FEC Hybrid Method Based on Reed–Solomon Codes for Transmitting Video Signal Over Limited Bandwidth and High Loss Rate Channels

Leonid V. Kaliuzhnyi, Viktor B. Rusalovskiy, Artem D. Ishchenko, and Vasyl M. Didenko

State University of Intelligent Technologies and Telecommunications, Odesa, Ukraine

https://doi.org/10.26636/jtit.2026.3.2637

Abstract — The paper considers the adaptive forward error correction (FEC) technique for real-time video transmission over a feedback-free channel with limited bandwidth and non-stationary interference. A hybrid Reed–Solomon coding method is proposed in GF(2^8), combining differentiated redundancy across specific frame types (I, P, and B), with adaptive code parameter control driven by physical layer channel state estimation. The channel is modeled by a three-state Markov process called Gilbert–Elliott–Jamming (GEJ), in which an explicit jamming state is added to the classical Gilbert–Elliott model. A closed-form expression for the block decoding failure probability under the erasure model is derived, and applicability bounds for representative RS configurations are determined. A redundancy adaptation function with a guaranteed upper bound on overhead is introduced together with a control algorithm. Experiments conducted under four scenarios (stable Wi-Fi, mobile 4G, tactical channel and active jamming) show that the mission success rate increases from 60% for fixed RS and 50% for SRT to 78% under sustained jamming, while the time to first frame decreases from 600 – 1200 ms to 200 ms. The method is implementable on the ESP32-S3 and ARM Cortex-M4 platforms with an encoding throughput of 4 to 8 Mbps.

Keywords — adaptive forward error correction, jamming-resistant video, Reed–Solomon codes, video streaming

1. Introduction

Video transmissions account for over 80% of global Internet traffic and continue to grow faster than the underlying infrastructure is capable of absorbing them [1], [2]. This growth is driven not only by entertainment streaming but also, increasingly, by interactive and mission-critical services such as remote surgery and telemedicine, disaster response and emergency communications, remote education, industrial teleoperation, as well as video links of unmanned aerial and robotic platforms. In such cases, an interruption of the picture is not merely a comfort issue, but a safety- and outcome-critical failure. Within this broader class, the most demanding applications are those in which the decoded picture is part of a control loop. These include the following: military and dual use unmanned aerial vehicles, tactical radio links, telemedicine, and critical infrastructure monitoring.

The efficiency of modern video delivery is based on predictive compression. Contemporary codecs such as H.264/AVC [3], H.265/HEVC [4] and the royalty-free AV1 solution [5] organize the stream into groups of pictures (GOPs) built from three frame types: intracoded I-frames, which are self-contained, as well as intercoded P- and B-frames, which are predicted from their neighbors. This creates a strong asymmetry in importance: the loss of an I frame invalidates the related entire GOP, whereas the loss of a single B-frame degrades only the frame itself. In this paper, a protection scheme that ignores this asymmetry and either wastes redundancy on the least important frames or leaves the most important ones exposed is developed.

In this paper, the impact of residual losses is ultimately judged in terms of perceived quality. Full-reference metrics such as peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) [6] quantify the fidelity against a reference, while the limitations of purely error-based measures are well documented [7]. Perceptual metrics, such as video multimethod assessment fusion (VMAF) [8], correlate more closely with subjective scores and have become a de-facto standard for streaming evaluation. These metrics provide an objective benchmark based on which the resilience of a transmission scheme is assessed.

Furthermore, combat experience gathered in 2022 – 2025 has shown that civilian real-time protocols, such as RTP, RTMP, HLS, and analog FPV channels operating in the 2.4 and 5.8 GHz bands, are not usable in zones of active electronic warfare (EW), because wideband jamming destroys both the video downlink and the UAV control uplink [9], [10]. Retransmission-based protocols, i.e., TCP and RTP/RTCP with ARQ, introduce end-to-end latency of 100 – 200 ms, which is unacceptable for real-time control and the loss of key frames (I frames) interferes with decoder synchronization.

Forward error correction (FEC) offers a good alternative. However, a fixed FEC scheme is not well matched to a non-stationary channel: at low loss rates it wastes bandwidth, while at high loss rates it does not deliver sufficient resilience.

From this dependency, we formulate the research problem as follows. The purpose of the research is to design an adaptive

FEC method, based on Reed–Solomon codes, that guarantees, in the feedback-free transmission mode, a specific target resilience level ($P_{fail} \leqslant \varepsilon$) while minimizing total redundancy ($O_{total} \rightarrow \min$) on a channel with non-stationary characteristics.

Several aspects of the problem still lack a solution in the literature. First, the classical Gilbert–Elliott model captures impulsive interference, but not long, periodic blackouts that characterize active jamming, and the analytical methodology for feedback-free channels with explicit jamming states has not been fully developed. Second, most published AFEC algorithms do not provide an analytical upper bound on total redundancy, which complicates deployment in bitrate-constrained channels. Third, the adaptive schemes usually apply a uniform FEC rate to all packets, ignoring the differing importance of I, P and B frames within the group-of-pictures (GOP) structure. Fourth, in one-way scenarios, for example, a UAV transmitting to a ground station with a jammed uplink, the algorithm has to rely only on a priori channel data and physical layer measurements.

The object of this research is the process of transmitting real-time video signals over a channel with limited bandwidth and strong interference, operating in feedback-free one-way mode. The purpose of the work is to develop an adaptive hybrid Reed–Solomon coding method that ensures $P_{fail} \leqslant 10^{-3}$ at packet loss rates of up to 20%, with a guaranteed upper bound on total redundancy equaling $O_{total} \leqslant 30\%$ and an end-to-end latency below 150 ms, with the solution remaining implementable on embedded microcontroller platforms of the ESP32-S3 and ARM Cortex-M4 class.

The research tasks set in support of this purpose are as follows:

- to derive a mathematical channel model that accounts for non-stationary jamming states,
- to derive a closed-form expression for the block decoding failure probability under the erasure model,
- to formalize the redundancy adaptation function and the differentiated frame-type FEC control algorithm,
- to compare the RS code against alternative FEC schemes,
- to implement and verify the method on an experimental software testbed.

2. Related Works

The theory of forward error correction for erasure channels is mature. Reed–Solomon codes [11], introduced as polynomial codes over finite fields, are maximum distance separable (MDS) and, therefore, optimal in their trade-off between redundancy and correction ability. Canonical treatments of their algebraic decoding are provided in [12] and [13]. For multimedia transport specifically, the low delay error protection design space – from block codes to rateless and streaming constructions – is surveyed in [14].

These foundations fix the achievable protection for a given overhead, but say nothing about how the overhead should be chosen when the channel is non-stationary, which is the main concern of the literature focusing on adaptive solutions.

As far as adaptive, channel state-driven FEC is concerned, the most recent work tunes the coding parameters to the estimated channel state. The ABRF scheme of [15] performs joint bitrate-FEC control and reports a 68 – 95% reduction in VMAF degradation compared to earlier adaptive FEC schemes. In [16], DeepRS, an LSTM network, predicts packet loss patterns and adjusts the RS parameters. Reinforcement learning controllers such as RL-AFEC [17] and R-FEC [18] select FEC parameters online from interaction with the network.

These methods report substantial quality gains, but they share two properties that matter here: they apply a single, uniform code rate to all packets of a frame, ignoring the GOP importance asymmetry, and they either rely on an accurate online loss estimate or on a learned predictor, without providing an analytical guarantee on the resulting redundancy.

In streaming, cross-layer, and neural approaches a parallel line of work targets low latency through streaming code – such as the Tambur method provided in [19] for videoconferencing – and the network adaptive error control presented in [20] and through cross-layer coordination between the application and transport layers, as in VOXEL [21]. A more radical direction replaces the codec itself with a neural solution.

The GRACE algorithm introduced in [22] reports a large reduction in unrecoverable frames and a gain in the mean opinion score at the cost of a learned encoder-decoder pair. These approaches improve resilience but move away from the lightweight, analytically tractable operation required on resource-constrained hardware.

On the channel-modeling side, the Gilbert–Elliott two-state model remains the standard abstraction of a bursty erasure channel, because it represents a single "bad" regime. It underestimates the long periodic blackouts produced by active jamming.

To the best of our knowledge, no existing methods address the combination of the three following elements jointly: an explicit jamming state inside the channel model, a closed-form analytical bound on FEC overhead, and an explicit use of GOP frame type differentiation in the adaptive control.

3. Proposed Method

3.1. Mathematical Channel Model

We modeled the channel as a three-state Markov GEJ (Gilbert–Elliott–Jamming) process with states $\{G, B, J\}$ and a transition probability matrix. In state G (good), the packet-loss probability is $p_G \ll 0.01$, in state B (bad) it lies in $p_B \in [0.05, 0.15]$ (impulsive interference), and in state J (jamming) it is $p_J \approx 1.0$ (a near-complete blackout). The transition matrix is parametrized by P_{GB}, P_{BG}, P_{GJ}, P_{JG}, P_{BJ} and P_{JB}, forming an irreducible ergodic chain in most realistic regimes. The problem of sweep jamming, which breaks stationarity, is treated separately.

Tab. 1. GEJ model parameters for typical operating scenarios.

Scenario	π_G	π_B	π_J	p_G	p_B	p_J	L_B (avg)	L_J (avg)
S1: Wi-Fi 5 GHz	0.97	0.03	0.00	0.001	0.05	—	50 pkts	—
S2: 4G/LTE	0.85	0.13	0.02	0.002	0.08	0.95	200 pkts	1000 pkts
S3: Tactical	0.60	0.30	0.10	0.005	0.15	0.99	500 pkts	5000 pkts
S4: Active EW	0.40	0.30	0.30	0.005	0.15	0.99	500 pkts	5000 pkts

The stationary distribution $\pi = (\pi_G, \pi_B, \pi_J)$ is the left principal eigenvector of the transition matrix, normalized by $\pi_G + \pi_B + \pi_J = 1$. The typical scenario parameters [10] are collected in Tab. 1. The steady-state packet loss probability is:

$$p_{\text{avg}} = \pi_G \cdot p_G + \pi_B \cdot p_B + \pi_J \cdot p_J \ . \tag{1}$$

The burst length distribution within each state follows an exponential law:

$$P(L > l) = e^{\frac{-l}{L_{avg}}} \ .$$

The model was calibrated by Baum–Welch estimation on two real, publicly available packet loss traces that bracket the regimes of interest: a benign link without jamming (the `CRAWDAD DUE` packet delivery dataset [10], overall loss $\approx$ 33.6%) and a link recorded under a periodic RF jammer (the `CRAWDAD vanetjamming2014` dataset [23], overall loss $\approx$ 77.9%). Under the Kolmogorov–Smirnov criterion, the three-state GEJ and the two-state Gilbert–Elliott models fit the benign trace comparably, achieving a KS error of 0.02 vs. 0.03 – without jamming the third state is largely redundant. Under active periodic jamming, the explicit jamming state becomes decisive (KS error 0.25 vs. 0.88), because a single bad state cannot simultaneously reproduce the short impulsive bursts and the long deliberate blackout.

3.2. Block Decoding Failure Probability

Consider a systematic Reed-Solomon code RS(n, k) over the finite Galois field GF(2^8), $n = 255$. Its error-correcting capability is:

$$t = \left\lfloor \frac{(n-k)}{2} \right\rfloor \ . \tag{2}$$

As a maximum distance separable (MDS) code, RS attains the Singleton bound $d_{\min} = n - k + 1$ and corrects up to t symbol errors with unknown positions or up to $2t$ erasures with known positions. In UDP transmission with packet sequence numbers, lost packets always have known positions, so the erasure regime applies, doubling the protection capability.

The probability of failing to decode a block of length n under independent Bernoulli losses with rate p is:

$$P_{\text{fail}}(p, n, t) = \sum_{i=2t+1}^{n} \binom{n}{i} p^i (1-p)^{n-i} \ . \tag{3}$$

Numerical values of Eq. (3) for representative RS configurations are provided in Tab. 2. From the table, we can read off applicability bounds: RS(255, 231, 12) keeps $P_{\text{fail}} \leqslant 10^{-3}$ as long as $p \leqslant 5\%$, and RS(255, 191, 32) does so up to $p \leqslant 10\%$. For p greater than 20%, the code alone is no longer enough and one has to add interleaving or a frequency hop.

Tab. 2. Block decoding failure probability for representative RS configurations (erasure model).

p [%]	Cyclic $g(x){=}x^{16}{+}x^{12}{+}x^3{+}1$	RS(255, 231, t=12)	RS(255, 191, t=32)
1	$9.2 \cdot 10^{-1}$	$3.0 \cdot 10^{-17}$	$6.4 \cdot 10^{-70}$
5	≈ 1	$1.1 \cdot 10^{-3}$	$7.8 \cdot 10^{-28}$
10	≈ 1	$5.7 \cdot 10^{-1}$	$1.2 \cdot 10^{-12}$
20	≈ 1	≈ 1	$1.9 \cdot 10^{-2}$
30	≈ 1	≈ 1	$9.5 \cdot 10^{-1}$

For burst losses of length L distributed exponentially, the block failure probability is approximated by:

$$P_{\text{loss}}^{\text{burst}} = e^{-\frac{2t}{L_{\text{avg}}}} \ . \tag{4}$$

An interleaving depth D distributes a burst of length L across D blocks, reducing the effective burst length by a factor of D:

$$P_{\text{loss}}^{\text{int}} = e^{-\frac{2tD}{L_{\text{avg}}}} \ . \tag{5}$$

Beyond the block-failure probability, the effective rate R_{eff} captures both the nominal code rate k/n and the full-block loss probability:

$$R_{\text{eff}}(p, n, t) = \frac{k}{n} \cdot \big(1 - P_{\text{fail}}(p, n, t)\big) \ . \tag{6}$$

The decoding complexity of Berlekamp–Massey + Forney over GF(2^8) is as follows:

$$C_{\text{dec}}(n, t) = O(n \cdot t + t^2) \ . \tag{7}$$

For RS(255, 223, 16) this amounts to roughly 12 000 elementary GF operations per block, which translates to 12 µs on an ARM Cortex-M4 clocked at 168 MHz or approximately 1 µs on an ESP32-S3 that features hardware GF acceleration. The effective channel throughput under a given code is:

$$C = \frac{k}{n} \cdot (1 - p) \cdot B \ , \tag{8}$$

where B is the physical radio-link capacity [bit/s].

3.3. Adaptive Redundancy Function and Differentiated Protection

The proposed adaptation function combines packet loss probability and the average burst length into a single scalar coefficient such as:

$$\kappa(p, L) = \min \left\{ 1 + 2p + \frac{L}{200}, \ 3 \right\} \ . \tag{9}$$

The cap $\kappa_{\max} = 3$ prevents unbounded overhead growth in catastrophically degraded channels. The redundancy for frame type $F \in \{I, P, B\}$ is as follows:

$$r_F = \min\left\{w_F \cdot r_{\text{base}} \cdot \kappa(p, L),\ 0.6\right\}, \tag{10}$$

where $w_I = 1.0$, $w_P = 0.3$, $w_B = 0.1$ are weights and $r_{\text{base}} \in [0.1, 0.4]$ constitutes the baseline rate driven by target VMAF.

The total redundancy for a GOP of period G with size ratio $\rho = S_I/S_P$ and $\rho \in [4, 10]$ is:

$$O_{\text{total}} = \frac{\rho \cdot r_I + (G-1) \cdot r_P}{\rho + (G-1)}. \tag{11}$$

For a typical H.264 stream with GOP $= 30$ and $\rho = 6$, Eq. (11) gives $O_{\text{total}} = 23.4\%$, against 40% for uniform protection at the same I-frame rate. This is a relative reduction of 41.5%.

3.4. RS Code Selection

Table 3 summarizes the comparative properties of the FEC schemes against key criteria. We selected the RS code as the baseline FEC algorithm, because it satisfies three requirements at once. It is MDS, i.e. it offers the optimal correction capability for a given redundancy, it tolerates burst losses up to $2\,t$ erasures, and its decoding cost remains moderate – about 12 000 GF operations per block at $t = 16$, which is well within the throughput capacity of an ESP32-S3 at 4 Mbps encoding.

The alternatives we considered fail to satisfy at least one of those criteria: RaptorQ does not provide a strict overhead bound, LDPC needs a more powerful platform than a Cortex-M4, and plain XOR-FEC offers very limited correction capability.

3.5. Parameter Sweep

To identify the best RS configuration for a given channel state, we ran a sweep over the (t, p) parametric plane. Five configurations with $n = 255$ and $k \in \{239, 223, 207, 191, 175\}$ so that $t \in \{8, 16, 24, 32, 40\}$ were tested at $p \in \{1, 5, 10, 15, 20, 25\}\%$. The test vector was the publicly available `mobile_cif.y4m` sequence (Xiph media, CIF 352×288), and the full RS encoding and decoding were performed on GF(2^8) using `reedsolo 1.7`.

This amounts to $5\times6\times30 = 900$ independent encode–decode runs with distinct seeds of random loss, requiring in total 65 min on an Intel CPU. At each of the 30 grid points, we computed P_{fail} from Eq. (3) and averaged the measured VMAF-proxy over 30 frames. UDP-baseline on the same 30 frames: $\text{VMAF}_{\text{udp}}(p) = \{71.8; 47.7; 37.9; 33.1; 28.3; 23.4\}$ for $p = \{1; 5; 10; 15; 20; 25\}$ [%]. The results are reported in the Tab. 4.

Each configuration has a clear applicability range, as one may notice in Tab. 4 and Fig. 1. The configuration $t = 8$ keeps VMAF $\geqslant 95$ up to $p = 5\%$ but then collapses at $p \geqslant 10\%$, $t = 16$ holds up to $p = 10\%$, $t = 24$ up to 15%, $t = 32$ up to 20%, and $t = 40$ covers the entire range up to 25%, at the price of 45.7% overhead. Therefore, a fixed configuration

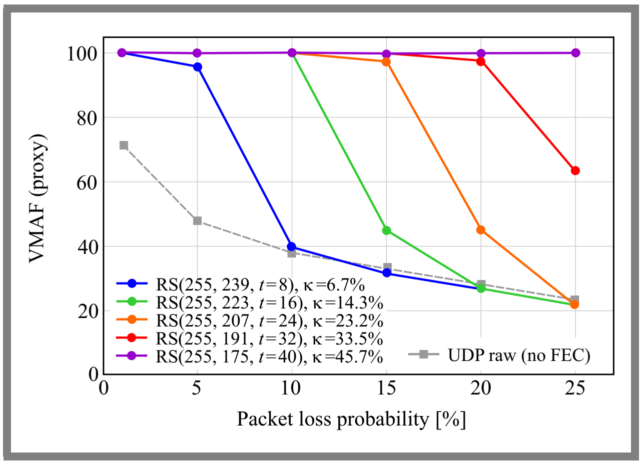

Fig. 1. Frame recovery quality for the family of RS configurations ($n = 255$, k varies; `mobile_cif` test vector).

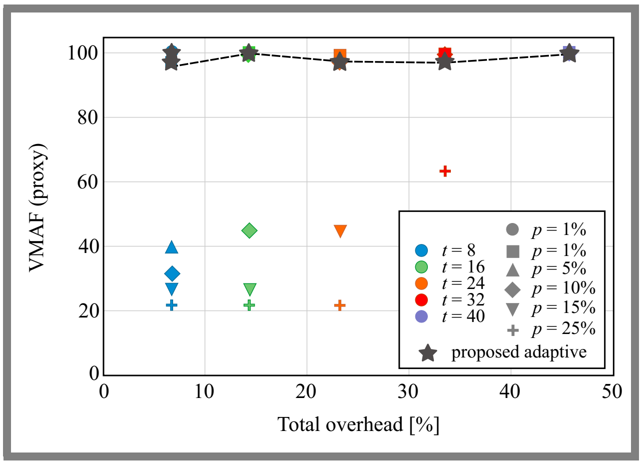

Fig. 2. Pareto chart overhead vs. quality for RS configurations.

sized for the worst case ($t = 40$) wastes 6.8× the necessary bandwidth in a stable channel at $p = 1\%$ (45.7% vs. 6.7%).

This is the main argument in favor of adaptive switching. The Pareto frontier in the overhead – VMAF plane (Fig. 2) gives the decision rule $t^*(p) = \arg\min\{\text{overhead} \mid \text{VMAF} \geqslant 95\}$: RS(255, 239, 8) at $p \in \{1, 5\}\%$; RS(255, 223, 16) at $p = 10\%$; RS(255, 207, 24) at $p = 15\%$; RS(255, 191, 32) at $p = 20\%$; RS(255, 175, 40) at $p = 25\%$.

3.6. Adaptive FEC Control Algorithm

Algorithm 1 estimates the channel state from physical layer metrics and selects RS code parameters before encoding each frame.

The complexity is $O(w + \log n)$, where w is the observation window size. For a typical $w = 30$, this is 30 comparison operations per call, which is negligible in addition to the cost of RS encoding itself.

3.7. Experimental Verification and Results

The method was verified on a software testbed written in Python 3.11, using the following libraries: `pyfinite` $\geqslant 1.9$ for GF arithmetic and `reedsolo` $\geqslant 1.7$ for systematic RS encoding. The testbed parameters are summarized in Tab. 5.

Tab. 3. Comparison of FEC schemes for real-time video transmission over channels with non-stationary interference.

Scheme	Correction capability	Decoding complexity	Burst resilience	Overhead at $t = 16$	MCU suitability
RS(255, 223)	16 (MDS)	$O(n\,t + t^2)$	High	14.3%	Yes
BCH(255, 215)	5	$O(n\,t)$	Medium	15.7%	Yes
Hamming(15, 11)	1	$O(n)$	Low	26.7%	Yes
RaptorQ	No hard bound	$O(n)$	High	Variable	Partial
LDPC	Near-Shannon	$O(n \log n)$	High	Variable	No
Plain XOR-FEC	1 per pair	$O(n)$	Low	100% (1:1)	Yes

Tab. 4. Sweep over RS(255, k, t) $\times$ p configurations.

t	k	Overhead [%]	VMAF at $p = 1\%$	VMAF at $p = 5\%$	VMAF at $p = 10\%$	VMAF at $p = 15\%$	VMAF at $p = 20\%$	VMAF at $p = 25\%$
8	239	6.7	100.0	95.9	39.7	31.5	26.7	21.7
16	223	14.3	100.0	100.0	100.0	44.9	26.7	21.7
24	207	23.2	100.0	100.0	100.0	97.4	45.1	21.7
32	191	33.5	100.0	100.0	100.0	100.0	97.4	63.3
40	175	45.7	100.0	100.0	100.0	100.0	100.0	100.0

Algorithm 1 Adaptive FEC controller

Input: packet history H (RSSI, SNR, AckRate over window W) of size w, current frame type $F \in \{I, P, B\}$, jamming-detector thresholds $\theta_{RSSI}, \theta_{SNR}$.

Output: RS code parameters (n, k, t) for encoding frame F

Start

$emp_p \leftarrow \big(1 - \text{mean}(AckRate, H)\big)$ ▷ empirical loss prob.

$med_{RSSI} \leftarrow \text{median}(RSSI, H)$

$var_{RSSI} \leftarrow \text{variance}(RSSI, H)$

$p_{hat} \leftarrow \alpha \cdot emp_p + (1 - \alpha) \cdot p_{prev}$ ▷ EWMA filter ($\alpha = 0.3$)

$L_{hat} \leftarrow \text{autocorr\_burst\_length}(H)$

$j_{flag} \leftarrow (med_{RSSI} < \theta_{RSSI}) \cup (var_{RSSI} > \theta_{SNR})$ ▷ jamming detector

$w_F \leftarrow \text{match } F : I \rightarrow 1.0\,; P \rightarrow 0.3\,; B \rightarrow 0.1$

$r_{base} \leftarrow \text{match } F : I \rightarrow 0.4\,; P \rightarrow 0.2\,; B \rightarrow 0.1$

$\kappa \leftarrow \min(1 + 2 \cdot p_{hat} + L_{hat}/200, 3)$ ▷ adaptation function

$r \leftarrow \text{clamp}(w_F \cdot r_{base} \cdot \kappa, 0.05, 0.6)$

if $j_{flag} = 1$ **then**

  $r \leftarrow \max(r, 0.5)$ ▷ jamming regime

end if

$n \leftarrow 255$ ▷ fixed for GF(256)

$t \leftarrow \lfloor r \cdot n/2 \rfloor$

$k \leftarrow n - 2 \cdot t$

return (n, k, t)

End

Each of the four scenarios S1 – S4 listed in Tab. 1 was run 30 times with different seeds from the random generator.

The testbed compares five solutions:

- UDP without FEC,
- Fixed RS(255, 223, 16),

Tab. 5. Software experimental testbed parameters.

Parameter	Value
Platform	Intel i7-12700H, 32 GB RAM, Linux Ubuntu 22.04
Language	Python 3.11
RS libraries	`pyfinite` $\geqslant$ 1.9 (GF(256)), `reedsolo` $\geqslant$ 1.7
Channel emulator	Custom GEJ implementation
Transport	UDP loopback with controlled latency / drop-rate
Test stream	Synthetic H.264-like, GOP = 30, $\rho = 6$, bitrate 0.5 – 10 Mbps
Quality metrics	VMAF (`libvmaf 2.3.1`), PSNR, SSIM
Integral metrics	MSR, TTFF, EFR
Runs per scenario	$K = 30$ with distinct seeds
Total sample size	4 scenarios $\times$ 5 solutions $\times$ 30 seeds = 600 runs

- SRT with default settings,
- WFB-NG-emulation with RS 12/8,
- The proposed method using Algorithm 1.

Aggregate results are reported in the Tab. 6.

Three observations follow from Tab. 6. In the stable channel S1, all five solutions deliver comparable quality (VMAF of approx. 91), but the proposed method reaches this with an overhead of 8 - 12%, while fixed RS uses 22 – 25%, a 2- to 3-fold saving in bandwidth. In S2 and S3, where the loss rate is 5 – 20%, the proposed algorithm is consistently 5 – 19 VMAF points above fixed RS and WFB-NG, and matches

Tab. 6. Aggregate experimental results (mean ± std over $K = 30$ channel emulator runs).

Metric	Scenario	UDP raw	Fixed RS	SRT	WFB-NG	Proposed
VMAF avg	S1	90 ±1	91±1	91 ±1	90 ±1	91 ±1
VMAF avg	S2	65 ±6	84±3	88 ±2	82 ±4	89 ±2
VMAF avg	S3	28 ±9	75±5	78 ±4	62 ±7	81 ±3
MSR [%]	S4	<10	40	50	60	78
Latency p95 [ms]	S2	60	80	350	90	95
Overhead [%]	S3	0	25	≈5 (ARQ)	33	30
TTFF [ms]	S4	—	800	1200	600	200

SRT in terms of quality, while running with 3.5 times lower latency (95 vs. 350 ms).

In S4, where the channel is under sustained EW, MSR increases from 60% with WFB-NG to 78% with the proposed method, and TTFF – after the channel exits a jamming state – drops from 600 – 1200 ms to 200 ms, due to the directed I-frame protection and fast resync mechanism.

A two-tailed Student's t-test at significance level of $\alpha = 0.05$ confirms a statistically significant separation between the proposed approach and each of the baselines in scenarios S2 – S4. The proposed method vs. fixed RS and the proposed approach vs. WFB-NG generate p-values below 0.01, and the proposed vs. SRT in S3 results in $p = 0.03$.

To illustrate the recovery effect, Fig. 3 shows the result of transmitting a test CIF frame from `mobile_cif.y4m` (352 × 288, fragmented into 149 UDP packets of 1024 bytes each) at three loss rates: $p = 5\%$, 15% and 25%.

At $p = 5\%$, UDP raw degrades to PSNR = 15.6 dB and SSIM = 0.85, leaving the frame visually unintelligible. The adaptive scheme switches to RS(255, 223, 16) at an overhead of 14.3% and recovers the frame in full (PSNR = 99 dB, SSIM = 1.00). At $p = 15\%$, the raw UDP drops to PSNR 11.9 dB, SSIM 0.76 and the adaptive scheme now selects RS(255, 191, 32) at an overhead of 33.5% and again returns PSNR 99 dB. At $p = 25\%$, the raw UDP collapses to PSNR 8.2 dB, SSIM 0.44, the adaptive scheme moves to RS(255, 175, 40) at an overhead of 45.7% and recovers the frame.

3.8. Comparison with Recent Adaptive-FEC Methods

Table 7 illustrates the proposed method against representative published adaptive FEC systems along the dimensions that matter for a resource-constrained one-way link.

The features of the present method are the explicit GOP frame type differentiation and, above all, the closed-form analytical bound on total redundancy (Sec. 3.3), which none of the learning-based systems provides and which is what makes the scheme deployable within the fixed compute and energy budget of a microcontroller class node.

The price of this simplicity is the controller's reactivity: instead of predicting variations, it relies on recent physical-layer measurements to estimate the channel. At the onset of a new interference regime, this estimate lags, causing frames transmitted during the transition to be protected according to the previous, milder state.

This is the principal limitation of the approach and the reason why a predictive channel state estimator in the light of DeepRS [16] is identified as the first line of further work (Sec. 4). We stress that the comparison provided in Tab. 7 is qualitative: the quantitative gains reported for [15], [16], and [22] were obtained under different channels, codecs, quality metrics, and are therefore not directly commensurable with the results reported here.

4. Conclusions

The work reported in this paper develops and verifies a hybrid forward error correction method based on Reed–Solomon codes with adaptive control of the coding parameters for a one-way channel with non-stationary characteristics. The main findings can be summarized as follows.

On a benign real link, the three-state GEJ and the classical two-state Gilbert–Elliott model fit the loss-burst distribution comparably (Kolmogorov–Smirnov error 0.02 vs. 0.03), whereas on a real trace recorded under active periodic jamming the explicit jamming state is decisive (0.25 vs. 0.88).

This supports including an explicit jamming state precisely when the channel is subject to deliberate interference and shows that the third state is not required under benign conditions.

The closed-form Eq. (3) for the block decoding failure probability under the erasure model gives applicability bounds for representative RS configurations: RS(255, 231, 12) keeps $P_{\text{fail}} \leqslant 10^{-3}$ as long as $p \leqslant 5\%$, and RS(255, 191, 32) does so up to $p \leqslant 10\%$. For p above 20%, interleaving or frequency hopping is needed in addition to the code itself.

The adaptation function $\kappa(p, L) = \min\{1 + 2p + L/200, 3\}$ packs the packet loss probability and the average burst length into a single scalar coefficient. Due to the cap $\kappa_{\max} = 3$, the overhead cannot grow without bound, no matter how degraded the channel becomes.

Frame-type-differentiated protection through Eq. (11) gives a total overhead $O_{\text{total}} = 23.4\%$, versus 40% for uniform protection at the same I-frame rate. This is a relative reduction of 41.5% for a typical H.264 stream with GOP = 30 and $\rho = 6$.

Fig. 3. Visual comparison of mobile CIF frame recovery under different loss rates: a) reference, b) raw UDP – horizontal stripes correspond to lost UDP packets, and c) proposed method with adaptive RS.

Tab. 7. Positioning of the proposed method relative to recent adaptive FEC systems.

Method	Adaptation driver	Rate granularity	Analytical overhead bound	Target platform	Reported gain
ABRF [15]	Joint bitrate FEC optimization	Uniform per frame	No	Server / GPU	Up to 68 – 95% less VMAF degradation
DeepRS [16]	LSTM loss predictor	Uniform per frame	No	Server / GPU	Proactive loss prediction
RL-AFEC [17]	Online reinforcement learning	Uniform per frame	No	Server	RL-tuned FEC in real time
R-FEC [18]	Online reinforcement learning	Uniform per frame	No	WebRTC endpoint	Improved QoE in WebRTC
GRACE [22]	Neural codec (loss-aware)	Codec level	No	GPU	≈95% fewer unrecoverable frames, +38% MOS
Proposed	PHY-layer estimate + GEJ state	Frame type differentiated (I/P/B)	Yes ($\leqslant$ 30%)	ESP32-S3 / ARM Cortex-M4	See Sec. 3.5 – 3.7

Experiments conducted under the four scenarios (mobile 4G, stable Wi-Fi, tactical channel, active EW), with $K = 30$ runs each, show a statistically significant advantage of the proposed method. The MSR climbs from 60% with fixed RS and 50%

with SRT to 78% with sustained jamming, TTFF falls from 600 – 1200 ms to 200 ms, and in stable channels the overhead is 2 – 3 × lower than for fixed RS.

The sweep over five configurations $t \in \{8, 16, 24, 32, 40\}$ and six loss rates $p \in \{1, 5, 10, 15, 20, 25\}$ [%] on the `mobile_cif.y4m` test vector pins down the applicability range of every configuration. It also shows that the adaptive scheme reaches VMAF $\geqslant 97$ at the minimum overhead needed in each case (between 6.7% and 45.7% depending on p), and avoids the $6.8\times$ bandwidth waste incurred by a fixed worst-case configuration.

In conclusion, the novel results of this work may be summarized as follows: (i) a closed-form expression for the block decoding failure probability under the UDP erasure model that links RS code parameters directly to channel characteristics, without requiring simulation, (ii) an adaptation function $\kappa(p, L)$ with a guaranteed upper bound, (iii) a frame-type-differentiated FEC control algorithm with a guaranteed 30% overhead bound that reduces total redundancy by 41.5% while keeping the I-frame protection at the same level, and (iv) an experimentally constructed Pareto chart of RS(255, k, t) configurations in the overhead – VMAF plane, from which five operating modes for the adaptive controller are read off.

The proposed solution applies to video surveillance of critical infrastructure, telemedicine, communication protocols for FPV UAVs and autonomous robotic systems, as well as videoconferencing services over unstable networks. The method runs on the ESP32-S3 (4 Mbps encoding) and ARM Cortex-M4 (1 – 3 Mbps) platforms, with a footprint of at most 1 KB of static data and 10 KB of dynamic memory.

Several lines of further work remain open. First, an LSTM channel-state predictor with a 5-step horizon, in the spirit of DeepRS [16], could be added on top of the present controller for proactive parameter adaptation. Second, the testbed can be extended to real H.265/AV1 streams with on-bench VMAF, PSNR and SSIM evaluation. Third, a cross-layer optimization following VOXEL [21], in which the encoder bitrate and the FEC parameters are jointly controlled, seems to be a natural next step for the future.

References

[1] Cisco, "Cisco Annual Internet Report (2018–2023)", White Paper, 2020 (`https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html`).

[2] AppLogic Networks (Sandvine), "Global Internet Phenomena Report 2024", 2024 (`https://www.sandvine.com/global-internet-phenomena-report-2024`).

[3] T. Wiegand, G.J. Sullivan, G. Bjøntegaard, and A. Luthra, "Overview of the H.264/AVC Video Coding Standard", *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 13, pp. 560–576, 2003 (`https://doi.org/10.1109/TCSVT.2003.815165`).

[4] G.J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand, "Overview of the High Efficiency Video Coding (HEVC) standard", *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 22, pp. 1649–1668, 2012 (`https://doi.org/10.1109/TCSVT.2012.2221191`).

[5] J. Han *et al.*, "A Technical Overview of AV1", *Proceedings of the IEEE*, vol. 109, pp. 1435–1462, 2021 (`https://doi.org/10.1109/JPROC.2021.3058584`).

[6] Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli, "Image Quality Assessment: from Error Visibility to Structural Similarity", *IEEE Transactions on Image Processing*, vol. 13, pp. 600–612, 2004 (`https://doi.org/10.1109/TIP.2003.819861`).

[7] Z. Wang and A.C. Bovik, "Mean Squared Error: Love It or Leave It? A New Look at Signal Fidelity Measures", *IEEE Signal Processing Magazine*, vol. 26, pp. 98–117, 2009 (`https://doi.org/10.1109/MSP.2008.930649`).

[8] R. Rassool, "VMAF Reproducibility: Validating a Perceptual Practical Video Quality Metric", *Proc. of IEEE Int. Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, Cagliari, Italy, 2017 (`https://doi.org/10.1109/BMSB.2017.7986143`).

[9] T. Stockhammer, M.M. Hannuksela, and T. Wiegand, "H.264/AVC in Wireless Environments", *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 13, pp. 657–673, 2003 (`https://doi.org/10.1109/TCSVT.2003.815167`).

[10] S. Fu and Y. Zhang, "CRAWDAD Due/packet-delivery Dataset: Packet-delivery and Delay Measurements of a Low-power Wireless Link", *IEEE DataPort*, 2015 (`https://ieee-dataport.org/collections/crawdad`).

[11] I.S. Reed and G. Solomon, "Polynomial Codes Over Certain Finite Fields", *Journal of the Society for Industrial and Applied Mathematics*, vol. 8, pp. 300–304, 1960 (`https://doi.org/10.1137/0108018`).

[12] S.B. Wicker, *Error Control Systems for Digital Communication and Storage*, Englewood Cliffs, NJ: Prentice Hall, 512 p., 1994 (ISBN 9780132008099).

[13] J.G. Proakis and M. Salehi, *Digital Communications*, 5th ed., New York: McGraw-Hill, 1150 p., 2007 (ISBN: 9780072957167).

[14] A. Badr, A. Khisti, W.-T. Tan, and J. Apostolopoulos, "Perfecting Protection for Interactive Multimedia: A Survey of Forward Error Correction for Low-delay Interactive Applications", *IEEE Signal Processing Magazine*, vol. 34, pp. 95–113, 2017 (`https://doi.org/10.1109/MSP.2016.2639062`).

[15] S. Cheng, H. Hu, and X. Zhang, "ABRF: Adaptive BitRate–FEC Joint Control for Real-time Video Streaming", *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 33, pp. 5212–5226, 2023 (`https://doi.org/10.1109/TCSVT.2023.3248164`).

[16] S. Cheng, H. Hu, X. Zhang, and Z. Guo, "DeepRS: Deep-learning Based Network-adaptive FEC for Real-time Video Communications", *2020 IEEE Int. Symposium on Circuits and Systems (ISCAS)*, Seville, Spain, 2020 (`https://doi.org/10.1109/ISCAS45731.2020.9180974`).

[17] K. Chen *et al.*, "RL-AFEC: Adaptive Forward Error Correction for Real-time Video Communication based on Reinforcement Learning", *Proc. of the 13th ACM Multimedia Systems Conference*, pp. 96–108, 2022 (`https://doi.org/10.1145/3524273.3528184`).

[18] I. Lee *et al.*, "R-FEC: RL-based FEC Adjustment for Better QoE in WebRTC", *Proc. of the 30th ACM International Conference on Multimedia*, pp. 2948–2956, 2022 (`https://doi.org/10.1145/3503161.3548370`).

[19] M. Rudow *et al.*, "Tambur: Efficient Loss Recovery for Videoconferencing via Streaming Codes", *Proc. of 20th USENIX Symposium on Networked Systems Design and Implementation*, pp. 953–971, 2023 (`https://www.usenix.org/conference/nsdi23/presentation/rudow`).

[20] S.L. Fong *et al.*, "Low-latency Network-adaptive Error Control for Interactive Streaming", *Proc. of the 27th ACM International Conference on Multimedia*, pp. 438–446, 2019 (`https://doi.org/10.1145/3343031.3350942`).

[21] M. Palmer *et al.*, "VOXEL: Cross-layer Optimization for Video Streaming with Imperfect Transmission", *Proc. of the 17th International Conference on emerging Networking EXperiments and Technologies*, pp. 359–374, 2021 (`https://doi.org/10.1145/3485983.3494864`).

[22] Y. Cheng *et al.*, "GRACE: Loss-resilient Real-time Video through Neural Codecs", *ArXiv*, 2024 (`https://doi.org/10.48550/arXiv.2305.12333`).

[23] O. Puñal, C. Pereira, A. Aguiar, and J. Gross, "CRAWDAD Dataset Uportorwthaachen/vanetjamming2014: IEEE 802.11p Packet Traces Under Constant, Periodic and Reactive RF Jamming", *IEEE DataPort*, 2014 (`https://doi.org/10.15783/C7Q306`).

Leonid V. Kaliuzhnyi, Ph.D. Candidate
Department of Intelligent
Technologies and Telecommunications
https://orcid.org/0009-0009-8445-1143
E-mail: leonid.kalyuzhnyy@gmail.com
State University of Intelligent Technologies and Telecommunications, Odesa, Ukraine
https://suitt.edu.ua

Viktor B. Rusalovskiy
Department of Intelligent
Technologies and Telecommunications
https://orcid.org/0009-0003-6134-531X
E-mail: it.ddo.od@gmail.com
State University of Intelligent Technologies and Telecommunications, Odesa, Ukraine
https://suitt.edu.ua

Artem D. Ishchenko
Department of Intelligent
Technologies and Telecommunications
https://orcid.org/0009-0000-5161-7818
E-mail: artem.ischchenko@gmail.com
State University of Intelligent Technologies and Telecommunications, Odesa, Ukraine
https://suitt.edu.ua

Vasyl M. Didenko
Department of Intelligent
Technologies and Telecommunications
https://orcid.org/0009-0008-9517-7963
E-mail: vahaod@gmail.com
State University of Intelligent Technologies and Telecommunications, Odesa, Ukraine
https://suitt.edu.ua