

Budgeted Twin Synchronization over V2X Uplinks for Electric Vehicle Routing

Abdelkader Kaddour and Lamri Sayad

University Mohamed Boudiaf of M'sila, M'sila, Algeria

 <https://doi.org/10.26636/jtit.2026.4.2820>

 Submitted: 18 Aug. 2026, Reviewed: 10 Sep. 2026, Accepted: 14 Sep. 2026

Abstract — Connected electric vehicle (EV) fleets depend on cloud digital twins that rely on telemetry over lossy Internet of Things (IoT) and vehicle-to-everything (V2X) uplinks. When updates are lost, stale or budget-limited, twin ages and dispatch plans can fail. This paper studies two coupled decisions: when to spend an uplink message and how to plan from a stale twin. We compare periodic age-of-information (AoI) threshold and deviation triggers in a full factorial with and without a conservative twin correction based on split-conformal prediction (CP) under three communication regimes. Within each trigger pair, attempts, receptions, and AoI are identical. Across 144 confirmatory cells, raw triggers that plan on uncorrected snapshots remain at 0.0% strict feasibility. With a six attempt uplink cap and correction, AoI threshold triggering feasibility reaches 38.2%, with 4.51 mean attempts, versus a score of 36.8% for deviation+CP and 35.4% for periodic+CP. Uncapped hybrid admission does not improve feasibility, but reduces mean twin age (0.111 versus 0.166). Conservative correction is therefore a mechanism that restores feasible plans; trigger choice mainly trades messages against freshness. The routing solver is a fixed adaptive large neighborhood search (ALNS) module used only as an application key performance indicator (KPI).

Keywords — age of information, conservative correction, digital twin, IoT/V2X uplink, message budget

1. Introduction

Intelligent transportation systems increasingly integrate wireless connectivity with cyberphysical decision layers [1]–[3]. In connected electric vehicle (EV) logistics, roadside units, onboard sensors, and charging station back-ends feed a cloud digital twin that aggregates traffic multipliers, energy consumption factors, and charger availability [4]–[6]. The twin is then used to compute or refresh routes before dispatch. Cooperative V2X applications showed that communication quality is not an abstract physical layer metric: the green light speed advisory depends on timely infrastructure messages [7], and platoon control degrades when onboard sensors are inaccurate, even if a radio link is present [8]. Similarly, allocation under incomplete vehicle information shows that ignoring or delaying data changes optimization [9].

A closely related control problem appears in wireless sensor networks, where adaptive holding times decide how long a topology remains active before a costly refresh [10]. Fleet twins face an analogous question: how long may twin fields

remain active, and when should the uplink spend a scarce synchronization message?

Age of information (AoI) was introduced precisely to quantify the freshness of status updates on a monitor [11]–[13]. In vehicular networks, AoI interacts with packet loss, persistence, and contention in dedicated short-range communications (DSRC), long-term evolution (LTE)-V2X, and new radio links (NR)-V2X [14]–[19]. The digital twin literature has begun to treat synchronization itself as a network control problem: on-demand updates, game-theoretic coordination, and AoI-error-response schemes [20]–[26].

However, most EV routing papers still optimize under perfect or statically uncertain information. The classic electric vehicle routing problem with time windows (EVRPTW) models capture battery limits, time windows, and recharging stations [27]–[30], while robust and stochastic variants address uncertainty of energy and travel time as random parameters rather than delayed telemetry [31]–[34]. Recent digital twin EV routing work often assumes reliable communication and lists communication uncertainty as future work [35], [36]. Consequently, uplink message counts, packet success, and twin AoI are rarely first-class experimental metrics.

This paper studies fleet twin maintenance as a resynchronization control problem. The controller decides not only whether to request telemetry, but also how to spend a scarce message budget under packet loss and observation lag conditions. The routing engine is kept fixed, so that the observed plan quality can be attributed to the freshness and correction policy rather than to a new metaheuristic. Existing EV routing studies focus primarily on improving the plan after the information state is assumed known. The present paper addresses a narrower timing issue. The holding time of twin fields is usually implicit in a fixed polling clock, so its effect on uplink cost and residual plan infeasibility is not explicitly controlled.

This work is organized around three technical elements. First, cloud twin maintenance is formulated as a budgeted synchronization problem with explicit message accounting, matched packet outcomes, and AoI in both normalized and physical time. Second, every time- or event-trigger is evaluated with and without conservative twin correction, so that freshness control is not confounded with planning from uncertainty. Third, confirmatory holdout cells, communication regimes, a validation slice inspired by the simulation of urban mobility (SUMO), and the AoI message Pareto front are reported.

The dominant feasibility effect is a conservative correction under a message cap: every raw trigger stays at 0.0%, whereas the same triggers with correction recover 35–38%. Among budgeted corrected policies, an AoI threshold uses fewer messages than a dense periodic schedule and has the highest observed feasibility, although the difference in trigger level is not statistically significant. Extra uncapped attempts mainly reduce twin age. Therefore, the results should be interpreted as a freshness overhead trade-off, not as dominance of one trigger over all routing metrics.

The remainder of the paper is organized as follows. Section 2 reviews V2X telemetry, AoI, digital twin synchronization, and EV routing under uncertainty. Section 3 defines the twin, uplink, AoI, and message budget objective. Section 4 presents the proposed synchronization pipeline. Section 5 describes the experimental protocol. Section 6 presents the results and limitations, while Section 7 concludes the paper.

2. Related Works

2.1. V2X Telemetry, AoI, and Refresh Timing

Dedicated short-range and cellular V2X technologies have been empirically compared for cooperative awareness [2], [3], [14], [19]. Application-centric evaluations interpret packet delivery through the service they enable, as in green light speed advisory [7] and platooning under sensor inaccuracy [8]. Incomplete vehicle information changes allocation quality [9]. Telemetry based on the message queueing telemetry transport protocol (MQTT) is a typical IoT implementation path [37]. Adaptive hold-time control in clustered wireless sensor networks (WSNs) decides not only whether to rebuild the state but also how long the selected state remains active [10].

Secure data aggregation in WSNs addresses energy-lifetime trade-offs under dynamic conditions [38], while blockchain architectures provide integrity for IoT data processing [39]. Localization under 6G semantic communication requirements of 6G motivates context-aware resource management [40]. Vehicular caching with mobility prediction reduces back-haul load in dynamic networks [41]. The detection in cloud infrastructures is based on flow log analysis [42].

In [11], status update freshness as age of information is formulated. Subsequent work studies when to update versus wait [12] and surveys AoI as a bridge between network performance and application utility [13]. Vehicular instantiations connect AoI to reliability and error-prone sidelinks [15]–[18]. Hybrid time- and event-triggered systems ask whether determinism or resource conservation matters more once resources are equalized [43]. We adopt twin field AoI as the freshness state and compare periodic, threshold, and deviation triggers under a hard message budget.

2.2. Digital Twin Synchronization and Conservative Planning

Vehicular twins require data synchronization under bandwidth limits [20], [22]–[25]. We keep the payload as a compact

disruption vector and study *when* to request it. Split conformal prediction provides distribution-free upper sets [44], [45] that we use as an enabling correction of stale fields, not as the primary theoretical contribution. EVRPTW with recharge stations is used only as a downstream application KPI [27], [46]. Recent digital twin EV routing papers generally assume reliable communication or list it as future work [35]. We keep ALNS fixed, so that observed plan quality can be attributed to the synchronization and correction layer.

3. System Model and Problem Formulation

Let the nominal fleet instance be I_0 , with customers, stations, time windows, cargo, and battery parameters as in [27]. A disruption state $d = (\tau, e, q, S)$ captures a traffic multiplier τ , an energy multiplier e , a charger queue multiplier q , and a set S of disabled stations. At the horizon fraction $t \in [0, 1]$, the physical world realizes a ramped disturbance $d^{\text{true}}(t)$ that interpolates from the nominal state to a final scenario (traffic grows non-linearly and station outages appear late). The cloud twin stores the newest successfully delivered observation d^{cloud} together with its source generation time g^{cloud} . At decision time t , its AoI is:

$$\Delta(t) = t - g^{\text{cloud}}, \quad (1)$$

remaining consistent with the definition of status update [11], [13]. If T denotes the instance horizon (depot due time, in minutes), the same age in minutes is $\Delta_{\min}(t) = \Delta(t)T$. We report both, so that the freshness is not left as a unitless fraction.

An uplink request may be lost with probability p_{loss} . If it is delivered at decision epoch t , the available backend snapshot describes $d^{\text{true}}(t - \ell)$, where ℓ is observation lag as a fraction of the horizon. Its source timestamp is therefore $g_{\text{new}} = t - \ell$, and the post-reception age is $\Delta(t^+) = t - g_{\text{new}} = \ell$, not zero.

This is an observation: it does not claim to simulate packet queuing or a causal physical layer propagation delay.

Sensor noise models the fact that even a delivered packet is not a perfect physical measurement [8]. We study three regimes inspired by the reported outdoor V2X reliability spread [15], [18], [19]: *normal* ($p_{\text{loss}} = 0.05$, $\ell = 0.02$), *sensor-noise factor* $\eta = 0.02$), *degraded* (0.25, 0.08, 0.07), and *severe* (0.45, 0.15, 0.12). All controllers share matched random streams, so packet outcomes at a given epoch are identical across policies.

Figure 1 summarizes the architecture. Vehicles and charging stations emit telemetry data. A lossy IoT/V2X uplink delivers a subset of requests. The cloud twin stores aged fields, a sync controller decides whether to spend a message, and a fixed ALNS planner consumes the planning world. The MQTT publish/subscribe is a natural implementation of the uplink [37], but we abstract it as a request-loss-lag channel, so that the contribution remains decision-layer control rather than protocol engineering.




Fig. 1. Communication architecture. The controller decides whether to spend an uplink attempt and a successfully delivered time-stamped sample updates the cloud twin. The raw or conformal planning world feeds the same ALNS module, while the hidden truth is used only for the final KPI evaluation.

Let

$$\mathcal{T}_{0.1} = \{0.10, 0.20, \dots, 0.90\}.$$

All controllers are evaluated on the same 11-epoch decision clock

$$\mathcal{T} = \mathcal{T}_{0.1} \cup \{0.25, 0.75\}.$$

The two quarter horizon epochs support the fixed periodic schedule

$$\mathcal{T}_P = \{0.10, 0.25, 0.40, 0.60, 0.75, 0.90\}.$$

Event controllers screen all epochs in $\mathcal{T}$. At each candidate epoch, the controller selects a refresh indicator $u(t) \in \{0, 1\}$. Let A be the number of attempted uplinks. The hard budget constraint for every method in the confirmatory factorial is:

$$A = \sum_{t \in \mathcal{T}} u(t) \leq B, \quad (2)$$

with $B = 6$.

Every factorial controller executes the same 11 candidate epochs: periodic sets $u(t) = 1$ only at its six scheduled epochs, while the event controllers screen their conditions at all 11 epochs. The cap admits at most six transmission attempts regardless of how many conditions are screened.

Six is a controlled experimental budget chosen to match the fixed periodic schedule. It is not a V2X standard or a claimed physical network limit. The supplementary uncapped hybrid uses $B = \infty$ instead and is not part of the factorial. Thus, the decision clock is fixed, and the budget constrains admission, not the number of screening opportunities. A successful update reduces Δ to the delivered snapshot's lag ℓ , analogous to resetting a holding time state while retaining non-zero source staleness [10].

4. Proposed Synchronization Controller

4.1. Overview of the Pipeline

The proposed pipeline has four blocks. The first block extracts compact twin descriptors: AoI, last cloud sample, and a dimensionless risk index. The second block selects $u(t)$ using a periodic age threshold, deviation, predictive admission (PAI), or hybrid AoI-risk rule. The third block attempts the uplink only when $u(t) = 1$ and $A < B$, then optionally builds a conservative planning world. The fourth block executes a matched ALNS re-solve and returns application KPIs at the hidden true disruption.

This separation keeps the role of each module explicit. The sync policy controls the refresh time, while the routing construction rule remains a fixed evaluation engine.

4.2. Comparison of Policies

The policies compared span the time-triggered versus event-triggered approaches, as discussed in hybrid information systems [10], [43]:

- Periodic raw @6 / periodic+CP @6: both request on the same six fixed epochs in $\mathcal{T}_P$ and obey $B = 6$.
- AoI raw @6 / AoI+CP @6: both request when the pre-decision age satisfies $\Delta(t) \geq \Delta_{th} = 0.25$, subject to $B = 6$. This is a classical age gate [11].
- Deviation raw @6 / deviation+CP @6: both request when the same local edge proxy changes by more than 10%, subject to $B = 6$.
- In each pair, the raw method plans on the received snapshot, whereas the split conformal prediction (CP) method applies conservative correction *after* reception. CP is not used as an input to any of the three trigger decisions. The resulting 3×2 factorial therefore changes only the planning world within each trigger pair.
- PAI @6: feasibility/risk admission with cooldown, conservative correction, capped at six attempts.
- PAI-hybrid uncapped: risk gate plus an AoI soft trigger without the hard cap, used only as a supplementary AoI message operating point.

Conservative correction uses a Mondrian split conformal upper world calibrated on independent trajectories [44], [45], [47]. For the three continuous fields $k \in \{\tau, e, q\}$, a causal forecast from the stale snapshot provides a location μ_k and scale σ_k at a fixed 0.18-horizon look-ahead. Its effective age is the snapshot age plus 0.18. The one-sided simultaneous nonconformity score is

$$r_k = \left[\frac{(y_k - \mu_k)}{\sigma_k} \right]_+ . \quad (3)$$

$$s_c = \max_{k \in \{\tau, e, q\}} r_k$$

For station outages, the separate score is the largest $1 - \pi_j$ among truly unavailable stations, where π_j is the forecast outage score. Finite sample quantiles use the order statistic $[(n+1)(1-\alpha)]$ within the communication regime and

Algorithm 1 Twin synchronization loop (pre-dispatch)

Require: Nominal instance I_0 , controller π , candidate set $\mathcal{T}$, message budget B

```

1: Plan initial routes on  $I_0$ ; set  $g^{\text{cloud}} \leftarrow 0$ ,  $d^{\text{cloud}} \leftarrow d_0$ 
2: for each candidate epoch  $t$  do
3:    $\Delta^- \leftarrow t - g^{\text{cloud}}$ 
4:   Decide trigger  $u \leftarrow \pi(t, \Delta^-, d^{\text{cloud}})$ 
5:   if  $u = 1$  and attempts  $< B$  then
6:     Increment attempts and request an uplink sample
7:     if the packet is received then
8:       Set  $d^{\text{cloud}} \leftarrow d^{\text{true}}(t - \ell) + \epsilon$ ,  $g^{\text{cloud}} \leftarrow t - \ell$ 
9:       Build  $d^{\text{plan}}$  (raw snapshot or CP-corrected world)
10:      Re-solve customer sequences with ALNS on  $d^{\text{plan}}$ 
11:      end if
12:    end if
13:    Record  $\Delta^+(t) \leftarrow t - g^{\text{cloud}}$ 
14:  end for
15: Evaluate final routes on hidden true disruption at  $t = 1$ 

```

effective age strata (strata with fewer than 25 records use the pooled quantile). We use $\alpha_c = 0.10$ for the simultaneous continuous upper rectangle and $\alpha_s = 0.05$ for the outage superset.

The locked model was fitted to 4 320 timestamped forecast-truth pairs from 480 trajectories (10 calibration instances, two severities, three regimes, and eight trajectories per combination). A further 864 pairs from 96 trajectories on two separate calibration-test instances yielded 91.9% continuous coverage against the 90% target and 100.0% outage coverage against the 95% target.

All 12 calibration instances are disjoint from the 24-instance confirmatory holdout. These are marginal synthetic generator coverage results under exchangeability, not a deployment certificate. The planning world replaces each field with its conformal upper bound before the fixed ALNS planner consumes it. In a deployed system, the calibration window would itself be refreshed over time, for example with sliding window or adaptively weighted conformal updates to track nonstationary traffic [48].

We treat this layer as an information technology (IT) mechanism for sensing uncertainty in stagnant fields, analogous to planning in prediction sets [49]. The planner is a fixed ALNS module [46] with matched iteration and time budgets. All synchronization occurs *before* dispatch. We do not claim reoptimization of the state of the vehicle en route, which would require a different dynamic vehicle routing problem protocol [50].

A dimensionless risk index

$$r(d) = 0.35 \max(\tau - 1, 0) + 0.40 \max(e - 1, 0) + 0.20 \max(q - 1, 0) + 0.12|S| \quad (4)$$

supports soft triggers.

Equation (4) is used only to interpret the refresh timing trade-off, and it is not a learned reward. PAI fires a *hard* trigger when the current plan is infeasible under the conservative world,

and a *soft* trigger when the risk increment exceeds a threshold selected on a disjoint C5 pilot. The hybrid mode also treats $\Delta \geq \Delta_{\text{th}}$ as a soft trigger, echoing adaptive holding-time refresh [10]. Soft triggers are suppressed by a cooldown and, under budgets, by a soft reserve that keeps remaining slots for later hard alarms. This reserve is a simple admission control analog of saving control energy for future topology rebuilds in sensor networks.

4.3. Metrics

Primary metrics follow the networked status update practice [13], [15]: mean and peak Δ (horizon fraction and minutes), attempted messages, received messages, and success rate received/attempted. The trigger reads the pre-decision age $\Delta^-(t)$. For reporting, each controller contributes exactly one post-decision sample $\Delta^+(t)$ at every common epoch $t \in \mathcal{T}$ and a loss leaves the age unchanged, while a successful delivery sets it to ℓ .

For example, if $\ell = 0.08$ and a sample is received at $t = 0.50$, its source timestamp is 0.42 and its post-reception age is $0.50 - 0.42 = 0.08$.

We report

$$\bar{\Delta} = |\mathcal{T}|^{-1} \sum_{t \in \mathcal{T}} \Delta^+(t)$$

and

$$\Delta_{\max} = \max_{t \in \mathcal{T}} \Delta^+(t),$$

then convert each cell to minutes using its own horizon T before aggregation. Hence $\bar{\Delta}$ is a sampled decision-epoch average, not a continuous-time age integral. The budget counts every transmission attempt, whether the packet is lost or received. The successful receptions and received/attempted are reported separately. Attempts are not airtime or bytes.

Secondary application KPIs are strict EVRPTW feasibility and late customers under the hidden true disruption at $t = 1$. A forecast message with perfect information and the same ALNS budget provides an information reference benchmark, not a mathematical optimization bound [35].

5. Experimental Design

5.1. Instances and Splits

We use Schneider EVRPTW instances [27] derived from Solomon families: clustered (C), random (R), and mixed (RC). Thresholds $\Delta_{\text{th}} = 0.25$, risk threshold 0.18, and cooldown 0.18 were selected on a disjoint C5 pilot that is never reused in confirmatory tables. The confirmatory set comprises 24 untouched 100-customer instances, two environment severities, and three communication regimes, yielding 144 paired cells. A complementary SUMO/V2X inspired slice uses the same controllers on a smaller instance set with bursty loss [18], [51].

5.2. Parameters and Reproducibility

Table 1 summarizes the key experimental parameters. All simulations use fixed random seeds per cell, and the calibra-

Tab. 1. Experimental parameters.

Parameter	Value
Instance set	Schneider EVRPTW, 24 instances
Customers per instance	100
Environment severities	2 (medium, high)
Communication regimes	3 (normal, degraded, severe)
Confirmatory cells	144 (24 inst. $\times$ 2 sev. $\times$ 3 reg.)
Common candidate epochs	11
Periodic request epochs	6
Message budget (all factorial methods)	6 attempts
AoI threshold Δ_{th}	0.25
CP look-ahead fraction	0.18
Risk threshold (PAI)	0.18
Cooldown (PAI)	0.18
Continuous/outage targets	0.90 / 0.95
Calibration fit/test pairs	4320 / 864
ALNS solve budget	800 iterations / 1.2 s
Random streams	1/cell matched, 3-seed check

tion set for conformal correction is strictly disjoint from the evaluation holdout.

The simulation code and result aggregation scripts are available upon request.

5.3. Fairness and Communication Matching

All controllers share identical scenario seeds and per-epoch communication outcomes. If two controllers request at the same epoch, they receive the same loss and noise realization. More importantly, the two members of each raw / CP confirmatory pair use the same candidate clock, trigger parameters, and $B = 6$ cap. CP is applied only after successful reception and is not an input to the periodic, AoI, or deviation trigger. Thus, within each raw/CP pair, communication exposure is held fixed by construction: the two methods request at the same epochs, face the same packet outcomes, receive the same observations, and produce the same AoI trajectory.

The only difference in treatment is whether the received snapshot is used directly or transformed into a conservative planning world. Automated validation checks these equalities before the tables are exported.

The primary factorial deliberately uses one matched communication stream per instance-severity regime cell, as its estimand is the within-cell CP-versus-raw effect under identical exposure, not unconstrained seed-to-seed variability. The separate three seed experiment is reported as a robustness check for its overlapping controllers. ALNS iteration and time budgets are matched. No controller receives the hidden severity label, and there is no final perfect information repair. These controls prevent correction from being confounded with communication opportunities or luckier packets [43].

5.4. Statistics

For each instance, we average feasibility differences across its cells, then apply an exact two-sided sign test on non-tied in-

stances and an instance cluster bootstrap (20 000 resamples) for 95% intervals. Clustering by instance avoids treating correlated regime/severity cells as independent. The exact sign test assesses directional balance among non-tied instance contrasts, whereas the bootstrap interval estimates the mean lift by resampling all instances, including ties. These procedures answer different questions and are reported as complementary. A bootstrap interval excluding zero does not override a non-significant sign test in our trigger ranking interpretation. We report lifts in percentage points, not relative percentages, because raw baselines are exactly zero in the confirmatory set. The three within-trigger the CP-versus-raw contrasts are primary. Comparisons among corrected triggers are secondary and are interpreted descriptively without a multiplicity-adjusted ranking claim.

6. Results and Discussion

6.1. Matched-factorial Freshness and Feasibility

Table 2 summarizes the full holdout. Raw periodic, AoI threshold, and deviation controllers all achieve 0.0% strict feasibility despite 4.51 – 6.00 uplink attempts and 71.8 – 76.1% packet success. This is the central negative result: exchanging telemetry is not sufficient if the planner consumes uncorrected snapshots under the loss and observation lag. The factorial with conservative correction isolates the mechanism. The same three triggers with CP and a six-try cap reach 35.4%, 38.2%, and 36.8% (periodic, AoI, deviation), respectively. Thus, the jump from 0% to the mid thirties occurs due to conservative planning, not different communication opportunities: every raw/CP pair has identical attempt, reception, and AoI values.

Among the capped corrected policies, AoI+CP achieved the highest observed feasibility and used the fewest attempts (4.51 versus 6.00 for periodic+CP), although the trigger-level difference was not statistically significant. The mean decision epoch age is 0.166 (198.2 min after cell-wise physical-time conversion) for AoI +CP and 0.165 (197.4 min) for periodic+CP. The uncapped hybrid controller reaches 34.7% with the lowest mean AoI (0.111, 132.6 min) at 10.99 attempts. Compared to the perfect information (PI) reference (56.2%), the 38.2% AoI+CP rate is approximately 68% as large. Late customers fall from 23.55 to 33.38 under raw policies to below 0.45 under corrected policies.



Figures 2 and 3 visualize the feasibility, while Fig. 4 shows the AoI-message trade-off. Each raw/CP matched pair occupies exactly the same communication point, so the raw-to-corrected annotation isolates the feasibility effect of planning correction. The differences between the periodic, AoI, and deviation points instead show the trigger-induced message freshness trade-off. The uncapped hybrid remains visually separate at a higher message count.

6.2. Statistical Contrasts


Against the matched raw trigger, periodic+CP, AoI+CP, and deviation+CP obtain instance level lifts of +35 to +38 percent-

Tab. 2. Full holdout (144 cells) freshness metrics and application KPIs. The first six rows form the matched factorial, PAI rows are supplementary.

Method	Feas. [%]	AoI	AoI [min]	Max AoI	Messages	Succ. [%]	Late
Periodic raw @6	0.0	0.165	197.4	0.332	6.00	74.4	23.55
Periodic+CP @6	35.4	0.165	197.4	0.332	6.00	74.4	0.44
AoI raw @6	0.0	0.166	198.2	0.312	4.51	76.1	29.18
AoI+CP @6	38.2	0.166	198.2	0.312	4.51	76.1	0.08
Deviation raw @6	0.0	0.185	212.1	0.387	5.28	71.8	33.38
Deviation+CP @6	36.8	0.185	212.1	0.387	5.28	71.8	0.08
PAI @6	31.2	0.229	266.0	0.528	6.00	73.3	1.10
PAI-hybrid uncapped	34.7	0.111	132.6	0.214	10.99	74.3	0.08
Oracle (PI)	56.2	—	—	—	—	—	0.00

**Fig. 2.** Strict feasibility in the full holdout. Hatched raw bars are zero, colored bars show the matched CP variants and supplementary controllers. The dashed line is the PI reference.**Fig. 3.** Feasibility by communication regime. Bars show the three corrected controllers; × markers denote their matched raw variants, all at 0%.

age points (pp) respectively, with sign-test $p = 4.9 \times 10^{-4}$ (12 positive, 0 negative, 12 tied instances) and bootstrap 95% intervals excluding zero (Tab. 3). AoI+CP versus periodic+CP is +2.8 pp (4/0/20, $p = 0.12$). The interval [0.7, 5.6] excludes zero, but most of the instances are tied, so we do not claim a significant trigger ranking under a matched cap. The interval

**Fig. 4.** AoI vs. mean attempted messages. Each arrow annotation gives the raw-to-CP feasibility change for a matched trigger. The pair occupies one communication point because its attempts, receptions, and age are identical. The uncapped hybrid is a supplementary operating point.

estimates a mean lift over all resampled instances, whereas the sign test uses directions among non-ties. Under the stated primary sign-test interpretation, the $p = 0.12$ result does not establish superiority. Deviation+CP versus periodic+CP is +1.4 pp, and AoI+CP versus deviation+CP is +1.4 pp. Both sign tests give $p = 1.00$ and intervals that span zero. Therefore, the dominant mechanism is a *conservative twin correction*. The trigger choice mainly reshapes the message and the age operating point. This echoes hybrid-triggered analyses, in which determinism can match event sampling once the resources are equalized [43].

6.3. Communication regimes

Figure 3 and Tab. 4 break the results by regime. Raw policies remain at 0% feasibility in normal, degraded, and severe channels. Under severe loss, AoI+CP remains at 37.5% feasibility with six attempts, while periodic+CP reaches 31.2% and deviation+CP reaches 33.3%. In the normal regime, all three corrected triggers reach 39.6% feasibility; AoI+CP uses 3.12 attempts, deviation+CP 4.94, and periodic+CP 6.00.

Tab. 3. Instance-clustered contrasts (full holdout).

Contrast	Lift (pp)	+ / - / =	p	95% CI
Periodic+CP @6 vs Periodic raw @6	+35.4	12 / 0 / 12	4.9×10^{-4}	[19.4, 52.1]
AoI+CP @6 vs AoI raw @6	+38.2	12 / 0 / 12	4.9×10^{-4}	[21.5, 54.9]
Deviation+CP @6 vs Deviation raw @6	+36.8	12 / 0 / 12	4.9×10^{-4}	[20.8, 53.5]
AoI+CP @6 vs Periodic+CP @6	+2.8	4 / 0 / 20	0.12	[0.7, 5.6]
Deviation+CP @6 vs Periodic+CP @6	+1.4	2 / 1 / 21	1.00	[-1.4, 4.9]
AoI+CP @6 vs Deviation+CP @6	+1.4	2 / 1 / 21	1.00	[-1.4, 4.9]

In the degraded regime, AoI+CP and deviation+CP reach 37.5%, versus 35.4% for periodic+CP. These patterns are consistent with the outdoor reliability of V2X that varies with load and geometry [16], [18], [19].

6.4. Multi-seed Confirmation

An existing 216 cell, three seed confirmation on 12 instances gives 37.0% feasibility for periodic+CP@6, 36.1% for the uncapped hybrid, and 56.9% for the oracle. We use it only as a stochastic check on these overlapping controllers. It is not used in the factorial contrasts, since that earlier suite did not contain all six matched raw/CP methods.

6.5. SUMO and Bursty-channel Slice


Table 5 reports a complementary 72 cell slice that replays SUMO speed profiles through a published V2X data set-inspired bursty sidelink model [18], [51]. Raw deviation and AoI-threshold triggers again yield 0.0% feasibility. Periodic+CP and a predictive CP controller reach 33.3%, with the predictive policy using more messages (10.4 versus 6.0) and fewer late customers. This slice is not a full city-scale co-simulation. It checks that the correction versus raw gap is not an artifact of i.i.d. Bernoulli loss.

6.6. Family and Severity Patterns

Figure 5 reports the feasibility of AoI+CP by instance family (C/R/RC) and the severity of the environment. Gains appear across families; high-severity cells remain harder. Clustered instances are not uniformly easier than random instances once communication impairment is present [27].

6.7. Operational Reading of the Message Budget

Under a hard cap of six attempts, AoI+CP is a practical default: 38.2% feasibility is achieved at 4.51 messages – the highest observed value among budgeted corrected policies, although pairwise differences with periodic+CP do not reach statistical significance. Periodic+CP remains a predictable provisioning choice at exactly six slots (35.4%). When operators can afford about 11 attempts, the hybrid gate buys substantially lower mean twin AoI (0.111, 132.6 min) with nearly the same feasibility (34.7%). This is a trade-off between freshness and control cost trade-off similar in spirit to adaptive holding time policies [10], [12].

**Fig. 5.** AoI+CP feasibility by family and severity.

6.8. Interpretation of the Negative Result

Periodic, AoI-threshold, and deviation controllers all exchange messages and achieve non-trivial packet success, yet they produce 0.0% strict feasibility. Reducing twin age is therefore insufficient if the planner consumes optimistic raw snapshots under loss and lag. The effects of sensor inaccuracy on cooperative applications have been documented for platooning [8]. The present results play a similar role for fleet twin planning fields. Robust electric vehicle routing problem (EVRP) models that treat energy as a static uncertainty set [31], [32] address a different failure mode. They do not explain why a freshly delivered but uncorrected snapshot still yields infeasible routes.

Under a hard six-try budget, AoI+CP and periodic+CP are both usable defaults. Event-triggered PAI is not better once correction is present. When the budget is relaxed, hybrid admission buys a lower mean AoI at nearly the same feasibility. Network operators who must provision update capacity – as in cellular V2X (C-V2X) application trials [7] – may prefer a fixed periodic or AoI schedule with correction, whereas operators who prioritize freshness may spend extra attempts. Since ALNS is fixed, feasibility lifts cannot be attributed to a new destroy/repair set [46]. Keeping the solver as an application KPI is consistent with application-centric V2X evaluation [7].

6.9. Scope and Limitations

The study is a synthetic pre-dispatch decision-layer evaluation. It does not claim physical layer (PHY) or medium access control (MAC) efficiency, real-city safety certification, or en-route

Tab. 4. Complete factorial by communication regime (full holdout).

Regime	Method	Feasibility [%]	AoI	AoI (min)	Messages
Normal	Periodic raw @6	0.0	0.065	72.0	6.00
	Periodic+CP @6	39.6	0.065	72.0	6.00
	AoI raw @6	0.0	0.110	128.5	3.12
	AoI+CP @6	39.6	0.110	128.5	3.12
	Deviation raw @6	0.0	0.102	113.8	4.94
	Deviation+CP @6	39.6	0.102	113.8	4.94
Degraded	Periodic raw @6	0.0	0.162	203.8	6.00
	Periodic+CP @6	35.4	0.162	203.8	6.00
	AoI raw @6	0.0	0.155	181.9	4.42
	AoI+CP @6	37.5	0.155	181.9	4.42
	Deviation raw @6	0.0	0.186	216.1	5.35
	Deviation+CP @6	37.5	0.186	216.1	5.35
Severe	Periodic raw @6	0.0	0.267	316.4	6.00
	Periodic+CP @6	31.2	0.267	316.4	6.00
	AoI raw @6	0.0	0.235	284.3	6.00
	AoI+CP @6	37.5	0.235	284.3	6.00
	Deviation raw @6	0.0	0.265	306.3	5.56
	Deviation+CP @6	33.3	0.265	306.3	5.56

Tab. 5. SUMO profile replay with a V2X dataset inspired burst loss model (72 cells).

Method	Feas. [%]	Messages	Late
Deviation (raw)	0.0	5.99	22.50
AoI-threshold	0.0	3.67	23.86
Periodic+CP @6	33.3	6.00	1.78
Predictive (CP)	33.3	10.36	0.15

replanning. Message counts are request attempts, not airtime or bytes. Reported AoI is a common clock, post-decision sample average, rather than a continuous-time integral. The three regime lags are deterministic observation-staleness parameters, and the 144-cell factorial uses one matched communication stream per instance-severity-regime cell. The three-seed check covers only the overlapping controllers noted above.

The matched design identifies the effect of paired treatment under fixed exposure but does not fully characterize the Monte Carlo variability across packet realizations. Solver wall clock time is constrained for fairness, but is not added to telemetry AoI. The three loss/lag/noise regimes are controlled experimental abstractions, not measurements from a deployed V2X network. They are inspired by published measurement trends [16], [18], [19], not a trace-driven MAC simulation. The conformal correction is calibrated under exchangeability assumptions that can be violated under non-stationary urban traffic [48]. Therefore, we report it as an enabling component rather than a coverage certificate for deployment. Microscopic traffic simulators, such as SUMO [51] and richer sidelink

traces remain complementary validation paths. Semantic compression of telemetry [52], [53] is out of scope – we study request timing, not payload coding.

7. Conclusions

This paper formulated fleet twin maintenance as an adaptive synchronization problem with explicit AoI, packet loss, and message budget accounting. The evaluation keeps the routing construction rule fixed and places the controller above it. On 144 holdout confirmatory cells, raw periodic, AoI threshold, and deviation sync yield 0.0% strict plan feasibility. The same triggers with conservative correction and a six-attempt cap reach 35.4 to 38.2%, with AoI +CP using the fewest attempts (4.51). Uncapped hybrid admission reaches 34.7% with a mean twin AoI of 0.111 (132.6 min). Compared to a perfect-information reference of 56.2%, the 38.2% AoI +CP rate is approximately two thirds as large. A 72-cell SUMO/bursty-channel slice reproduces the raw versus correction gap. The method is a freshness / uplink-overhead control mechanism and not a claim of dominance over all routing metrics. Budgeted AoI +CP or periodic+CP is the default under a hard message cap. Extra attempts are useful when twin age, not extra feasible plans, is the operating goal.

Future work includes MQTT-style telemetry stacks [37], packet-level sidelink replay [18], and online request arrivals after dispatch [50].

References

- [1] “White Paper: Communication Technologies for Intelligent Transportation Systems: From Railways to UAVs and Beyond”, S. Rizkalla and A. Kliks, Eds., *Journal of Telecommunications and Information Technology*, 2025 (<https://doi.org/10.26636/jtit.2025.COST-CA20120-VT2.2385>).
- [2] K. Abboud, H.A. Omar, and W. Zhuang, “Interworking of DSRC and Cellular Network Technologies for V2X Communications: A Survey”, *IEEE Transactions on Vehicular Technology*, vol. 65, pp. 9457–9470, 2016 (<https://doi.org/10.1109/TVT.2016.2591558>).
- [3] S. Chen *et al.*, “Vehicle-to-everything (V2X) Services Supported by LTE-based Systems and 5G”, *IEEE Communications Standards Magazine*, vol. 1, pp. 70–76, 2017 (<https://doi.org/10.1109/MCOMSTD.2017.1700015>).
- [4] M. Grieves and J. Vickers, “Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems”, in: *Transdisciplinary Perspectives on Complex Systems*, pp. 85–113, 2017 (https://doi.org/10.1007/978-3-319-38756-7_4).
- [5] F. Tao, H. Zhang, A. Liu, and A.Y.C. Nee, “Digital Twin in Industry: State-of-the-art”, *IEEE Transactions on Industrial Informatics*, vol. 15, pp. 2405–2415, 2019 (<https://doi.org/10.1109/TII.2018.2873186>).
- [6] W.A. Ali, M.P. Fanti, M. Roccotelli, and L. Ranieri, “A Review of Digital Twin Technology for Electric and Autonomous Vehicles”, *Applied Sciences*, vol. 13, art. no. 5871, 2023 (<https://doi.org/10.3390/app13105871>).
- [7] I.P. Chochliouros *et al.*, “C-V2X Communications for the Support of a Green Light Optimized Speed Advisory (GLOSA) Use Case”, *Journal of Telecommunications and Information Technology*, vol. 84, pp. 93–106, 2021 (<https://doi.org/10.26636/jtit.2021.152321>).
- [8] M. Sybis, P. Sroka, and P. Kryszkiewicz, “V2X Communications for Platooning: Sensor Inaccuracy Aspects”, *Journal of Telecommunications and Information Technology*, vol. 80, pp. 78–85, 2020 (<https://doi.org/10.26636/jtit.2020.142820>).
- [9] F. Marzuk, A. Vejar, and P. Cholda, “k-Anonymity in Resource Allocation for Vehicle-to-everything (V2X) Systems”, *Journal of Telecommunications and Information Technology* (Special Issue), pp. 1–4, 2025 (<https://doi.org/10.26636/jtit.2025.FITCE2024.1998>).
- [10] T.-M. Hoang *et al.*, “A Lightweight Adaptive Holding-time Policy for Clustered Wireless Sensor Networks”, *Journal of Telecommunications and Information Technology*, vol. 104, pp. 56–66, 2026 (<https://doi.org/10.26636/jtit.2026.2.2598>).
- [11] S.K. Kaul, R.D. Yates, and M. Gruteser, “Real-time Status: How Often Should One Update?”, *Proc. of IEEE INFOCOM*, pp. 2731–2735, 2012 (<https://doi.org/10.1109/INFCOM.2012.6195689>).
- [12] Y. Sun *et al.*, “Update or Wait: How to Keep Your Data Fresh”, *IEEE Transactions on Information Theory*, vol. 63, pp. 7492–7508, 2017 (<https://doi.org/10.1109/TIT.2017.2735804>).
- [13] R.D. Yates *et al.*, “Age of Information: An Introduction and Survey”, *IEEE Journal on Selected Areas in Communications*, vol. 39, pp. 1183–1210, 2021 (<https://doi.org/10.1109/JSAC.2021.3065072>).
- [14] J.B. Kenney, “Dedicated Short-range Communications (DSRC) Standards in the United States”, *Proceedings of the IEEE*, vol. 99, pp. 1162–1182, 2011 (<https://doi.org/10.1109/JPROC.2011.2132790>).
- [15] C. Guo, X. Wang, L. Liang, and G.Y. Li, “Age of Information, Latency, and Reliability in Intelligent Vehicular Networks”, *IEEE Network*, vol. 37, pp. 109–116, 2023 (<https://doi.org/10.1109/MNET.124.2200132>).
- [16] A. Rolich, I. Turcanu, A. Vinel, and A. Baiocchi, “Understanding the Impact of Persistence and Propagation on the Age of Information of Broadcast Traffic in 5G NR-V2X Sidelink Communications”, *Computer Networks*, vol. 248, art. no. 110503, 2024 (<https://doi.org/10.1016/j.comnet.2024.110503>).
- [17] C. Zhang *et al.*, “Optimizing Age of Information in Internet of Vehicles over Error-prone Channels”, *Sensors*, vol. 24, art. no. 7888, 2024 (<https://doi.org/10.3390/s24247888>).
- [18] R. Mo, B. Wu, Z. Tan, and H. Qiu, “SEE-V2X: C-V2X Direct Communication Dataset: An Application-centric Approach”, *Proc. of the ACM Conference on Embedded Networked Sensor Systems*, pp. 305–311, 2025 (<https://doi.org/10.1145/3715014.3722077>).
- [19] E. Moradi-Pari *et al.*, “DSRC Versus LTE-V2X: Empirical Performance Analysis of Direct Vehicular Communication Technologies”, *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, pp. 4889–4903, 2023 (<https://doi.org/10.1109/TITS.2023.3247339>).
- [20] X. Yang *et al.*, “Data Synchronization for Vehicular Digital Twin Network”, *IEEE Global Communications Conference*, Rio de Janeiro, Brazil, 2022 (<https://doi.org/10.1109/GLOBECOM48099.2022.10000831>).
- [21] J. Zheng *et al.*, “Data Synchronization in Vehicular Digital Twin Network: A Game Theoretic Approach”, *IEEE Transactions on Wireless Communications*, vol. 22, pp. 7635–7647, 2023 (<https://doi.org/10.1109/TWC.2023.3254158>).
- [22] Y. Li *et al.*, “On-demand Data Synchronization for Differentiated Digital Twins in Vehicular Networks”, *IEEE/CIC International Conference on Communications in China*, Hangzhou, China, 2024 (<https://doi.org/10.1109/ICCC62479.2024.10681749>).
- [23] K. Li *et al.*, “AoI-error-aware Data Synchronization for Vehicular Digital Twin”, *GLOBECOM 2025 IEEE Global Communications Conference*, Taipei, Taiwan, 2025 (<https://doi.org/10.1109/GLOBECOM59602.2025.11432664>).
- [24] A. Loubany, M. Itani, and S. Sharafeddine, “From Age of Information to Age of Digital Twin: A Review on Synchronization Metrics for IoT Networks”, *IEEE Access*, vol. 13, pp. 131102–131119, 2025 (<https://doi.org/10.1109/ACCESS.2025.3591589>).
- [25] L.V. Kadir *et al.*, “Intelligent Digital Twin Communication Framework for Addressing Accuracy and Timeliness Tradeoff in Resource-constrained Networks”, *IEEE Transactions on Cognitive Communications and Networking*, vol. 11, pp. 1954–1965, 2025 (<https://doi.org/10.1109/TCCN.2024.3469234>).
- [26] K.-S. Kim, J.-W. Kim, and H.-S. Lee, “Prediction-Based Scheduling for Digital Twin Synchronization with AoI and Spatial Urgency”, *IEEE Transactions on Industrial Informatics*, vol. 22, pp. 9280–9291, 2026 (<https://doi.org/10.1109/TII.2026.3698880>).
- [27] M. Schneider, A. Stenger, and D. Goeke, “The Electric Vehicle-routing Problem with Time Windows and Recharging Stations”, *Transportation Science*, vol. 48, pp. 465–694, 2014 (<https://doi.org/10.1287/trsc.2013.0490>).
- [28] G. Desaulniers, F. Errico, S. Irnich, and M. Schneider, “Exact Algorithms for Electric Vehicle-routing Problems with Time Windows”, *Operations Research*, vol. 64, pp. 1388–1405, 2016 (<https://doi.org/10.1287/opre.2016.1535>).
- [29] M. Keskin and B. Çatay, “Partial Recharge Strategies for the Electric Vehicle Routing Problem with Time Windows”, *Transportation Research Part C: Emerging Technologies*, vol. 65, pp. 111–127, 2016 (<https://doi.org/10.1016/j.trc.2016.01.013>).
- [30] G. Hiermann, J. Puchinger, S. Ropke, and R.F. Hartl, “The Electric Fleet Size and Mix Vehicle Routing Problem with Time Windows and Recharging Stations”, *European Journal of Operational Research*, vol. 252, pp. 995–1018, 2016 (<https://doi.org/10.1016/j.ejor.2016.01.038>).
- [31] S. Pelletier, O. Jabali, and G. Laporte, “The Electric Vehicle Routing Problem with Energy Consumption Uncertainty”, *Transportation Research Part B: Methodological*, vol. 126, pp. 225–255, 2019 (<https://doi.org/10.1016/j.trb.2019.06.006>).
- [32] J. Jeong, B. Ghaddar, N. Zufferey, and J. Nathwani, “Adaptive Robust Electric Vehicle Routing Under Energy Consumption Uncertainty”, *Transportation Research Part C: Emerging Technologies*, vol. 160, art. no. 104529, 2024 (<https://doi.org/10.1016/j.trc.2024.104529>).
- [33] D. Wang, W. Zheng, and H. Zhou, “Robust Optimization for Electric Vehicle Routing Problem Considering Time Windows Under Energy Consumption Uncertainty”, *Applied Sciences*, vol. 15, art. no. 761, 2025 (<https://doi.org/10.3390/app15020761>).
- [34] B. Han, Y. Yan, T. Chi, and Y. Park, “The Electric Vehicle Routing Problem with Travel Time and Energy Consumption Uncertainty”,

- Transportation Research Part E: Logistics and Transportation Review*, vol. 202, art. no. 104211, 2025 (<https://doi.org/10.1016/j.tre.2025.104211>).
- [35] Y. Xie *et al.*, “A Digital Twin-based Approach for Dynamic Traffic-aware Routing and Charging of Electric Vehicles”, *Expert Systems with Applications*, vol. 312, art. no. 131380, 2026 (<https://doi.org/10.1016/j.eswa.2026.131380>).
- [36] T. Kaarlela *et al.*, “Digital Twin Framework for Electric Truck Route and Infrastructure Planning in the Nordic Region”, *Transportation Engineering*, vol. 23, art. no. 100426, 2026 (<https://doi.org/10.1016/j.treng.2026.100426>).
- [37] A. Banks and R. Gupta, “MQTT Version 3.1.1”, OASIS Standard, 2014 (<https://www.oasis-open.org/standard/mqttv3-1-1/>).
- [38] O. Semenova *et al.*, “Intelligent Secure Data Aggregation in WSNs”, *Journal of Telecommunications and Information Technology*, vol. 101, pp. 95–104, 2025 (<https://doi.org/10.26636/jtit.2025.3.2220>).
- [39] S. Das and U.K. Mondal, “Blockchain-implied Architecture for Secure and Energy Efficient Processing of IoT Data in Pervasive WSNs”, *Journal of Telecommunications and Information Technology*, vol. 102, pp. 50–60, 2025 (<https://doi.org/10.26636/jtit.2025.4.2194>).
- [40] S. Rabhi, “Aware Node Localization in Wireless Sensor Networks Using Harris Hawks Optimization”, *Journal of Telecommunications and Information Technology*, vol. 102, pp. 1–10, 2025 (<https://doi.org/10.26636/jtit.2025.4.2285>).
- [41] R. Baghiani, L. Guezouli, and A. Korichi, “TS-based Mobility-aware Multi-hierarchical Caching Model with Vehicle Clustering and Content Popularity Prediction”, *Journal of Telecommunications and Information Technology*, vol. 97, pp. 65–78, 2024 (<https://doi.org/10.26636/jtit.2024.3.1616>).
- [42] H. Wójcik, M. Roszkowski, and A. Mycek, “Network Threat Detection in IaaS Environments Based on Flow Log Data”, *Journal of Telecommunications and Information Technology*, vol. 104, pp. 33–41, 2026 (<https://doi.org/10.26636/jtit.2026.3.2628>).
- [43] M. Noroozi and M. Fidler, “Age and Deviation-of-information of Hybrid Time- and Event-triggered Systems: What Matters More, Determinism or Resource Conservation?”, *Performance Evaluation*, vol. 164, art. no. 102412, 2024 (<https://doi.org/10.1016/j.peva.2024.102412>).
- [44] J. Lei *et al.*, “Distribution-free Predictive Inference for Regression”, *Journal of the American Statistical Association*, vol. 113, pp. 1094–1111, 2018 (<https://doi.org/10.1080/01621459.2017.1307116>).
- [45] A.N. Angelopoulos and S. Bates, “Conformal Prediction: A Gentle Introduction”, *Foundations and Trends in Machine Learning*, vol. 16, pp. 494–591, 2023 (<https://doi.org/10.1561/22000000101>).
- [46] S. Ropke and D. Pisinger, “An Adaptive Large Neighborhood Search Heuristic for the Pickup and Delivery Problem with Time Windows”, *Transportation Science*, vol. 40, pp. 455–472, 2006 (<https://doi.org/10.1287/trsc.1050.0135>).
- [47] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*, Springer, 324 p., 2005 (<https://doi.org/10.1007/b106715>).
- [48] I. Gibbs and E.J. Candès, “Adaptive Conformal Inference Under Distribution Shift”, *arXiv*, 2021 (<https://doi.org/10.48550/arXiv.2106.00170>).
- [49] L. Lindemann, M. Cleaveland, G. Shim, and G.J. Pappas, “Safe Planning in Dynamic Environments Using Conformal Prediction”, *IEEE Robotics and Automation Letters*, vol. 8, pp. 5116–5123, 2023 (<https://doi.org/10.1109/LRA.2023.3292071>).
- [50] V. Pillac, M. Gendreau, C. Guéret, and A.L. Medaglia, “A Review of Dynamic Vehicle Routing Problems”, *European Journal of Operational Research*, vol. 225, pp. 1–11, 2013 (<https://doi.org/10.1016/j.ejor.2012.08.015>).
- [51] P. Álvarez López *et al.*, “Microscopic Traffic Simulation Using SUMO”, *IEEE International Conference on Intelligent Transportation Systems*, Maui, USA, 2018 (<https://doi.org/10.1109/ITSC.2018.8569938>).
- [52] S. Ye, Q. Wu, P. Fan, and Q. Fan, “A Survey on Semantic Communications in Internet of Vehicles”, *Entropy*, vol. 27, art. no. 445, 2025 (<https://doi.org/10.3390/e27040445>).
- [53] Y. Feng *et al.*, “Semantic Communication for Edge Intelligence Enabled Autonomous Driving System”, *IEEE Network*, vol. 39, pp. 149–157, 2025 (<https://doi.org/10.1109/MNET.2024.3468328>).

Abdelkader Kaddour, Ph.D.

Laboratory of Informatics and its Applications (LIAM)

 <https://orcid.org/0009-0006-9566-5559>

E-mail: abdelkader.kaddour@univ-msila.dz

University Mohamed Boudiaf of M'sila, M'sila, Algeria

<https://www.univ-msila.dz>

Lamri Sayad, Professor

Laboratory of Informatics and its Applications (LIAM)

 <https://orcid.org/0000-0002-0648-1859>

E-mail: lamri.sayad@univ-msila.dz

University Mohamed Boudiaf of M'sila, M'sila, Algeria

<https://www.univ-msila.dz>