Journal of
Marine Science and Engineering

MDPI

Article

Channel-Adaptive Joint Selection of FEC Scheme, Rate, and Segment Size for Short-Block Underwater Acoustic Communication

Seunggyu Kim, Saeyong Park and Taeho Im *

Division of Information and Communication Engineering, Hoseo University, Asan 31499, Republic of Korea; 20255412@vision.hoseo.edu (S.K.); 20245812@vision.hoseo.edu (S.P.)
* Correspondence: taehoim@hoseo.edu; Tel.: +82-54-540-9642

Abstract

Underwater acoustic (UWA) sensor and control links carry mostly short messages (below ∼10 kB) over time-varying, low-SNR multipath channels, which is a regime where forward error correction (FEC) operates on short, finite blocks where a single static configuration is inefficient. Adaptive schemes for these links typically adjust the modulation order and code rate; the payload segment (block) size—which, together with the code rate, sets the coded block length that governs the finite-blocklength penalty for short messages—is seldom adapted per transmission jointly with the choice of FEC scheme on a like-for-like footing. We propose a per-transmission controller that jointly selects the FEC scheme, code rate, and segment size from a prediction of the near-term channel state, which is paired with a like-for-like short-block benchmark of LDPC, list-decoded polar, BCH, Reed–Solomon, convolutional, and turbo codes. No single code dominates: under a unified ARQ goodput metric, the reliability–throughput frontier has a crossover that shifts with the channel, so the optimal FEC choice is channel-dependent. Across our experiments, the segment-size degree of freedom is the dominant throughput lever, capturing essentially all of the adaptation gain at low-to-mid SNR and over fading; switching the FEC family adds a further, bounded gain only where the frontier crosses between families (up to 10% at high-SNR AWGN, polar to RS). The joint controller essentially matches a fair single-family adaptive baseline off the crossover (to within a negligible prediction-overhead margin) and exceeds it at that crossover, beats a fixed, no-CSI code by up to 16%, and captures 91–99% of an oracle; a lightweight persistence predictor matches a learned LSTM for the first-order channel-state model studied. A statistics-driven replay using measured-channel parameters, and a recorded-channel replay over the public Watermark benchmark, preserve the same family ordering.

Keywords: underwater acoustic communication; forward error correction; adaptive coding; finite blocklength; payload segmentation

Check for updates

Academic Editor: Weicheng Cui

Received: 9 June 2026
Revised: 26 June 2026
Accepted: 1 July 2026
Published: 3 July 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Underwater acoustic (UWA) links must operate over time-varying multipath channels with non-negligible Doppler spread and low signal-to-noise ratios, conditions under which forward error correction (FEC) is indispensable for reliable delivery [1]. Much of the traffic on underwater sensor networks and autonomous-underwater-vehicle control links consists of short command, telemetry, and status messages rather than bulk data, so FEC operates on

short, finite-length blocks where the finite-blocklength penalty is severe and a single static code configuration is inefficient across the rapidly changing channel [2,3]. To track these variations, adaptive modulation and coding (AMC) schemes adjust the modulation order and code rate to the estimated channel state [4–6]; recent schemes are increasingly driven by machine learning and designed to tolerate the outdated channel-state information that long propagation delays impose. A complementary line of work adapts redundancy at the packet layer through feedback-driven packet coding [7], while a separate body of studies characterizes individual short-block FEC families in the underwater and finite-blocklength regimes [8–11]. The block (or codeword) length is not entirely fixed in prior art—rate-compatible and rateless constructions and energy-aware code-size optimization vary the effective codeword length within a single family [12,13], and packet-layer schemes vary the packet size [7,14]—but the payload segment (block) size is not adapted *jointly with a per-transmission choice of FEC scheme* on a like-for-like short-block footing, even though the segment size sets the FEC block length that governs the finite-blocklength penalty for the short messages that dominate underwater links.

We argue that this leaves a substantial degree of freedom unexploited: jointly choosing the segment size together with the FEC scheme and code rate, conditioned on the predicted channel state, should yield reliability and efficiency gains beyond single-family rate adaptation wherever the optimal code family changes with the channel. To this end, we propose a channel-state-aware controller that, from a prediction of the near-term UWA channel state (received signal-to-noise ratio [SNR], delay spread, and Doppler), uses a lightweight policy to jointly select, per transmission, the FEC scheme, the code rate, and the payload segment size before encoding. We accompany this controller with a controlled, like-for-like short-block FEC benchmark for messages below 10 kB.

1.1. Research Questions

We organize this paper around three questions. RQ1: How do candidate short-block FEC families trade reliability against throughput, and does the best family depend on the channel state? RQ2: Does jointly adapting the FEC scheme, code rate, and segment size outperform single-axis adaptation (a fixed FEC family with its rate/segment adapted), and where does the advantage come from? RQ3: How does the joint gain depend on feedback delay and prediction error, and is a learned channel-state predictor needed over a lightweight one?

1.2. Contributions

(i) A per-transmission, channel-state-aware controller jointly selects the FEC scheme, code rate, and payload segment size from a predicted near-term channel state, subsuming rate-only adaptive modulation and coding. We find the segment-size degree of freedom to be the dominant throughput lever, with the FEC-*family* switch contributing a bounded additional gain only where the reliability–throughput frontier crosses between families. (ii) Using a controlled, like-for-like short-block FEC benchmark of six families (LDPC, cyclic redundancy check [CRC]-aided list-decoded polar, BCH, Reed–Solomon, convolutional, turbo) on a common information-bit axis over AWGN with the four fading-robust families (LDPC, polar, BCH, RS)—and, for completeness, the soft convolutional code—additionally evaluated over an orthogonal frequency-division multiplexing (OFDM) link over a tapped-delay-line (TDL) fading channel, we establish that no single code dominates and that the reliability–throughput crossover is channel-dependent (RQ1). (iii) Our policy-evaluation focus shows that joint adaptation matches a fair single-family adaptive baseline and exceeds it where the optimal action is to switch FEC family (RQ2), capturing 91–99% of an oracle, together with a robustness analysis (RQ3) of feedback delay and prediction error in which

a lightweight persistence predictor matches a learned long short-term memory (LSTM) network for the first-order channel-state model studied.

The remainder of this paper reviews related work, then formalizes the channel, segmentation, and decision problem (Section 3), presents the channel-adaptive joint controller (Section 4), details the benchmark and evaluation methodology (Section 5), reports results for RQ1–RQ3 (Section 6), and discusses limitations and deployment before concluding.

2. Related Work

We organize prior work along four axes—adaptive modulation and coding, adaptive packet-layer coding, short-block FEC characterization, and channel prediction with finite-blocklength theory—and locate our contribution against each. Table 1 summarizes the comparison.

2.1. Adaptive Modulation and Coding (AMC)

The dominant approach to time-varying underwater links is to adapt the modulation order and the code rate to the estimated channel state, increasingly with machine-learning or reinforcement-learning controllers [4,5,15,16]. Several of these explicitly target the underwater channel's defining difficulty—long propagation delay that renders the feedback channel state outdated—by learning policies robust to that staleness [6,17]. These methods typically hold the FEC *block length* (equivalently, the payload segment size) fixed and adapt only the modulation and rate. Our controller subsumes rate adaptation but adds the segment size and the FEC *scheme* as jointly selected variables. Section 6.6 shows that the segment-size lever carries most of the adaptation gain and that the additional FEC-*family* degree of freedom yields goodput beyond what single-family adaptation reaches where the reliability–throughput frontier crosses families while matching it elsewhere.

2.2. Adaptive Packet-Layer Coding

A complementary line adapts redundancy above the physical layer, tuning packet-coding rates or packet sizes through feedback to improve reliability and throughput [7,14,18]. This operates at the packet/network layer and is largely orthogonal to the finite-blocklength behavior of the underlying FEC block. We instead act at the FEC layer, where the segment size sets the FEC block length that governs the finite-blocklength penalty, and we couple it with the code choice rather than with packet-level redundancy.

2.3. Short-Block FEC Characterization

A separate body of work characterizes individual FEC families in the short-block and underwater regimes—short LDPC and raptor-like LDPC [9,11], polar codes for underwater/OFDM links [10,19,20], and comparative studies of LDPC, turbo, and convolutional codes [8]. These studies typically fix a single code family and quantify its performance. We provide instead a unified, like-for-like benchmark across families on a common information-bit axis (Sections 6.1–6.5) and make the per-state best code the *output* of a state-conditioned selection policy (an interpretable benchmark-table argmax optionally distilled into a compact classifier) rather than a design-time decision.

2.4. Channel Prediction and Finite-Blocklength Theory

Underwater channel prediction—forecasting the time-varying channel or its link-quality features—has been pursued with classical and deep-learning predictors [21–24], and our controller relies on such a prediction to act ahead of the feedback delay. Our RQ3 study (Section 6.7), however, finds that for the channel-state process considered here, a lightweight persistence predictor is as effective as a learned LSTM, so we keep the predictor simple. The finite-blocklength penalty that motivates segment-size adaptation is grounded

in normal-approximation theory for short packets [2,3,25] and in recent channel-coding surveys [26,27]; we use these as the theoretical anchor for why short payloads make the FEC choice channel-dependent.

Table 1. Positioning against representative prior work. "Length-adaptive" indicates whether a block/codeword/packet length is adapted, and at which layer; "Cross-family joint" indicates whether the FEC *scheme* (family) is selected jointly with rate and segment size per transmission. The gap we target is the joint, cross-family, per-transmission selection at the FEC-block layer.

Cluster	Representative Work	Layer	Adapts	Length-Adaptive	Cross-Fam. Joint
AMC (ML/RL)	[4–6]	physical	mod. + rate	no (fixed block)	no
Adaptive packet coding	[7,14]	packet	packet rate/size	packet layer	no
Rateless/energy-aware	[12,13]	physical (FEC)	codeword length	within family	no
Short-block FEC studies	[8,10,11]	physical	(fixed code)	no	no
AMC surveys/benchmarks	[26,27]	—	survey	—	—
This work	—	physical (FEC)	**scheme + rate + segment**	**FEC block**	**yes**

Bold marks this work—the proposed joint, cross-family, per-transmission selection at the FEC-block layer—to contrast it with the prior clusters.

2.5. Positioning

To our knowledge, the joint, channel-state-conditioned adaptation of payload segment size together with the FEC scheme and rate—evaluated on a unified short-block (<10 kB) underwater benchmark—is not claimed by any of these clusters individually. In particular, the authors' prior work applied channel prediction to adaptive polar–raptor coded OFDM telemetry [28] within a single code construction; this paper differs in treating the FEC *family* and the payload *segment size*—not only the coded rate—as joint, per-transmission adaptation variables, and in grounding that choice on a like-for-like multi-family short-block benchmark.

3. System Model and Problem Formulation

3.1. Underwater Acoustic Channel

We consider a point-to-point underwater acoustic (UWA) link whose passband is modeled, after baseband conversion, as a time-varying multipath channel. Over a transmission of duration short relative to the channel's coherence time, the channel is represented by a tapped delay line (TDL) with L resolvable paths,

$$h(\tau, t) = \sum_{\ell=1}^{L} a_\ell(t)\, \delta(\tau - \tau_\ell), \tag{1}$$

where τ_ℓ represents the path delays spanning a delay spread τ_{rms} and $a_\ell(t)$ are the time-varying complex path gains whose average powers follow an exponential power-delay profile (PDP) and each of which follows a Rayleigh or Rician distribution with a Doppler power spectrum of one-sided spread f_d. This delay–Doppler description matches measured UWA channels in which both the multipath delay spread and the Doppler caused by platform and surface motion vary on the time scale of seconds [21,29,30]. The two modulation arms (single-carrier and OFDM) and the receiver front-end that operates on this channel—including the genie-CSI assumption and the per-subcarrier frequency-domain processing—are specified once in Section 5. We summarize the instantaneous channel by a state vector

$$\mathbf{c}(t) = \big(\gamma(t),\ \tau_{\mathrm{rms}}(t),\ f_d(t)\big), \tag{2}$$

i.e., the received signal-to-noise ratio (SNR) γ, the delay spread, and the Doppler spread, which together govern the achievable reliability of a coded block. The long propagation delay of the acoustic channel (on the order of seconds per kilometer) means that any channel-state information (CSI) report fed back to the transmitter is necessarily outdated, which is a constraint that we make explicit in Section 4 and quantify in the RQ3 study.

3.2. Transmitter Chain and Payload Segmentation

A message of B payload bytes—with $B < 10$ kB, the regime that dominates underwater sensor and control traffic—is partitioned into $N_s = \lceil 8B/s \rceil$ segments of at most s information bits each. Each segment is the information word of one FEC block: a code $\mathcal{C}$ of rate $R = k/n$ maps $k = s$ information bits to n coded bits, which are mapped by a modulator (single-carrier phase-shift keying [PSK] or OFDM with PSK subcarriers, per the frames of Table 2) to channel symbols and transmitted. The receiver forms per-bit log-likelihood ratios (LLRs) from the demodulated symbols and decodes each block independently. The segment size $s = k$ is therefore the information block length seen by the FEC—which, at code rate R, fixes the coded block length $n = k/R$ that sets the finite-blocklength penalty—and it is the lever this paper treats as a first-class adaptation variable rather than a fixed constant.

Table 2. System parameters (OFDM frames from the authors' experiments; single-carrier arm parallel at half-bandwidth).

Parameter	F1	F2	F3	F4
Center freq. [kHz]	14	14	30	30
OFDM bandwidth [kHz]	12	12	16	4
FFT size	1024	2048	1024	512
Subcarrier spacing [Hz]	11.71	5.85	15.62	7.81
CP length [ms]	21.4	21.4	22.6	22.6
Useful symbol [ms]	85.3	170.7	64.0	128.0
Sampling rate [kHz]	200	200	200	200
Pilot grid (t, f)	(3, 3)	(3, 3)	(3, 3)	(3, 3)
Subcarrier mod.	BPSK/QPSK	BPSK/QPSK	BPSK/QPSK	BPSK/QPSK
Single-carrier BW [kHz]	6	6	8	2

Segmentation is not free: every segment carries a header and a CRC, and the OFDM waveform additionally spends a fixed fraction of its resource elements on the $(3, 3)$ time–frequency pilot grid of Table 2. Writing o for the per-segment header-and-CRC overhead in bits—the pilot cost is charged separately, once, through the usable-data fraction φ of Equation (5), and is not double-counted in o—the effective rate delivered to the application for a payload split into N_s segments is

$$R_{\text{eff}} = \frac{8B}{\sum_{i=1}^{N_s} (n_i + o)}, \tag{3}$$

so that smaller segments cut last-segment padding and retransmission granularity but pay a higher relative overhead and—because they shorten the coded block $n = k/R$—a larger finite-blocklength penalty. This tension is the reason segment size cannot be optimized in isolation from the code.

3.3. Finite-Blocklength Penalty

At the short block lengths induced by sub-10 kB payloads, the maximum coding rate achievable at a target error probability lies well below the channel capacity by a gap that grows as the block length shrinks—the finite-blocklength penalty [2,3,25]. For a block of

length n at block-error probability ϵ over an AWGN channel, the normal approximation gives the achievable rate as

$$R^{\star}(n,\epsilon) \approx C - \sqrt{\tfrac{V}{n}}\, Q^{-1}(\epsilon) + \mathcal{O}\Big(\tfrac{\log n}{n}\Big), \tag{4}$$

with C representing the capacity, V the channel dispersion, and Q^{-1} the inverse Gaussian tail function; over fading the penalty is larger and depends on $\mathbf{c}(t)$ [2,26]. The practical consequence, confirmed empirically in Section 6, is that the ranking of FEC families and the best $(\mathcal{C}, R, s)$ choice both depend on the operating point, which is precisely what a fixed configuration cannot track.

3.4. Objective and Decision Problem

Let the goodput G of a transmission be the expected number of successfully delivered information bits per channel use (channel symbol). Under the segmentation of Section 3.2, a B-byte payload occupies $N_s = \lceil 8B/k \rceil$ coded blocks, each carrying $n + o$ coded and overhead bits—equivalently $(n+o)/(\varphi \log_2 M)$ channel uses after modulation and pilot insertion—and a block that fails its CRC is retransmitted—a stop-and-wait/selective-repeat automatic repeat request (ARQ) appropriate to a reliable control link—so on a channel with block-error rate BLER, a block costs $1/(1-\mathrm{BLER})$ transmissions on average and *every* block is eventually delivered (no message-level collapse to zero). The resulting ARQ effective goodput is

$$G = \frac{8B\,\varphi\,(1-\mathrm{BLER})\,\log_2 M}{N_s\,(n+o)}, \tag{5}$$

where $R = k/n$ is the code rate, M the modulation order, BLER the block-error rate at the operating channel state, and φ the usable data fraction of the waveform ($\varphi = 1$ on the single-carrier arm; $\varphi = 8/9$ for the OFDM $(3,3)$ pilot grid of Table 2). Equation (5) is the single goodput metric reported throughout Section 6: it charges each transmission its true per-segment overhead o and its retransmission cost, and it penalizes a payload that does not fill its last segment (since $N_s k \geq 8B$). As BLER $\to 0$ for an overhead-free, segment-aligned payload, it reduces to the familiar block-level spectral efficiency $R\,(1-\mathrm{BLER})\,\log_2 M$, which we retain only as a high-SNR reference; the overhead-charged ceiling at BLER $= 0$ is $R_{\mathrm{eff}}\,\varphi\,\log_2 M$ with R_{eff} of Equation (3). Where end-to-end timing matters—the latency budget below—the corresponding throughput in bits per second is $G_t = (1-\mathrm{BLER})\,8B/T_{\mathrm{att}}$, where T_{att} is the time of a *single* transmission attempt (codeword plus per-segment overhead) and the retransmission cost is carried by the $(1-\mathrm{BLER})$ factor, so the expected delivery time is $T_{\mathrm{att}}/(1-\mathrm{BLER})$; this timing form is used only for latency-feasibility checks—not for the goodput reported in Section 6. We also track the energy per delivered bit, $E_b^{\mathrm{del}} = E_{\mathrm{tx}}/[(1-\mathrm{BLER})\,8B]$ [31], and the end-to-end latency (encoding + propagation + decoding + retransmission), whose decoder component differs markedly across families (Section 4).

The decision problem is, per transmission, to choose an action $\mathbf{x} = (\mathcal{C}, R, s) \in \mathcal{X}$—the FEC scheme, its code rate, and the segment size—so as to maximize the expected goodput subject to a latency budget $T_{\max}$ and conditioned on the channel state that will hold when the packet is received:

$$\mathbf{x}^{\star}(\mathbf{c}) = \arg\max_{\mathbf{x}\in\mathcal{X}} \mathbb{E}\big[G \mid \mathbf{c}, \mathbf{x}\big] \quad \text{s.t.} \quad T_{\mathrm{dec}}(\mathbf{x}) + T_{\mathrm{att}}(\mathbf{x}) \leq T_{\max}. \tag{6}$$

Because the true $\mathbf{c}$ at reception is unknown at transmit time (outdated feedback), the controller of Section 4 acts on a *predicted* state $\hat{\mathbf{c}}$. The action space $\mathcal{X}$ spans the FEC families benchmarked in Section 6, the code rates from $R = \frac{1}{3}$ (turbo) through $\frac{1}{2}$ and $\frac{2}{3}$ (the strong

codes) up to the high-rate algebraic variants ($R = 0.71$–0.78 for BCH and 0.73–0.81 for RS), within each family's feasible range, and a payload segment (information-block) size $s = k$ that the action set realizes over ≈44–256 information bits (the k values of Table 3); the segment size is thus selected jointly with the family and rate rather than from an independent grid.

Table 3. FEC action space (decoder and soft/hard decision per family). All (n, k) are in bits except Reed–Solomon, in GF(2^4)/GF(2^5) symbols [$(15, 11)$, $(31, 25)$, i.e., $(60, 44)$, $(155, 125)$ bits]. Convolutional/turbo list the mother-code rate. Algebraic codes use hard-decision decoding here; their soft-decision (Chase-II) variant is studied separately (Section 6.5).

Family	Configs (n, k)	Rates	Decoder	Decision	Channels
Polar CA-SCL (GA)	(128, 64), (256, 128), (512, 256), (128, 86), (256, 171)	$\frac{1}{2}$, $\frac{2}{3}$	CRC-aided SCL	soft	AWGN/TDL
Short LDPC	(96, 50), (192, 98)	0.52, 0.51	belief propagation	soft	AWGN/TDL
BCH	(63, 45), (63, 51), (127, 99)	0.71–0.78	bounded-distance	hard	AWGN/TDL
Reed–Solomon	(15, 11), (31, 25)	0.73–0.81	bounded-distance	hard	AWGN/TDL
Convolutional $K = 7$	$k = 128, 256$	$\frac{1}{2}$	unquantized Viterbi	soft	AWGN
Turbo	$k = 256$	$\frac{1}{3}$	iterative BCJR	soft	AWGN

4. Proposed Method—Channel-Adaptive Joint Controller

We propose a per-transmission controller that predicts the near-term underwater channel state and then jointly selects the FEC scheme, code rate, and payload segment size that maximize the expected goodput under the latency budget. The controller has two parts: a channel-state predictor (Section 4.1) and a lightweight selection policy built from the E1 benchmark (Section 4.2) coupled by an outdated-CSI margin (Section 4.3). The per-transmission loop is illustrated in Figure 1 and specified as Algorithm 1 in Section 4.5.

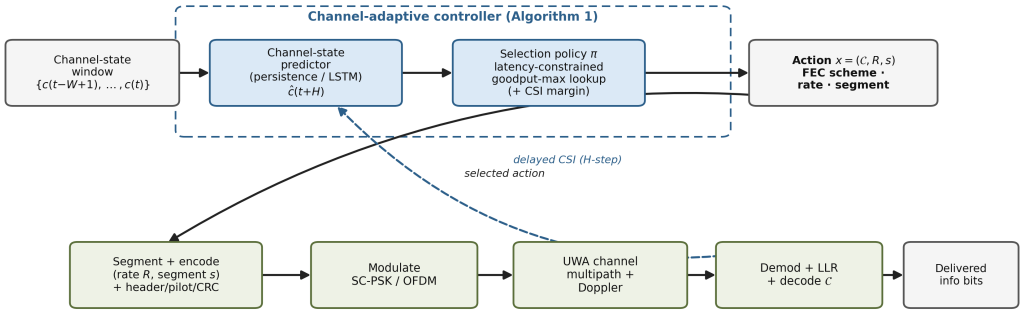


Figure 1. The channel-adaptive joint controller (Algorithm 1). From recent channel-state observations, the predictor forms an H-step estimate $\hat{\mathbf{c}}(t+H)$, and the policy π returns the latency-constrained, goodput-maximizing action $x = (\mathcal{C}, R, s)$ (FEC scheme, rate, segment size) from the E1 table, which is tempered by an outdated-CSI margin.

4.1. Channel-State Prediction

The long propagation delay of the acoustic link makes the most recent channel-state report stale by the time a packet arrives, so the controller acts on a prediction of the state H transmissions ahead, $\hat{\mathbf{c}}(t + H)$, formed from a window of the W most recent observations $(\gamma, \tau_{\mathrm{rms}}, f_d)$. Our operating choice is a lightweight persistence predictor, $\hat{\mathbf{c}}(t + H) = \mathbf{c}(t)$—the most recent observation—which for a temporally correlated channel state is close to the conditional-mean predictor and costs nothing beyond a register read, which is an attractive property for a low-power underwater modem. We also evaluated a learned alternative: a two-layer long short-term memory (LSTM) network (hidden size 64) with a fully connected head over the three state features, $\hat{\mathbf{c}}(t + H) = f_{\boldsymbol{\theta}}(\mathbf{c}(t-W+1), \ldots, \mathbf{c}(t))$, with window

$W = 8$ and trained by supervised MSE regression (Adam, learning rate 10^{-2}, 40 epochs) on channel-state series with a strict train/test split across channel realizations [17,22,23]. It did not improve goodput over persistence for the channel-state process studied here, so we adopt persistence (Section 6.7). We sweep the horizon $H \in \{0, 1, 2, 5\}$ and an explicit prediction-error level to characterize robustness to outdated CSI (Section 6.7).

4.2. Joint Selection Policy

Given the predicted state $\hat{\mathbf{c}}$, the policy returns the action $\pi(\hat{\mathbf{c}}) = \mathbf{x} = (\mathcal{C}, R, s)$ that solves the per-transmission problem of Section 3.4. The policy is built directly from the controlled E1 benchmark: E1 measures, for every action $\mathbf{x} \in \mathcal{X}$ and a grid of channel states, the resulting BLER, goodput, decoder latency, and energy, yielding a state-to-action map

$$\pi^{\star}(\mathbf{c}) = \arg\max_{\mathbf{x} \in \mathcal{X}} \widehat{G}(\mathbf{c}, \mathbf{x}) \quad \text{s.t.} \quad \widehat{T}_{\text{dec}}(\mathbf{x}) + T_{\text{att}}(\mathbf{x}) \leq T_{\max}, \tag{7}$$

where $\widehat{G}$ and $\widehat{T}_{\text{dec}}$ are the goodput and latency interpolated from the E1 table. In its first-pass form, the policy is this argmax lookup (optionally distilled into a small supervised classifier for compactness); contextual-bandit or lightweight reinforcement-learning refinements are left as a comparison point. This construction makes the policy interpretable and cheap to evaluate on an embedded modem, and it lets the controller move all three degrees of freedom at once, which is in contrast to data-driven AMC schemes that adapt only the modulation order and code rate while holding the block length fixed [1,4,6]. Because the E1 ordering of families is channel-dependent (Section 6, RQ1), the resulting map is itself channel-dependent: it selects rate-$\frac{1}{2}$ strong codes (e.g., CRC-aided successive-cancellation list [SCL] polar or short LDPC) at low SNR and high-rate algebraic codes (BCH/RS) once the channel is good enough, and it adjusts the segment size to balance the finite-blocklength penalty against per-segment overhead.

4.3. Outdated-CSI Handling and Overhead Accounting

The controller evaluation includes two conservative safeguards. First, prediction error and feedback delay are absorbed by acting on $\hat{\mathbf{c}}(t + H)$ and, optionally, by selecting on a downward-adjusted SNR estimate $\hat{\gamma}_{\text{safe}} = \hat{\gamma} - \alpha\, \hat{\sigma}_{\gamma}$, where $\hat{\sigma}_{\gamma}$ is the predictor's recent error standard deviation and $\alpha \geq 0$ is a robustness knob that trades a little throughput for a conservative, lower-rate choice. Setting $\alpha = 0$ recovers the plain goodput-maximizing policy and is the value used for *all* reported RQ2/RQ3 results, so the gains below are not inflated by margin tuning; a positive α is a deployment safeguard motivated by the RQ3 sensitivity. That study reports how the joint gain degrades with prediction error and feedback delay and what minimum horizon preserves it. Second, every action's goodput is computed with its true cost—the segment header and CRC charged through the per-segment overhead o, and the OFDM $(3, 3)$ pilot grid of Table 2 charged once through the usable-data fraction φ (Section 3.2)—so that the apparent benefit of small segments is charged its real cost.

4.4. Latency Budget

The latency budget $T_{\max}$ is a first-class constraint because the decoder cost varies by orders of magnitude across the action space. CRC-aided SCL polar decoding scales with the list size L, and turbo decoding scales with its iteration count, whereas algebraic BCH/RS decoding and the short-LDPC belief-propagation decoder are comparatively cheap; we therefore record a decoder latency proxy per action and report it alongside goodput. Where decoding latency would otherwise dominate the budget, the literature offers concrete reductions—fast simplified SCL decoding for polar codes [32] and low-latency short-blocklength LDPC decoder hardware [33]—which we treat as evidence that

the latency cost of the strong codes is implementable rather than prohibitive while still charging it in $T_{\max}$.

4.5. Per-Transmission Algorithm

The controller adds one predictor evaluation (a single state read for the persistence predictor) and one table lookup per transmission. These are both negligible against the acoustic propagation delay, so its computational overhead does not enter the latency budget in any practically relevant way.

Algorithm 1: Per-transmission joint adaptation

1. Observe the recent channel-state window $\{\mathbf{c}(t-W+1), \ldots, \mathbf{c}(t)\}$.
2. Predict $\hat{\mathbf{c}}(t+H)$ with the channel-state predictor (Section 4.1; persistence by default).
3. Select $\mathbf{x} = (\mathcal{C}, R, s) = \pi(\hat{\mathbf{c}})$ by the latency-constrained goodput-maximizing lookup (Section 4.2), applying the optional CSI margin of Section 4.3 ($\alpha = 0$ in the reported runs).
4. Segment the B-byte payload into $N_s = \lceil 8B/s \rceil$ blocks; encode each with $\mathcal{C}$ at rate R; append the segment header and CRC; insert OFDM pilots according to the frame grid.
5. Modulate and transmit; on the return path, update the channel-state observations for the next prediction.

5. Materials and Methods

5.1. Channel and Modulation

We evaluate on two channels; the channel model itself is specified once in Section 3 and not repeated here. The first is an additive white Gaussian noise (AWGN) single-carrier PSK link, which is used to isolate code behavior from fading. The second is an OFDM link over the time-varying TDL channel of Section 3, with its exponential PDP, per-tap Rayleigh/Rician fading, and first-order Doppler process. The receiver front-end—common to both arms and held identical across all FEC families—is described below and summarized in Figure 2. The TDL parameters are swept over the delay spread, Doppler, and Rician factor; the synthetic baseline is an eight-tap exponential-PDP channel at Rician $K = +5$ dB (Section 6.8) with the full swept grid specified in the configuration described in this section.

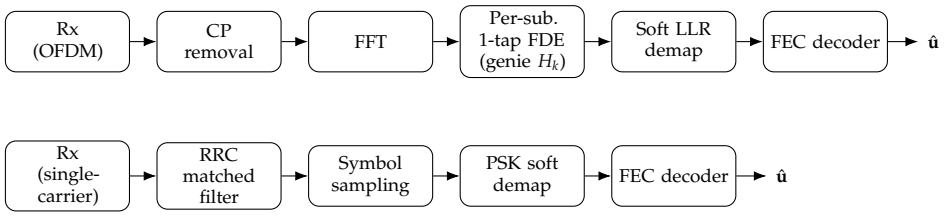


Figure 2. Receiver front-end, identical across all FEC families (top: OFDM arm; bottom: single-carrier arm). The OFDM arm performs CP removal, FFT, and per-subcarrier single-tap FDE with *genie* CSI ($\mathrm{conj}(H_k)Y_k$); no pilot estimation, Doppler/CFO compensation, or turbo equalization (Section 7.2).

The OFDM parameters (Table 2) follow the authors' prior underwater acoustic OFDM systems [28]. The four frame configurations are chosen to span the operating envelope of those field deployments along the two axes that matter most for a UWA OFDM link—center frequency and the bandwidth/symbol-duration trade-off. F1 and F2 share the 14 kHz band and 12 kHz bandwidth but differ in FFT size (1024 vs. 2048), trading subcarrier spacing against symbol duration: F2's longer symbol improves resilience to multipath delay spread

but is more sensitive to Doppler. F3 and F4 move to the 30 kHz band at opposite bandwidth extremes—a wide 16 kHz, short-symbol frame (F3) for high throughput, and a narrow 4 kHz, long-symbol frame (F4) that resists multipath and Doppler in poor channels at the cost of rate. The set thus brackets the center-frequency, bandwidth, and symbol-duration choices a deployed link must make, so the results obtained across all four are not tied to a single frame design. The single-carrier arm mirrors these frames at half the OFDM bandwidth with an RRC pulse ($\beta = 0.25$).

Receiver Front-End (What It Does, and Does Not, Do)

Figure 2 shows the receiver chain, which is held *fixed and identical across all FEC families* so that the benchmark isolates the coding axis. On the OFDM arm, the receiver removes the cyclic prefix, takes the FFT, and performs a per-subcarrier single-tap frequency-domain equalization (FDE)—a coherent matched filter $\mathrm{conj}(H_k)\, Y_k$ (matched-filter combining, not ZF/MMSE amplitude inversion) using the genie channel response H_k, so deeply faded subcarriers are down-weighted rather than noise-enhanced; then, it forms per-bit LLRs for the decoder. On the single-carrier arm, an RRC-matched filter and symbol-rate sampling precede PSK soft demapping and the same per-block decode. What the receiver *does not* do is equally important for a fair reading: it assumes *genie* CSI (no pilot-based estimation—the $(3, 3)$ pilot grid is charged in the goodput accounting but not used to estimate the channel); it performs *no* explicit Doppler or carrier-frequency-offset compensation (the channel taps are held constant within an OFDM symbol, so intra-symbol inter-carrier interference is assumed negligible, which is justified for the studied regime by the small normalized Doppler $f_d T_{\mathrm{sym}}$); and it uses neither a time-domain/decision-feedback equalizer nor iterative (turbo) equalization. These are deliberate idealizations: a stronger front-end (pilot-based estimation, Doppler compensation, turbo equalization) would improve *all* families together and is complementary to—not in competition with—the FEC-scheme and segment-size selection studied here, which is why it is held fixed; its realism cost is revisited in Section 7.2.

Payload sizes sweep $\{128\,\mathrm{B}, 512\,\mathrm{B}, 1, 2, 4, 8\,\mathrm{kB}\}$ (all < 10 kB). The segment (information-block) size is the per-configuration information length k of Table 3; across the action space, it ranges over ≈44–256 information bits, and the controller selects it jointly with the FEC family and rate (Section 3.4).

5.2. FEC Action Space and Fair Comparison

The benchmark spans six FEC families—short LDPC, CRC-aided list-decoded polar (with a Gaussian-approximation frozen set), BCH, Reed–Solomon, convolutional (soft Viterbi), and turbo—instantiated as the valid configurations of Table 3 (15 actions covering segment sizes ≈ 44–256 information bits and code rates 0.25–0.81). All families are compared like-for-like: identical information words, identical channel realizations, identical modulation, and a common per-information-bit axis, so a rate-R code is charged for its redundancy via $E_s/N_0 = E_b/N_0 + 10\log_{10} R$ [34,35]. Decoder settings are fixed across the benchmark: polar CA-SCL uses list size $L = 8$ with an 8-bit CRC and a per-operating-point Gaussian-approximation (GA) frozen set; LDPC is a regular ($d_v = 3$, $d_c = 6$) belief-propagation decoder capped at 30 iterations; convolutional is $K = 7$ unquantized Viterbi; turbo runs 8 BCJR iterations. These settings were chosen at the diminishing-returns knee of a sensitivity sweep at 2 dB AWGN (`results/decoder_sensitivity_e10.csv`): polar BLER falls $0.33\,(L = 1) \to 0.16\,(2) \to 0.11\,(4) \to 0.064\,(8) \to 0.034\,(16)$, LDPC $0.40\,(5\,\mathrm{it}) \to 0.22\,(10) \to 0.20\,(30) \to 0.17\,(50)$, and turbo $0.91\,(2\,\mathrm{it}) \to 0.63\,(4) \to 0.45\,(8) \to 0.41\,(12)$, so $L = 8$, 30, and 8 iterations sit where further effort yields little, which is consistent with the values used for short-block decoders in the literature [8,10]. We report decoder cost via

a transparent complexity-order latency proxy (SCL list size $\times\, n \log n$, LDPC iterations $\times\, n$, Viterbi trellis states $\times\, n$, algebraic $\approx n$) rather than wall-clock time.

The polar frozen set is matched to each operating point (design SNR $= E_b/N_0 + 10\log_{10} R$). The finite-blocklength penalty that makes the optimal code channel-dependent is the theoretical anchor for the benchmark [2,3,26].

5.3. Controller Evaluation (E2/E3)

The action-space sweep (E1 extended over the action set $\times$ a grid of channel states) yields a measured state-to-performance table. The controller experiments are then policy evaluation over that measured table: along a channel-state trajectory; each controller selects an action from its (predicted) state, and the realized goodput/energy/BLER for the *true* state is read back from the table, so no Monte Carlo is re-run, and every reported number is a table-based policy-evaluation result. Channel-state trajectories are a stated design model—a first-order autoregressive walk over the sweep grid (AR(1) coefficient $\rho = 0.9$; received SNR mean-reverting to the nominal operating point; delay spread and Doppler wandering across their grids)—with a strict train/test split across trajectories, and metrics are reported as the mean over 100 test trajectories with 95% confidence intervals. The controller acts on a predicted state from a persistence predictor at feedback delay H; E3 sweeps H and an additive prediction-error level (Section 6.7). Baselines are a fair single-family adaptive AMC controller (best FEC family with its rate and segment adapted), a fixed-best static action (no CSI), and an oracle given the true state (Section 6.6).

Off-grid validation. Because the trajectory states are continuous while the sweep is on a grid, each read snaps the true state to the nearest grid point. We verified this nearest-neighbor read against direct simulation at held-out operating points that lie *between* grid points, using a fresh, independent Monte-Carlo (a different RNG seed from the table build, so the table and the validation never share realizations). At AWGN 3/5/7 dB and at OFDM-TDL 6/10 dB, the directly simulated goodput lies between the two bracketing grid entries at every SNR-interior point, so the table behaves as a faithful monotone interpolant along the SNR axis; an additional interior delay-spread/Doppler point (off the τ_{rms}/Doppler grid as well) is reproduced to within the grid resolution but is not strictly bracketed. The value the controller actually reads (the nearer grid point) differs from the direct simulation by at most 0.24 and on average 0.05 information bits per channel use—the half-grid-step resolution at a worst-case midpoint, which is driven by the high-rate RS action whose goodput rises steeply through the waterfall—and, because the identical snap is applied to all four controllers, this resolution is common to the evaluation protocol rather than biasing any one controller; the held-out simulations did not change the controller ordering at the checked points, and the finite grid resolution is reported as part of the approximation error underlying the RQ2 and RQ3 conclusions.

5.4. Metrics

The primary metric is the ARQ effective goodput of Equation (5) (successfully delivered information bits per channel use), with every action charged its true per-segment overhead (header and CRC, through o), the OFDM pilot grid (once, through φ), and its retransmission cost, so that small segments pay their real cost. We also report energy per delivered bit [31], block-error rate, and the decoder latency proxy. Across the action space and operating points evaluated here, the latency proxy never exceeds the budget—no action is pruned by $T_{\max}$, so the reported goodput is not shaped by the latency constraint; $T_{\max}$ is retained as the explicit lever by which a tighter deployment budget would prune the more expensive strong-code decoders.

5.5. Experiment Matrix and Reproducibility

Table 4 maps the experiments (E1–E4) to the research questions and their outputs. Monte Carlo runs use the *full* stopping rule ($\geq$200 block errors and $\geq$2000 blocks per point, capped at 2×10^5 blocks) for the reported E1 curves and tables with 95% confidence intervals shown on the BLER figures. All runs are seeded per condition for reproducibility, and the simulator is implemented in Python 3.12.3/NumPy 1.26.4 with galois 0.4.11, scikit-commpy 0.8.0, and pyldpc 0.7.9 for the codecs and PyTorch 2.12.0 for the optional learned predictor.

Table 4. Experiment matrix.

Exp.	Question	Sweep	Output
E1	RQ1	6 families over AWGN; 5 over TDL (4 fading-robust + soft conv; turbo AWGN-only) ($\times$ rate $\times$ block length $\times$ SNR)	Figure 3 BLER, Figure 4 goodput, crossover (Tables 5–8)
E2	RQ2	{joint, single-family AMC, fixed-best, oracle} $\times$ SNR $\times$ channel	controller goodput and energy (Section 6.6)
E3	RQ3	proposed controller $\times$ feedback delay H $\times$ prediction error σ	retained-gain maps (Section 6.7)
E4	RQ1 ext.	{4 fading families} $\times$ SNR $\times$ {synthetic, 3 measured-statistic channels}	measured-statistics replay (Section 6.8)

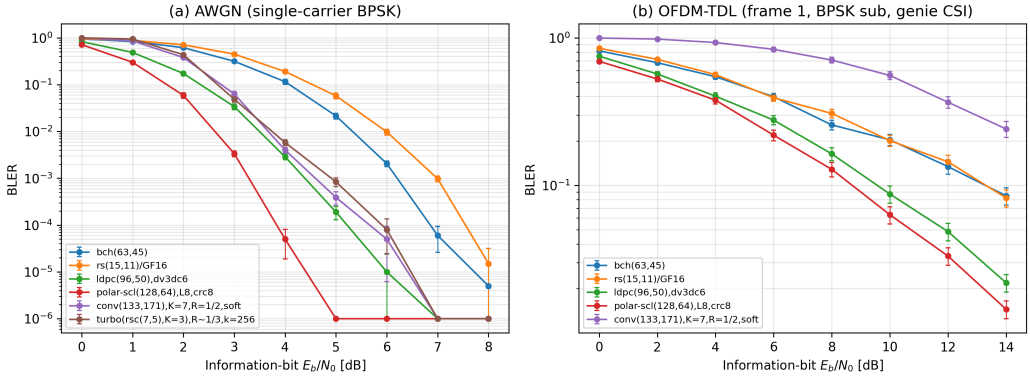


Figure 3. Block-error rate versus per-information-bit E_b/N_0. (**a**) AWGN single-carrier BPSK, six families; (**b**) OFDM-TDL fading (frame 1, BPSK subcarriers, genie CSI), the fading-capable families including the soft convolutional code (uppermost curve, dominated at every SNR). The information-bit axis sets $E_s/N_0 = E_b/N_0 + 10\log_{10} R$.

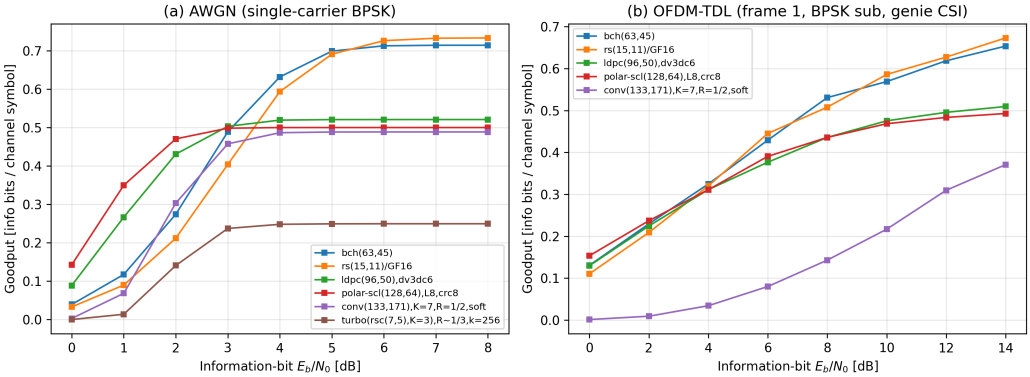


Figure 4. ARQ effective goodput (Equation (5); $B = 256$ B, $o = 32$ bits) versus E_b/N_0. (**a**) AWGN three-stage frontier (polar < 3 dB, soft convolutional 3–5 dB, RS/BCH > 6 dB); (**b**) OFDM-TDL, where the high-rate codes overtake polar only by $\approx$10 dB and the soft convolutional code is dominated at every SNR (turbo AWGN-only).

Table 5. AWGN single-carrier BPSK, six-family like-for-like benchmark (measured BLER, 95% Wilson CIs $\leq \pm 0.022$). $R_{\text{eff}} = k/(n+o)$ is the overhead-charged rate ceiling and the last column the ARQ effective goodput at 8 dB ($B = 256$ B, $o = 32$ bits, Equation (5)); raw rate R for reference.

FEC Family	R	BLER @1 dB	BLER @2 dB	BLER @3 dB	R_{eff}	Goodput @8 dB
Polar CA-SCL, GA (128,64)	0.50	**0.300**	**0.059**	**0.003**	0.40	0.400
LDPC (96, 50)	0.52	0.488	0.173	0.034	0.39	0.390
Convolutional $K = 7$, R$\frac{1}{2}$ (soft)	0.49	0.86	0.38	0.063	0.46	0.460
Turbo, R$\frac{1}{3}$	0.25	0.945	0.435	0.049	0.24	0.242
BCH (63, 45)	0.71	0.836	0.616	0.316	0.47	0.469
RS (15, 11)	0.73	0.877	0.711	0.449	0.48	**0.474**

Bold marks the best value in each column (lowest BLER; highest goodput).

Table 6. Construction sensitivity at fixed $(n, k) = (128, 64)$ (polar) or $k = 256$ (conv): the FEC *construction*—frozen-set design, decoder decision, fading design SNR—moves BLER by ∼1 order of magnitude, as much as switching family (`results/construction_e9.csv`; OFDM-TDL row: frame 1, BPSK, 8-tap, $K = +5$ dB).

Construction Axis	Operating Point	Weak	Strong
Polar frozen set: BEC vs. GA (matched)	AWGN, 2 dB	0.842 (BEC)	0.049 (GA)
Convolutional decision: hard vs. soft Viterbi	AWGN, 2 dB	0.983 (hard)	0.388 (soft)
Polar fading design: mean-matched vs. fade-robust	OFDM-TDL, 14 dB	0.087 (mean)	0.012 (robust)

Table 7. OFDM over the TDL fading channel (frame 1, BPSK subcarriers, genie CSI; measured BLER, 95% Wilson CIs $\leq \pm 0.021$, convolutional row $\leq \pm 0.035$). $R_{\text{eff}} = k/(n+o)$; goodput at 14 dB for $B = 256$ B, $o = 32$ bits, $\varphi = 8/9$. The soft convolutional code is *dominated at every fading SNR* (long $n = 524$ codeword, not bit-interleaved).

FEC Family	R	BLER @0 dB	BLER @6 dB	BLER @14 dB	R_{eff}	Goodput @14 dB
BCH (63, 45)	0.71	0.817	0.399	0.085	0.47	0.381
RS (15, 11)	0.73	0.850	0.393	0.082	0.48	**0.386**
LDPC (96, 50)	0.52	0.751	0.277	0.022	0.39	0.339
Polar CA-SCL, GA (128, 64)	0.50	**0.693**	**0.219**	**0.014**	0.40	0.350
Convolutional (133, 171), soft	0.49	0.998	0.836	0.241	0.46	0.311

Bold marks the best value in each column (lowest BLER; highest goodput).

Table 8. Goodput frontier by channel under the ARQ effective goodput ($B = 256$ B). The crossover is the E_b/N_0 where the high-rate algebraic codes overtake the rate-$\frac{1}{2}$ codes; on fading, the soft convolutional code (the AWGN mid-SNR stage) is dominated at every SNR and turbo is AWGN-only, so the fading frontier has no mid-SNR stage.

Channel	Low SNR	Mid SNR	High SNR	Crossover
AWGN (BPSK)	Polar CA-SCL (<3 dB)	conv, soft (3–5 dB)	RS/BCH (≥6 dB)	≈6 dB
OFDM-TDL fading (frame 1)	Polar CA-SCL (0–8 dB)	conv (dominated)	RS/BCH (≥10 dB)	≈10 dB

6. Results and Analysis

We first establish, through a controlled like-for-like benchmark (E1), how the candidate short-block FEC families trade reliability against throughput and how that trade-off moves with the channel (RQ1). All families are compared on a common information-bit axis: results are plotted against the per-information-bit E_b/N_0, so that the physical channel sees $E_s/N_0 = E_b/N_0 + 10\log_{10} R$ and a low-rate code is charged for the extra energy its redundancy costs. Each family encodes the same information bits, over the same channel realizations, with the same modulation, and we report the ARQ effective goodput

of Equation (5)—delivered information bits per channel use, charged the per-segment overhead and retransmission cost at a headline payload of $B = 256$ B—alongside BLER so that reliability and rate are weighed jointly rather than in isolation. Two channels are evaluated: an additive white Gaussian noise (AWGN) single-carrier BPSK link, which isolates code behavior from fading, and an OFDM link over a time-varying tapped-delay-line (TDL) multipath-and-Doppler channel (frame 1 of Table 2, BPSK subcarriers, frequency-domain genie channel-state information), which exposes the families to the frequency-selective fading characteristic of the underwater channel.

6.1. Reliability and Throughput over AWGN

Table 5 reports BLER at three low-SNR operating points and the high-SNR goodput for the six FEC families on the AWGN link; Figure 3a plots the full BLER curves. Three findings stand out. First, at this short blocklength, the CRC-aided SCL polar code with a GA frozen set is the reliability champion, attaining the lowest BLER at every low-SNR point—0.059 at 2 dB versus 0.173 for the rate-$\frac{1}{2}$ LDPC code and 0.38–0.44 for the convolutional and turbo codes—consistent with the strength of CRC-aided list decoding of polar codes reported for short underwater and short-control-channel blocks [10,19,26]. Second, code construction matters as much as the code family (Table 6): replacing the SNR-agnostic Bhattacharyya/BEC frozen set with the channel-matched GA density-evolution design lowers the polar code's 2 dB BLER from 0.84 to 0.05—roughly a 17-fold improvement at a fixed $(n, k) = (128, 64)$—and switching the convolutional decoder from hard decision to unquantized (soft) Viterbi lowers its 2 dB BLER from 0.98 to 0.39. Third, once reliability is achieved, the throughput ordering is governed by the overhead-charged effective rate rather than the raw code rate. Under the ARQ effective goodput of Equation (5) (256-byte payload, 32-bit per-segment overhead), the high-SNR (8 dB) ordering is RS 0.474 $\approx$ BCH 0.469 $>$ conv 0.460 $>$ polar 0.400 $>$ LDPC 0.390 $>$ turbo 0.242. The high-rate algebraic codes still lead, but their margin over the rate-$\frac{1}{2}$ codes shrinks from the raw-rate gap (0.71–0.73 vs. 0.50–0.52) to a few hundredths, because their short blocks ($n = 60$–63) pay the 32-bit overhead most heavily; the long-block soft convolutional code ($n = 524$) dilutes that overhead and rises to near-parity despite its lower rate (0.49), while the short LDPC code falls below it.

The construction effects quoted above are made explicit in Table 6, which shows a fresh, reproducible run that isolates each construction axis at the operating point the prose cites. The frozen-set design, the decoder decision type, and the fading design SNR each move the BLER by roughly an order of magnitude at a *fixed* (n, k)—i.e., the construction is as decisive as the family name, which is the point that Sections 6.2 and 6.5 develop.

The consequence for goodput (Figure 4a) is a three-stage frontier. Below $\approx$3 dB, the short rate-$\frac{1}{2}$ polar CA-SCL code leads, because reliability dominates and it fails least often. Between $\approx$3 and 5 dB, the soft convolutional code leads: it has cleared the waterfall and its long block dilutes the per-segment overhead, so its effective goodput ($\approx$0.46) exceeds that of the short rate-$\frac{1}{2}$ codes. Above $\approx$6 dB, the high-rate algebraic BCH and RS codes overtake it ($\approx$0.47), spending the least of the channel on redundancy. No single family is best across the operating range, and—because the ordering turns on block length as well as code rate—the result already implicates the segment size, not only the family and rate, as an adaptation axis (Section 6.4).

6.2. Reliability and Throughput over the OFDM-TDL Fading Channel

Table 7 and Figures 3b and 4b repeat the comparison over the fading channel for the fading-capable families—LDPC and polar CA-SCL (soft-decision), BCH and RS (hard-decision bounded-distance), and the soft convolutional code; the turbo decoder remains

AWGN-only here (Section 6.5). The convolutional code, which wins the AWGN mid-SNR stage, is by contrast *dominated at every fading SNR* (Table 7, lowest curve in Figure 4b): its long codeword is not bit-interleaved across the independent OFDM-symbol fades, so the channel's deep-fade tail produces burst errors that its Viterbi decoder cannot correct. It is therefore retained in the RQ1 benchmark for completeness; because it is dominated at every fading SNR, it is never goodput-optimal over fading and so does not enter the controller's fading decisions (and a realistic latency budget prunes it regardless; see Section 7.3), leaving the RQ2/RQ3 results unchanged. The AWGN picture survives but shifts to a higher SNR. The GA frozen set again rescues the polar code provided it is designed for the fading channel: matching it to the per-subcarrier *mean* E_s/N_0 leaves a high-SNR error floor (short-block reliability over fading is set by the deep-fade tail, not the mean), which a fixed fade-robust design removes, dropping the polar code's 14 dB BLER from 0.087 to 0.012 (Table 6). Polar and rate-$\frac{1}{2}$ LDPC then form the reliability pair (polar slightly ahead at $\geq$6 dB), which are both clearly beating the high-rate BCH and RS codes at a low-to-mid SNR. The fading penalty is severe: even at 14 dB, the strong codes sit at BLER $\approx$ 0.01–0.02 and the high-rate codes sit near 0.08, against BLER $\approx$ 0 by 5 dB on AWGN, which is the finite-blocklength behavior expected for short packets over multipath fading [2,3].

The goodput crossover (Figure 4b) is delayed by fading. Through the waterfall, the four fading-robust families stay close on effective goodput (the convolutional code trailing well below them), and the high-rate BCH and RS codes overtake the rate-$\frac{1}{2}$ polar and LDPC codes only by $\approx$10 dB—several decibels later than the $\approx$6 dB crossover on AWGN. At 14 dB, the algebraic codes reach an ARQ effective goodput of $\approx$0.38–0.39 (256-byte payload) against $\approx$0.34–0.35 for polar and LDPC—a narrower high-SNR margin than the raw code rates (0.71–0.73 vs. 0.50–0.52) would suggest for the same per-segment-overhead reason as on AWGN (Section 6.1): the algebraic codes' short blocks pay that overhead most heavily. The turbo code is absent here (its probability-domain BCJR decoder is AWGN-only; see Section 6.5), and the soft convolutional code—present but dominated at every SNR (Table 7)—does not open a mid-SNR stage as it does on AWGN, so the fading frontier is a genuine two-stage trade-off: the long-block convolutional code that wins the AWGN mid-SNR stage does not survive frequency-selective fading.

6.3. The Reliability–Throughput Frontier Is Channel-Dependent (RQ1)

Tables 5, 7 and 8 together answer RQ1: no single short-block FEC configuration dominates the reliability–throughput frontier, and the location of the optimum depends on the channel state. Under the unified ARQ effective goodput, the AWGN frontier is a *three-stage* trade-off—the short rate-$\frac{1}{2}$ polar CA-SCL code at low SNR ($<$3 dB), the long-block soft convolutional code at mid SNR (3–5 dB), and the high-rate algebraic BCH/RS codes at high SNR ($\geq$6 dB)—while over the fading channel, it is a *two-stage* trade-off (rate-$\frac{1}{2}$ polar at low-to-mid SNR, high-rate BCH/RS above $\approx$10 dB): the soft convolutional code that forms the AWGN mid-SNR stage is dominated at every fading SNR (Table 7), and the turbo decoder is AWGN only. The SNR at which the high-rate codes overtake the strong codes on goodput moves with the channel from $\approx$6 dB on AWGN to $\approx$10 dB over the TDL fading channel (Table 8). Because the ordering depends on block length as well as code rate (Section 6.4), the optimum is a function of the payload size as well as the channel state. Because the underwater channel state varies within a session faster than a fixed configuration can be re-provisioned, this channel- and payload-dependent optimum is direct evidence that the best FEC choice is a function of the (predicted) channel state—the premise of the adaptive controller proposed in Section 4.

One scope caveat once bounded this frontier: the BCH and RS codes used bounded-distance *hard*-decision algebraic decoding, whereas the polar, LDPC, and convolutional

codes are soft decision. We have since added soft-decision (Chase-II) algebraic decoding (Section 6.5). It sharply improves the algebraic codes' low-SNR reliability—at 3 dB AWGN, BCH$(63, 45)$ drops from BLER 0.32 to 0.06 and RS$(15, 11)$ from 0.45 to 0.11—and moves this algebraic-versus-polar goodput crossover (the E5 BCH/RS-versus-polar pairing at goodput 0.40, distinct from the $\approx$ 6 dB three-stage frontier crossover of Table 8) *earlier*, from $\approx$4 to $\approx$3 dB, but it does *not* erase it: the rate-$\frac{1}{2}$ polar code still attains the lowest BLER at the lowest SNR, and the high-rate algebraic codes still require a sufficiently good channel before their throughput advantage materializes. The existence and channel-dependence of the crossover—on which the adaptive premise rests—therefore hold under soft-decision decoding, which is now as a *measured* result rather than a conservative hard-decision bound. (Ordered-statistics decoding, a still-stronger soft variant, would move the crossover at most marginally further and is left as future work.)

A methodological corollary follows from Sections 6.1 and 6.2 and is reported here because it bears directly on the controller's action space: the **code *construction*—the polar frozen-set design and soft- versus hard-decision decoding—is as decisive as the code family**. A polar code is the weakest or the strongest family in this benchmark depending only on whether its frozen set is matched to the channel; an earlier finding that polar "underperformed over fading" reflected the SNR-agnostic BEC construction rather than a family-level limitation, so that construction is not used as the representative polar design in the remainder of this paper. We therefore report and adapt the construction (the GA design SNR), not merely the family, throughout the remainder of this paper.

6.4. Payload Size and the Overhead–Padding Trade-Off

The frontier above is reported at a single payload ($B = 256$ B); Figure 5 isolates the segmentation trade-off by sweeping B at a representative high-SNR point. Since BLER is independent of the payload size, the variation of each curve with B reflects only the per-segment overhead and last-segment padding. Two effects compete. A short payload cannot fill a large segment, so codes with a long block (the convolutional and turbo codes, $k = 256$) waste a large padding fraction on tiny payloads—collapsing to a goodput of 0.23 at $B = 16$ B, where their 256-bit segment is half empty—while the short-block codes (BCH, RS, polar, $k \leq 64$) pay little padding and lead there. As B grows past a code's segment size, the padding amortizes and the per-segment overhead o dominates instead, now favoring the longer blocks: the convolutional code recovers its full 0.46 as soon as the payload fills its segment ($B = 32$ B = 256 bits), matching the algebraic codes to within a few hundredths. The goodput-optimal code therefore changes with the payload size as well as the channel state—which is a second axis that the joint controller of Section 4 exploits through its segment-size degree of freedom and that a single fixed configuration cannot. This is the empirical basis for treating the segment size, and not only the rate and family, as an adaptation variable.

Segment size varied independently of the family and rate. To confirm that segment size is a first-class lever and not merely a by-product of the family/rate choice, Table 9 holds the FEC family (polar CA-SCL) *and* the code rate ($R = \frac{1}{2}$) fixed and varies only the information-block (segment) size $k \in \{64, 128, 256\}$ (i.e., $(n, k) = (128, 64), (256, 128), (512, 256)$). With family and rate constant, the ARQ goodput at a 256-byte payload rises from 0.40 to 0.47 (AWGN, 8 dB) purely because a larger segment amortizes the per-segment header/CRC overhead—a +18% swing that is *larger* than the high-SNR cross-family gain (+10.4%, Section 6.6). Segment-size adaptation is therefore independent and, across most of the operating range, the dominant degree of freedom.

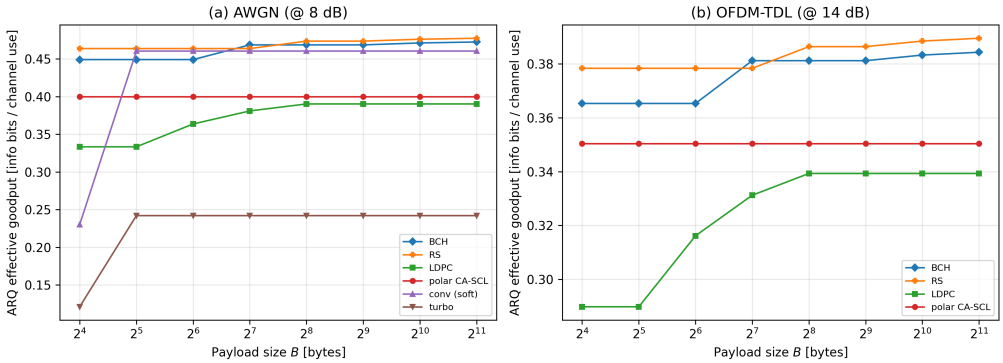


Figure 5. Segmentation trade-off: ARQ effective goodput versus payload size B at high SNR, (**a**) AWGN 8 dB and (**b**) OFDM-TDL 14 dB. Since BLER is independent of B, the variation is pure overhead/padding: short-block codes ($k \leq 64$) lead on tiny payloads and yield to longer blocks as B grows.

Table 9. Segment size varied with the FEC family (polar CA-SCL) and code rate ($R = \frac{1}{2}$) held fixed: ARQ effective goodput ($B = 256$ B) rises monotonically with the segment size k as the per-segment overhead is amortized, isolating the segment-size lever from family and rate (from `results/action_sweep_awgn.csv`).

Polar $R = \frac{1}{2}$, AWGN	$k = 64$	$k = 128$	$k = 256$
Goodput @ 2 dB	0.367	0.433	0.467
Goodput @ 8 dB	0.400	0.444	0.471

A complementary decomposition of the controller gain (from the E2 baselines of Section 6.6) attributes the joint advantage over a fixed code to its two sources. Adapting the rate and segment size *within* the best single family raises goodput over the fixed-best code at low SNR (AWGN 0 dB 0.310→0.333; TDL 0 dB 0.145→0.170) and accounts for essentially all of the gain over fading and at low-to-mid AWGN; the additional cross-family switch contributes only at the high-SNR AWGN crossover (0.593→0.655 at 8 dB). The two analyses agree: the segment-size (with rate) lever carries the gain almost everywhere, and family switching is a bounded, high-SNR-AWGN-specific bonus.

6.5. Validity and Scope of the E1 Benchmark

Several caveats bound these E1 results (expanded in Section 7.2). The BLER numbers use the full Monte-Carlo profile (≥200 block errors per point) with 95% confidence intervals (Figure 3), and the fading-channel polar code uses a fixed fade-robust GA design (Section 6.2). The fading link assumes genie CSI; only frame 1 with BPSK subcarriers is reported over fading; the soft convolutional code is included on fading but is dominated there (Table 7) and the turbo decoder is AWGN-only in this build; and BCH and RS are evaluated with both hard-decision and soft-decision (Chase-II) algebraic decoding (Table 10). None of these affects the RQ1 conclusion, which rests on the *relative* family ordering and the *existence and channel-dependence* of the goodput crossover, which are both robust across the operating range and under soft-decision algebraic decoding.

Robustness of the polar construction to design-SNR mismatch. Because the polar GA frozen set is matched to the operating point using the *predicted* SNR, a natural concern is how sensitive the construction is when that prediction is wrong. Figure 6 sweeps the design SNR of polar$(128, 64)$ while holding the operating E_b/N_0 fixed with this paper's matched design (the operating E_s/N_0) marked. The construction is robust within a wide window: across roughly ±4 dB of design–SNR mismatch around the matched point, the BLER stays essentially flat at its floor—at a 2 dB operating SNR, it remains near 0.05 for

design SNRs from -4 to $+3$ dB and at 3 dB near 0.003 for design SNRs from -4 to $+6$ dB. A 4–10$\times$ rise in BLER appears only beyond about ± 5 dB of mismatch, and it is worst when the frozen set is designed for an SNR well *above* the operating point (over-designing freezes bit-channels that are in fact reliable): at 2 dB, the BLER rises to 0.21 for a -5 dB mismatch and 0.30 for a $+9$ dB mismatch versus 0.05 at the match. The matched design therefore sits in the middle of a broad valley, so the per-operating-point construction is not fragile to the SNR-prediction error; moreover, a conservative CSI margin ($\alpha > 0$, Section 4.3) would bias the design toward *lower* SNR, i.e., onto the flatter, safer side of that valley.

Table 10. Soft-decision (Chase-II) vs. hard-decision algebraic decoding (E5; `results/e5_soft_algebraic.csv`). Soft decoding sharply lowers the algebraic codes' low-SNR BLER and moves the AWGN goodput crossover (vs. the rate-$\frac{1}{2}$ polar code, goodput 0.40) from $\approx$4 to $\approx$3 dB without erasing it (polar still leads at 2 dB).

	BLER @ 3 dB		ARQ Goodput @ 4 dB	
Code	**Hard**	**Soft**	**Hard**	**Soft**
BCH $(63, 45)$	0.319	0.055	0.413	0.466
BCH $(63, 51)$	0.418	0.116	0.447	0.518
RS $(15, 11)$	0.454	0.107	0.385	0.468
RS $(31, 25)$	0.775	0.483	0.382	0.590

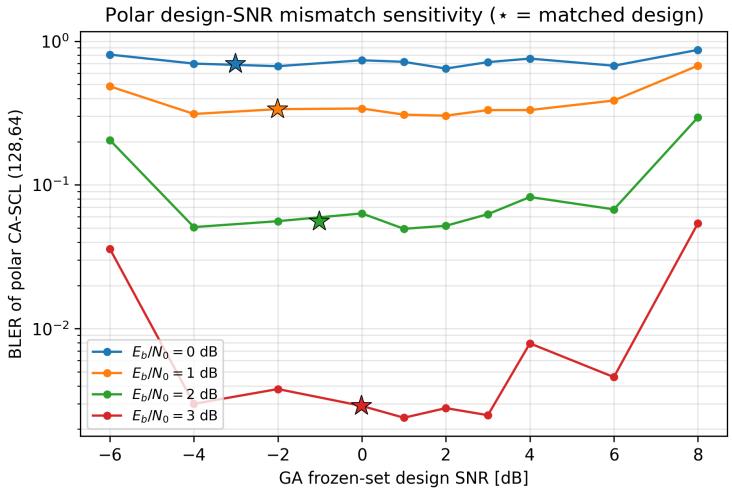


Figure 6. Polar design–SNR mismatch sensitivity: BLER of polar CA-SCL$(128, 64)$ versus the GA frozen-set design SNR, one line per operating E_b/N_0 ($\star$ marks the matched design). Each curve has a broad flat valley (± 4 dB tolerance) and degrades by 4–10$\times$ only under large mismatch—most steeply when over-designed.

Separating the transmission scheme from the channel. The benchmark pairs each channel with one transmission scheme—single-carrier on AWGN, OFDM on fading—which raises the question of whether the family ordering reflects the codes or the waveform. To separate the two, we re-ran the four fading-capable families over OFDM on a *flat* (deterministic unit-gain) channel with AWGN—i.e., OFDM-on-AWGN—and compared their BLER to the single-carrier AWGN benchmark (Table 11). The two agree to within the Monte-Carlo confidence interval at every operating point (mean $|\Delta\text{BLER}| = 0.010$, max 0.026), and the family ordering is identical (polar $<$ LDPC $<$ BCH $<$ RS at 2 dB on both arms). The reliability ordering is therefore a property of the codes, not of the single-carrier-versus-OFDM waveform, so pairing each channel with one scheme does not bias the RQ1 comparison.

Table 11. Transmission scheme vs. channel: BLER of the four fading-capable families over single-carrier (SC) AWGN and over OFDM on a flat unit-gain AWGN channel agree within Monte-Carlo error (`results/ofdm_awgn_e6.csv`), so the family ordering is a code property rather than an OFDM/single-carrier artifact.

	2 dB		4 dB		6 dB	
Family	**SC**	**OFDM**	**SC**	**OFDM**	**SC**	**OFDM**
Polar CA-SCL	0.053	0.061	0.000	0.000	0.000	0.000
LDPC	0.240	0.217	0.003	0.019	0.008	0.000
BCH	0.607	0.603	0.106	0.122	0.001	0.002
RS	0.723	0.697	0.199	0.183	0.012	0.010

Imperfect-CSI bound on the fading results. The fading benchmark uses genie CSI, which is optimistic for reliability. To bound that optimism, we replaced the genie channel response in the per-subcarrier matched filter with an imperfect estimate $\hat{H} = H + e$, $e \sim \mathcal{CN}(0, N_0)$—the error of a least-squares estimate from a single unit-power pilot at the data E_s/N_0 (the $(3,3)$ pilot grid's per-subcarrier error before any interpolation gain), which is a parametric sensitivity rather than a full pilot-LS interpolation. Table 12 reports the BLER, genie versus imperfect, for the four families over the synthetic TDL. An imperfect CSI roughly doubles the BLER (and lowers the ARQ goodput by ≈8% at a high SNR to ≈50% at a low SNR, where doubling an already-high BLER halves the success probability), but—critically—the family ordering is *preserved* at every operating point (polar and LDPC remain the reliability pair; BCH and RS remain throughput-oriented), so the genie idealization is optimistic in absolute terms yet does not change the RQ1 conclusion. This is also the front-end-invariance evidence anticipated in Section 5: the FEC-selection-relevant ordering survives a realistic degradation of the receiver front-end, so the controller's family and segment decisions are robust to the front-end idealizations rather than artifacts of them.

Table 12. Imperfect-CSI sensitivity: BLER of the four fading-robust families over the synthetic OFDM-TDL channel (frame 1, BPSK, 8-tap, $K = +5$ dB) under genie CSI vs. an LS-at-pilot estimate $\hat{H} = H + e$, $e \sim \mathcal{CN}(0, N_0)$ (`results/imperfect_csi_e13.csv`). The family ordering is preserved.

	4 dB		8 dB		12 dB	
Family	**Genie**	**Imperf.**	**Genie**	**Imperf.**	**Genie**	**Imperf.**
Polar CA-SCL	0.356	0.672	0.131	0.347	0.037	0.118
LDPC	0.399	0.699	0.190	0.373	0.039	0.135
BCH	0.524	0.721	0.270	0.397	0.135	0.209
RS	0.559	0.728	0.280	0.439	0.135	0.217

6.6. Joint Channel-Adaptive Selection Beats Single-Axis Adaptation (RQ2)

We now evaluate the proposed controller against its baselines (E2) as a table-based policy evaluation over the E1 sweep (Section 5): each controller selects an action from the predicted state, and the realized goodput for the true state is read from the sweep. All values use the ARQ effective-goodput metric of Equation (5), and the action space spans the full Table 3, including higher-rate algebraic variants beyond the six configurations of Table 5. The algebraic actions use the hard-decision bounded-distance decoders of Table 3 for direct comparability with the E1 sweep the policy reads; the soft-decision (Chase-II) algebraic variant is an auxiliary sensitivity study (Table 10, Section 6.5) and is not part of the E2/E3 policy evaluation. Because Chase-II shifts the AWGN goodput crossover only modestly (from ≈4 to ≈3 dB) without erasing it, including it in the action space would not change the controller's family or segment ordering. The high-SNR AWGN action is therefore RS$(31, 25)$ ($k = 125$, $R = 0.81$), whose longer block amortizes the per-segment

overhead far better than the representative RS$(15, 11)$ row of Table 5 ($k = 44$); this is why the joint controller's 8 dB goodput (0.655) exceeds that table's RS entry (0.474).

Choosing a *fair* single-axis baseline requires care in this short-block action space. A classical AMC controller fixes the FEC scheme and adapts the code rate, but here, each FEC family realizes a given rate only at a specific block length, so "fix the segment, adapt scheme and rate" degenerates to a single fixed code (Table 3 has at most one (scheme, rate) per segment size). We therefore compare against the strongest fair single-axis baseline: a single-family adaptive AMC controller that fixes the FEC family to the best one—selected by the mean goodput over the benchmark grid (polar CA-SCL on both channels), independently of the test trajectories—and adapts its (n, k) configuration (rate and block length together) to the channel. We also report a fixed-best static code (no CSI) and an oracle given the true state (the gain ceiling). The proposed joint controller additionally adapts the FEC *family*, so the joint-over-single-family margin isolates the value of multi-family adaptation [4]. Each controller acts on a persistence predictor at feedback delay $H = 1$ (AWGN)/$H = 2$ (TDL); we report the mean over 100 independent test trajectories with 95% confidence intervals. Section 6.7 studies the sensitivity to the prediction.

Table 13 and Figure 7 report goodput with 95% confidence intervals, and Figure 8 the corresponding energy per delivered bit. Because all controllers are evaluated on the *same* 100 trajectories, we test the joint-versus-single-family difference at each operating point with a paired Wilcoxon signed-rank test and report Cohen's d_z effect size rather than inferring "ties" from confidence-interval overlap. The joint controller's advantage over the fair single-family baseline is concentrated exactly where the RQ1 frontier crosses between families. On AWGN, the two are statistically indistinguishable at 4 dB (0.547 vs. 0.545; paired $p = 0.88$), and the joint controller gains $+10.4\%$ at 8 dB (0.655 vs. 0.593; paired $p < 10^{-17}$, $d_z = 4.7$, a large effect) by switching from polar to the high-rate RS code once the channel is good enough—exactly the AWGN goodput crossover of RQ1 (Section 6.1). Off the crossover, the extra family degree of freedom does *not* help: at 0 dB AWGN and across the entire fading range, the joint controller stays within $\pm 0.8\%$ of the single-family AMC, but the paired test resolves a *small, consistent deficit* ($\Delta \approx -0.001$ to -0.002 information bits per channel use; $|d_z| \approx 0.5$–0.9, $p < 10^{-3}$)—the prediction overhead of carrying an unused family choice where no frontier crossing exists (polar is near optimal at every fading operating point, and the convolutional/turbo codes that create AWGN's mid-SNR stage are unavailable here). This deficit is statistically resolvable only because it is so consistent across trajectories; in absolute terms, it is negligible, and it *confirms* rather than contradicts the central finding—there is no family-switch headroom off the crossover, so the family degree of freedom should be exercised only there. Against the fixed-best, no-CSI code, the gains are larger where the rate adaptation matters—$+7\%$ at 0 dB on AWGN and $+16\%$ at 0 dB over fading (both paired $p < 10^{-16}$, $d_z \geq 1.3$)—because both the single-family AMC and the joint controller can drop to a lower-rate configuration in a poor channel that a static code cannot. The joint controller captures 91–99% of the oracle goodput on both channels. Joint multi-family adaptation thus matches the best single-family AMC off the crossover (to within a negligible prediction-overhead margin) and adds real value precisely at the family crossover the RQ1 benchmark predicts (high-SNR AWGN, polar→RS) without manufacturing a gain where the frontier does not cross.

The fading-channel headroom for family adaptation over a well-chosen single family is negligible (oracle$-$fixed ≈ 0.001–0.007), so the one regime with a real multi-family gain, high-SNR AWGN, is also the most sensitive to stale CSI, which motivates the RQ3 study.

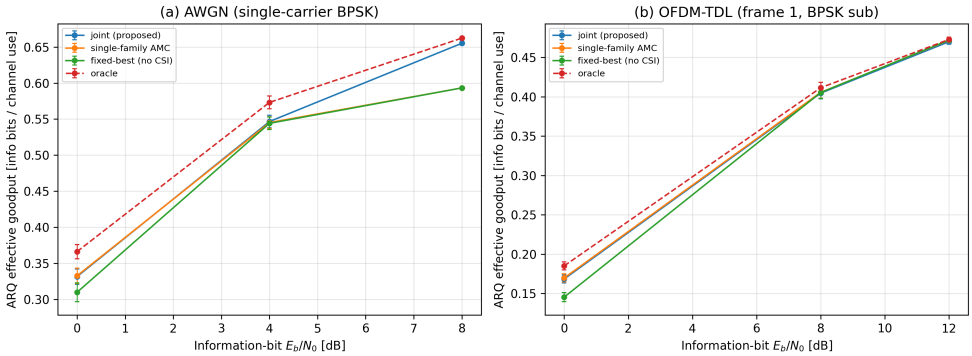


Figure 7. ARQ effective goodput versus E_b/N_0 for the joint controller, the fair single-family AMC baseline, the fixed-best (no-CSI) code, and the oracle, (**a**) AWGN and (**b**) OFDM-TDL (persistence predictor; error bars are 95% CIs over 100 trajectories). The joint controller separates from the baseline only at the high-SNR AWGN where it switches family (polar→RS).

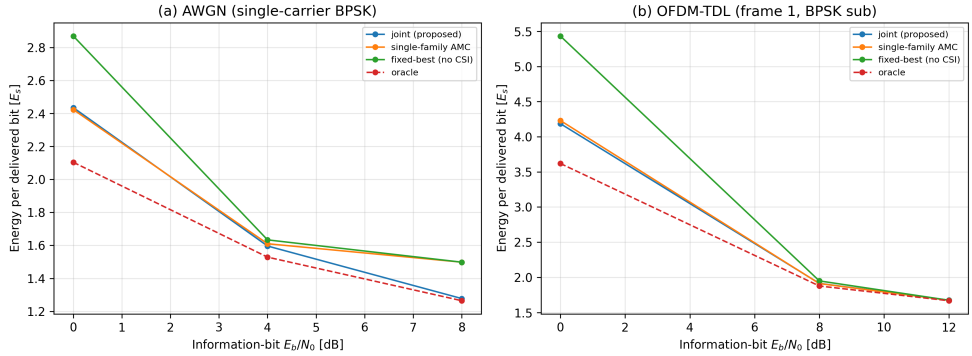


Figure 8. Energy per delivered information bit (in units of E_s) versus E_b/N_0 for the same controllers. The joint controller matches the single-family AMC baseline, and both undercut the fixed no-CSI code at low SNR (15–23% less); the high-SNR AWGN family switch (polar→RS) lowers it further (1.28 vs. 1.50 E_s at 8 dB).

Table 13. ARQ effective goodput by controller and operating point (persistence predictor, $H = 1$ AWGN/$H = 2$ TDL; mean over 100 trajectories, 95% CI $\leq \pm 0.011$). Joint vs. single-family AMC is assessed by a paired Wilcoxon signed-rank test: a large gain at 8 dB AWGN (polar→RS; $p < 10^{-17}$, $d_z = 4.7$) and a negligible difference off the crossover ($|\Delta| \leq 0.002$).

	AWGN (BPSK)			OFDM-TDL		
Controller	**0 dB**	**4 dB**	**8 dB**	**0 dB**	**8 dB**	**12 dB**
joint (proposed)	0.332	0.547	**0.655**	0.168	0.405	0.470
single-family AMC (polar)	0.333	0.545	0.593	0.170	0.406	0.472
fixed-best (no CSI)	0.310	0.544	0.593	0.145	0.405	0.472
oracle (true CSI)	0.366	0.573	0.662	0.185	0.412	0.473

Bold marks the joint controller's goodput at 8 dB AWGN—the family-crossover operating point (polar→RS) where the joint controller's gain over the single-family baseline is large and statistically significant.

Figure 9 makes the controller's policy explicit by mapping the selected action—FEC family, code rate, and segment size—onto the channel state. On AWGN, (panel a) the joint controller walks the RQ1 frontier exactly: it holds the rate-$\frac{1}{2}$ polar code at 0 dB, raises the polar rate to $\frac{2}{3}$ at 2–4 dB, and switches to the high-rate RS code ($R = 0.81$, $k = 125$) above 6 dB—a coordinated move of family, rate, *and* block length rather than rate alone. Over fading (panel b), the policy is dominated by the polar family, whose rate climbs from $\frac{1}{2}$ to $\frac{2}{3}$ as the SNR rises, and the high-rate RS code surfaces only in the high-SNR, low-delay-spread, low-Doppler corner; the single low-SNR LDPC cell reflects the polar/LDPC

reliability tie of RQ1 (Section 6.2). The map confirms that the joint controller exercises the multi-family action space where the frontier rewards it (high-SNR AWGN) and otherwise reduces to single-family rate-and-segment adaptation, which is consistent with the goodput separation of Table 13.

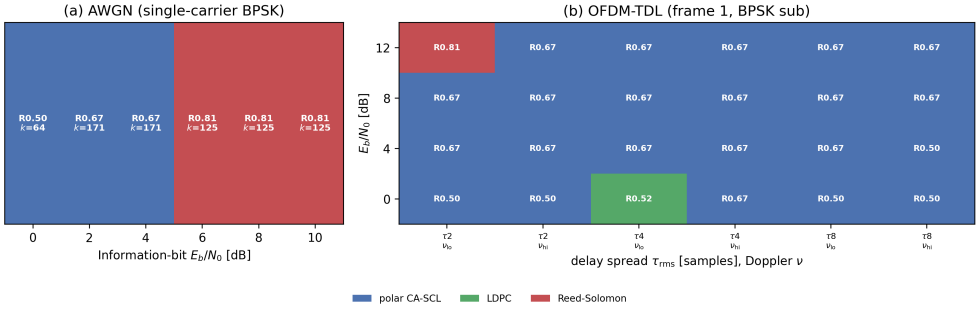


Figure 9. Selected-action map of the joint controller: FEC family (cell color), code rate (R), and block length k versus channel state. (**a**) AWGN: polar at low SNR, switching to RS above 6 dB. (**b**) OFDM-TDL: polar dominates, RS appears only in the high-SNR, low-delay, low-Doppler corner.

6.7. The Gain Needs Fresh CSI, and a Lightweight Predictor Suffices (RQ3)

E3 sweeps the feedback delay/prediction horizon H and an additive prediction-error standard deviation σ (as a fraction of each state axis's range), and it summarizes each operating point by the retained gain (joint−fixed-best)/(oracle−fixed-best): the fraction of the oracle-over-static adaptation headroom that survives that prediction quality (1 = full gain, 0 = collapses to the static code, negative = adapting on bad CSI is worse than not adapting). Table 14 reports these findings.

Table 14. Retained adaptation gain (E_b/N_0-averaged) versus feedback delay H and prediction error σ; persistence predictor.

	AWGN		**OFDM-TDL**	
H	$\sigma = 0$	$\sigma = 0.25$	$\sigma = 0$	$\sigma = 0.25$
0	1.00	−0.09	1.00	0.22
1	0.56	−0.23	0.55	0.02
2	0.40	−0.31	0.43	−0.05
5	0.17	−0.43	0.27	−0.14

Two messages follow. First, the joint gain requires reasonably fresh CSI: the retained gain falls monotonically with both the feedback delay (AWGN 1.00 → 0.56 → 0.40 → 0.17 for $H = 0, 1, 2, 5$) and the prediction error. It is largely retained only for roughly $H \leq 1$ and $\sigma \leq 0.1$, and it collapses beyond that—going *negative* at large H or σ, where adapting on stale or noisy state estimates underperforms the static code (Figures 10 and 11). This quantifies the cost of the long acoustic propagation delay and outdated feedback [6], and it justifies the optional conservative margin of the controller (Section 4.3): when a recent prediction error is high, a positive α makes the controller adapt less aggressively. Reassuringly, degradation is graceful—the controller's *absolute* goodput (SNR averaged) moves by under 10% across the H sweep at $\sigma = 0$ (AWGN 0.534 → 0.491, TDL 0.357 → 0.345), falling back toward the fixed-best code (0.482/0.341) rather than below it for moderate prediction quality. Because the fading-channel retained-gain denominator (oracle − fixed-best) is small (≈0.001–0.007, Section 6.6), individual TDL retained-gain entries near zero are ill-conditioned—a small ratio of small quantities—rather than statistically robust; the paired analysis of Section 6.6 shows the underlying joint−single-family differences on fading are tiny though consistent (a sub-0.8% prediction-overhead deficit), so we read the retained-

gain qualitatively—by sign and trend—rather than as precise values, and the robust signal is the monotone decay with H and σ that holds on both channels.

Second, and importantly for a low-power underwater modem, a learned predictor is not required here: a two-layer LSTM channel-state predictor gave no consistent improvement over the lightweight persistence predictor on either channel, because for the first-order autoregressive state model of Section 5, the most-recent observation is already the near-optimal (conditional-mean) predictor by construction (e.g., AWGN $H = 2, \sigma = 0$: retained gain 0.40 for persistence vs. 0.36 for the LSTM, within Monte-Carlo noise). This is specific to the Markov state process used here and is not a claim that learned prediction cannot help underwater. To probe that scope directly, we measured the one-step-ahead prediction RMSE on (a) real field SNR time series from two West Sea deployments (108 runs, 6294 frames; `results/west_sea_e11.csv`) and (b) a higher-order AR(2) synthetic process, comparing persistence, a 1-tap, and a 2-tap (AR(2)) least-squares predictor. The reduction in RMSE from the second tap measures the exploitable structure beyond lag one: it is essentially nil for the AR(1) process (-0.2%, persistence suffices, as above) but 13.5% for the AR(2) process and 3.2% on the real West Sea traces—i.e., the real channel state *does* carry a modest higher-order structure that a richer (or learned) predictor could exploit, which the AR(1) model by construction does not. The "persistence is enough" result is therefore a property of the first-order trajectory model, not a general claim; a richer or measured channel-state process is the regime where learned prediction could pay off. The joint-adaptation contribution (RQ2) does not depend on the predictor choice, and we adopt persistence throughout for the trajectory model studied.

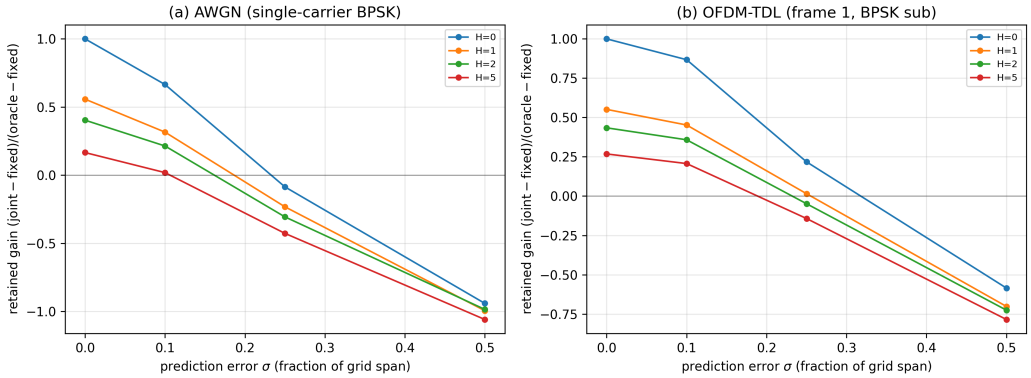


Figure 10. Retained adaptation gain (joint−fixed)/(oracle−fixed) versus prediction error σ, one line per feedback delay H, (**a**) AWGN and (**b**) OFDM-TDL. The gain is recovered for $\sigma \lesssim 0.1$ and collapses (turning negative) as σ grows.

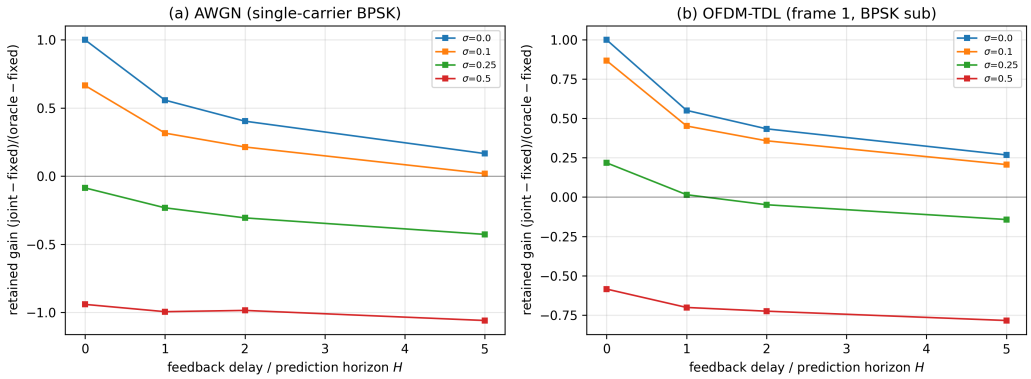


Figure 11. Retained adaptation gain versus feedback delay H, one line per prediction error σ. The gain falls monotonically with H, going negative for $\sigma \geq 0.25$ by $H = 5$ over fading; the persistence and LSTM predictors overlap, so persistence is used.

6.8. External Validity: Replay over Measured-Channel Statistics

The RQ1–RQ3 results are obtained over the synthetic statistical TDL channel of Section 5 (Rician $K = +5$ dB). To test whether they hold on a real underwater multipath, we replay the channel from the statistics of measured channel impulse responses from the authors' prior field campaigns—395 quality-controlled CIR measurements at $f_c = 30$ kHz from East Sea and South Sea deployments, with RMS delay spread 0.05–18 ms, Rician K-factor -21 to -2 dB, and 4–389 resolved paths. The measured channels are markedly *harsher* than the synthetic model on one axis and *richer* on another: their K-factor is uniformly negative (median ≈ -5 dB—the channel is diffuse/Rayleigh-dominated rather than the specular $+5$ dB synthetic model), while their multipath spans many resolved paths over a wide delay window. We map the dominant measured marginals—K-factor, RMS delay spread (onto the exponential power-delay-profile decay), and path richness (onto the tap count)—onto the same validated frequency-domain CP-OFDM TDL model, and we re-run the four fading-robust families over the synthetic baseline and three representative measured operating points: a near-range South Sea channel ($K = -3.6$ dB), a longer-delay East Sea channel ($K = -8$ dB), and a harsh deep-Rayleigh tail ($K = -12$ dB, long delay). This is a *statistics-driven* replay—the measured channel-parameter distribution drives the model—rather than a raw tap-level convolution of each recorded response; we complement it below with a *recorded-channel* replay over the public Watermark benchmark, which drives the per-subcarrier channel directly from a measured time-varying impulse response. Recorded UWA links to rank FEC families have been replayed directly on in situ signals [36]; our replay is complementary in intent, driving the validated TDL model with the measured parameter *distribution* to test whether the controller's benchmark generalizes beyond the synthetic channel rather than scoring codes over a single recorded link.

Table 15 and Figure 12 report the comparison, and the RQ1 family ordering survives: polar CA-SCL is the highest-goodput family on every channel, with LDPC the usual runner-up on the measured channels and BCH taking second only on the milder synthetic channel at 8 dB, while the high-rate RS and BCH codes stay throughput-oriented and competitive only at higher SNR, and the crossover persists. Two effects appear with opposite signs. At a low SNR, the measured channels punish the weak high-rate codes, as the negative K-factor removes the specular component: at 4 dB, the East Sea channel makes RS fail outright (BLER ≈ 1.0, goodput 0.00 vs. 0.08 synthetic) and roughly halves BCH's goodput. At mid-to-high SNRs, the rich frequency diversity instead aids the strong codes, whose codewords span many subcarriers: polar reaches full goodput (0.53) by 8 dB over the East Sea and harsh channels against 0.44 over the eight-tap synthetic channel. The measured statistics thus reinforce the conclusion: the strong channel-matched codes are even more clearly preferred and the high-rate codes even more clearly high-SNR-only, so the RQ1 frontier and the joint-adaptation rationale hold beyond the synthetic model.

Table 15. ARQ effective goodput ($B = 256$ B) of the four fading-robust families over the synthetic TDL ($K = +5$ dB) and three measured-statistic channels (South Sea -3.6, East Sea -8, harsh -12 dB), at 4 and 8 dB. Polar CA-SCL leads on every channel; LDPC and BCH trade second (RQ1 ordering preserved). Best per column in bold.

	Goodput @ 4 dB				Goodput @ 8 dB			
FEC Family	**Synth**	**S. Sea**	**E. Sea**	**Harsh**	**Synth**	**S. Sea**	**E. Sea**	**Harsh**
Polar CA-SCL	**0.251**	**0.260**	0.236	0.224	**0.444**	**0.477**	**0.526**	**0.526**
LDPC	0.198	0.203	**0.244**	**0.238**	0.289	0.308	0.344	0.344
BCH	0.194	0.153	0.065	0.051	0.310	0.293	0.262	0.276
Reed–Solomon	0.076	0.039	0.001	0.000	0.186	0.155	0.056	0.055

Recorded-channel replay over the Watermark benchmark. To validate the family ordering on an *actual recorded* channel rather than a parameter-mapped one, we replay the four fading-robust families over the NOF1 channel of the public Watermark benchmark [37] (Norway–Oslofjord, $f_c = 14$ kHz, ≈16 kHz band—numerology close to OFDM frame 1). For each OFDM symbol, we take one snapshot of the measured time-varying impulse response $h(\tau, t)$ (258 delay taps × 2048 snapshots per file, 60 files), transform it to the per-subcarrier frequency response, normalize to unit average power so E_s/N_0 is comparable, and run the standard frequency-domain matched-filter detection; a fresh recorded snapshot per OFDM symbol reproduces the measured Doppler across blocks. (A numerology caveat applies: the subcarrier grid is matched to the recording and the per-snapshot power is normalized, so the replay isolates the measured *small-scale* frequency-selective and time-varying fading.) Table 16 reports the result. The RQ1 family ordering is preserved on the real channel at every SNR—polar CA-SCL and LDPC are the reliability pair (polar best), and the high-rate BCH and RS codes trail and recover only at high SNR—and the goodput crossover is pushed to a *higher* SNR than on the synthetic model (the recorded NOF1 channel is harsher: even at 12 dB, BCH's goodput, 0.36, has not yet overtaken polar's 0.39). The recorded-channel replay therefore confirms, on a real measured multipath, both the RQ1 ordering and the channel-dependence of the crossover. This validates the RQ1 family ordering on a recorded channel; it does not re-run the closed-loop E2/E3 controller over recorded trajectories, which remains future work (Section 7.2).

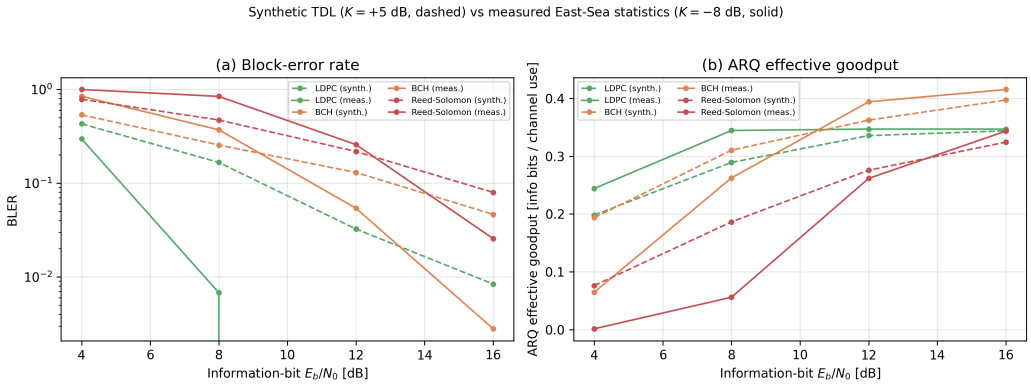


Figure 12. Replay over measured-channel statistics: (**a**) BLER and (**b**) ARQ goodput versus E_b/N_0 for the four fading-robust families over the synthetic TDL ($K = +5$ dB, dashed) and a measured East-Sea channel ($K = -8$ dB, solid). The RQ1 ordering is preserved.

Table 16. Recorded-channel replay over the Watermark NOF1 channel [37] (BLER and ARQ goodput, $B = 256$ B; `results/watermark_e12.csv`). The RQ1 family ordering (polar < LDPC < BCH < RS in BLER) holds on the real recorded channel; the high-rate algebraic codes remain high-SNR-only with the crossover pushed beyond 12 dB on this harsh channel.

	4 dB		8 dB		12 dB	
Family	**BLER**	**GP**	**BLER**	**GP**	**BLER**	**GP**
Polar CA-SCL	**0.397**	**0.241**	**0.043**	**0.383**	**0.018**	**0.393**
LDPC	0.665	0.131	0.121	0.343	0.049	0.371
BCH	0.963	0.018	0.661	0.159	0.224	0.364
RS	0.979	0.010	0.779	0.105	0.346	0.310

Bold marks the best value in each column (lowest BLER; highest goodput).

7. Discussion

7.1. What the Results Imply

The results carry three design implications. The benchmark settles a question of principle: since the goodput-optimal code moves with both the channel state and the payload size (per-segment overhead penalizing short blocks, Figure 5), per-state code selection is an independent design variable with measurable goodput impact. The controller study then bounds what that lever buys. Multi-family adaptation helps only where the reliability–throughput frontier crosses between families—there is a high-SNR AWGN here, a $+10\%$ gain from the polar-to-RS switch—and collapses to a single-family rate and segment adaptation where one family is near-optimal throughout, as over the fading channel. We state that scope limit plainly: the joint controller matches, but does not beat, a well-chosen single-family baseline when the frontier does not cross. The gain can also be captured with low computational overhead: an argmax lookup over the measured table with a persistence predictor reaches 91–99% of the oracle, and a learned LSTM gave no improvement for our channel-state model, so a small embeddable controller is enough.

A methodological point recurs throughout: the FEC *construction* matters as much as the family. The same polar code is the weakest or the strongest family in the benchmark depending only on whether its frozen set is channel-matched (a Gaussian-approximation rather than an SNR-agnostic design), and soft decoding likewise reshapes the convolutional results. We therefore report and adapt the construction rather than the family name alone.

7.2. Limitations

Several limitations bound the results presented in this paper. (i) The fading link assumes genie CSI, while the $(3,3)$ pilot grid is nonetheless charged in the goodput accounting (through $\varphi = 8/9$) without being used for estimation: the comparison is therefore *conservative* for throughput (the pilot overhead is paid for but not exploited) yet *optimistic* for reliability (the channel estimate is idealized), which are two idealizations that push the goodput crossover in opposite directions. A parametric imperfect-CSI sensitivity (Section 6.5, Table 12) bounds the reliability optimism—an LS-at-pilot estimate roughly doubles BLER but leaves the family ordering unchanged—while a full pilot-LS interpolation on the $(3,3)$ grid remains the next realism step and is expected to shift the fading results toward lower goodput. (ii) The fading channel is a statistical TDL model; we validated the RQ1 ordering both against the statistics of 395 measured channel impulse responses (Section 6.8) and against an actual recorded-channel replay over the public Watermark NOF1 benchmark (Table 16); a raw tap-level replay of the authors' own field CIRs remains future work. (iii) BCH and RS are reported with both bounded-distance hard-decision and soft-decision (Chase-II) decoding (Table 10); ordered-statistics decoding (OSD), a still-stronger soft variant, remains future work [34]. (iv) The soft convolutional code is now evaluated over fading (Table 7) but is dominated at every SNR—its long ($n = 524$) codeword is not bit-interleaved across the independent OFDM-symbol fades, so the deep-fade tail produces burst errors the Viterbi decoder cannot correct; the turbo decoder remains AWGN-only here (its commpy probability-domain BCJR needs a scalar noise variance, invalid per-subcarrier over frequency-selective fading; a max-log-MAP port is future work). Only frame 1 with BPSK is reported over fading. (v) The channel-state trajectory is a design model (a first-order autoregressive walk), and the negative LSTM result is specific to this process and a limited training budget rather than a claim that learned prediction never helps underwater.

7.3. Latency and Deployment

Latency enters the controller formulation through a decoder-cost term. Beyond the complexity-order proxy used in the action space, we measured the actual per-block decode

time of each family (median over many blocks on a single CPU thread; Table 17). The measured *ordering* matches the proxy: hard-decision BCH/RS are cheapest (≈0.14 ms), the short LDPC and the soft-decision (Chase-II) algebraic and CA-SCL polar decoders are intermediate (≈2–7 ms), and the long-block convolutional and turbo decoders are one to two orders of magnitude more expensive (≈114–150 ms). The Python wall-clock is library-overhead-dominated and is *not* a deployment latency; we report it only as a relative ordering and take absolute feasibility from cited fixed-point/FPGA implementations—fast simplified successive-cancellation-list decoding for polar codes [32] and low-latency short-blocklength LDPC decoder hardware [33]. The measured ordering nonetheless makes the budget concrete: a per-message decode budget set to a fraction of the OFDM useful-symbol duration (64–171 ms, Table 2) is non-binding for the strong short-block codes but prunes the long-block convolutional and turbo actions first—e.g., a 32 ms budget excludes both conv and turbo while retaining polar, LDPC, and the algebraic codes. (This finite-budget pruning is an added sensitivity; the headline RQ2/RQ3 results use an unconstrained budget, so pruning does not reshape them.) The controller itself adds only a state read and a table lookup per transmission, which are negligible against the acoustic propagation delay.

Table 17. Measured per-block decode time (median; single CPU thread, realistic LLRs at 4 dB; `results/latency_e10.csv`). A *relative* ordering only—Python wall-clock is not deployment latency; absolute feasibility is from cited HW decoders [32,33]. A budget at a fraction of the 64–171 ms OFDM symbol prunes conv/turbo first.

	RS Hard	BCH Hard	LDPC	RS Soft	Polar L8	BCH Soft	Turbo	Conv
median ms/block	0.14	0.14	1.8	4.8	5.7	6.5	114	150

8. Conclusions and Future Work

We addressed forward error correction for the short messages (below ∼10 kB) that dominate underwater acoustic sensor and control links, where FEC operates on short, finite-length blocks over a time-varying multipath-and-Doppler channel. We argued that the payload segment (block) size is an under-used degree of freedom—it sets the FEC block length and hence the finite-blocklength penalty—and we proposed a per-transmission controller that jointly selects the FEC scheme, the code rate, and the segment size from a prediction of the near-term channel state, paired with a controlled, like-for-like short-block FEC benchmark.

Three findings stand out. First, no single FEC configuration dominates: under the unified ARQ effective-goodput metric, the reliability–throughput frontier crosses to the high-rate codes at an SNR that shifts with the channel (≈6 dB on AWGN, ≈10 dB over fading) and at an optimum that also depends on the payload size, so the best code is channel- and payload-dependent; the code construction (a channel-matched polar frozen set, soft decoding) matters as much as the family. Second, the segment-size degree of freedom is the dominant adaptation lever: across both channels at low-to-mid SNR values, it captures essentially all of the achievable adaptation gain, and joint adaptation additionally switches the FEC *family* only where the reliability–throughput frontier crosses between families—a bounded extra gain of up to 10% at high-SNR AWGN (polar to RS), 16% over a fixed no-CSI code, and 91–99% of an oracle; where one family is near-optimal throughout, as over fading, the controller collapses to single-family rate-and-segment adaptation and matches that baseline rather than beating it. Third, the controller can be simple: an argmax lookup fed by a persistence predictor recovers nearly all of this gain, a learned LSTM gave no improvement for the first-order channel-state model studied, and a robustness analysis shows the gain eroding with feedback delay and prediction error, which motivates a conservative margin. Replaying the link over the statistics of 395 measured

channel impulse responses preserves the family ordering and the crossover, favoring the strong channel-matched codes even more clearly on a real multipath (Section 6.8), so the conclusions are not specific to the baseline synthetic TDL setting.

Future work follows directly from the study's remaining limitations: a raw tap-level replay of the authors' own recorded channel impulse responses—extending the Watermark recorded-channel replay and the statistics-driven measured-channel replay reported here—together with a full pilot-LS channel estimator to replace the genie CSI (beyond the parametric imperfect-CSI bound of Table 12); ordered-statistics decoding for the algebraic codes (beyond the Chase-II soft decoding added here) and a fading-aware turbo/convolutional decoder to bring those families onto the fading axis; a richer or less predictable channel-state process, the regime—suggested by the higher-order structure in the West Sea traces and the AR(2) test (Section 6.7)—in which a learned predictor could justify its cost over persistence; and an on-device implementation of the lookup controller on an embedded acoustic modem, which is consistent with its lightweight design.

Author Contributions: Conceptualization, T.I. and S.K.; methodology, S.K. and S.P.; software, S.K. and S.P.; validation, S.K. and T.I.; formal analysis, S.K.; investigation, S.K. and S.P.; data curation, S.P.; writing—original draft preparation, S.K.; writing—review and editing, T.I.; visualization, S.K.; supervision, T.I.; project administration, T.I.; funding acquisition, T.I. All authors have read and agreed to the published version of the manuscript.

Funding: This paper was supported by the Korea Research Institute for defense Technology Planning and Advancement (KRIT)—grant funded by the Korea government (DAPA [Defense Acquisition Program Administration]) (No. KRIT-CT-23-026, Integrated Underwater Surveillance Research Center for Adapting Future Technologies, 2026).

Data Availability Statement: The Watermark benchmark dataset analyzed in Section 6.8 is publicly available from the Norwegian Defence Research Establishment (FFI) at https://www.ffi.no/en/research/watermark (accessed on 26 June 2026) [37]. The simulation code and the generated result files (`results/*.csv`) that support the reported tables and figures are available from the corresponding author upon reasonable request. The measured channel impulse responses from the authors' own field campaigns (Section 6.8) are summarized in prior work and available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AMC	adaptive modulation and coding
ARQ	automatic repeat request
AWGN	additive white Gaussian noise
BCH	Bose–Chaudhuri–Hocquenghem code
BCJR	Bahl–Cocke–Jelinek–Raviv algorithm
BLER	block-error rate
CA-SCL	CRC-aided successive-cancellation list
CIR	channel impulse response
CRC	cyclic redundancy check
CSI	channel-state information
FBL	finite blocklength
FEC	forward error correction
GA	Gaussian approximation
LDPC	low-density parity-check code
LLR	log-likelihood ratio
LSTM	long short-term memory

OFDM	orthogonal frequency-division multiplexing
PDP	power-delay profile
RS	Reed–Solomon code
SCL	successive-cancellation list
SNR	signal-to-noise ratio
TDL	tapped-delay-line
UWA	underwater acoustic

References

1. Diamant, R.; Lampe, L. Adaptive Error-Correction Coding Scheme for Underwater Acoustic Communication Networks. *IEEE J. Ocean. Eng.* **2015**, *40*, 104–114. [CrossRef]
2. Durisi, G.; Koch, T.; Östman, J.; Polyanskiy, Y.; Yang, W. Short-Packet Communications over Multiple-Antenna Rayleigh-Fading Channels. *IEEE Trans. Commun.* **2016**, *64*, 618–629. [CrossRef]
3. Liu, Y.; Effros, M. Finite-Blocklength and Error-Exponent Analyses for LDPC Codes in Point-to-Point and Multiple Access Communication. In *2020 IEEE International Symposium on Information Theory (ISIT)*; IEEE Press: New York, NY, USA, 2020. [CrossRef]
4. Jing, L.; Dong, C.; He, C.; Shi, W.; Yin, H. Adaptive Modulation and Coding for Underwater Acoustic Communications Based on Data-Driven Learning Algorithm. *Remote Sens.* **2022**, *14*, 5959. [CrossRef]
5. Zhou, J.; Sun, X. Adaptive Modulation and Coding for Underwater Acoustic. In Proceedings of the 2023 IEEE 7th Information Technology and Mechatronics Engineering Conference (ITOEC), Chongqing, China, 15–17 September 2023. [CrossRef]
6. Zhang, Y.; Zhu, J.; Wang, H.; Shen, X.; Wang, B.; Dong, Y. Deep Reinforcement Learning-Based Adaptive Modulation for Underwater Acoustic Communication with Outdated Channel State Information. *Remote Sens.* **2022**, *14*, 3947. [CrossRef]
7. Jing, L.; Tang, Y.; He, C.; Yin, H. Adaptive Packet Coding for Reliable Underwater Acoustic Communications. *Remote Sens.* **2022**, *14*, 4712. [CrossRef]
8. Falk, M.; Bauch, G.; Nissen, I. On Channel Codes for Short Underwater Messages. *Information* **2020**, *11*, 58. [CrossRef]
9. Al-Kadhimi, A.M.; Abdelkareem, A.E.; Tsimenidis, C. P-LDPC coded image transmission with OFDM over underwater acoustic channel. *Acta Polytech.* **2024**, *64*, 68–76. [CrossRef]
10. Lidström, V. Evaluation of Polar-Coded Noncoherent Acoustic Underwater Communication. *IEEE J. Ocean. Eng.* **2025**, *50*, 448–461. [CrossRef]
11. Hyla, J.; Sułek, W. Short Blocklength Nonbinary Raptor-Like LDPC Coding Systems Design and Simulation. *IEEE Access* **2025**, *13*, 3849–3863. [CrossRef]
12. Barreto, G.; Simão, D.H.; Pellenz, M.E.; Souza, R.D.; Jamhour, E.; Penna, M.C.; Brante, G.; Chang, B.S. Energy-Efficient Channel Coding Strategy for Underwater Acoustic Networks. *Sensors* **2017**, *17*, 728. [CrossRef] [PubMed]
13. Agarwal, S.; Hagedorn, A.; Trachtenberg, A. Adaptive rateless coding under partial information. In Proceedings of the 2008 Information Theory and Applications Workshop, San Diego, CA, USA, 27 January–1 February 2008. [CrossRef]
14. Wang, X.; Chen, W.; Cao, Z. Throughput-Efficient Rateless Coding with Packet Length Optimization for Practical Wireless Communication Systems. In Proceedings of the 2010 IEEE Global Telecommunications Conference GLOBECOM 2010, Miami, FL, USA, 6–10 December 2010. [CrossRef]
15. Huang, L.; Zhang, Q.; Tan, W.; Wang, Y.; Zhang, L.; He, C.; Tian, Z. Adaptive modulation and coding in underwater acoustic communications: a machine learning perspective. *EURASIP J. Wirel. Commun. Netw.* **2020**, *2020*, 203. [CrossRef]
16. Fu, Q.; Song, A. Adaptive Modulation for Underwater Acoustic Communications Based on Reinforcement Learning. In Proceedings of the OCEANS 2018 MTS/IEEE Charleston, Charleston, SC, USA, 22–25 October 2018. [CrossRef]
17. Zhang, Y.; Zhu, J.; Liu, Y.; Wang, B. Underwater Acoustic Adaptive Modulation with Reinforcement Learning and Channel Prediction. In Proceedings of the 15th International Conference on Underwater Networks & Systems, Shenzhen, China, 22–24 November 2021. [CrossRef]
18. Ahmed, R.; Stojanovic, M. Reliable communication using packet coding for underwater acoustic channels. In Proceedings of the 2016 IEEE 17th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC), Edinburgh, UK, 3–6 July 2016. [CrossRef]
19. Zhai, Y.; Li, J.; Feng, H. Research on Polar Coding Application for Underwater Acoustic OFDM Communication System. In Proceedings of the 2020 IEEE 6th International Conference on Computer and Communications (ICCC), Chengdu, China, 11–14 December 2020. [CrossRef]
20. Choi, G.; Lee, N. Deep-Polar Codes: An Efficient Error Correction Code for Short Blocklength Transmission. In Proceedings of the 2023 IEEE Globecom Workshops (GC Wkshps), Kuala Lumpur, Malaysia, 4–8 December 2023. [CrossRef]

21. Radosevic, A.; Duman, T.M.; Proakis, J.G.; Stojanovic, M. Channel prediction for adaptive modulation in underwater acoustic communications. In Proceedings of the OCEANS 2011 IEEE, Santander, Spain, 6–9 June 2011. [CrossRef]
22. Zhu, Z.; Tong, F.; Zhou, Y.; Zhang, Z.; Zhang, F. Deep Learning Prediction of Time-Varying Underwater Acoustic Channel Based on LSTM with Attention Mechanism. *J. Mar. Sci. Appl.* **2023**, *22*, 650–658. [CrossRef]
23. Liu, L.; Cai, L.; Ma, L.; Qiao, G. Channel State Information Prediction for Adaptive Underwater Acoustic Downlink OFDMA System: Deep Neural Networks Based Approach. *IEEE Trans. Veh. Technol.* **2021**, *70*, 9063–9076. [CrossRef]
24. Liu, H.; Ma, L.; Wang, Z.; Qiao, G. Channel Prediction for Underwater Acoustic Communication: A Review and Performance Evaluation of Algorithms. *Remote Sens.* **2024**, *16*, 1546. [CrossRef]
25. Polyanskiy, Y.; Poor, H.V.; Verdú, S. Channel Coding Rate in the Finite Blocklength Regime. *IEEE Trans. Inf. Theory* **2010**, *56*, 2307–2359. [CrossRef]
26. Miao, S.; Kestel, C.; Johannsen, L.; Geiselhart, M.; Schmalen, L.; Balatsoukas-Stimming, A.; Liva, G.; Wehn, N.; Brink, S. Trends in Channel Coding for 6G. *Proc. IEEE* **2024**, *112*, 653–675. [CrossRef]
27. Cooper-Baldock, Z.; Panteli, E.; Santos, P.E. A state-of-the-art survey and benchmarking of Adaptive Modulation and Coding for Underwater Acoustic Communications. *Appl. Ocean. Res.* **2026**, *170*, 105018. [CrossRef]
28. Park, S.; Kim, S.; Lee, H.; Im, T. Reliable Underwater Acoustic Telemetry for Ocean Remote Sensing Platforms: Channel-Prediction-Based Adaptive Polar–Raptor Coded OFDM. *Remote Sens.* **2026**, *18*, 1747. [CrossRef]
29. Li, B.; Tong, F.; Zheng, S.; Chen, D. Bit-error rate gradient descent Doppler estimation for underwater acoustic OFDM communication. *Appl. Acoust.* **2021**, *171*, 107557. [CrossRef]
30. Wang, S.; Li, W.; Guo, R. Underwater Acoustic Channel Prediction in delay-Doppler Domain based on Reinforcement Learning. In Proceedings of the OCEANS 2022 Hampton Roads, Hampton Roads, VA, USA, 17–20 October 2022. [CrossRef]
31. Liu, Y.; Lee, J.; Sun, C.; Han, Y.; Chen, W. Energy Efficient Scheduling for Short Packet Communications With Finite Blocklength Coding. *IEEE Trans. Green Commun. Netw.* **2024**, *8*, 1645–1660. [CrossRef]
32. Hashemi, S.A.; Condo, C.; Gross, W.J. Fast Simplified Successive-Cancellation List Decoding of Polar Codes. In Proceedings of the 2017 IEEE Wireless Communications and Networking Conference Workshops (WCNCW), San Francisco, CA, USA, 19–22 March 2017. [CrossRef]
33. Nonaca, D.; Guichemerre, J.; Wiesmayr, R.; Tunali, N.E.; Studer, C. A 14ns-Latency 9Gb/s 0.44mm^2 62pJ/b Short-Blocklength LDPC Decoder ASIC in 22FDX. In Proceedings of the 2025 IEEE European Solid-State Electronics Research Conference (ESSERC), Munich, Germany, 8–11 September 2025. [CrossRef]
34. Cai, J.; Tomlinson, M.; Tjhai, C.; Ambroze, M.; Ahmed, M. Comparison of Concatenated Reed-Solomon Coding with Hard-decision to Soft-decision LDPC Coding for the Rayleigh Fading Channel. In Proceedings of the 2006 IEEE Information Theory Workshop—ITW '06, Chengdu, China, 22–26 October 2006. [CrossRef]
35. Wulandari, P.S.; Utami, E.Y.D.; Timotius, I.K. Bit Error Rate Performance Comparison of Low-Density Parity-Check Code and Polar Code. In Proceedings of the 2023 29th International Conference on Telecommunications (ICT), Toba, Indonesia, 8–9 November 2023. [CrossRef]
36. Pelekanakis, K.; Paglierani, P.; Alvarez, A.; Alves, J. Comparison of error correction codes via optimal channel replay of high north underwater acoustic channels. *Comput. Netw.* **2024**, *242*, 110270. [CrossRef]
37. van Walree, P.A.; Socheleau, F.X.; Otnes, R.; Jenserud, T. The Watermark Benchmark for Underwater Acoustic Modulation Schemes. *IEEE J. Ocean. Eng.* **2017**, *42*, 1007–1018. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.