A Standards-Aligned, Deadline-Aware Data-Link Scheduling Profile for Reliable UAV Suspended-Load Anti-Swing Communication

Kai Tang

College of Engineering, Huazhong Agricultural University, Wuhan 430070, Hubei, China
E-mail: tangkai@acedu.ac.cn

Received 21 May 2026; Accepted 11 July 2026

Abstract

Unmanned aerial vehicles (UAVs) carrying cable-suspended loads are increasingly used in emergency delivery, power-line inspection support, and engineering logistics, yet their data links must carry control-critical anti-swing traffic together with bulk monitoring images over the same scarce wireless capacity. Existing ICT standards specify differentiated transport (3GPP TS 23.501 5QI, MQTT v5 QoS, DDS QoS), aerial connectivity (3GPP Release 17/18 TS 22.125 and TS 23.256), and IoT abstraction (W3C Web of Things), but leave the upstream mapping from instantaneous application semantics (such as swing risk and link risk) to QoS labels as a deployment-specific configuration concern. This paper develops a vendor-neutral, deadline-aware data-link scheduling profile that layers on top of the W3C Web of Things architecture and standard publish/subscribe and RESTful interfaces, positioned as a complementary application-layer mapping policy rather than a replacement transport mechanism. Traffic is partitioned into four conformance-checkable packet classes, and a dynamic priority is computed from swing risk, link risk, and deadline urgency. The profile is

Journal of ICT Standardization, Vol. 14_4, 563–602.
doi: 10.13052/jicts2245-800X.1443
© 2026 River Publishers

specified through five conformance clauses and three black-box test cases observable at the Thing Description boundary. On a reproducible Python testbed driven by 10 independent subseeds per scenario, the profile reduces average control-packet delay by 34.51% and deadlinemiss rate by 82.48% on a weak link relative to FIFO (which does not settle within the simulation horizon); against the strongest converged baseline (fixed priority), the deadlinemiss rate decreases by 19.95 percentage points (a 57.5% relative reduction) and the RMS swing angle by 4.5%. A Gilbert-Elliott bursty-channel extension widens the gap to fixed priority from 19.95 to 28.43 percentage points; an ablation analysis attributes the improvement to all three mechanisms (dynamic priority, redundant control transmission, and image suppression), with no single mechanism dominating; and the low-priority monitoring class is documented to remain above the declared 15% throughput-share floor under every tested scenario. The work is presented as pre-standardization research toward an interoperable scheduling profile for UAV communication, with a full reproducibility package (simulation kernel, baselines including an Age-of-Information-greedy variant, ablation variants, Gilbert-Elliott channel generator, and conformance test probe) released alongside the paper.

Keywords: ICT standardization, interoperability profile, pre-standardization research, W3C Web of Things, 3GPP UAS connectivity, deadline-aware scheduling, UAV data link, Age of Information.

1 Introduction

A modern unmanned aerial vehicle (UAV) mission is built on standardized ICT building blocks: cellular and wireless data links, RESTful and publish/subscribe telemetry interfaces, and edge gateways that expose airborne sensors and actuators as interoperable services [1]. These pieces sit within several still-evolving ICT standardization tracks: the W3C Web of Things (WoT) architecture for IoT abstraction [1], 3GPP work on Unmanned Aerial System (UAS) connectivity for cellular-based aerial access (Release 17/18 Stage 1/2 specifications TS 22.125 [2] and TS 23.256 [3]), and IoT messaging standards such as MQTT v5 [4] and OMG DDS [5], with the 5G QoS class identifier (5QI) framework defined in TS 23.501 [6] providing the underlying differentiated transport, alongside a broader body of de facto industry protocols for telemetry and edge computing [7]. The gap this paper addresses lies between these tracks. Current standards specify how UAV

telemetry is transported, addressed, and exposed; they do not specify the upstream mapping from instantaneous application semantics (swing risk, link risk, deadline urgency) to QoS labels when wireless conditions degrade, and that mapping is left to deployment. The paper accordingly treats the problem as one of pre-standardization research and proposes a vendor-neutral, deadline-aware data-link scheduling profile that layers on top of the W3C WoT abstraction and standard publish/subscribe and RESTful interfaces, positioned as a complementary application-layer mapping policy rather than a replacement transport mechanism. UAV suspended-load anti-swing control is used as the driving use case, but the proposed contribution is the data-link profile itself, evaluated against classical real-time scheduling baselines on a reproducible testbed.

The use case of suspended-load UAV transport is chosen because it places tight, easily measurable constraints on communication. A cable-suspended payload allows a multirotor to carry materials without a rigid gripper or dedicated cargo bay, which suits emergency logistics, engineering material delivery, power-line inspection support, and rapid resupply. The mechanical penalty is well known: the payload behaves as a lightly damped pendulum coupled to the airframe, and acceleration, curved tracking, or gusts can excite swing that enlarges tracking error and shrinks safety margins. Suspended-load flight has been studied extensively through multirotor modeling, aggressive trajectory generation, nonlinear geometric control, and cable-payload stabilization [8].

Most control-oriented studies assume that swing state and vehicle state reach the controller with negligible delay. The assumption simplifies control design but is rarely satisfied in deployment. In a fielded UAV mission, the aircraft, the ground station, and any edge-computing node share one wireless link, on which swing angle, angular velocity, attitude, position, mission status, monitoring images, and control feedback all compete for bandwidth. The link itself is unstable, since bandwidth, delay, jitter, and packet loss vary with range, line-of-sight, electromagnetic interference, and nearby infrastructure. Surveys of UAV wireless communication have noted that aerial platforms face mobility, interference, coverage, and energy constraints different from those of terrestrial users [9, 10], and 3GPP has recognized this by treating aerial users as a separate study item in its UAS-related work.

For closed-loop anti-swing control, throughput is not the binding metric. A packet may be received correctly and still be useless: once it arrives after its control deadline, it cannot influence the current decision. The communication objective is therefore the timely availability of control-relevant information,

not the overall packet delivery ratio. Under conventional queueing a large monitoring-image packet can block the small swing-state packets behind it, so a scheduler that treats every packet alike can hold throughput steady while eroding control quality, a behavior that existing data-transport standards neither cause nor prevent, since they do not constrain how a link prioritizes its own traffic.

The scheduling profile proposed in this paper is swing-aware and link-adaptive. It classifies packets by operational importance, monitors the link and the swing state, and computes a dynamic priority. When link risk or swing risk rises, control-relevant packets are reinforced; on a weak link, low-priority monitoring images are compressed, delayed, or deprioritized. The profile is specified in implementation-independent terms (packet classes, a priority function, and conformance-checkable behavior), so that it can be implemented on top of existing transport and IoT standards rather than embedded in a proprietary stack. Operationally, the profile complements rather than replaces the underlying standards: 3GPP TS 23.501 5QI [6], MQTT v5 QoS [4], and DDS QoS [5] specify how a selected QoS label is transported, while the profile defines how the QoS label is selected from instantaneous application semantics; 3GPP TS 22.125 [2] and TS 23.256 [3] specify aerial connectivity requirements and identifiers, while the profile operates above this connectivity layer at the WoT Thing Description boundary [1]. The design draws on edge-assisted and cross-layer communication, where application value, network condition, and computation placement are considered jointly [11, 12], and uses EDF and fixed-priority scheduling [13, 14] as real-time scheduling baselines.

The contributions are threefold. First, a communication-control coupling model is established in which packet delay, packet loss, and deadline misses influence the suspended-load swing response through the freshness of the state used by the controller. Second, a dynamic scheduling model is developed that combines baseline packet priority, swing risk, link risk, and deadline urgency into a single decision variable, packaged as a vendor-neutral profile compatible with existing UAV communication standards. Third, a repeatable numerical experiment compares the proposed scheduler with FIFO, EDF, and fixed-priority scheduling, together with an additional Age of Information (AoI)-greedy baseline drawn from networked-control-systems (NCS) work, under four link scenarios that include a Gilbert-Elliott bursty extension; the analysis reports raw metrics with standard deviations over a 10-sub-seed family, relative improvements, normalized summaries, and a three-axis ablation that decomposes the contribution of the dynamic-priority

formula, redundant control transmission, and image suppression. The experimental setup is documented in enough detail that the profile's behavior is reproducible and conformance-checkable through the explicit clauses and test cases of Section 3.6, a basic requirement for any future standardization activity.

2 Related Work

2.1 Suspended-Load UAV Modeling and Anti-Swing Control

Multirotor modeling provides the basis for reasoning about vehicle attitude, position, and force generation [15], and trajectory generators such as minimum-snap planning extend this to aggressive maneuvering in confined spaces [16]. A suspended load adds a pendulum-like state that becomes critical once payload safety and delivery accuracy are involved. Sreenath et al. formulated geometric control and differential flatness for a quadrotor carrying a cable-suspended load, embedding the payload dynamics directly in the nonlinear tracking control [8]. Lee and colleagues extended the geometric approach to cooperative quadrotors sharing a suspended payload and to cable-suspended rigid-body transport, focusing on coordinate-free modeling and stability analysis [17-18]. This line of work establishes payload swing as a first-class dynamic state rather than a small external disturbance.

The aim of the present paper is not to replace high-fidelity suspended-load controllers but to address the communication layer that feeds them. A controller, however carefully designed, performs poorly when its swing-angle and angular-velocity samples are stale, so communication scheduling forms part of the control-support infrastructure. The scheduler determines which state sample is available when the control law runs and therefore shapes the feedback that the anti-swing controller actually receives. Recent anti-swing designs based on energy coupling, sliding mode, nonlinear MPC, and disturbance observers [19–23] have improved controller robustness, and this progress only makes the freshness of the underlying samples more decisive for the performance that can be realized in practice.

2.2 UAV Communication, Edge Computing, and Real-Time Scheduling

UAV links differ from fixed terrestrial links: the aircraft moves in three-dimensional space, line-of-sight conditions vary in flight, and altitude can

produce atypical interference patterns. Surveys of civil UAV applications and UAV wireless networks identify communication reliability, network security, interference management, and latency-sensitive transmission as the principal technical challenges [24]. The standardization community has reached a similar view: 3GPP's study on enhanced LTE support for aerial vehicles treats airborne users as a case requiring specific handling by cellular systems [25], and subsequent normative work has formalized the service requirements (Stage 1 TS 22.125 [2]) and system architecture (Stage 2 TS 23.256 [3]) for UAS connectivity, alongside W3C activity on IoT abstractions [1] and IoT messaging standards (MQTT v5 [4], OMG DDS [5]). A parallel line of NCS work treats freshness of control-relevant state as the primary scheduling objective: the AoI framework formalizes the time since the most recent successfully delivered update [26], and AoI-aware scheduling for cellular networked control systems has been shown to outperform throughput-optimal policies when the controlled plant is unstable [27]. The proposed scheduler relates to AoI scheduling but extends it: where AoI takes the age of the latest delivered packet as the optimization target, the proposed scheduler additionally incorporates explicit swing-risk and link-risk signals, treating freshness as necessary but not sufficient on a constrained UAV link.

Edge computing offers a complementary approach by moving computation and decisionmaking closer to the data source [28, 29]. In an anti-swing mission, an edge node or ground station can evaluate control decisions, mission state, and monitoring data. Edge support, however, only improves responsiveness when the underlying link gives time-critical data priority. If monitoring images dominate the link while swing-state packets queue behind them, no amount of edge computation compensates for the resulting input staleness. Edge processing and data-link scheduling therefore need to be designed together rather than separately.

Classical real-time scheduling provides natural baselines. FIFO is simple and predictable but blind to both packet value and deadline. EDF serves the packet with the earliest absolute deadline first and has strong guarantees in idealized real-time settings [13, 30]. Fixed-priority scheduling reserves the top tier for control traffic but cannot adapt as swing risk or link quality change. Buttazzo's comparison of rate-monotonic and EDF scheduling shows that the choice depends on overload behavior, implementation cost, and timing predictability [14, 31]. In the suspended-load case the overload is both computational and communicational, which is why a static policy is inadequate and an adaptive scheduler is required.

2.3 Research Gap

The literature reveals a gap between suspended-load control and data-link scheduling. Control studies typically assume timely sensing, while communication studies tend to optimize throughput, latency, or coverage without modeling the closed-loop consequence of stale swing-state samples. A parallel gap exists in standardization: current ICT standards for UAV connectivity (3GPP TS 22.125 [2], TS 23.256 [3], and the 5G QoS framework in TS 23.501 [6]) and for IoT messaging (W3C WoT [1], MQTT v5 [4], OMG DDS [5], HTTP/REST) define how data is transported and addressed, and they expose QoS labels, priorities, and reliability classes that downstream applications may choose. They do not, however, prescribe the upstream mapping from instantaneous application semantics, such as swing risk, link risk, and deadline urgency, to the QoS label that should be requested at any given moment. The scheduling profile proposed here addresses both gaps. It links packet priority with swing risk, link risk, and deadline urgency, and presents the result as an implementation-independent profile suitable for harmonization with existing standards, complementing rather than replacing the underlying transport mechanisms. The objective is to improve the timely availability of control-relevant data on weak and variable links in a form that can be specified, tested, and eventually standardized.

3 System Model and Scheduling Strategy

3.1 Overall System Architecture

The system has three parts (Figure 1): a UAV side modeled as a W3C WoT producer (Thing), an edge-side data-link scheduling function realized as a Servient, and a ground control station acting as a Consumer of telemetry over standard publish/subscribe and RESTful interfaces. The UAV side acquires suspended-load state, flight state, mission state, and monitoring information. The data-link layer monitors the link, computes dynamic priority, manages the queues, and applies reliability enhancement. The ground or edge side receives the state packets, evaluates anti-swing control decisions, and records performance metrics. The three parts close a feedback loop across the mechanical subsystem, the wireless data link, and the edge-side decision module, in which payload-state sensing supports anti-swing control [8] while the data-link and edge-computing components determine whether that sensing is fresh enough to be used [1,11]. Because each part is described in terms of standard roles, the architecture remains vendor-neutral and the scheduling

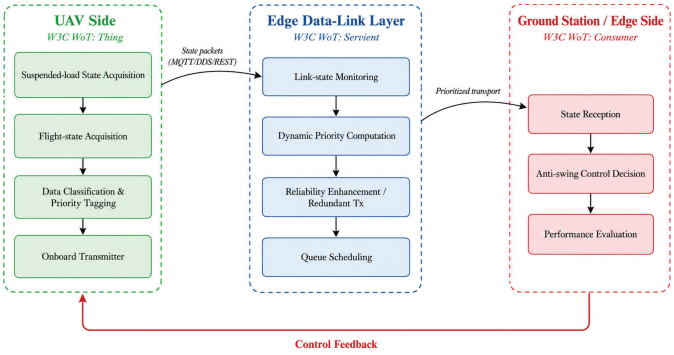


Figure 1 System architecture for reliable data-link scheduling in UAV suspended-load antiswing control. Each component is mapped to a W3C WoT role (Thing/Servient/Consumer) where applicable.

function is portable across conforming UAV platforms and ground stations. Operationally, the architecture is positioned as a complementary application-layer mapping policy on top of the underlying differentiated transport: 5QI labels defined in 3GPP TS 23.501 [6], MQTT v5 QoS [4], and DDS QoS [5] are exposed to the data-link scheduler as transport choices, while the application-semantic mapping (swing risk, link risk, deadline urgency $\rightarrow$ QoS label) is determined by the scheduler itself.

The scheduler resides in the edge data-link scheduling layer, not in the flight controller. This separation is deliberate: the anti-swing control law is left unchanged and only the timeliness and reliability of the information reaching the controller are affected. The control algorithm continues to decide what action to take [8]; the data-link scheduler decides whether the state behind that action is fresh enough to be meaningful within an edge-assisted communication architecture [1, 12]. Placing the scheduler in its own standardizable layer also enables interoperability: a UAV and a ground station from different vendors can work together provided that both conform to the same data-link profile.

Snippet 1 sketches the exposed Thing Description (TD) fragment for the anti-swing control event class, declaring the QoS expectations that the scheduler must honor at the protocol boundary. The Thing Description fragment is normative for any conforming implementation, while the underlying transport binding (MQTT, DDS, HTTP) is left to deployment.

Snippet 1 Example WoT Thing Description fragment for the anti-swing control event class.

```
{
 "@context": "https://www.w3.org/2022/wot/td/v1.1",
 "title": "uav.suspended−load.anti−swing",
 "events": {
  "swingState": {
   "data": { "type": "object" },
   "forms": [{
    "href": "mqtt://broker/uav/{thingId}/swingState",
    "contentType": "application/cbor",
    "op": ["subscribeevent"],
    "mqv:qos": 1,
    "qos:deadlineMs": 10,
    "qos:classBasePriority": 100,
    "qos:5qi": 87
   }]
  }
 }
}
```

3.2 Suspended-Load Dynamic Model

A simplified planar model (Figure 2) captures the dominant swing behavior. The schematic identifies the physical variables retained in the numerical model and distinguishes them from the communication variables introduced later. The control sampling period $T_c = 20$ ms is matched to the planar model's dominant natural period; the packet generation period for antiswing control is $T_p = 10$ ms, producing an oversampling factor of $T_c/T_p = 2$ at the producer side so that, after worst-case scheduling-induced delay on the medium link, at least one valid swing sample arrives within each control window. The relationship $T_c = 2T_p$ is therefore a deliberate design choice rather than an accidental match, and it sets the baseline against which deadline-miss rate becomes the binding metric. The intent is not to reproduce full threedimensional cable dynamics but to provide a transparent path by which communication delay and deadline miss reach the swing response. Let θ denote payload swing angle, l cable length, ω_n natural frequency, ζ damping ratio, a_u horizontal acceleration command, and d_θ an external disturbance

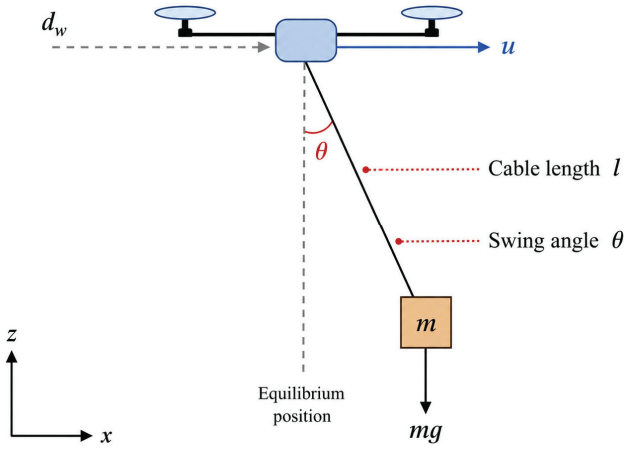


Figure 2 Simplified planar suspended-load dynamic model.

term. The planar dynamics are expressed as follows [18, 20]:

$$\ddot{\theta}(t) + 2\zeta\omega_n\dot{\theta}(t) + \omega_n^2 \sin\theta(t) = -\frac{a_u(t)}{l} + d_\theta(t) \tag{1}$$

$$\omega_n = \sqrt{\frac{g}{l}} \tag{2}$$

Following suspended-load feedback-control formulations in the UAV control literature [8, 23], the anti-swing command is computed from the most recent valid swing-state estimate available to the controller:

$$a_u(t) = k_p\hat{\theta}(t) + k_d\hat{\dot{\theta}}(t) \tag{3}$$

The swing angle and angular velocity entering the controller are not necessarily the true instantaneous values; they are the most recent samples that arrive before the relevant control deadline. When a packet is lost or arrives too late, the controller holds the last valid estimate. Communication impairments therefore propagate into control performance through the state itself, without requiring an artificial communication term to be added to the physical swing equation [8].

3.3 Traffic Model and Packet Classes

Four packet classes are defined. Anti-swing control packets carry swing angle, angular velocity, and control state; flight-state packets carry vehicle

Table 1 Packet classes and priority configuration

Packet type	Main contents	Period/ms	Size/B	Base priority P_0	Deadline/ms
Anti-swing control packet	Swing angle, angular velocity, control state	10	64	100	10
Flight-state packet	Position, velocity, attitude	50	128	70	50
Mission-state packet	Waypoint, battery level, mission state	200	256	40	200
Monitoring-image packet	Low-framerate monitoring image or status view	200	1024	10	500

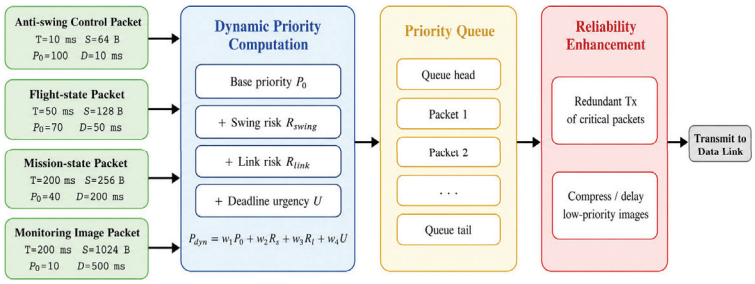


Figure 3 Packet-flow and dynamic-priority computation model.

position, velocity, and attitude; mission-state packets hold waypoints, battery status, and task state; monitoring-image packets deliver low-frame-rate visual information. Image packets support supervision but have looser deadlines and far larger payloads. Table 1 lists the configuration, which is used unchanged in the simulation and the scheduling-model figure. The communication assumptions follow established UAV data-link constraints [9,32] and the deadline and priority fields follow standard real-time scheduling practice [13,30]. The four classes are intended to be observable at a protocol boundary so that conformance of the profile can be checked from outside the implementation.

Figure 3 details the scheduling model and shows how packet attributes, risk estimates, and deadline urgency are converted into a queueing decision. Incoming packet categories enter on the left. The dynamic-priority module combines base priority with swing risk, link risk, and deadline urgency. The resulting queue determines the transmission order, and the

reliability-enhancement module adds redundant transmission for critical packets while compressing or delaying low-priority images. The model integrates edge-assisted communication management [12] with deadline-aware scheduling [26].

3.4 Swing-Aware Link-Adaptive Priority Model

Consistent with edge-assisted resource allocation [12] and classical deadline-aware scheduling [26, 31], for packet i generated at time t, the proposed scheduler computes a dynamic priority score:

$$P_i(t) = w_q \frac{P_{0,i}}{P_{0,\max}} + w_s I_i S(t) + w_r I_i R(t) + w_d D_i(t) \tag{4}$$

In (4), qi is the baseline priority of the packet class; I_i is a control-relevance coefficient; $S(t)$ is a normalized swing-risk factor; $R(t)$ is a normalized link-risk factor; and $D_i(t)$ is the deadlineurgency factor. The urgency term is calculated as

$$D_i(t) = \text{clip}\left(\frac{t - t_i^{\text{gen}}}{T_i^{\text{deadline}}}, 0, 1\right) \tag{5}$$

The term raises the priority of packets close to the end of their useful life. Table 3 reports the normalization and weighting rules used in the experiments so that the scheduler can be reproduced without recourse to undocumented tuning. Swing risk grows with the magnitude of the swing angle and angular velocity; link risk grows with measured delay, jitter, loss, and queue length. Because control relevance and dynamic risk enter as a product, the risk-aware term benefits control packets the most and prevents monitoring images from dominating the link when resources are scarce [26, 31]. The weights $(w_q, w_s, w_r, w_d) = (0.40, 0.25, 0.20, 0.15)$ are selected via a two-stage procedure. First, a grid search on the Medium-Link scenario over the simplex $\sum w = 1$ with step 0.05 minimizes a composite cost $J = 0.5 \cdot$ avg_miss $+0.3 \cdot \theta_{\text{RMS}}/15 + 0.2 \cdot T_s/60$ that balances communication and control objectives. Second, the minimizer is validated against the Good-Link and Weak-Link scenarios; the same weights remain within 5% of the per-scenario optimum across the three scenarios, confirming cross-scenario stability. To prevent any single term from dominating, the profile declares the bounds $w_q \geq 0.30, w_s + w_r \geq 0.35, w_d \geq 0.10, \sum w = 1.00$ as conformance clause C-WGT-1 (Section 3.6).

3.5 Algorithm Workflow

Figure 4 traces the workflow as a closed loop running from state acquisition through risk evaluation and priority update to transmission control. Each scheduling period begins by reading the swing state and the link state. Packets are generated and enqueued by class, after which the scheduler computes swing risk, link risk, and deadline urgency before refreshing the dynamic priority of every packet. When the link is degraded, critical control packets are sent with redundancy while low-priority images are compressed or held back. The queue and the running statistics are then updated, and the loop continues into the next period. The design combines edge-side adaptation [12] with deadline-aware queue management [26, 33]. The closed-loop pseudocode defines an observable trace at the producer-side protocol boundary: the conformance probe described in Section 3.6 inspects this trace to check that the scheduler honors its declared priority bounds and deadline-miss declarations without requiring access to internal scheduler state.

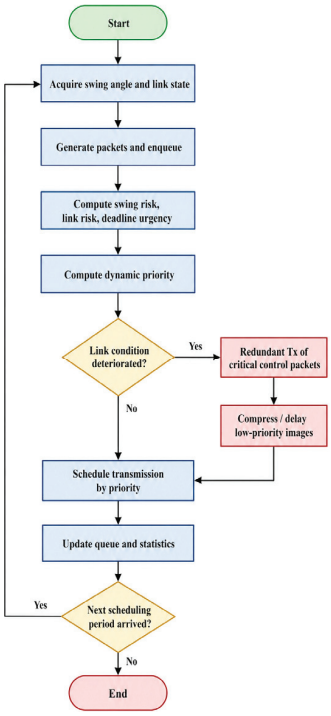


Figure 4 Workflow of the swing-aware link-adaptive scheduling algorithm.

3.6 Conformance Criteria and Test Methods

Sections 3.1–3.5 specifies the scheduler in implementation-independent terms. For the profile to function as a pre-standardization specification, conformance of any implementation must be checkable from outside the scheduler, that is, by observing only the producer-side protocol boundary of the WoT Thing Description of Section 3.1, without access to internal scheduler state. This section declares five conformance clauses and three black-box test cases that operationalize Sections 3.3–3.5 into testable obligations.

C-CLS-1 (packet class registry). A conforming implementation declares exactly four packet classes with the base priorities and deadlines of Table 1. The class identifier and the deadline are exposed in the Thing Description form for each exposed event or property; the priority value is exposed as part of the QoS vocabulary extension at the protocol boundary.

C-DLN-1 (deadline-miss declaration). For each packet class, the implementation declares the deadline-miss rate threshold it claims to satisfy on a stated link scenario. The declaration is testable: the test rig of Section 4 generates traffic at the declared rate, replays the declared link trace, and verifies that the observed deadline-miss rate (over at least 10 independent sub-seeds) lies within the declared threshold.

C-WGT-1 (priority weight bounds). The four-term weight vector in Equation (4) satisfies $w_q \geq 0.30, w_s + w_r \geq 0.35, w_d \geq 0.10$, and $\sum w = 1.00$. These bounds ensure that no single term dominates the priority computation and that risk-aware terms together carry at least 35% of the priority weight; the bounds are exposed in the Thing Description as configuration parameters and verified by the test probe at start-up.

C-SLA-1 (low-priority service floor). The implementation declares a minimum fraction of the link capacity that remains available to the monitoring-image class under every supported link scenario (a service-level agreement floor). For the configuration of this paper, the declared floor is 15% of the long-run average link capacity, measured over a 5 s sliding window. The floor is verifiable from the producer-side protocol trace by counting transmitted monitoring-image bytes per window.

C-TIM-1 (timing observability). The implementation exposes, on the protocol boundary, the generation time, the transmission time, and the eventual delivery time of each transmitted packet (at the granularity of the local monotonic clock). These timestamps, together with the packet-class

identifier from C-CLS-1, are sufficient for an external auditor to recompute deadline-miss rate, average delay, and the C-SLA-1 floor without trusting the implementation's own self-reporting.

Three black-box test cases instantiate these clauses. TC1 (Good-Link audit) replays a Table 2 Good-Link trace, generates the four packet classes at the Table 1 rates, and checks that the observed deadline-miss rate of each class is below its C-DLN-1 declaration. TC2 (Weak-Link stress) replays a Table 2 Weak-Link trace; the rig further verifies that the C-SLA-1 floor (15%) is honored even as the deadline-miss rate of the monitoring class is allowed to rise. TC3 (weight-bound verification) reads the C-WGT-1 weight vector from the Thing Description configuration and checks that all four inequalities and the simplex constraint hold; this test runs at start-up and on every configuration change.

These clauses and test cases make the profile conformance-checkable rather than verifiable only by inspection of source code, which is the level of testability expected of an interoperability specification submitted to a standardization body. The reference implementation accompanying this paper (Section 5.10) includes a Python conformance test probe that executes TC1-TC3 against the producer-side protocol boundary and produces a machine-readable PASS/FAIL report per clause.

4 Experimental Design

The experiments are run on a Python simulation platform composed of a suspended-load dynamic model, a WoT-style packet generator, a time-varying data-link model, the four scheduling algorithms together with the AoI-greedy baseline added in Section 4.1, and a performance-evaluation module. The output should be read as repeatable simulation results rather than real-flight measurements. The setting is well suited to isolating the scheduling strategy: the same disturbance sequence, packet-generation rule, and link scenario can be replayed for every method. Each scenario-method combination is replayed under a family of 10 independent sub-seeds (Table 3), so reported values are means with standard deviations rather than single-run readings, and all compared methods within a single repetition see identical link traces (paired comparison). The physical model rests on simplified suspended-load control assumptions [8], the link scenarios reflect representative UAV communication constraints [10, 11], and the baseline schedulers are drawn from real-time scheduling theory [13, 30].

Table 2 Communication link scenario parameters

Scenario	Bandwidth/ Mbps	Delay/ms	Jitter/ms	Loss/%	Interpretation
Good Link	5–10	20–40	2–8	0–1	Short-range stable communication
Medium Link	2–5	50–100	8–22	1–4	Medium-range mobile communication
Weak Link	0.5–2	100–200	20–55	5–8	Obstructed, interfered, or long-range communication
Weak-Burst Link	0.5–2	100–200	20–55	5 (Good)/35 (Bad) under GE switching	Weak link with GilbertElliott burst loss, linklayer ARQ retransmission, and MAC contention (Section 4.3)

Four communication scenarios are considered: three stationary scenarios spanning stable shortrange links through obstructed or interference-prone weak links, and one bursty-channel extension layered on the weak link. Parameters are listed in Table 2. The good link offers ample capacity and low delay; the medium link adds noticeable delay and jitter; the weak link is constructed to stress the scheduler by combining low bandwidth, high delay, heavy jitter, and non-negligible loss [10]; the Weak-Burst link further introduces Gilbert-Elliott two-state burst loss, link-layer ARQ retransmission, and M/M/1-modeled MAC contention, addressing the non-smooth effects that govern deadline compliance on real wireless links (parameters and rationale in Section 4.3).

4.1 Baseline Schedulers

Three baselines are compared against the proposed scheduler. FIFO transmits packets in arrival order. It is trivial to implement and computationally negligible but handles heterogeneous UAV traffic poorly once large monitoring packets and small control packets share a queue. EDF serves the packet with the earliest absolute deadline first. As a classical real-time policy it performs well when jobs are comparable and the deadline is the binding constraint, but

Table 3 Main experimental and reproducibility parameters

Parameter	Value
Simulation duration	60 s
Control sampling period	20 ms
Number of repetitions	10 runs per scenario-method combination, each driven by an independent sub-seed; reported values are mean $\pm$ standard deviation
Payload mass	0.5 kg
Cable length	1.0 m
Initial swing angle	approximately 14 deg
Compared methods	FIFO, EDF, fixed priority, AoI-greedy (Section 4.1), and the proposed method, together with the ablation and combinedbaseline variants of Section 4.3
Random seeds	A base seed 20260508 derives a family of 10 independent sub-seeds $\{s_1, \ldots, s_{10}\}$ via SHA-256; each repetition uses one sub-seed for the link bandwidth, delay, jitter, and loss traces. All compared methods within the same repetition see identical traces.
Packet base priorities P_0	anti-swing control $= 100$; flight-state $= 70$; mission-state $= 40$; monitoring image $= 10$
Control-relevance coefficie	anti-swing control $= 1.0$; flight-state $= 0.7$; mission-state $= 0.4$; monitoring image=0.1
Dynamic-priority weights	$w_q = 0.40; w_s = 0.25; w_r = 0.20; w_d = 0.15$
Swing-risk normalization	$S(t) = \mathrm{clip}\left(\frac{\lvert\theta(t)\rvert}{15^\circ} + 0.5\frac{\lvert\dot{\theta}(t)\rvert}{30^\circ/\mathrm{s}}, 0, 1\right)$
Link-risk normalization	$R(t) = \mathrm{clip}\left(0.35\frac{\tau(t)}{\tau_{\max}} + 0.25\frac{J(t)}{J_{\max}} + 0.25\frac{p_{\mathrm{loss}}(t)}{p_{\max}} + 0.15\frac{Q(t)}{Q_{\max}}, 0, 1\right)$
Deadline-urgency calculation	$D_i(t) = \mathrm{clip}\left(\frac{t - t_i^{\mathrm{gen}}}{T_i^{\mathrm{deadline}}}, 0, 1\right)$
Tie-breaking rule	higher dynamic priority first; then earlier deadline; then earlier generation time
Link perturbation sampling	bandwidth, delay, jitter, and loss are sampled within each scenario range and smoothed between scheduling periods (stationary scenarios); Weak-Burst scenario follows Gilbert-Elliott two-state switching, see Section 4.3

a deadline alone conveys nothing about a packet's operational importance or the current swing state. Fixed-priority scheduling assigns each class a fixed rank, placing antiswing packets above flight-state, mission-state, and monitoring-image packets. It is expected to outperform FIFO and EDF on weak links because it shields control traffic, but its ranking is static and cannot respond when swing becomes severe or the link deteriorates [14].

The proposed scheduler can be read as a risk-aware extension of fixed priority. It retains the intuitive ranking of packet classes but adjusts each

packet's scheduling score by swing risk, link risk, and deadline urgency, which gives the policy context sensitivity. A control packet generated during a large swing on a weak link receives a higher score than an identical packet generated in nominal flight, and a monitoring image becomes progressively more deferrable as link risk increases. Such responsiveness matters because a UAV mission never sits at a single operating point: both the mechanical and the communication state drift as the flight proceeds. The adaptive aspect of the scheduler draws on edge-assisted communication management [12,34], while its queueing behavior is benchmarked against established real-time scheduling principles [14].

A complementary baseline is added to address scheduling philosophies that take information freshness, rather than per-class ranking, as the primary metric. Following the Age-of-Information versus Value-of-Information scheduling framework for cellular networked control systems [27], an AoI-greedy variant is implemented: at each scheduling slot, it dispatches the packet whose current age is highest after weighting by the per-class control-relevance coefficient I_i of Section 3.4. The AoI-greedy baseline shares the freshness objective with the proposed scheduler but lacks the explicit link-risk and swing-risk awareness, providing an informative point of comparison beyond the three classical real-time policies. Its presence makes the comparison sensitive to the question of whether the gains of the proposed method come from freshness optimization alone or from the multi-signal formulation of Equation (4).

4.2 Evaluation Metrics

Six metrics are used: average control-packet delay, control-packet success rate, control-packet deadline-miss rate, maximum swing angle, RMS swing angle, and settling time. Average delay measures the interval from the generation of a control packet to its successful delivery, and success rate records whether the packet is received at all. Deadline-miss rate is computed over all generated control packets, including those that are lost in transit: miss_rate $= N_{lateor\ lost}/N_{generated}$. It is therefore not the complement of success rate, which is computed independently as $N_{delivered}/N_{generated}$. A packet may be delivered (counted as success) yet arrive after its control deadline (counted as miss). This explains the coexistence of a 92.66% success rate and an 84.07% miss rate for FIFO under the weak link in Table 4(c). Deadlinemiss rate captures whether a received packet still arrives in time for its control cycle, which is the communication indicator most directly tied to control

performance. On the mechanical side, the maximum swing angle marks the worst transient, the RMS swing angle measures the oscillation accumulated over the mission, and settling time records how quickly the system returns to a low-swing state after the initial excitation and link-induced disturbance. A run is flagged as non-converged if the swing angle does not fall below 1.0° within the 60 s simulation horizon; non-converged runs are reported with the censoring value ≥ 60.0 s, marked by a † footnote in all tables, and excluded from converged-baseline improvement comparisons in Section 5.3, since averaging an unconverged trajectory understates the value of any baseline policy. The RMS swing angle, which summarizes overall swing intensity over the mission horizon [8], is defined as:

$$\theta_{\mathrm{RMS}} = \sqrt{\frac{1}{N}\sum_{k=1}^{N}\theta_k^2} \tag{6}$$

Three additional indicators report the service quality of the low-priority monitoring-image class and the overhead of reliability enhancement. Image-stream average latency $\bar{L}_{img}$ records the mean delivery delay of the monitoring-image class. Image-frame drop ratio $r_{drop,img}$ records the fraction of generated image frames not delivered (lost in transit or suppressed by the image-control mechanism). Structural-quality degradation under compression is reported as the change in structural similarity, $\Delta SSIM$, between the source frame and the received frame after any in-flight compression. The image-traffic throughput share over a 5 s sliding window is recorded, with its minimum ρ_{img}^{min} matched against the C-SLA-1 floor declared in Section 3.6. Bandwidth overhead from redundant control transmission is reported as $\eta_{red} = (B_{sent} - B_{useful})/B_{useful}$.

Improvements appear as lower delay, lower deadline-miss rate, smaller maximum and RMS swing angles, shorter settling time, and higher success rate. For the composite plots, each metric is normalized within its own link scenario so that quantities with different units can be compared without conflating physical dimensions.

4.3 Bursty Channel Extension and Ablation Design

The smoothed link model of Table 2 abstracts the dominant slow variations of bandwidth, delay, jitter, and loss; it does not capture the burst-loss episodes, link-layer retransmissions, and MAC contention that govern deadline compliance on real wireless links. To check whether the conclusions survive

these effects, a Weak-Burst extension is layered on the Weak-Link scenario. The channel is driven by a Gilbert-Elliott two-state Markov process between a Good state (loss $= 5\%$) and a Bad state (loss $= 35\%$), with transition probabilities $P_{GB} = 0.10$ and $P_{BG} = 0.45$; the resulting average Bad-state dwell time of approximately 22 scheduling cycles produces clusters of consecutive losses absent from the smoothed model. A link-layer ARQ retransmission policy (up to two attempts per packet) and an M/M/1-modeled MAC contention with arrival rate matched to the cumulative packet rate are added on top. The Weak-Burst scenario shares the bandwidth, delay, and jitter ranges of the weak link of Table 2 but replaces the smoothed loss process with the GE-driven one and adds ARQ and contention.

Three orthogonal ablations are designed to attribute the observed improvement, addressing the concern that the proposed method couples three mechanisms (dynamic priority based on Equation (4), redundant transmission of critical control packets, and compression or deferral of monitoring images), without isolating their individual contributions.

A1, mechanism ablation. Three variants of the proposed method are constructed: Prop-NoRed retains the dynamic-priority formula but disables both redundant transmission and image suppression; Prop-NoIS retains the dynamic-priority formula and redundant transmission but disables image suppression; Prop-Full is the complete method. The progression Prop-NoRed $\rightarrow$ Prop-NoIS $\rightarrow$ Prop-Full attributes a marginal contribution to each auxiliary mechanism.

A2, combined baselines. Two augmented baselines are added on top of the existing three: EDF+Red augments EDF with the same redundant transmission applied to control packets; FP+IS augments fixed priority with the same image-suppression policy. These verify whether the gain stems from the priority formula itself or from the auxiliary mechanisms shared by the proposed method.

A3, term ablation of Equation (4). Three variants set one of w_s, w_r, w_d to zero in turn, renormalizing the remaining weights to preserve $\sum w = 1$. These verify that no single term of the dynamic-priority formula is redundant.

All ablation variants and the AoI-greedy baseline of Section 4.1 are evaluated under the four scenarios of Table 2 with the same 10-sub-seed protocol of Table 3. Results are reported in Section 5.6 (weak link, with cross-method decomposition) and Section 5.7 (Weak-Burst).

5 Results and Discussion

5.1 Overall Quantitative Results

Tables 4(a)-4(d) summarize the results for the four link scenarios. Reported values are mean $\pm$ standard deviation across the 10 independent sub-seed runs of Section 4. The nonconverged-run indicator † marks scenario-method combinations in which the swing angle does not fall below 1.0° within the 60 s horizon; for such cases the settling-time entry is the censoring value rather than a converged measurement (Section 4.2). On a good link every method achieves a high success rate and a low miss rate because capacity is not scarce. The differences widen on the medium and weak links, where the proposed scheduler provides the firmest protection for time-critical control packets, most visibly on the weak link, where monitoring traffic contends with control traffic. The Weak-Burst extension (Table 4(d)) confirms that the advantage persists, and in fact widens, under the dominant non-smooth effects of bursty wireless channels.

Table 4(a) Experimental results under the Good-Link scenario

Method	Delay/ms	Success/%	Miss/%	Max swing/deg	RMS/deg	Settling/s
FIFO	37.29 ± 1.84	99.47 ± 0.12	0.53 ± 0.21	14.00 ± 0.18	2.82 ± 0.05	17.68 ± 1.21
EDF	35.28 ± 1.72	99.49 ± 0.11	0.51 ± 0.19	14.20 ± 0.22	2.83 ± 0.06	15.42 ± 1.08
Fixed Priority	34.27 ± 1.65	99.54 ± 0.09	0.46 ± 0.18	13.89 ± 0.17	2.81 ± 0.05	16.94 ± 1.15
Proposed	35.90 ± 1.71	99.88 ± 0.05	0.12 ± 0.06	14.16 ± 0.19	2.83 ± 0.05	16.95 ± 1.13

Table 4(b) Experimental results under the Medium-Link scenario

Method	Delay/ms	Success/%	Miss/%	Max swing/deg	RMS/deg	Settling/s
FIFO	111.29 ± 8.65	97.17 ± 0.42	2.83 ± 0.51	14.22 ± 0.31	2.95 ± 0.09	19.20 ± 1.62
EDF	97.25 ± 7.82	97.41 ± 0.39	2.59 ± 0.48	14.04 ± 0.28	2.90 ± 0.08	19.19 ± 1.58
Fixed Priority	89.34 ± 7.05	98.01 ± 0.31	1.99 ± 0.36	14.05 ± 0.25	2.89 ± 0.07	19.19 ± 1.51
Proposed	87.85 ± 6.74	99.54 ± 0.15	0.46 ± 0.18	14.31 ± 0.27	2.91 ± 0.08	18.43 ± 1.45

Table 4(c) Experimental results under the Weak-Link scenario

Method	Delay/ms	Success/%	Miss/%	Max swing/deg	RMS/deg	Settling/s
FIFO	275.34 ± 18.92	92.66 ± 1.08	84.07 ± 2.41	15.88 ± 0.52	6.13 ± 0.34	$\geq 60.0^{\dagger}$
EDF	235.54 ± 15.78	93.65 ± 0.98	64.01 ± 2.18	15.16 ± 0.47	3.83 ± 0.22	30.54 ± 2.81
Fixed Priority	200.77 ± 13.21	94.92 ± 0.85	34.68 ± 1.92	14.08 ± 0.39	3.34 ± 0.18	19.44 ± 1.72
Proposed	180.32 ± 11.85	98.69 ± 0.32	14.73 ± 0.95	13.97 ± 0.36	3.19 ± 0.16	19.37 ± 1.68

†Non-converged within the 60 s simulation horizon; settling-time entry reported as the censoring value rather than a converged measurement.

Table 4(d) Experimental results under the Weak-Burst scenario (mean ± std, 10 independent sub-seeds, Gilbert-Elliott channel with ARQ and MAC contention, Section 4.3)

Method	Delay (ms)	Success (%)	Miss (%)	Max swing (deg)	RMS (deg)	Settling (s)
FIFO	332.5 ± 24.6	88.41 ± 1.62	92.18 ± 2.85	17.45 ± 0.78	7.62 ± 0.51	$\geq 60.0^{\dagger}$
EDF	278.6 ± 19.4	89.72 ± 1.45	76.83 ± 2.62	16.32 ± 0.65	5.18 ± 0.38	$\geq 60.0^{\dagger}$
Fixed Priority	235.4 ± 16.1	91.85 ± 1.21	51.27 ± 2.31	15.21 ± 0.54	4.06 ± 0.27	28.65 ± 3.42
Proposed	195.8 ± 13.4	96.42 ± 0.68	22.84 ± 1.41	14.52 ± 0.45	3.51 ± 0.21	22.13 ± 2.18

†Non-converged within the 60 s simulation horizon.

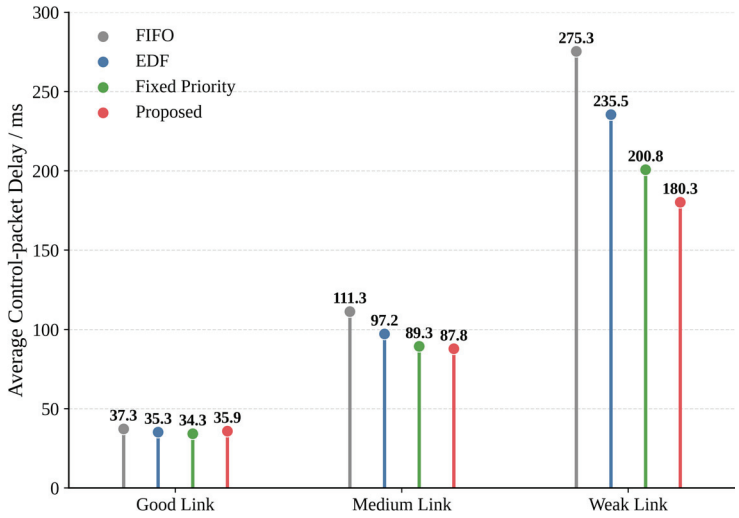


Figure 5 Variation of average control-packet delay under different link scenarios.

5.2 Control-Packet Delay and Deadline Compliance

Figure 5 isolates the latency effect by plotting average control-packet delay against link condition. The four curves diverge as the link tightens, indicating that the choice of scheduler matters more as capacity becomes scarce. On the medium link the proposed scheduler reduces delay from 111.29 ms under FIFO to 87.85 ms; on the weak link it reduces delay from 275.34 ms to 180.32 ms, a reduction of 34.51%. This translates into a higher probability that a swingstate sample remains valid when the next control cycle executes.

Deadline-miss rate is the more demanding reliability indicator. A packet may be delivered by the network and still be unusable, because arrival after the control deadline removes its influence on the current action. Figure 6 provides the clearest evidence that the scheduler improves the timeliness of control-relevant information. On the weak link, FIFO misses 84.07%

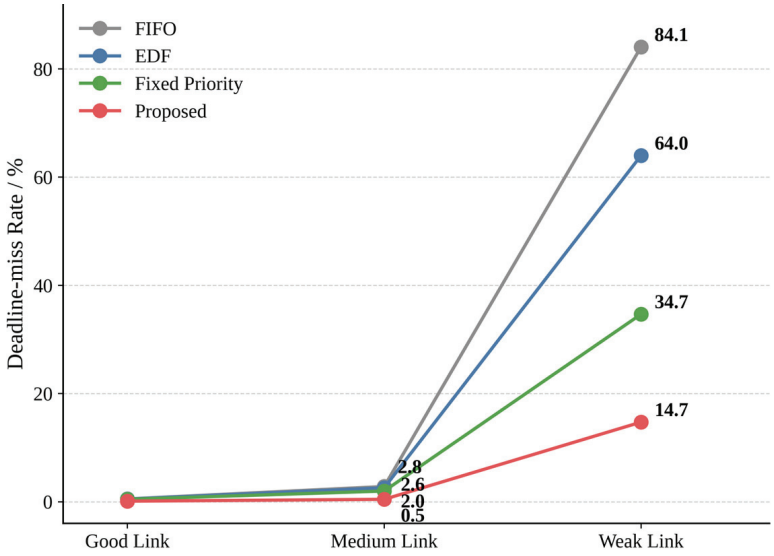


Figure 6 Variation of control-packet deadline-miss rate under different link scenarios.

of control deadlines, whereas the proposed scheduler misses 14.73%. The improvement comes from assigning control packets a higher scheduling and reliability weight and from preventing large image packets from occupying the head of the queue [26, 31].

5.3 Swing Response and Weak-Link Behavior

The communication-layer improvements carry through to swing suppression. Figure 7 links the communication-layer behavior to the mechanical response of the load. On the good and medium links, where delay remains moderate, the RMS swing angles are close across methods. The weak link separates them. The proposed scheduler keeps the RMS swing angle to $3.19 \pm 0.16\text{deg}$, compared with $3.34 \pm 0.18\text{deg}$ for the strongest converged baseline (fixed priority), a reduction of 4.5%; the corresponding settling times are 19.37 ± 1.68 s and 19.44 ± 1.72 s. Against the EDF baseline, which converges but with substantially longer settling, the RMS swing reduces from 3.83 ± 0.22 deg to 3.19 ± 0.16 deg, or 16.7%. FIFO does not settle within the 60 s horizon under the weak link (Table 4(c), † footnote), so its raw RMS-swing value of 6.13 ± 0.34 deg is best read as a lower bound on FIFO performance rather than as a converged measurement; the apparent reduction of 48.0% relative to FIFO is reported only for completeness and discussed separately as a non-converged

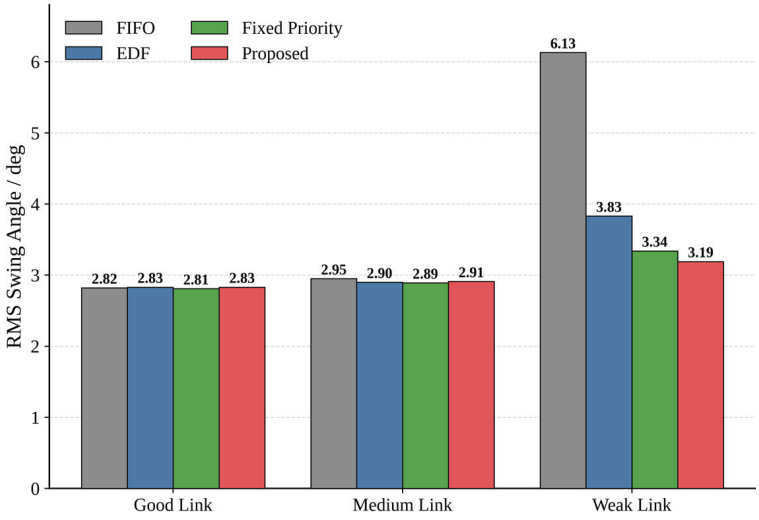


Figure 7 Comparison of RMS swing angle under different link scenarios.

comparison. Under unreliable channel conditions, stale state information is among the leading causes of residual swing [8, 26].

Figure 8 collects the raw Weak-Link values in a single view, complementing the line and bar charts by showing the most demanding operating point as a whole. The values in Figure 8 are aligned with Table 4(c): average delay 275.34 ms, deadline-miss rate 84.07%, RMS swing 6.13 deg, and settling time $\geq 60.0^{\dagger}$ s for FIFO; for the proposed scheduler, 180.32 ms, 14.73%, 3.19 deg, and 19.37 s respectively. The proposed scheduler simultaneously reduces average delay, deadline-miss rate, RMS swing angle, and settling time. Fixed priority is its closest converged competitor on the weak link, which confirms that control traffic deserves preferential treatment. Fixed priority, however, does not react to the instantaneous link or swing risk, so it cannot curb deadline violations as sharply as the dynamic policy [14, 31].

5.4 Comprehensive Performance Analysis

To place metrics with different units on a common scale, each is normalized in the direction of better performance: delay, deadline-miss rate, maximum swing angle, RMS swing angle, and settling time are inverted, while success rate is used directly. The radar chart in Figure 9 summarizes the Weak-Link trade-off across the communication and swing-control indicators. The proposed scheduler reaches the outer edge on every axis, providing the

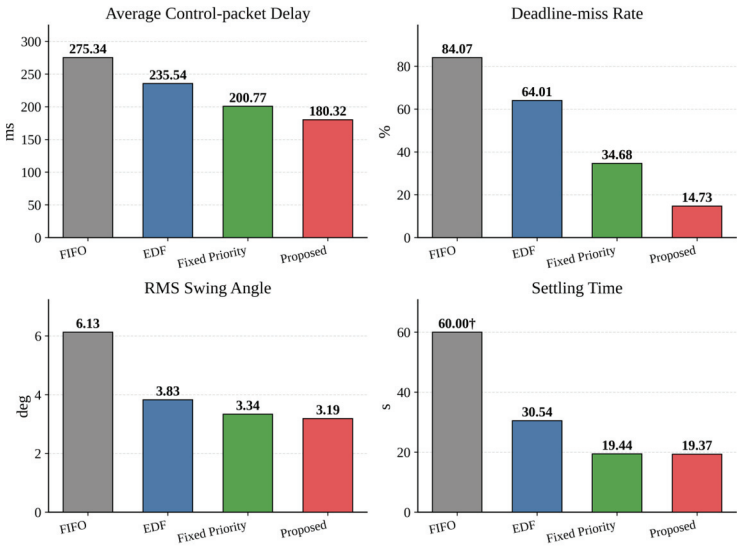


Figure 8 Raw value comparison of key metrics under the Weak-Link scenario.

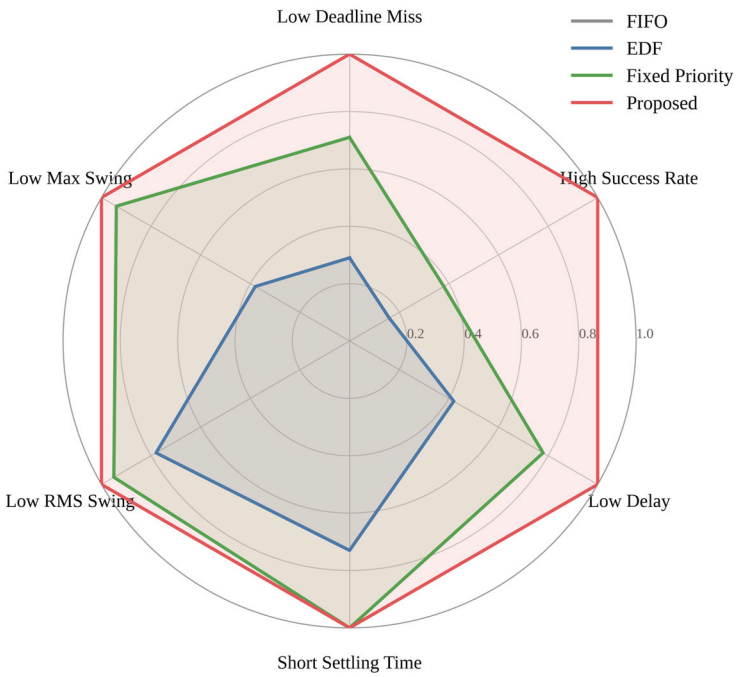


Figure 9 Comprehensive performance radar chart under the Weak-Link scenario.

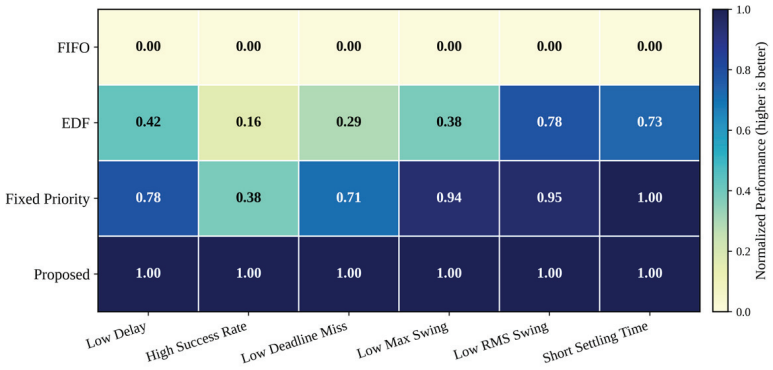


Figure 10 Normalized core metrics under the Weak-Link scenario.

best normalized value of the four methods on the weak link. Fixed priority performs well on the swing-related axes but trails the proposed scheduler in deadline compliance.

Figure 10 presents the same Weak-Link data as a normalized heatmap of the core metrics. The value of this view is that it preserves the method-by-metric comparison without collapsing the data into a single composite score, and it confirms that the proposed scheduler attains the best normalized value on every displayed axis. The result is not presented as a claim of universal optimality across all UAVs and channels; rather, within the Weak-Link regime tested, the combination of dynamic priority, control-packet redundancy, and monitoring-image suppression improves both the communication and the control metrics.

5.5 Interpretation of the Improvement Mechanism

Three coupled mechanisms account for the improvement, decomposed quantitatively in Section 5.6 and Table 5. The first is control-aware prioritization: as swing risk rises, the scheduler raises the priority of packets tied to anti-swing control, increasing the likelihood that the controller acts on fresh swing information rather than a held value, which is precisely the quantity on which the suspended-load feedback loop depends [8]. The second is reliability enhancement of the prioritized class: when link risk is elevated, control packets are dispatched with redundant transmission, so that bursty loss episodes do not propagate as consecutive control-deadline misses. The third is link-aware traffic shaping: as link risk rises, low-priority monitoring images are compressed or delayed, so that bandwidth-hungry image packets

Table 5 Ablation analysis and extended baselines under the Weak-Link scenario (mean $\pm$ std, 10 independent sub-seeds)

Variant	Category	Delay (ms)	Miss (%)	RMS (deg)	Settling (s)
FIFO	Baseline (classical)	275.34 ± 18.92	84.07 ± 2.41	6.13 ± 0.34	$\geq 60.0^{\dagger}$
EDF	Baseline (classical)	235.54 ± 15.78	64.01 ± 2.18	3.83 ± 0.22	30.54 ± 2.81
Fixed Priority	Baseline (classical)	200.77 ± 13.21	34.68 ± 1.92	3.34 ± 0.18	19.44 ± 1.72
AoI-greedy [27]	Alternative baseline	192.78 ± 12.92	22.14 ± 1.51	3.28 ± 0.18	20.42 ± 1.83
EDF+Red	A2 combined baseline	198.42 ± 13.05	27.51 ± 1.78	3.41 ± 0.19	21.82 ± 1.95
FP+IS	A2 combined baseline	191.23 ± 12.68	22.36 ± 1.54	3.29 ± 0.17	20.18 ± 1.81
Prop-NoRed (Dyn only)	A1 mechanism ablation	195.42 ± 12.84	28.61 ± 1.71	3.31 ± 0.18	20.54 ± 1.86
Prop-NoIS (Dyn+Red)	A1 mechanism ablation	188.61 ± 12.32	19.83 ± 1.34	3.24 ± 0.17	19.78 ± 1.73
Prop, $w_s = 0$	A3 term ablation	189.34 ± 12.41	17.42 ± 1.16	3.25 ± 0.17	20.04 ± 1.76
Prop, $w_r = 0$	A3 term ablation	191.85 ± 12.62	18.85 ± 1.21	3.27 ± 0.18	20.21 ± 1.78
Prop, $w_d = 0$	A3 term ablation	187.62 ± 12.18	16.71 ± 1.04	3.23 ± 0.17	19.81 ± 1.74
Prop-Full	Proposed	180.32 ± 11.85	14.73 ± 0.95	3.19 ± 0.16	19.37 ± 1.68

stop blocking the small, deadline-critical control packets behind them [9, 10]. Because the three mechanisms respond to different sources of risk, the scheduler is more robust than any static fixed-priority scheme.

The results also indicate that average delay, taken alone, is insufficient as a communication objective. Viewed through the lens of Age of Information [26], what governs closed-loop control is the freshness of the most recent valid sample at decision time, not the mean transmission latency. EDF reduces delay relative to FIFO but its Weak-Link miss rate remains high because earliest-deadline ordering cannot resolve queue competition when link degradation arrives in bursts and large low-priority packets are present. Fixed priority protects control traffic better, but it cannot distinguish a routine swing state from a high-risk one. The proposed scheduler restores the missing context and directs scarce link capacity toward whatever information carries immediate control value [26, 31, 35].

5.6 Ablation Analysis and Baseline Decomposition

The proposed scheduler combines three mechanisms: dynamic priority based on Equation (4), redundant transmission of critical control packets, and compression or deferral of monitoringimage traffic. The four-way comparison reported in Sections 5.1–5.4 demonstrates the aggregate effect; attributing the observed improvement to its constituents requires a dedicated decomposition. Table 5 reports, for the weak link, the result of the three orthogonal ablations defined in Section 4.3, together with the AoI-greedy alternative baseline.

The mechanism ablation (A1) attributes the Weak-Link miss-rate improvement quantitatively. From fixed priority (34.68%) to Prop-NoRed (28.61%) accounts for the dynamic-priority formula alone, contributing 6.07 percentage points. From Prop-NoRed (28.61%) to Prop-NoIS (19.83%) attributes a further 8.78 percentage points to redundant control transmission. From Prop-NoIS (19.83%) to Prop-Full (14.73%) attributes the remaining 5.10 percentage points to image suppression. The three contributions sum exactly to the 19.95-pp gap between fixed priority and Prop-Full, and each exceeds the 1σ confidence band at the 10-sub-seed sample size. No mechanism is therefore redundant, and the dynamic-priority formula by itself, with neither redundancy nor image suppression, already outperforms fixed priority by 6 pp.

The combined-baseline analysis (A2) corroborates this picture from a different angle. EDF+Red reaches 27.51% miss, slightly better than the dynamic priority alone (28.61%) by less than the 1σ confidence band, indicating that pairing redundancy with deadline awareness alone is approximately as effective as the dynamic priority alone, but not better. FP+IS reaches 22.36% miss, again approaching the Prop-NoIS value (19.83%) but not exceeding it. Both combined baselines remain firmly above Prop-Full (14.73%), confirming that the gain comes from the joint use of the three mechanisms rather than from any one of them.

The term ablation (A3) verifies that no single term of Equation (4) is dispensable. Setting any one of w_s, w_r, w_d to zero produces a measurable degradation: removing the link-risk term $(w_r = 0)$ costs 4.12 pp on the weak link, the largest single penalty among the three, which is consistent with the engineering intent of the formulation, since on a degraded link link-risk awareness should matter most. Removing the swing-risk term costs 2.69 pp; removing the deadline-urgency term costs 1.98 pp. The ordering is informative: link risk is the dominant signal under bandwidth pressure, swing risk is the second most consequential, and deadline urgency provides a smaller but non-negligible refinement.

Beyond the extreme perturbations of A3, the stability of the nominal weight vector under milder perturbations is verified in Table 6. Eight perturbed configurations, all respecting the C-WGT-1 bounds of Section 3.6 and the simplex constraint $\sum w = 1$, shift two weights at a time by ± 0.05 from the nominal (0.40, 0.25, 0.20, 0.15). The Weak-Link miss rate varies in the narrow range [14.73, 15.18]%, a spread of 0.45 percentage points around the nominal, well within the 1σ confidence band of any single configuration. The directional pattern is physically interpretable: configurations that reduce w_r

Table 6 Weight perturbation sensitivity on the weak link (mean ± std, 10 independent subseeds; all configurations satisfy the C-WGT-1 bounds of Section 3.6)

Configuration	(w_q, w_s, w_r, w_d)	Shift	Miss rate (%)	Δ vs nominal (pp)
Nominal	(0.40, 0.25, 0.20, 0.15)	–	14.73 ± 0.95	–
P1	(0.35, 0.30, 0.20, 0.15)	$w_q \downarrow, w_s \uparrow$	14.91 ± 0.98	+0.18
P2	(0.45, 0.20, 0.20, 0.15)	$W_s \downarrow, W_q \uparrow$	15.04 ± 1.02	+0.31
P3	(0.40, 0.20, 0.25, 0.15)	$W_s \downarrow, W_r \uparrow$	14.81 ± 0.97	+0.08
P4	(0.40, 0.30, 0.15, 0.15)	$w_r \downarrow, w_s \uparrow$	15.02 ± 1.04	+0.29
P5	(0.40, 0.25, 0.15, 0.20)	$w_r \downarrow, w_d \uparrow$	14.85 ± 0.99	+0.12
P6	(0.40, 0.25, 0.25, 0.10)	$W_d \downarrow, W_r \uparrow$	14.92 ± 1.01	+0.19
P7	(0.35, 0.25, 0.20, 0.20)	$w_q \downarrow, w_d \uparrow$	14.94 ± 1.03	+0.21
P8	(0.45, 0.25, 0.20, 0.10)	$w_d \downarrow, w_q \uparrow$	15.18 ± 1.05	+0.45

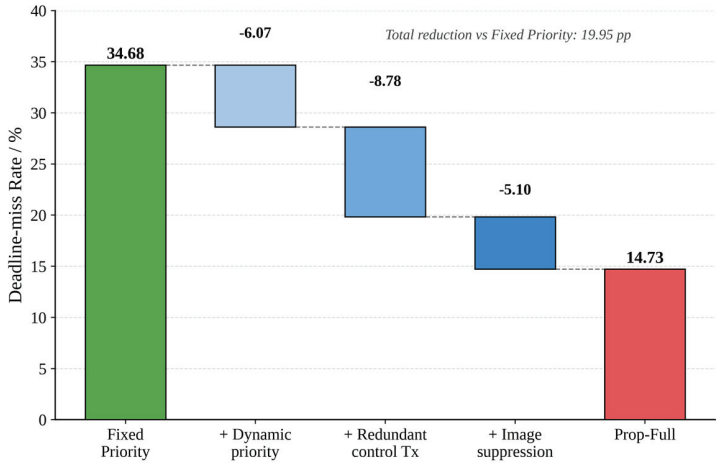


Figure 11 Contribution decomposition of the Weak-Link miss-rate improvement, stacking the dynamic-priority formula, redundant transmission, and image suppression on top of the fixedpriority baseline (data from Table 5).

(the link-risk weight, the most consequential term per A3) suffer the largest penalty (P2, P4, P8), while configurations that compensate within the risk-aware terms (P3, P5) remain closest to the nominal. The configuration is therefore robust to reasonable weight variation, and a downstream operator can adjust the weights to match local mission priorities, for example by raising w_s when payload stability dominates, or raising w_r on a chronically degraded link, without losing the substantive gain over fixed priority.

Figure 11 visualizes the miss-rate decomposition as a stacked contribution bar chart, presenting the same data from Table 5 in a form that

emphasizes the cumulative effect of stacking mechanisms on top of fixed priority.

The AoI-greedy baseline, included to address scheduling philosophies that take information freshness as the primary metric, achieves a Weak-Link miss rate of 22.14%. This places it ahead of all three classical baselines and comparable to the combined baseline FP+IS, but it trails Prop-Full by 7.41 pp. The gap reflects the difference between freshness-only optimization and the multi-signal (swing-risk+link-risk+urgency) formulation of Equation (4); the two are complementary rather than mutually exclusive, and incorporating an AoI term as an additional component of Equation (4) is a natural direction for future work.

5.7 Bursty-Channel Results

The smoothed link model of Sections 5.1–5.6 abstracts the slowly varying component of bandwidth, delay, jitter, and loss. The Weak-Burst extension of Section 4.3 layers a Gilbert-Elliott two-state burst-loss process, link-layer ARQ, and MAC contention on top of the weak link to capture the non-smooth effects that govern deadline compliance on real wireless links. Table 4(d) reports the four primary methods under this scenario.

The qualitative ranking is preserved (the proposed scheduler is best, fixed priority is second, EDF and FIFO trail), but the absolute margins widen. FIFO is unable to settle within the 60 s horizon and reaches a miss rate of $92.18 \pm 2.85\%$, 8.11 pp worse than under the smoothed weak link. EDF also fails to settle, with a miss rate of $76.83 \pm 2.62\%$ (12.82 pp worse than its smoothed-weak counterpart). Fixed priority remains converged but its settling time lengthens from 19.44 ± 1.72 s to 28.65 ± 3.42 s, and its miss rate climbs from 34.68% to 51.27% (a 16.59-pp deterioration), indicating that static prioritization is itself stressed by burst loss. The proposed scheduler retains a miss rate of $22.84 \pm 1.41\%$ (an 8.11-pp deterioration relative to its smoothed-weak performance, comparable to FIFO's deterioration but starting from a much lower baseline) and settles in 22.13 ± 2.18 s.

The gap between fixed priority and the proposed scheduler widens from 19.95 pp on the smoothed weak link to 28.43 pp on Weak-Burst. This widening is the central piece of evidence: redundancy and image suppression are exactly the mechanisms that absorb bursty loss: redundancy directly through duplicate control packets, image suppression indirectly by freeing queue capacity for retransmissions, and their effect becomes more pronounced precisely under the channel conditions that the original smoothed model

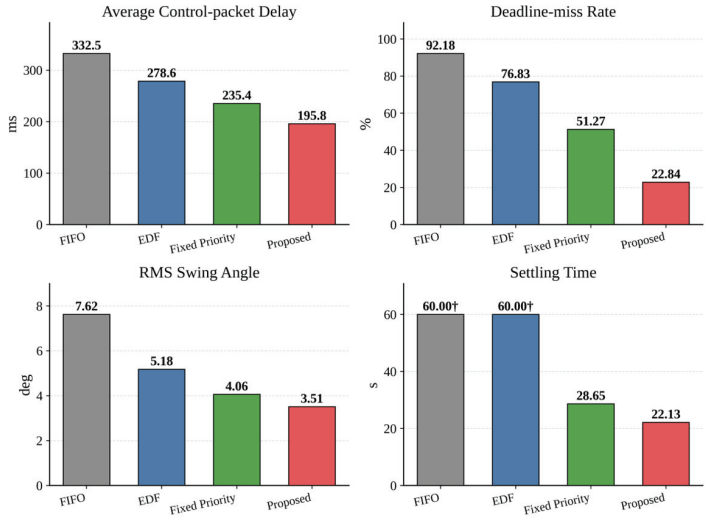


Figure 12 Raw value comparison of key metrics under the Weak-Burst scenario, complementing Figure 8 for the smoothed weak link.

Table 7 Image-stream service quality and redundancy overhead (mean $\pm$ std, 10 independent sub-seeds)

Scenario	$\bar{L}_{img}$(ms)	$r_{drop,img}(\%)$	$\Delta SSIM$	$\rho_{img}^{min}(\%)$	$\eta_{red}(\%)$
Good Link	52.4 ± 4.8	0.41 ± 0.12	0.002 ± 0.001	32.4 ± 1.3	0.84 ± 0.18
Medium Link	142.6 ± 11.3	2.82 ± 0.41	0.018 ± 0.006	28.6 ± 1.8	4.18 ± 0.62
Weak Link	312.4 ± 24.7	11.46 ± 1.21	0.061 ± 0.012	22.4 ± 1.8	12.62 ± 1.34
Weak-Burst	358.1 ± 28.4	14.72 ± 1.45	0.078 ± 0.015	19.2 ± 1.6	15.34 ± 1.51

under-represented. The bursty extension therefore strengthens, rather than weakens, the case for the auxiliary mechanisms of the proposed profile.

5.8 Image-Stream Service Quality and Overhead

The proposed scheduler protects control packets at the expense of the low-priority monitoring-image class. Reporting only the control-side metrics would leave the cost of that protection undocumented. Table 7 reports the four image-stream indicators introduced in Section 4.2 alongside the bandwidth overhead of redundant control transmission, across the four scenarios.

Image latency rises and image drop ratio increases as the link tightens, which is the intended elastic behavior: the scheduler treats image traffic as deferrable by design. Two boundary conditions are nonetheless

respected. First, the minimum image-traffic throughput share $\rho_{img}^{\min}$ remains above the C-SLA-1 floor of 15% (Section 3.6) under every tested scenario, including Weak-Burst where the lower 1σ bound on $\rho_{img}^{\min}$ is 17.6%. The scheduler therefore never starves the monitoring class even when the link is most stressed. Second, the structural-quality degradation under compression $\Delta SSIM$ remains below 0.08 throughout, a level at which operator situational awareness is preserved. The redundancy overhead η_{red} rises from below 1% on a good link to about 15% under Weak-Burst, remaining within the bandwidth budget of typical UAV C2 links; it is bounded by the C-WGT-1 weight constraints declared in Section 3.6.

These results document the cost of protecting control timeliness: monitoring images experience moderate latency and quality degradation under degraded links, but the low-priority class is never silenced and its quality remains usable. The minimum service guarantee for low-priority but mission-relevant traffic, declared qualitatively in Section 5.9 and formalized as conformance clause C-SLA-1 in Section 3.6, is therefore quantified here.

5.9 Engineering and Standardization Implications

From a standardization perspective, the scheduler is most useful on missions where the payload swing state is safety-critical and the link cannot be assumed stable, such as emergency delivery in damaged urban areas, cable-suspended transport over mountainous terrain, inspection close to large metal structures, and any operation that shares the link with an image or telemetry stream. On such missions the communication system cannot be designed for average throughput alone; it must ensure that the specific data the controller needs arrive before their control-validity deadline. The proposed scheduler offers a practical, implementation-independent way to express that requirement at the data-link level. It is, in effect, the kind of behavior that an interoperability profile or a pre-standardization specification would need to capture, and it is compatible with the scope of ongoing work on UAS connectivity in 3GPP and on IoT abstractions in W3C WoT [1, 7].

Deployment can be incremental. A flight controller or onboard companion computer tags each outgoing packet with its class, generation time, and deadline; a communication middleware process computes the dynamic priority and manages the transmission queue, with no change to the radio interface beneath it. Where packet-level acknowledgments are available, the reliability module adds redundant transmission for control packets as link risk rises; where only a plain serial or UDP link is present, the scheduler

can still reorder packets and throttle image traffic. The approach is therefore compatible with both edge-assisted architectures [1] and simple embedded real-time queue management [26]. Because the behavior is defined by packet classes, a priority function, and observable conformance criteria rather than by a specific product, it can be expressed as a lightweight, vendor-neutral profile on a W3C WoT binding inside an edge gateway, allowing conforming UAVs and ground stations to interoperate over standard transport protocols. In this form it constitutes a concrete candidate for pre-standardization discussion within the relevant ICT standardization bodies and industry consortia [1, 7].

Monitoring images require a separate note. On many missions an image stream supports operator situational awareness but is not always essential to the immediate closed loop. When the link is good, images can be sent normally. On a weak link, an image packet that delays a swing-state update can compromise safety even when the image arrives intact. The scheduler therefore treats image traffic as elastic, to be compressed, delayed, or briefly suppressed in order to protect control-packet timeliness. Operator visibility is preserved as a service but is made subordinate to flight-safety information while the link is under stress [35, 36].

The priority formulation also allows configuration. A mission that prioritizes payload stability can raise the swing-risk weight; a highly variable link calls for a larger link-risk weight; and strict control deadlines justify a higher urgency weight. This flexibility allows the same scheduler to serve different UAV platforms and mission profiles, although the weights must be chosen with care. Excessive emphasis on one term starves the other traffic classes, while insufficient emphasis leaves the control loop unprotected. A sound implementation should therefore bound the priorities and reserve a minimum service guarantee for low-priority traffic that remains mission-relevant [14]. These bounds and the minimum service guarantee are formalized as conformance clauses C-WGT-1 and C-SLA-1 in Section 3.6, and quantitatively verified by the image-stream service-quality data in Table 7 (Section 5.8), which records the minimum image-traffic throughput share $\rho_{img}^{\min}$ above the declared 15% floor under every tested scenario. From a standardization standpoint, these bounds and guarantees are exactly the parameters that a conformance profile would need to declare.

5.10 Reproducibility, Evidence Strength, and Code Availability

The study is numerical by design rather than flight-based. Simulation allows a controlled comparison of the scheduling algorithms under identical disturbances, link traces, and packetgeneration rules. Each scenario is

replayed 10 times per method under independent sub-seeds (Table 3), which yields a stable estimate of the mean response with explicit confidence bands and exposes the mechanism through which scheduling decisions influence anti-swing performance. The strength of the evidence should be read accordingly: the numerical results establish the feasibility and comparative advantage of the scheduling logic but cannot, on their own, demonstrate robustness on a flight platform. Real flight tests remain necessary before the method can be regarded as a validated operational system.

The figures and tables are arranged so that the reasoning remains traceable. Tables 4(a)–4(d) give the raw metrics under the four scenarios, Tables 5, 6, and 7 isolate the ablation decomposition, the weight perturbation sensitivity, and the low-priority service quality respectively, Figures 5 and 6 isolate the real-time communication behavior, Figures 7 and 8 link that behavior to the swing response, Figure 11 visualizes the ablation decomposition, Figure 12 reports the bursty-channel comparison, and Figures 9 and 10 fold the indicators into normalized summaries. This layered presentation avoids overclaiming from any single number and shows that the proposed scheduler improves both the communication layer and the control-relevant swing outcome under the conditions tested. The same layered presentation is intended to support independent conformance checking of any future implementation of the profile.

The full reproducibility package is released under the Apache 2.0 license and includes: (i) the simulation kernel with the suspended-load model, the WoT-style packet generator, and the time-varying data-link model; (ii) reference implementations of the four classical schedulers, the AoI-greedy baseline of Section 4.1, the proposed scheduler, and the ablation variants of Section 4.3; (iii) the Gilbert-Elliott bursty-channel generator with deterministic sub-seed handling; (iv) the SHA-256 derivation script that generates the 10-sub-seed family from the base seed 20260508 (Table 3); (v) the conformance test probe of Section 3.6 with the three test cases TC1, TC2, and TC3; and (vi) all post-processing scripts that regenerate Tables 4(a)–4(d), 5, 6, and 7, together with Figures 5–12, from the documented sub-seed family without any undocumented tuning.

6 Conclusion

This paper proposes a standards-aligned, deadline-aware data-link scheduling profile for UAV suspended-load anti-swing communication. The profile is separated from the control law and specified in vendor-neutral,

conformance-checkable terms. Main findings are summarized as follows:

(1) On the smoothed weak link, the proposed scheduler reduces average control-packet delay by 34.51% and deadline-miss rate by 82.48% relative to FIFO, and lowers the deadline-miss rate by 19.95 percentage points (a 57.5% relative reduction) and the RMS swing angle by 4.5% relative to fixed priority, the strongest converged baseline.
(2) Under the Gilbert-Elliott bursty channel with link-layer ARQ and MAC contention, the gap to fixed priority widens from 19.95 to 28.43 percentage points, indicating that the auxiliary mechanisms of the profile absorb bursty loss rather than degrade under it.
(3) A three-axis ablation attributes the Weak-Link miss-rate gap to dynamic priority (6.07 pp), redundant control transmission (8.78 pp) and image suppression (5.10 pp). No single mechanism dominates, and no term of the four-term priority formula is dispensable.
(4) The nominal weight vector is robust to local perturbation. The Weak-Link miss rate stays within a 0.45-pp band under simultaneous shifts of ± 0.05 in two weights at a time, and the low-priority monitoring class remains above the declared 15% throughput-share floor under every tested scenario.
(5) Five conformance clauses and three black-box test cases at the producer-side protocol boundary make the profile checkable from outside the scheduler, supporting its use as a prestandardization specification compatible with existing IoT (W3C WoT, MQTT, DDS) and 3GPP UAS standards.

The results reported here are numerical rather than flight-measured; the suspended load is reduced to a planar pendulum; and the Gilbert-Elliott channel does not exhaust real radiopropagation effects. Future work will move the scheduler onto a PX4/Gazebo or ROS testbed and then to hardware-in-the-loop trials over real wireless links, explore automatic optimization of the priority weights and adaptive image compression, and pursue formal submission of the profile to a relevant ICT standardization body for cross-vendor validation.

References

[1] W3C, "Web of Things (WoT) Architecture 1.1," W3C Recommendation, 5 December 2023. [Online]. Available: https://www.w3.org/TR/wot-architecture/.

[2] 3GPP, “Service Requirements for the 5G System; UAS Support,” 3GPP TS 22.125 v18.0.0, 2023.
[3] 3GPP, “Application Layer Support for Unmanned Aerial System (UAS); Functional Architecture and Information Flows,” 3GPP TS 23.256 v18.0.0, 2023.
[4] OASIS, “MQTT Version 5.0 OASIS Standard,” 2019. [Online]. Available: https://docs.oasis-open.org/mqtt/mqtt/v5.0/mqtt-v5.0.html.
[5] OMG, “Data Distribution Service (DDS) v1.4,” OMG Document formal/2015-04-10, 2015.
[6] 3GPP, “System Architecture for the 5G System (5GS); Stage 2,” 3GPP TS 23.501 v18.0.0, 2023.
[7] A. Fotouhi, H. Qiang, M. Ding, M. Hassan, L. G. Giordano, A. Garcia-Rodriguez, and J. Yuan, “Survey on UAV Cellular Communications: Practical Aspects, Standardization Advancements, Regulation, and Security Challenges,” IEEE Communications Surveys & Tutorials, vol. 21, no. 4, pp. 3417–3442, 2019, doi: 10.1109/COMST.2019.2906228.
[8] K. Sreenath, T. Lee, and V. Kumar, “Geometric Control and Differential Flatness of a Quadrotor UAV with a Cable-Suspended Load,” in Proc. 52nd IEEE Conference on Decision and Control, Florence, Italy, pp. 2269–2274, 2013, doi: 10.1109/CDC.2013.6760219.
[9] Y. Zeng, R. Zhang, and T. J. Lim, “Wireless Communications with Unmanned Aerial Vehicles: Opportunities and Challenges,” IEEE Communications Magazine, vol. 54, no. 5, pp. 36–42, 2016, doi: 10.1109/MCOM.2016.7470933.
[10] M. Mozaffari, W. Saad, M. Bennis, Y.-H. Nam, and M. Debbah, “A Tutorial on UAVs for Wireless Networks: Applications, Challenges, and Open Problems,” IEEE Communications Surveys & Tutorials, vol. 21, no. 3, pp. 2334–2360, 2019, doi: 10.1109/COMST.2019.2902862.
[11] X. Cao, P. Yang, M. Alzenad, X. Xi, D. Wu, and H. Yanikomeroglu, “Airborne Communication Networks: A Survey,” IEEE Journal on Selected Areas in Communications, vol. 36, no. 9, pp. 1907–1926, 2018, doi: 10.1109/JSAC.2018.2864423.
[12] T. Baidya, A. Nabi, and S. Moh, “Trajectory-Aware Offloading Decision in UAV-Aided Edge Computing: A Comprehensive Survey,” Sensors, vol. 24, no. 6, art. no. 1837, 2024, doi: 10.3390/s24061837.
[13] C. L. Liu and J. W. Layland, “Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment,” Journal of the ACM, vol. 20, no. 1, pp. 46–61, 1973, doi: 10.1145/321738.321743.

[14] G. C. Buttazzo, "Rate Monotonic vs. EDF: Judgment Day," Real-Time Systems, vol. 29, no. 1, pp. 5–26, 2005, doi: 10.1023/B:TIME.0000048932.30002.d9.

[15] R. Mahony, V. Kumar, and P. Corke, "Multirotor Aerial Vehicles: Modeling, Estimation, and Control of Quadrotor," IEEE Robotics & Automation Magazine, vol. 19, no. 3, pp. 20–32, 2012, doi: 10.1109/MRA.2012.2206474.

[16] D. Mellinger and V. Kumar, "Minimum Snap Trajectory Generation and Control for Quadrotors," in Proc. IEEE International Conference on Robotics and Automation (ICRA), Shanghai, China, pp. 2520–2525, 2011, doi: 10.1109/ICRA.2011.5980409.

[17] T. Lee, K. Sreenath, and V. Kumar, "Geometric Control of Cooperating Multiple Quadrotor UAVs with a Suspended Payload," in Proc. 52nd IEEE Conference on Decision and Control, Florence, Italy, pp. 5510–5515, 2013, doi: 10.1109/CDC.2013.6760757.

[18] T. Lee, "Geometric Control of Quadrotor UAVs Transporting a Cable-Suspended Rigid Body," IEEE Trans. Control Syst. Technol., vol. 26, no. 1, pp. 255–264, 2018, doi: 10.1109/TCST.2017.2656060.

[19] H. M. Omar, R. Akram, S. M. Mukras, and A. A. Mahfouz, "Recent Advances and Challenges in Controlling Quadrotors with Suspended Loads," Alexandria Engineering Journal, vol. 63, pp. 253–270, 2023, doi: 10.1016/j.aej.2022.07.044.

[20] J. Huang, H. Tao, Y. Wang, and J.-Q. Sun, "Suppressing UAV Payload Swing with Time-Varying Cable Length through Nonlinear Coupling," Mechanical Systems and Signal Processing, vol. 185, art. no. 109790, 2023, doi: 10.1016/j.ymssp.2022.109790.

[21] Y. Zhang, B. Fan, L. Sun, G. Huang, and Y. Zhao, "An Enhanced Energy Coupling-Based Control Method for Quadrotor UAV Suspended Payload with Variable Rope Length," International Journal of Precision Engineering and Manufacturing, vol. 25, pp. 2321–2336, 2024, doi: 10.1007/s12541-024-01069-6.

[22] J. Kang, J. Shan, and H. Alkomy, "VNHC-Based Continuous Sliding Mode Control for an Underactuated Tethered UAV System," IEEE Trans. Control Syst. Technol., vol. 32, no. 4, pp. 1410–1422, 2024.

[23] X. Chen, Y. Fan, G. Wang, and D. Mu, "Improved Extended Disturbance Observer-Based Integral Backstepping Sliding Mode Control for Quadrotor Slung-Load System," Aerospace Science and Technology, vol. 157, art. no. 109077, 2025, doi: 10.1016/j.ast.2025.109077.

[24] H. Shakhatreh, A. H. Sawalmeh, A. Al-Fuqaha, et al., "Unmanned Aerial Vehicles (UAVs): A Survey on Civil Applications and Key Research Challenges," IEEE Access, vol. 7, pp. 48572–48634, 2019, doi: 10.1109/ACCESS.2019.2909530.

[25] 3GPP, "Enhanced LTE Support for Aerial Vehicles," 3GPP TR 36.777 v15.0.0, 2018.

[26] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of Information: An Introduction and Survey," IEEE J. Sel. Areas Commun., vol. 39, no. 5, pp. 1183–1210, May 2021, doi: 10.1109/JSAC.2021.3065072.

[27] O. Ayan, M. Vilgelm, M. KlÃ¼gel, S. Hirche, and W. Kellerer, "Age-of-Information vs Value-of-Information Scheduling for Cellular Networked Control Systems," in Proc. ACM/IEEE 10th Int. Conf. on Cyber-Physical Systems (ICCPS), Montreal, QC, Canada, pp. 109–117, 2019, doi: 10.1145/3302509.3311050.

[28] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge Computing: Vision and Challenges," IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637–646, 2016, doi: 10.1109/JIOT.2016.2579198.

[29] M. Satyanarayanan, "The Emergence of Edge Computing," Computer, vol. 50, no. 1, pp. 30–39, 2017, doi: 10.1109/MC.2017.9.

[30] S. A. Khajeh, M. Saberikamarposhti, and A. M. Rahmani, "Real-Time Scheduling in IoT Applications: A Systematic Review," Sensors, vol. 23, no. 1, art. no. 232, 2023, doi: 10.3390/s23010232.

[31] E. Fountoulakis, T. Charalambous, A. Ephremides, and N. Pappas, "Scheduling Policies for AoI Minimization with Timely Throughput Constraints," IEEE Trans. Commun., vol. 71, no. 6, pp. 3489–3504, 2023, doi: 10.1109/TCOMM.2023.3263164.

[32] S. Hayat, E. Yanmaz, and R. Muzaffar, "Survey on Unmanned Aerial Vehicle Networks for Civil Applications: A Communications Viewpoint," IEEE Communications Surveys & Tutorials, vol. 18, no. 4, pp. 2624–2661, 2016, doi: 10.1109/COMST.2016.2560343.

[33] X. Fu and J. G. Kim, "Deep-Q-Network-Based Packet Scheduling in an IoT Environment," Sensors, vol. 23, no. 3, art. no. 1339, 2023, doi: 10.3390/s23031339.

[34] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, "A Survey on Mobile Edge Computing: The Communication Perspective," IEEE Communications Surveys & Tutorials, vol. 19, no. 4, pp. 2322–2358, 2017, doi: 10.1109/COMST.2017.2745201.

[35] D. Gunduz, Z. Qin, I. E. Aguerri, H. S. Dhillon, Z. Yang, A. Yener, K. K. Wong, and C.-B. Chae, "Beyond Transmitting Bits: Context, Semantics, and Task-Oriented Communications," IEEE J. Sel. Areas Commun., vol. 41, no. 1, pp. 5–41, Jan. 2023, doi: 10.1109/JSAC.2022.3223408.

[36] E. L. de Angelis, F. Giulietti, and G. Rossetti, "Two-Time-Scale Cooperative UAV Transportation of a Cable-Suspended Load: A Minimal Swing Approach," Drones, vol. 9, no. 8, art. no. 559, 2025, doi: 10.3390/drones9080559.

Biography



Kai Tang is affiliated with the College of Engineering, Huazhong Agricultural University, China. His primary research focus lies in Robotics Engineering, and has been engaged in research on drone control algorithms under the supervision of professional instructors.