin overestimates the number of remaining retransmission opportunities, leading to an insufficient number of FEC packets. Consequently, the best-performing instance of Hairpin



(a) Percent of video spent frozen vs. (b) Bandwidth cost vs. packet loss rates in the bad state of the GE model.

Fig. 20. The performance of the AutoRec and Tambur under varying packet loss rates in the bad state of Gilbert-Elliott model.

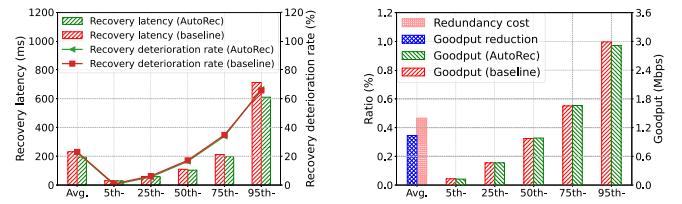
(Hairpin_{HC}) only reduces the deadline miss rate to 0.35% with a bandwidth cost of 1.85%. In comparison, the least effective instance of AutoRec (AutoRec_{LC}) achieves a lower deadline miss rate of 0.33% with a bandwidth cost of 4.56%. Among the WebRTC versions, WebRTC_{NOW} achieves the best result, yet its deadline miss rate and bandwidth cost are both worse than those of AutoRec.

These results demonstrate AutoRec’s effectiveness in RTC scenarios. Under low packet loss conditions, AutoRec outperforms these FEC mechanisms, achieving comparable stall reduction with similar or lower bandwidth cost. However, as packet loss rates increase, AutoRec’s effectiveness diminishes relative to these FEC mechanisms. To match their stall reduction at higher loss rates, AutoRec incurs greater bandwidth costs that increasingly exceed those of the FEC mechanisms.

Second, we evaluate the percent of video spent frozen and bandwidth cost of AutoRec, AutoRec_{LC}, AutoRec_{HC}, Tambur, Tambur_{LC}, and Tambur_{HC} in Gilbert-Elliott loss model [39] using Ringmaster (an extensible videoconferencing research platform for FEC benchmarking, replacing WebRTC for controlled experiments) [40].

The Gilbert-Elliott (GE) model is a two-state Markov chain widely used to simulate burst packet losses in networks, where the channel alternates between good (low-loss) and bad (high-loss) states. We implement a Gilbert-Elliott loss model with state transitions every 30 ms, fixed transition probabilities (good to bad: 20%; bad to good: 80%), fixed good state loss rate (1%), and varying bad state loss rates (10%, 20%, 40%, 60%, 80%) to evaluate AutoRec’s performance under varying burst loss intensities. Ringmaster utilizes a constant bitrate transmission scheme without congestion control. It does not employ a timeout-based termination mechanism. Instead, if the first unacknowledged datagram remains outstanding for more than one second, the sender forces the next frame to be a key frame and continues transmission.

Fig. 20 illustrates the experimental results. At low bad state loss rate (10%), AutoRec, AutoRec_{LC}, and AutoRec_{HC} can reduce the percent of video spent frozen to 0%, with AutoRec_{LC} operating under the strictest redundancy cost constraint and incurring only 4.93% bandwidth cost. Meanwhile, Tambur, Tambur_{LC}, and Tambur_{HC} achieve 0% of video spent frozen. Tambur_{LC} (the one with the smallest bandwidth cost among Tambur, Tambur_{LC}, and Tambur_{HC}) requires over 13.81% bandwidth cost, significantly exceeding AutoRec’s bandwidth cost. At moderate bad state loss rates (20% to 40%), the performance in reducing the percent of video spent frozen of both AutoRec and Tambur across configurations



(a) The average and quantile of the benefits.

(b) The average overhead and the quantile of the goodput.

Fig. 21. AutoRec’s benefits and overhead in the real network.

consistently correlates with their bandwidth cost rankings. Under high bad state loss rates (60% and 80%), Tambur achieves 3.88% and 11.91% of video spent frozen at a bandwidth cost of ~28%; by contrast, even AutoRec_{HC} (the best-performing one for freeze reduction among AutoRec, AutoRec_{LC}, and AutoRec_{HC}) achieves only 3.13% and 12.38% of video spent frozen at a bandwidth cost of ~53%.

These results demonstrate that AutoRec delivers strong performance even under burst packet loss conditions. While AutoRec outperforms Tambur at low burst loss rates, its effectiveness in reducing video freezes diminishes as the burst loss rate increases. Under high burst loss rates, AutoRec falls short of Tambur’s performance.

In summary, AutoRec effectively reduces video freezes within redundancy cost constraints for RTC applications. It consistently mitigates freezes under all tested packet loss conditions, including varying loss magnitudes (low to high) and types (random and burst), though its effectiveness decreases as loss rates increase. Crucially, AutoRec outperforms multiple FEC schemes at low loss rates, matches their performance at moderate loss rates, but is surpassed by multiple FEC mechanisms in high-loss scenarios.

B. Real-Network Evaluation

To further explore AutoRec’s performance, we deploy AutoRec prototype in our CDN proxy and evaluate video freezing metrics, loss recovery benefits, and reinjection overhead.

Parameter configuration. To facilitate the deployment of AutoRec in real-world networks, the user-customizable parameters in the Redundancy Adapter (§ V-A), namely recovery latency tolerance (α), redundancy cost tolerance (β), and goodput reduction tolerance (γ), are set to 200 ms, 50%, and 50%, respectively.

AutoRec’s loss recovery benefits and overhead. Fig. 21(a) demonstrates the performance of AutoRec in the real network, showing consistent improvements in recovery latency and recovery deterioration rate. The average recovery latency is reduced by 15.8%, from 231.2 ms to 194.6 ms, while the average recovery deterioration rate decreases by 1.7%, from 23.0% to 22.6%. Notably, the 95th percentile recovery latency is reduced by 102.2 ms. This means that for the worst-affected 5% of receivers, the waiting time for recovering from the loss of a first retransmission is cut by 102.2 ms. These results indicate that AutoRec effectively reduces the recovery latency below the recovery latency tolerance, which satisfies the recovery latency constraints.

The overhead of AutoRec under the evaluated parameter configuration is shown in Fig. 21(b), which demonstrates

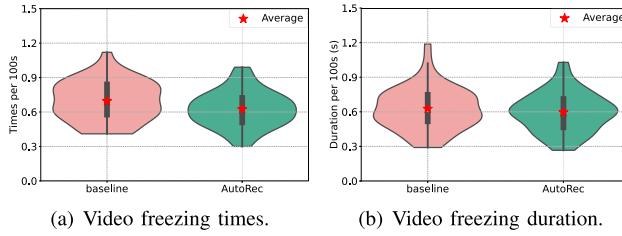


Fig. 22. The video freezing changes in the real network.

that the introduced overhead is negligible. Specifically, the average redundancy cost of AutoRec is only 0.47%, indicating that the total traffic increases by at most 0.47%. Meanwhile, the average goodput reduction of AutoRec is just 0.34%. Compared to the baseline, the reduction in goodput is more pronounced at higher percentiles, with a decrease of 2.58% at the 95th percentile, from 2.99 Mbps to 2.91 Mbps. From the above results, it is evident that AutoRec's redundancy cost and goodput reduction are both strictly within the redundancy cost tolerance and goodput reduction tolerance, meeting the constraints on overhead.

In summary, in the real network, the deployed AutoRec accelerates loss recovery to meet user demands whenever possible while keeping the overhead controllable.

Video freezing. The AutoRec performance can be further evaluated by the observed client-side video freezing, including its frequency and duration. As Fig. 22(a) shows, AutoRec reduces the average freezing times (per 100s) from 0.69 to 0.62, representing an improvement of 10.10%. The minimum and maximum values of video freezing times (per 100s) are reduced by 26.83% and 11.61%, respectively. Furthermore, AutoRec also reduces freezing duration (per 100s) by 4.74%, in which the maximum value is reduced by 159.03 ms (with the ratio of 13.39%), as Fig. 22(b) shows.

In summary, the results demonstrate that AutoRec effectively enhances loss recovery timeliness, leading to a measurable reduction in both the frequency and duration of client-side video freezing.

VIII. DISCUSSION

A. What About the Overhead of AutoRec?

The computational and storage overhead of the AutoRec remains minimal even under extreme conditions.

Computational overhead. The computational overhead of AutoRec consists of the following two aspects: (i) the calculation of the redundancy level in the Redundancy Adapter; (ii) the computational operations of insertion, deletion, and movement of the reinjection queue in the Reinjection Controller.

For aspect (i), the redundancy level is periodically calculated at each decision interval. The time complexity of this calculation is $O(\log(n))$, where n is the maximum allowable value of K_θ (§ V-A), typically set to 10. Given the small value of n and the logarithmic time complexity, the computational overhead is negligible.

For aspect (ii), since the insertion, deletion, and movement operations do not affect the order that the entire reinjection queue should maintain, i.e., an order where the time of their last transmission is from the earliest at the front to the latest at

the back, there is no need for sorting operations. Insertion and movement operations do not require traversing the reinjection queue. Instead, they only access the head and tail of the reinjection queue, which means no excessive computation is required. However, deletion requires traversing the reinjection queue to find the acknowledged packets and then delete them. The time complexity here is $O(n)$, where n is the number of packets in the reinjection queue. In the worst-case scenario, namely that all sent packets are lost, the reinjection queue will contain all the packets sent within one RTT. According to our measurement results, the average bitrate of live streams is 984.5 kbps, the average SRTT is 95.5 ms, and the amount of data carried by each packet is approximately 1300 bytes. Therefore, in the worst-case scenario, the average number of packets in the reinjection queue of each connection is approximately 9. So, the overhead required to traverse the reinjection queue once is extremely small.

Storage overhead. The storage overhead of AutoRec mainly consists of the following two aspects: (i) the storage overhead required for measuring the loss detection time in the Redundancy Adapter; (ii) the storage overhead needed to implement the reinjection queue in the Reinjection Controller.

For aspect (i), we need to store the first transmission time of each currently unacknowledged packet so that the loss detection time of the packet can be obtained when it is first retransmitted. As analyzed in § VIII-A, the average number of packets in the reinjection queue of each connection is 9, therefore, the storage overhead of storing the sending time of each currently unacknowledged packet will not be large either.

For aspect (ii), we need to store each packet that is currently being retransmitted but has not been acknowledged and whose number of injection times is less than K_θ (§ V-A) in the reinjection queue. Additionally, we have to record some information for these packets. Evidently, compared with the overhead of storing packets, the overhead of recording information for these packets can be neglected. As analyzed in § VIII-A, the average number of packets in reinjection queue at any given moment is 9. Therefore, in the worst-case scenario, the overhead of storing packets in reinjection queue will not be significant either.

B. Will AutoRec Exacerbate Congestion?

In fact, the number of packets injected by AutoRec will not be so large as to exacerbate network congestion. Firstly, AutoRec can adaptively adjust the number of injections for each packet. When there is congestion, the packet loss rate will increase. At this time, AutoRec will adaptively reduce the number of injections for each lost packet in the next decision interval according to the goodput reduction tolerance and redundancy cost tolerance, so that the total number of injected packets will not exceed a certain threshold. In addition, as analyzed in § VIII-A, the average number of packets that are sent but not acknowledged is very small. As shown in Fig. 4, the average packet loss rate in the real network is about 4%, that is, the lost packets account for a very small part of the total number of sent packets. And we only inject packets that are retransmitted but not acknowledged, so the number of injected packets will not be large to exacerbate congestion.

C. How AutoRec Integrates With QUIC's Existing Loss-Recovery Mechanisms?

QUIC's native loss-recovery mechanism exclusively tracks the last transmission attempt (whether initial transmission, retransmission, or reinjection) for each packet. It utilizes this final attempt's information to determine packet loss and initiate retractions, thereby ensuring reliable data delivery. AutoRec, conversely, performs reinjection of previously retransmitted or injected packets only under specific conditions. Notably, both reinjection and retransmission operations require first renumbering the original data packet and then sending the renumbered packet. Crucially, AutoRec fundamentally advances retransmission timing by executing these packet retractions before loss occurs. This process operates without interfering with QUIC's native loss-recovery mechanism or compromising reliability.

Illustrative example (AutoRec's redundancy level = 2): When QUIC's native loss-recovery mechanism detects that packet P_1 (carrying application data D) is lost, it retransmits P_1 by renumbering it as P_2 and sending P_2 . QUIC then stops monitoring P_1 for loss and begins monitoring P_2 for loss. If AutoRec's conditions are met, it reinjects P_2 by renumbering P_2 as P_3 and sending P_3 . QUIC consequently stops monitoring P_2 for loss and begins monitoring P_3 for loss. When AutoRec's conditions are met again, it reinjects P_3 by renumbering P_3 as P_4 and sending P_4 . QUIC then stops monitoring P_3 for loss and begins monitoring P_4 for loss. Upon a third trigger, AutoRec skips reinjection of P_4 because the packet carrying data D has been reinjected twice, reaching the maximum allowed by the redundancy level. Should P_4 be lost, QUIC's native loss-recovery mechanism detects this loss, retransmits P_4 by renumbering P_4 as P_5 and sending P_5 , then stops monitoring P_4 for loss and begins monitoring P_5 for loss. Thereafter, QUIC's existing loss-recovery mechanism ensures data D is successfully delivered to the receiver through this persistent cycle of loss detection and retransmission. Notably, if any packet from P_1 to P_5 is ACKed, the loss detection for P_5 is terminated to prevent infinite retransmission loops.

IX. RELATED WORK

Adaption to on-off traffic pattern. To the best of our knowledge, almost all prior works [41], [42], [43], [44], [45], [46], [47], [48], [49], [50], [51], [52], [53] fall into the category of adapting congestion control to on-off traffic patterns. For example, Zhang et al. [46] proposed a TCP variant to overcome the challenges, in which the on-off traffic pattern disturbs the increase of the TCP congestion window and triggers packet loss at the beginning of the on-mode. This paper does not focus on the congestion control issues of live streaming. Instead, in a manner analogous to how TADOC [54] leverages repetitive textual features to optimize text analytics, we leverage the on-off pattern to accelerate packet loss recovery.

Loss tolerance control for live streaming. Many studies have been proposed to enhance the data timeliness of live streaming to achieve efficient loss recovery and optimize client-side video freezes. The key ideas of these works include injecting supplement data (e.g., FEC [5], [6], [7], [8] and multi-path retractions [9], [10], [55]), ignoring some non-critical losses (e.g., application-level controls [13],

[14], [15], semi-reliable transmissions [11], [12], [56]). For instance, Tambur [7] employs streaming codes, a specialized FEC variant optimized for burst losses. In contrast, AutoRec optimizes loss recovery for diverse loss patterns by selectively reinjecting only lost packets, significantly reducing bandwidth overhead under low-loss conditions. Whereas GRACE [56] preserves QoE across varying packet losses via a novel neural video codec implemented through specialized application-layer frame encoding and decoding, AutoRec operates at the transport layer by reinjecting lost packets. The mechanism most close to AutoRec is Hairpin [8], which adaptively applies FEC encoding to lost packets specifically for edge-based interactive streaming. In fact, Hairpin can be viewed as an advanced variant of AutoRec designed for scenarios where modifying both client and server endpoints is feasible, since under these conditions, FEC is more efficient than simple packet reinjection. However, in commercial large-scale live-streaming product networks, the CDN vendor is mainly responsible for optimizing the end-to-end transmission performance and has control rights only on the sender side. Therefore, the above arts that apply client-side modification cannot meet the requirements of real-world deployment. In this paper, we propose a sender-side approach from the perspective of CDN vendors.

Advancements upon ARQ. Modern CDN vendors simply apply the ARQ for loss tolerance control in commercial live-streaming services. However, the legacy ARQ-based loss recovery is far from satisfactory. There also exist many studies [29], [33], [57], [58], [59], [60], [61], [62] that improve the performance of ARQ by introducing redundancy to loss recovery. These works, however, suffer from obviously deteriorated goodput since the inserted extra packets occupy the sender-side and in-network resources. This paper overcomes the above challenges by taking full advantage of the off-mode of live streaming.

Upstream path. The data transmission paths for live streaming can generally be simply divided into the upstream path from the live streamer to the server and the downstream path from the server to the user. Currently, many research studies [4], [63], [64], [65] are dedicated to improving the transmission quality of the upstream path. In contrast to these efforts, AutoRec primarily focuses on the optimization of the downstream path performance, with particular emphasis on reducing packet loss recovery latency. Since the downstream path is directly relevant to the users, optimizing the packet loss recovery latency using AutoRec on the downstream path rather than the upstream path is more likely to result in an improvement of users' Quality of Experience.

X. CONCLUSION

Slow loss recovery is not primarily caused by the retransmission process itself, but rather by the loss of retransmitted packets, which constitutes the fundamental challenge. To address this issue, we propose AutoRec, a mechanism that accelerates loss recovery by enabling senders to reinject lost packet duplicates in a strategic and controlled manner. AutoRec leverages the on-off mode switching prevalent in live streaming to improve recovery timeliness without disrupting ongoing data transmission. The effectiveness of AutoRec is

validated through extensive testbed experiments and real-world deployment. It has been integrated into Tencent's global CDN [66] and its EdgeOne platform [67], serving billions of live-streaming users worldwide.

REFERENCES

- [1] Technavio. (2023). *Live Streaming Market By Product, End-User, and Geography-Forecast and Analysis 2023-2027*. [Online]. Available: <https://www.technavio.com/report/live-streaming-market-industry-analysis>
- [2] S. Hatchet. (2022). *2022 Yearly Video Game Live Streaming Trends Report*. [Online]. Available: <https://streamhatchet.com/blog>
- [3] J. Li et al., "LiveNet: A low-latency video transport network for large-scale live streaming," in *Proc. ACM SIGCOMM Conf.*, Aug. 2022, pp. 812–825.
- [4] J. Kim, Y. Jung, H. Yeo, J. Ye, and D. Han, "Neural-enhanced live streaming: Improving live video ingest via online learning," in *Proc. Annu. Conf. ACM Special Interest Group Data Commun. Appl., Technol., Archit., Protocols Comput. Commun.*, Jul. 2020, pp. 107–125.
- [5] S.-H.-G. Chan, X. Zheng, Q. Zhang, W.-W. Zhu, and Y.-Q. Zhang, "Video loss recovery with FEC and stream replication," *IEEE Trans. Multimedia*, vol. 8, no. 2, pp. 370–381, Apr. 2006.
- [6] F. Michel, Q. De Coninck, and O. Bonaventure, "QUIC-FEC: Bringing the benefits of forward erasure correction to QUIC," in *Proc. IFIP Netw. Conf.*, May 2019, pp. 1–9.
- [7] M. Rudow, F. Y. Yan, A. Kumar, G. Ananthanarayanan, M. Ellis, and K. Rashmi, "Tambur: Efficient loss recovery for videoconferencing via streaming codes," in *Proc. USENIX NSDI*, Apr. 2023, pp. 953–971.
- [8] Z. Meng et al., "Hairpin: Rethinking packet loss recovery in edge-based interactive video streaming," in *Proc. USENIX NSDI*, Apr. 2024, pp. 907–926.
- [9] Z. Zheng et al., "XLINK: QoE-driven multi-path QUIC transport in large-scale video services," in *Proc. ACM SIGCOMM Conf.*, Aug. 2021, pp. 418–432.
- [10] G. Chen et al., "FUSO: Fast multi-path loss recovery for data center networks," *IEEE/ACM Trans. Netw.*, vol. 26, no. 3, pp. 1376–1389, Jun. 2018.
- [11] M. Palmer, M. Appel, K. Spiteri, B. Chandrasekaran, A. Feldmann, and R. K. Sitaraman, "VOXEL: Cross-layer optimization for video streaming with imperfect transmission," in *Proc. 17th Int. Conf. Emerg. Netw. Experiments Technol.*, Dec. 2021, pp. 359–374.
- [12] T. Pauly, E. Kinnear, and D. Schinazi, "An unreliable datagram extension to QUIC," in *Internet Engineering Task Force Draft-Pauly-Quicdatagram-00*. USA: Internet Engineering Task Force (IETF), Sep. 2018.
- [13] B. Vamanan, J. Hasan, and T. Vijaykumar, "Deadline-aware datacenter TCP (D²TCP)," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 4, pp. 115–126, 2012.
- [14] C. Zhou, W. Wu, D. Yang, T. Huang, L. Guo, and B. Yu, "Deadline and priority-aware congestion control for delay-sensitive multimedia streaming," in *Proc. 29th ACM Int. Conf. Multimedia*, Oct. 2021, pp. 4740–4744.
- [15] H. Zhang, X. Shi, X. Yin, F. Ren, and Z. Wang, "More load, more differentiation—A design principle for deadline-aware congestion control," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Apr. 2015, pp. 127–135.
- [16] Y. Arda and J.-C. Hennet, "Inventory control in a multi-supplier system," *Int. J. Prod. Econ.*, vol. 104, no. 2, pp. 249–259, Dec. 2006.
- [17] D. Bertsekas and R. Gallager, *Gallager, Data Networks*, vol. 1. Upper Saddle River, NJ, USA: Prentice-Hall, 1992, p. 2.
- [18] T. Li, K. Zheng, and K. Xu, "Acknowledgment on demand for transport control," *IEEE Internet Comput.*, vol. 25, no. 2, pp. 109–115, Mar. 2021.
- [19] T. Benson, A. Anand, A. Akella, and M. Zhang, "Understanding data center traffic characteristics," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 1, pp. 92–99, Jan. 2010.
- [20] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proc. 10th ACM SIGCOMM Conf. Internet Meas.*, Nov. 2010, pp. 267–280.
- [21] S. H. Yeganeh, Y. Cheng, N. Cardwell, and V. Jacobson, "Delivery rate estimation," in *IETF Internet-Draft Draft-Cheng-Iccrg-Delivery-Rate-Estimation-02*. USA: Internet Engineering Task Force (IETF), 2017.
- [22] A. Sivaraman, K. Winstein, P. Thaker, and H. Balakrishnan, "An experimental study of the learnability of congestion control," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 4, pp. 479–490, Feb. 2015.
- [23] S. Rajasekaran, M. Ghobadi, G. Kumar, and A. Akella, "Congestion control in machine learning clusters," in *Proc. 21st ACM Workshop Hot Topics Netw.*, Nov. 2022, pp. 235–242.
- [24] S. Floyd, M. Handley, J. Padhye, and J. Widmer, "Equation-based congestion control for unicast applications," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 30, no. 4, pp. 43–56, Oct. 2000.
- [25] J. Iyengar and I. Swett, *QUIC Loss Detection and Congestion Control*, document IETF RFC 9002, 2021.
- [26] LiteSpeed Tech. *LiteSpeed QUIC and HTTP/3 Library*. Accessed: Oct. 19, 2025. [Online]. Available: <https://github.com/litespeedtech/lsquic>
- [27] I. Ovchinnikov and K. Pavlov. *Nginx*. [Online]. Available: <https://github.com/nginx>
- [28] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson, "BBR: Congestion-based congestion control," *Queue*, vol. 14, no. 5, pp. 20–53, Oct. 2016.
- [29] T. Li et al., "ART: Adaptive retransmission for wide-area loss recovery in the wild," in *Proc. IEEE 31st Int. Conf. Netw. Protocols (ICNP)*, Oct. 2023, pp. 1–11.
- [30] S. Holmer, M. Shemer, and M. Paniconi, "Handling packet loss in WebRTC," in *Proc. IEEE Int. Conf. Image Process.*, Sep. 2013, pp. 1860–1864.
- [31] (2020). *PSA: WebRTC M88 Release Notes*. [Online]. Available: <https://groups.google.com/g/discuss-webrtc/c/A0FjOcTW2c0/n/UAvveyPCAAJ>
- [32] Y. Cheng, N. Cardwell, N. Dukkipati, and P. Jha, *The RACK-TLP Loss Detection Algorithm for TCP*, document RFC 8985, 2021.
- [33] V. Paxson, M. Allman, J. Chu, and M. Sargent, *Computing TCP's Retransmission Timer*, document RFC 6298, 2011.
- [34] M. Mathis and J. Mahdavi, "Forward acknowledgement: Refining TCP congestion control," in *Proc. Conf. Appl., Technol., Archit., Protocols Comput. Commun.*, Aug. 1996, pp. 281–291.
- [35] *WebRTC*. [Online]. Available: <https://webrtc.org/>
- [36] A. Di Domenico, G. Perna, M. Trevisan, L. Vassio, and D. Giordano, "A network analysis on cloud gaming: Stadia, GeForce now and PSNow," *Network*, vol. 1, no. 3, pp. 247–260, Oct. 2021.
- [37] brianhu. (2021). *Google Meet Troubleshooting Playbook—Network and Hardware Troubleshooting*. [Online]. Available: <https://www.googlecloudcommunity.com/gc/Workspace-Product-Articles/Google-Meet-Troubleshooting-Playbook-Network-and-Hardware/ta-p/165810>
- [38] S. Zhang. (2020). *Soonyangzhang/Webrtc-GCC-NS3: Test Google Congestion Control on NS3*. [Online]. Available: <https://github.com/SoonyangZhang/webrtc-gcc-ns3>
- [39] E. O. Elliott, "Estimates of error rates for codes on burst-noise channels," *Bell Syst. Tech. J.*, vol. 42, no. 5, pp. 1977–1997, Sep. 1963.
- [40] M. Corporation. (2023). *Ringmaster*. [Online]. Available: <https://github.com/microsoft/ringmaster>
- [41] A. Wierman and T. Osogami, "A unified framework for modeling TCP-vegas, TCP-SACK, and TCP-reno," in *Proc. 11th IEEE/ACM Int. Symp. Modeling, Anal. Simulation Comput. Telecommun. Syst. (MASCOTS)*, Oct. 2003, pp. 269–278.
- [42] A. Wierman, T. Osogami, and J. Olsén, "Modeling TCP-vegas under on/off traffic," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 31, no. 2, pp. 6–8, Sep. 2003.
- [43] J. Esteban, S. A. Benno, A. Beck, Y. Guo, V. Hilt, and I. Rimac, "Interactions between HTTP adaptive streaming and TCP," in *Proc. 22nd Int. Workshop Netw. Operating Syst. Support Digit. Audio Video*, Jun. 2012, pp. 21–26.
- [44] T. Kupka, P. Halvorsen, and C. Griwodz, "Performance of on-off traffic stemming from live adaptive segmented HTTP video streaming," in *Proc. 37th Annu. IEEE Conf. Local Comput. Netw.*, Oct. 2012, pp. 401–409.
- [45] L. De Cicco, V. Caldaralo, V. Palmisano, and S. Mascolo, "ELASTIC: A client-side controller for dynamic adaptive streaming over HTTP (DASH)," in *Proc. 20th Int. Packet Video Workshop*, Dec. 2013, pp. 1–8.
- [46] T. Zhang, J. Wang, J. Huang, J. Chen, Y. Pan, and G. Min, "Tuning the aggressive TCP behavior for highly concurrent HTTP connections in intra-datacenter," *IEEE/ACM Trans. Netw.*, vol. 25, no. 6, pp. 3808–3822, Dec. 2017.
- [47] Y. Zhao, B. Zhang, C. Li, and C. Chen, "ON/OFF traffic shaping in the internet: Motivation, challenges, and solutions," *IEEE Netw.*, vol. 31, no. 2, pp. 48–57, Mar. 2017.

- [48] V. Arun and H. Balakrishnan, "Copa: Practical delay-based congestion control for the internet," in *Proc. USENIX Symp. Netw. Syst. Design Implement.*, 2018, pp. 329–342.
- [49] T. Li et al., "TACK: Improving wireless transport performance by taming acknowledgments," in *Proc. Annu. Conf. ACM Special Interest Group Data Commun. Appl., Technol., Archit., Protocols Comput. Commun.*, Jul. 2020, pp. 15–30.
- [50] M. Yanev, S. McQuistin, and C. Perkins, "Does TCP new congestion window validation improve HTTP adaptive streaming performance?," in *Proc. ACM NOSSDAV*, 2022, pp. 29–35.
- [51] B. Wu, T. Li, C. Luo, C. Ouyang, X. Du, and F. Wang, "AutoPlex: Inter-session multiplexing congestion control for large-scale live video services," in *Proc. ACM SIGCOMM Workshop Netw.-Appl. Integr.*, Aug. 2022, pp. 1–6.
- [52] X. Yan et al., "Poster: TOO: Accelerating loss recovery by taming on-off traffic patterns," in *Proc. ACM SIGCOMM Conf.*, Sep. 2023, pp. 1147–1149.
- [53] B. Wu et al., "Reducing first-frame delay of live streaming by simultaneously initializing window and rate," in *Proc. IEEE 44th Int. Conf. Distrib. Comput. Syst. (ICDCS)*, Jul. 2024, pp. 232–242.
- [54] F. Zhang et al., "TADOC: Text analytics directly on compression," *VLDB J.*, vol. 30, no. 2, pp. 163–188, Mar. 2021.
- [55] L. Li et al., "A measurement study on multi-path TCP with multiple cellular carriers on high speed rails," in *Proc. Conf. ACM Special Interest Group Data Commun.*, Aug. 2018, pp. 161–175.
- [56] Y. Cheng et al., "GRACE: Loss-resilient real-time video through neural codecs," in *Proc. USENIX NSDI*, 2023.
- [57] B. S. Bakshi, P. Krishna, N. H. Vaidya, and D. K. Pradhan, "Improving performance of TCP over wireless networks," in *Proc. 17th Int. Conf. Distrib. Comput. Syst.*, May 1997, pp. 365–373.
- [58] D. Barman, I. Maitta, E. Altman, and R. El-Azouzi, "TCP optimization through FEC, ARQ and transmission power tradeoffs," in *Proc. WWIC*, Cham, Switzerland: Springer, 2009, pp. 87–98.
- [59] J. Chen, W. Tan, L. Liu, X. Hu, and F. Xu, "Towards zero loss for TCP in wireless networks," in *Proc. IEEE 28th Int. Perform. Comput. Commun. Conf.*, Dec. 2009, pp. 65–70.
- [60] H. Xie and T. Li, "Revisiting loss recovery for high-speed transmission," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Apr. 2022, pp. 1987–1992.
- [61] K. Liu and J. Y. B. Lee, "Improving TCP performance over mobile data networks with opportunistic retransmission," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Apr. 2013, pp. 1992–1997.
- [62] S. Liu, J. Huang, W. Jiang, J. Wang, and T. He, "Reducing flow completion time with replaceable redundant packets in data center networks," in *Proc. IEEE 39th Int. Conf. Distrib. Comput. Syst. (ICDCS)*, Jul. 2019, pp. 46–56.
- [63] J. Li, Z. Li, Q. Wu, and G. Tyson, "On uploading behavior and optimizations of a mobile live streaming service," in *Proc. IEEE INFOCOM Conf. Comput. Commun.*, May 2022, pp. 1299–1308.
- [64] D. Ray, J. Kosaian, K. V. Rashmi, and S. Seshan, "Vantage: Optimizing video upload for time-shifted viewing of social live streams," in *Proc. ACM Special Interest Group Data Commun.*, Aug. 2019, pp. 380–393.
- [65] M. Siekkinen, E. Masala, and J. K. Nurminen, "Optimized upload strategies for live scalable video transmission from mobile devices," *IEEE Trans. Mobile Comput.*, vol. 16, no. 4, pp. 1059–1072, Apr. 2017.
- [66] *Tencent Cloud CDN*. [Online]. Available: <https://www.tencentcloud.com/products/cdn>
- [67] *Tencent Cloud Edgeone*. [Online]. Available: <https://www.tencentcloud.com/products/teo>

Tong Li (Member, IEEE) received the Ph.D. degree from the Department of Computer Science and Technology, Tsinghua University, China, in 2017. He was the Chief Engineer with Huawei in 2022. Currently, he is as an Associate Professor with Renmin University of China. His research interests include network protocols, distributed systems, and network security.

Xu Yan received the B.S. degree from the Department of Computer Science and Technology, Shandong University, Qingdao, China, in 2023. He is currently pursuing the M.S. degree with Renmin University of China, Beijing, China. His research interests include network transport protocols and congestion control.

Bo Wu received the Ph.D. degree from the Department of Computer Science and Technology, Tsinghua University, Beijing, China, in 2019. Currently, he is a Research Fellow with the Department of Cloud Architecture and Platform, Tencent Technologies. His research interests include next-generation Internet, network security, network transport, and congestion control.

Cheng Luo received the M.S. degree from Zhejiang University, Hangzhou, China, in 2015. Currently, he leads the Application Framework Group with the Department of Cloud Architecture and Platform, Tencent Technologies. His research interests include the Internet architecture and transmission optimization.

Fuyu Wang received the M.S. degree in electromagnetic field and microwave technology from UESTC, China, in 2013. Currently, he is the Chief Engineer with the Department of Cloud Architecture and Platform, Tencent Technologies. His research interests include the Internet architecture, network transport protocols, and congestion control.

Jiuxiang Zhu received the B.S. degree from the Department of Computer Science and Engineering, Central South University, Changsha, China, in 2024. He is currently pursuing the M.S. degree with Renmin University of China, Beijing, China. His research interests include high-performance web proxy and QUIC optimization.

Haoyi Fang received the B.S. degree from Renmin University of China, Beijing, China, where he is currently pursuing the M.S. degree. His research interests include networking and big data.

Xinle Du received the B.E. degree from the Department of Computer Science and Technology, Xidian University, Xi'an, China, in 2018, and the Ph.D. degree from the Department of Computer Science and Technology, Tsinghua University, Beijing, China, in 2023. He has been the Chief Engineer with the Computer Network and Protocol Laboratory, Huawei Technologies, since 2023. His research interests include networking and LLM systems.

Ke Xu (Fellow, IEEE) received the Ph.D. degree from Tsinghua University, Beijing, China. He is currently a Full Professor with the Department of Computer Science and Technology, Tsinghua University. He has published more than 200 technical articles in the research areas of next-generation Internet, blockchain systems, and network security. He received the IWQoS 24 Best Paper Award and the Distinguished Paper Award from USENIX Security 23/24, CCS 25, and NDSS 25.