

Goal-Oriented Communication and Control Co-Design via Semantic Push-Pull in Industrial IoT

Muhammad Azeem Khan^{*†}, Yuriy Zacchia Lun[†], Aamir Mahmood^{*}, Piergiuseppe Di Marco[†], Mikael Gidlund^{*}, and Fortunato Santucci[†]

^{*}Department of Computer and Electrical Engineering, Mid Sweden University, 851 70 Sundsvall, Sweden

[†]Department of Information Engineering, Computer Science and Mathematics, University of L'Aquila, 67100 L'Aquila, Italy

[†]Email: muhammadazeem.khan@graduate.univaq.it, yuriy.zacchialun@univaq.it, piergiuseppe.dimarco@univaq.it, fortunato.santucci@univaq.it ^{*}Email: aamir.mahmood@miun.se, mikael.gidlund@miun.se

Abstract—Emerging 6G industrial IoT architectures require wireless networked control systems capable of stabilizing diverse control loops over tightly constrained radio resources. Conventional periodic and Age-of-Information (AoI) based scheduling guarantees bounded staleness at the cost of persistent channel saturation. Conversely, pure event-triggered (PureET) strategies minimize transmissions but risk catastrophic silent deterioration when local sensor-side thresholds fail to reflect critical state evolution. To bridge this gap, we propose a communication-control co-design framework governed by a 6G Semantic Layer that independently arbitrates uplink and downlink resources. Instead of relying on freshness, our architecture evaluates the actual control impact of a packet using the state-to-error ratio (SER). We unify this control confidence with channel reliability in terms of signal-to-noise ratio (SNR) to orchestrate a threshold-based sensor push and a state-aware controller pull mechanism. To ensure equitable resource allocation across dynamically heterogeneous plants, the proposed framework explicitly scales actuation deadbands according to local plant dynamics. Simulations over Rayleigh-faded channels demonstrate that this approach fundamentally shifts the Pareto frontier between transmission rate and control quality. The proposed scheme achieves tracking accuracy comparable to periodic schedulers at a reduced communication overhead, while mitigating the estimation errors characteristic of PureET.

Index Terms—Wireless Networked Control Systems, semantic communication, goal-oriented communication, event-triggered control, Age of Information, State-to-Error Ratio, 6G.

I. INTRODUCTION

Industrial-IoT (IIoT) deployments and large-scale digital twins require stabilizing physical processes over shared wireless infrastructure [1]. Decoupling wired loops into wireless networked control systems (WNCS) forces sensors and actuators to compete for finite radio resources, challenging 6G ultra-reliable low-latency communication (uRLLC) targets [2]. In dense deployments, control loops compete for scarce radio resources against diverse traffic, making intelligent resource allocation critical.

However, this scarcity reframes the scheduling problem: a finite resource pool must stabilize heterogeneous loops without causing network collapse. Conventional throughput-centric schedulers fail for control workloads because measurements from plants near equilibrium are highly predictable and largely redundant, while updates from disturbed plants are critical to prevent closed-loop instability. Bandwidth and

information value rarely align, and schedulers blind to this mismatch waste resources while starving critical processes.

Periodic and Age of Information (AoI) approaches prioritize transmissions based on data staleness [3], [4]. While analytical and robust, freshness is only a proxy for relevance, and periodic schemes amplify this inefficiency by demanding full channel utilization to bound staleness regardless of instantaneous need. Pure event-triggered (PureET) architectures shift the decision to the sensor, transmitting only when locally observed innovations exceed a threshold [5], [6]. This enables transmission suppression during steady-state operation but introduces a critical vulnerability. If triggering conditions miss relevant state evolution due to unobserved drift, process disturbances, or packet losses, the controller misses corrective updates and estimation degradation proceeds unchecked, risking loop destabilization. PureET and periodic scheduling thus represent opposing extremes: bandwidth-efficient but exposed to silent deterioration on one side, robust but resource-intensive on the other, revealing that information utility is inherently state-dependent.

A defining shift in recent 6G cyber-physical research is the move toward goal-oriented semantic communication-control co-design. Goal-oriented communication indicates a paradigm where transmission decisions are governed strictly by their contribution to the specific control objective, such as stabilizing an unstable plant, rather than merely fulfilling network-centric requirements. Complementary to this, the semantic approach refers to a framework where the value of information is quantified by its utility in reducing closed-loop estimation error, to prioritize updates that carry the highest functional relevance to the task at hand [7]. Recent frameworks operationalize this through AoI and Value of Information metrics [8], [9], context-aware estimation [10], [11], and joint digital twin orchestration [12], [13]. Unifying abstractions like the Semantic Aggregation Layer [14] enable context-aware scheduling, yet limitations persist within these foundational models.

Existing semantic-scheduling frameworks frequently rely on assumptions that limit practical applicability. Most focus exclusively on the sensor-to-controller uplink, abstracting away controller-to-actuator downlink contention [4]. This asymmetric modeling reflects a broader divide that isolates

control confidence from channel reliability, missing a fundamental structural duality between them. To operationalize control confidence, we adopt the State-to-Error Ratio (SER), a metric used in networked control systems to quantify a controller's confidence by measuring the ratio of the estimated state magnitude to the estimation error. A trustworthy predictor (high SER) and a robust channel (high SNR) serve identical purposes: both guarantee sufficient quality to safely defer transmissions [15]. Standard event-triggered architectures also struggle with heterogeneous plants, as uniform thresholds induce fairness violations, with high-gain control loops ignoring major drifts while low-gain loops overreact to minor noise. Resolving these intertwined challenges requires a unified architecture that jointly manages duplex capacity, exploits SER-SNR duality, and normalizes deadbands to individual dynamics.

Addressing these gaps, we propose a semantic-aware communication-control co-design framework, adopting semantic communication as a paradigm where transmission value is quantified exclusively by its capacity to reduce closed-loop estimation error. The framework bridges event-triggered efficiency at equilibrium with periodic robustness under stress, operationalizing the SER-SNR uniformity principle for broader theoretical impact. Our core technical contributions are fourfold:

- Arbitrate independent uplink and downlink capacity budgets under a unified semantic prioritization policy.
- Encode control-side predictive confidence and communication-side channel reliability into a single, interpretable transmission-suppression rule.
- Neutralize plant heterogeneity within the actuation deadband to guarantee equitable evaluation across diverse physical dynamics.
- Counteract silent degradation in PureET paradigms via an AoI-driven pull-request mechanism empowering controllers to solicit updates when estimates deteriorate.

The remainder of this paper is organized as follows. Section II presents the system model. Section III defines the semantic metrics and estimation architecture. Section IV details the scheduling and control framework. Section V reports the simulation results, and Section VI concludes this study.

II. SYSTEM MODEL

We consider a WNCS where a central edge controller stabilizes N spatially distributed unstable plants (e.g., robotic arms or automated guided vehicles) over a shared 5G/6G access point. Each plant is equipped with a local sensor and actuator that communicate with a centralized edge controller over bandwidth-constrained bidirectional wireless links. Uplink and downlink transmissions are allocated over orthogonal FDMA resource blocks, with separate capacity budgets for sensor updates and actuator commands. This setup reflects industrial IoT deployments where numerous high-rate control loops compete for limited radio resources. The architecture is depicted in Fig. 1.




Fig. 1. Semantic push-pull architecture for the proposed communication-control co-design framework. The 6G Semantic Layer coordinates sensor-initiated push updates, controller-initiated pull requests, and downlink control actions over shared wireless resources.

The plants, indexed by $\mathcal{N} = \{1, 2, \dots, N\}$, evolve in discrete time with a sampling period T_s and index $k \in \{0, 1, \dots, K-1\}$. At each step, the edge controller arbitrates limited uplink (UL) and downlink (DL) channels. The UL carries sensor-initiated *push* updates (triggered by local innovation thresholds) or controller-initiated *pull* requests (issued upon estimator degradation). The DL delivers resulting control commands to the actuators.

A. Plant Dynamics

Each plant $i \in \mathcal{N}$ is modeled as an inverted pendulum on a cart, a benchmark for networked control protocols due to its open-loop instability. The state vector $\mathbf{x}_i(k) = [p_i(k) \ \dot{p}_i(k) \ \theta_i(k) \ \dot{\theta}_i(k)]^T$ captures position, velocity, pendulum angle, and angular velocity, respectively. Linearizing the continuous-time dynamics around the upright equilibrium and applying a zero-order hold at sampling period T_s yields the discrete-time evolution,

$$\mathbf{x}_i(k+1) = \mathbf{A}_i \mathbf{x}_i(k) + \mathbf{B}_i u_i(k) + \mathbf{w}_i(k), \quad (1)$$

where $\mathbf{A}_i \in \mathbb{R}^{4 \times 4}$, $\mathbf{B}_i \in \mathbb{R}^{4 \times 1}$ are the state-transition and input matrices, $u_i(k) \in \mathbb{R}$ is the horizontal control force, and $\mathbf{w}_i(k) \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_w)$ models process disturbances with covariance $\mathbf{Q}_w \in \mathbb{R}^{4 \times 4}$. Sensors observe the complete state via an additive Gaussian model,

$$\mathbf{y}_i(k) = \mathbf{C} \mathbf{x}_i(k) + \mathbf{v}_i(k), \quad \mathbf{v}_i(k) \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_v), \quad (2)$$

with $\mathbf{C} = \mathbf{I}_4$ and measurement noise covariance $\mathbf{R}_v$. The controller employs a discrete-time Linear Quadratic Regulator (LQR) law $\mathbf{K}_i = -\text{dlqr}(\mathbf{A}_i, \mathbf{B}_i, \mathbf{Q}, \mathbf{R})$, where $\mathbf{Q}$ and $\mathbf{R}$ are the weighting matrices penalizing state deviations and control effort, respectively, prioritizing angular stability. LQR is adopted for analytical tractability and clean separation of control and communication layers.

B. Shared Wireless Network

At each step k , the scheduler grants UL/DL access to at most $M_{\text{UL}}, M_{\text{DL}} \ll N$ plants, respectively. Letting $\mathcal{G}^{\text{UL}}(k), \mathcal{G}^{\text{DL}}(k) \subseteq \mathcal{N}$ denote the scheduled sets,

$$|\mathcal{G}^{\text{UL}}(k)| \leq M_{\text{UL}}, \quad |\mathcal{G}^{\text{DL}}(k)| \leq M_{\text{DL}}. \quad (3)$$

Channel impairment is isolated to the UL, reflecting power-constrained sensor nodes in IIoT. The UL is modeled as a Rayleigh block-fading channel with outage-based reception. The instantaneous received SNR is $\gamma_i(k) = \bar{\gamma} |h_i(k)|^2$, yielding the closed-form packet error rate $P_{\text{out}} = \Pr[\gamma < \gamma_{\text{th}}] = 1 - \exp(-\gamma_{\text{th}}/\bar{\gamma})$, where γ_{th} is the minimum decoding threshold. With $S_i(k) \in \{0, 1\}$ denoting the success indicator, the effective reception sets are

$$\mathcal{R}^{\text{UL}}(k) = \{i \in \mathcal{G}^{\text{UL}}(k) : S_i(k) = 1\}, \quad (4)$$

while the downlink is treated as capacity-limited but error-free,

$$\mathcal{R}^{\text{DL}}(k) = \mathcal{G}^{\text{DL}}(k). \quad (5)$$

This asymmetric assumption reflects cell-edge IIoT realities, where the base station enjoys higher transmit power and richer antenna provisioning than the sensor nodes. In our framework, the sensor only checks local innovation, while the controller or 6G Semantic Layer (6G SL) uses channel-quality information for scheduling.

III. SEMANTIC METRICS AND ESTIMATION ARCHITECTURE

Translating raw physical signals into semantic scheduling decisions requires tracking data freshness and predictive confidence. This section defines the core metrics and estimation architecture supporting our framework.

A. Information Freshness

Freshness is tracked via the AoI abstraction. Two separate AoI counters are maintained per plant

$$\Delta_i^c(k+1) = \begin{cases} 1, & i \in \mathcal{R}^{\text{UL}}(k), \\ \Delta_i^c(k) + 1, & \text{otherwise}, \end{cases} \quad (6)$$

$$\Delta_i^a(k+1) = \begin{cases} 1, & i \in \mathcal{R}^{\text{DL}}(k), \\ \Delta_i^a(k) + 1, & \text{otherwise}. \end{cases} \quad (7)$$

The controller-side AoI $\Delta_i^c(k)$ quantifies the time since the controller last received a fresh measurement, while the actuator-side AoI $\Delta_i^a(k)$ measures actuator staleness. The proposed scheduler exploits this two-sided view to gate pull requests and downlink priority boosts.

B. Controller-Side Predictor

The edge controller maintains a two-stage update/propagation predictor for each plant. With $\hat{\mathbf{x}}_i(k|k)$ the filtered estimate at step k (using data up to step k) and

$\hat{\mathbf{x}}_i(k|k-1)$ the predicted estimate at step k (using data only up to $k-1$),

$$\hat{\mathbf{x}}_i(k|k) = \begin{cases} \mathbf{y}_i(k), & i \in \mathcal{R}^{\text{UL}}(k), \\ \hat{\mathbf{x}}_i(k|k-1), & \text{otherwise}, \end{cases} \quad (8a)$$

$$\hat{\mathbf{x}}_i(k+1|k) = \mathbf{A}_i \hat{\mathbf{x}}_i(k|k) + \mathbf{B}_i u_i(k). \quad (8b)$$

Eq. (8a) updates the estimate upon a successful uplink or retains the prior prediction otherwise, while (8b) propagates it through the plant dynamics. As $\Delta_i^c(k)$ increases, accumulated noise and linearization error degrade $\hat{\mathbf{x}}_i$, a drift that triggers controller pull requests.

C. State-to-Error Ratio: A Semantic Confidence Metric

When a fresh measurement $\mathbf{y}_i(k)$ arrives, comparing it against the pre-update prediction $\hat{\mathbf{x}}_i^-(k) := \mathbf{A}_i \hat{\mathbf{x}}_i(k-1) + \mathbf{B}_i u_i(k-1)$ yields a local indicator of predictor fidelity. We define the instantaneous State-to-Error Ratio (SER) as

$$\text{SER}_i(k) = \frac{\|\hat{\mathbf{x}}_i^-(k)\|^2 + \epsilon}{\|\mathbf{y}_i(k) - \hat{\mathbf{x}}_i^-(k)\|^2 + \epsilon}, \quad (9)$$

with $\epsilon > 0$ a small regularizer ensuring the ratio remains well-defined near equilibrium. A large SER_i signals that the innovation $\|\mathbf{y}_i - \hat{\mathbf{x}}_i^-\|$ is small relative to the estimated state, so the predictor can be trusted, while a small value indicates open-loop extrapolation is unsafe. Since no innovation is available between measurements, we propagate the last computed SER_i with a monotone decay tied to sensor AoI as

$$\text{SER}_i^{\text{eff}}(k) = \min\left(\text{SER}_{\text{max}}, \frac{\text{SER}_i}{\min(\Delta_i^c(k), \Delta_{\text{decay}})}\right). \quad (10)$$

The saturation floor Δ_{decay} prevents unbounded decay during prolonged outages, and SER_{max} caps the ratio near equilibrium. The effective SER in (10) serves as the control-side semantic confidence signal.

D. Performance Objective

Our design goal is to stabilize all plants while minimizing the system-wide number of transmissions. The *control cost*, expressed as the time-averaged root-mean-square (RMS) angular deviation

$$J_{\theta}^{\text{RMS}} = \sqrt{\frac{1}{KN} \sum_{k=0}^{K-1} \sum_{i=1}^N \theta_i^2(k)}, \quad (11)$$

penalizes poor stabilization. The *communication cost*, defined as the total cumulative count of successful uplink transmissions across the network over the time horizon K ,

$$J_{\text{tx}} = \sum_{k=0}^{K-1} \sum_{i=1}^N \mathbb{1}\{i \in \mathcal{R}^{\text{UL}}(k)\}, \quad (12)$$

penalizes wasteful bandwidth use. A preferable scheduler minimizes both quantities simultaneously.

IV. PROPOSED FRAMEWORK FOR SEMANTIC-AWARE SCHEDULING AND CONTROL

Our proposed framework builds on a single organizing principle: a channel slot is worth spending only when the transmission's content meaningfully alters the closed-loop outcome. We implement this via decentralized event-triggering, 6G SL arbitration, and edge controller logic.

A. Decentralized Semantic Event-Triggering at the Sensor

Each sensor i maintains a local memory of its last successfully transmitted measurement $(\mathbf{y}_i^{\text{last}}, \theta_i^{\text{last}})$. It evaluates a local innovation score $\eta_i^s(k) = |\theta_i(k) - \theta_i^{\text{last}}|$, and generates an event-driven measurement data unit (MDU) only when $\eta_i^s(k) > \delta_i^s$. The threshold δ_i^s is tailored to plant dynamics. When the trigger fires, the sensor submits a *push* MDU carrying the current measurement, its generation age $\Delta_i^c(k)$, and a priority

$$\rho_i^{\text{push}}(k) = \frac{\eta_i^s(k)}{\delta_i^s}, \quad (13)$$

which normalizes urgency across plant types. Silent sensors contribute no traffic. The locality is important: the sensor never needs to know the global plant population, channel state, or scheduler queue, which makes the trigger implementable at the extreme edge of the network.

B. 6G Semantic Layer and Duplex Arbitration

We propose the 6G SL between the application logic and radio resource scheduler as the AI-native semantic arbitration function envisioned for 6G RAN architectures [2], rather than a generic traffic scheduler retrofitted with priority tags. Every communication event carries semantic metadata (type, priority, age-at-generation) that the scheduler can act upon, and the layer plays three roles:

- *Event aggregation*: merging sensor-side push events and controller-side pull requests into a single priority queue.
- *Semantic prioritization*: ranking queue entries by a composite score that combines AoI, innovation magnitude $\|\mathbf{y}_i(k) - \hat{\mathbf{x}}_i(k)\|$ (the discrepancy between a fresh measurement and the predictor's estimate, capturing how much new information the update carries), and downlink actuator staleness.
- *Duplex arbitration*: granting the top M_{UL} uplink entries and the top M_{DL} pending control updates independently within their capacity budgets.

The uplink queue $\mathcal{Q}^{\text{UL}}(k)$ aggregates two sources: sensor-originated push MDUs that cleared their local triggers at step k , and controller-initiated pull MDUs corresponding to plants whose pull flag $\pi_i(k)$ has been raised by the *Condition C* in Sec. IV-C. Each pull inherits a priority that scales with controller-side AoI as

$$\rho_i^{\text{pull}}(k) = \frac{\Delta_i^c(k)}{\max(\tau^c, 1)}, \quad (14)$$

escalating monotonically as the estimate ages. The DL queue $\mathcal{Q}^{\text{DL}}(k)$ collects control actions carrying a base priority $\rho_i^{\text{base}}(k)$ proportional to the planned change magnitude, with

a staleness-aware boost keeping long-frozen actuators competitive even when their raw priority is modest

$$\rho_i^{\text{DL}}(k) = \rho_i^{\text{base}}(k) + \beta \mathbb{1}\{\Delta_i^a(k) \geq \tau^a\} \frac{\Delta_i^a(k)}{\tau^a}, \quad (15)$$

with $\beta > 0$ controlling the boost strength.

Both queues are sorted in descending priority and deduplicated, ensuring that no plant occupies more than one slot per direction in a given step. Letting $\tilde{\mathcal{Q}}^{\text{UL}}(k)$ and $\tilde{\mathcal{Q}}^{\text{DL}}(k)$ denote the resulting sorted lists, the grant sets follow as

$$\begin{aligned} \mathcal{G}^{\text{UL}}(k) &= \text{TopM}(\tilde{\mathcal{Q}}^{\text{UL}}(k), M_{\text{UL}}), \\ \mathcal{G}^{\text{DL}}(k) &= \text{TopM}(\tilde{\mathcal{Q}}^{\text{DL}}(k), M_{\text{DL}}), \end{aligned} \quad (16)$$

where $\text{TopM}(\cdot, m)$ extracts the first m elements. When fewer than M_{UL} uplink entries are queued, the remaining slots are deliberately left *idle*. The 6G SL never fabricates traffic to fill the channel, operationalizing the principle that silence itself carries information. Under the uplink-only fading model, downlink winners receive their commands reliably; capacity truncation in (16) is the sole source of downlink failure. Plants outside $\mathcal{G}^{\text{DL}}(k)$ retain their previous input.

C. Edge Controller Logic

The edge controller executes once per slot on a per-plant basis, partitioning every (plant, slot) pair into one of three mutually exclusive branches (formalized in Algorithm 1):

Algorithm 1 Per-slot edge controller logic for plant i

Require: $\mathcal{R}^{\text{UL}}(k)$, $\mathbf{y}_i(k)$, $\hat{\mathbf{x}}_i(k)$, $u_i(k)$, $\Delta_i^c(k)$, $\text{SER}_i^{\text{eff}}(k)$, $\gamma_i^{\text{UL}}(k)$

- 1: **Constants:** SER_{th} , γ_{th} , λ , τ^c , $\Delta\theta_{\text{base}}$
- 2: **if** $i \in \mathcal{R}^{\text{UL}}(k)$ **then**
- 3: **Condition A: fresh measurement**
- 4: $\hat{\mathbf{x}}_i(k) \leftarrow \mathbf{y}_i(k)$
- 5: $u_i^*(k) \leftarrow \mathbf{K}_i \hat{\mathbf{x}}_i(k)$
- 6: $\rho_i^{\text{base}}(k) \leftarrow |u_i^*(k) - u_i(k)| + 1.0$
- 7: Queue control action with priority $\rho_i^{\text{base}}(k)$
- 8: $\pi_i(k) \leftarrow 0$
- 9: **else**
- 10: $U_i(k) \leftarrow \frac{\text{SER}_i^{\text{eff}}(k)}{\text{SER}_{\text{th}}} + \lambda \frac{\gamma_{\text{th}}}{\gamma_i^{\text{UL}}(k)}$
- 11: **if** $U_i(k) \geq 1$ **then**
- 12: **Condition B: silence**
- 13: no action queued
- 14: **else**
- 15: **Condition C: stale estimate**
- 16: $u_i^*(k) \leftarrow \mathbf{K}_i \hat{\mathbf{x}}_i(k)$
- 17: $\Delta_i^u \leftarrow |K_{i,3}| \cdot \Delta\theta_{\text{base}}$
- 18: **if** $|u_i^*(k) - u_i(k)| > \Delta_i^u$ **then**
- 19: Queue control action with priority $|u_i^*(k) - u_i(k)|$
- 20: **end if**
- 21: $\pi_i(k) \leftarrow \mathbb{1}\{\Delta_i^c(k) \geq \tau^c\}$
- 22: **end if**
- 23: **end if**

- *Condition A (Fresh Measurement)*: When $i \in \mathcal{R}^{\text{UL}}(k)$, the predictor synchronizes, generating a fresh command queued with high priority, clearing the pull flag.
- *Condition B (Channel-Aware Silence)*: The controller evaluates a semantic uniformity score combining predictive confidence (SER) and channel quality (γ^{UL}). This uniformity score integrates SER and SNR as complementary indicators governing transmission suppression. A high SER directly enforces silence by confirming predictor reliability, rendering fresh updates semantically unnecessary. The SNR term functions as a conditional penalty: a degraded link inflates the score to block likely-to-fail transmissions, whereas a strong channel neutralizes this penalty. A high SNR does not inherently force a transmission; it merely allows the scheduler to update the plant if the control side requires correction. Ultimately, these metrics suppress low-yield traffic across two distinct axes: SER evaluates whether new information is *needed*, while SNR dictates whether it can be reliably *delivered*.
- *Condition C (Stale Estimate)*: Operating on open-loop extrapolation, the controller tests candidate commands against an angle-equivalent deadband $\Delta_i^u = |K_{i,3}| \cdot \Delta\theta_{\text{base}}$. Negligible updates are discarded. If controller-side AoI exceeds threshold τ^c , the pull flag π_i is raised to solicit fresh data.

V. RESULTS AND DISCUSSION

We evaluate the proposed framework against four baselines on a heterogeneous population of $N = 30$ inverted pendulums sharing a duplex-constrained 6G uplink under Rayleigh block fading. To capture diverse control requirements, the plants are evenly cycled across three distinct instability classes: highly unstable ($L = 0.2$ m), moderately unstable ($L = 0.5$ m), and slightly unstable ($L = 1.0$ m) dynamics. Sharing $M_{\text{UL}} = 10$ resource blocks among these 30 plants establishes a strict 3 : 1 contention ratio, explicitly chosen to stress the network and represent a congested industrial workcell. The proposed architecture scales efficiently to larger populations due to its asymmetric design: sensor-side complexity remains strictly $\mathcal{O}(1)$, offloading the $\mathcal{O}(N)$ uniformity evaluations to the resource-rich edge controller. Furthermore, as N grows, the SER-SNR rule inherently suppresses excess traffic, gracefully degrading individual update rates to prevent total channel collapse under extreme plant densities. The evaluated baselines (sharing identical physical and channel parameters) are configured as follows:

- *PureET*: transmits only when the local angular innovation crosses the per-plant threshold δ_i^s , with no controller-initiated pulls and no SER-SNR uniformity test.
- *Round-Robin*: cyclically grants the M_{UL} uplink slots to plants in fixed order, ignoring plant state and channel quality.
- *Max-AoI*: greedily allocates each slot to the M_{UL} plants with the largest controller-side AoI $\Delta_i^c(k)$, prioritizing freshness without semantic context.

TABLE I
SIMULATION PARAMETERS AND PER-METHOD PERFORMANCE.
EVALUATED METRICS: $u_{\text{Energy}} = \sum u^2 T_s$ THE CUMULATIVE ACTUATION ENERGY; $J_{\text{cc}} = (J_{\text{tx}}^{\text{UL}} + J_{\text{tx}}^{\text{DL}}) \cdot (\theta_{\text{RMS}})^2$ THE JOINT COMMUNICATION-CONTROL COST; JAIN = PER-SEGMENT JAIN FAIRNESS INDEX OVER 1 S WINDOWS.

Configuration: $N=30$, $T_s=10$ ms, $T_{\text{sim}}=60$ s, $\lambda=1.0$, $\text{SER}_{\text{th}}=50$, $\gamma_{\text{th}}=0$ dB, $M_{\text{UL}}=M_{\text{DL}}=10$, $P_{\text{out}} \in [0.05, 0.50]$, $\bar{\gamma}=9.77$ dB, $\tau^c=\tau^a=10$, $\beta=10$, $\Delta\theta_{\text{base}}=0.3^\circ$				
Method	θ_{RMS} (deg)	θ_{max} (deg)	u_{Energy} (N ² s)	J_{cc}
Proposed (6G SL)	0.80	6.02	484	2969
PureET	1.95	9.56	28867	43728
Round-Robin	0.49	6.62	188	10728
Max-AoI	0.47	5.93	150	9881
Random	0.59	14.67	816	17280

Method	UL Util. (%)	UL Tx (pkts)	DL Tx (pkts)	Jain (1 s)
Proposed (6G SL)	7.7	4610	21369	0.967
PureET	19.1	11443	11443	0.982
Round-Robin	75.7	45445	45445	1.000
Max-AoI	75.7	45445	45445	1.000
Random	75.7	45432	45432	0.999




Fig. 2. Empirical Pareto frontier in the (transmission rate, angular tracking error) plane. Operating points in the lower-left dominate.

- *Random*: selects M_{UL} plants uniformly at random per slot, serving as a stateless reference.

Table I consolidates the simulation configuration and the per-method results.

Fig. 2 maps each scheduler onto the bandwidth-versus-quality plane. Round-Robin and Max-AoI cluster in the high-bandwidth region at near-saturated channel use (75% UL utilization, ~ 757 packets/s on the UL) and deliver the lowest θ_{RMS} values (0.49° , 0.47°), but only because every plant is being served whether or not it needs attention. PureET pushes toward opposite extreme, transmitting at ~ 191 packets/s but paying a substantial control penalty ($\theta_{\text{RMS}}=1.95^\circ$, peak 9.56°). Random collapses entirely under the heterogeneous-plant configuration, with peak excursions exceeding 14.67° on the worst-served loops, confirming that stateless allocation cannot meet the demands of unstable dynamics. The proposed scheme uniquely occupies the lower-left region ($\theta_{\text{RMS}}=0.80^\circ$ at ~ 77 packets/s and 7.7% UL utilization), cutting UL traffic by $9.9\times$ over periodic baselines while staying within $\sim 0.33^\circ$ of their tracking accuracy.

Fig. 3 consolidates the comparison across four normalized cost dimensions. The proposed scheme is the only method




Fig. 3. Normalized strategy comparison across four cost dimensions: Angular RMS Error, Actuation Switching Energy, Control Transmission Rate, and Joint Communication-Control Cost. Bars are normalized within each metric to the worst-performing method.

whose bars stay below ~ 0.7 on every axis simultaneously. PureET wins on the transmission-rate axis but its switching energy u_{Energy} exceeds the proposed scheme's by $60\times$, exposing how scarce transmissions translate into large corrective inputs after silent drift. The proposed scheme's actuation energy reflects the per-plant gain-normalized deadband in Table I: aggressively-gained plants are not over-actuated on noise, and softly-gained plants do not require large remedial inputs because they are correctly tracked. The joint communication-control cost places the proposed scheme roughly $14.7\times$ below PureET, $5.8\times$ below Random, and $3.3\text{--}3.6\times$ below the periodic baselines, despite using a fraction of their bandwidth.

Condition C accounts for most transmission suppressions, showing that the framework primarily silences transmissions when control utility of a fresh update is negligible (falling within actuation deadband Δ_i^u), rather than due to poor link quality. This behavior matches the evaluated high-SNR regime ($\bar{\gamma}=9.77$ dB), where channel penalties only dictate silence during rare outages. Lower SNRs naturally shift suppression toward Condition B. Thus, the uniformity rule dynamically restricts traffic according to the active bottleneck, balancing control confidence and channel reliability. Concentrating bandwidth on active transients slightly lowers Jain's fairness index to 0.967 (vs. 1.0 for baselines) without starving plants. Through Pareto dominance, the framework absorbs stochastic channel impairments into its transmission-silence decisions, preventing closed-loop control failures.

VI. CONCLUSION

We presented a semantic-aware communication-control co-design framework for 6G WNCs, governed by a 6G Semantic Layer that arbitrates uplink and downlink resources independently, encodes SER and SNR into a unified silence rule, neu-

tralizes plant heterogeneity through a gain-normalized actuation deadband, and closes the silent-deterioration loop of pure event-triggering via an AoI-driven controller-initiated pull. Across a heterogeneous population of inverted pendulums under Rayleigh-faded uplinks, the framework approaches the tracking accuracy of periodic schedulers at a fraction of their transmission rate while eliminating the divergent estimation excursions characteristic of PureET. Future work will extend the framework along three directions: formal stability and convergence guarantees under stochastic outages, downlink fading with PHY-layer link adaptation, and a learning-assisted Model Predictive Control layer that integrates cleanly above the existing communication substrate without altering the scheduling logic.

ACKNOWLEDGMENTS

This work is supported by the Swedish Knowledge Foundation via the IRS TransTech, Mid Sweden University, and by the Center of Excellence Ex-Emerge at the University of L'Aquila, funded by the Italian Government under CIPE Resolution 70/2017.

REFERENCES

- [1] P. Park *et al.*, "Wireless network design for control systems: A survey," *IEEE Commun. Surv. Tutor.*, vol. 20, no. 2, pp. 978–1013, 2017.
- [2] W. Saad *et al.*, "A vision of 6G wireless systems: Applications, trends, technologies, and open research problems," *IEEE Netw.*, vol. 34, no. 3, pp. 134–142, 2020.
- [3] R. D. Yates *et al.*, "Age of information: An introduction and survey," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, pp. 1183–1210, 2021.
- [4] O. Ayan *et al.*, "Age-of-information vs. value-of-information scheduling for cellular networked control systems," in *ACM/IEEE ICCPS*, 2019, pp. 109–117.
- [5] S. Trimpe *et al.*, "Event-based state estimation with variance-based triggering," *IEEE Trans. Automat. Contr.*, vol. 59, no. 12, pp. 3266–3281, 2014.
- [6] X. Wang *et al.*, "Event-triggering in distributed networked control systems," *IEEE Trans. Automat. Contr.*, vol. 56, no. 3, pp. 586–601, 2010.
- [7] D. Gündüz *et al.*, "Beyond transmitting bits: Context, semantics, and task-oriented communications," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 1, pp. 5–41, 2022.
- [8] W. Wu *et al.*, "Goal-oriented semantic communications for robotic waypoint transmission: The value and age of information approach," *IEEE Trans. Wireless Commun.*, vol. 23, no. 12, pp. 18903–18915, 2024.
- [9] Y. Wang *et al.*, "Goal-oriented transmission scheduling for energy-efficient wireless networked control in SAGIN: An aoi-thresholding mechanism," *IEEE Trans. Veh. Technol.*, vol. 73, no. 8, pp. 11503–11517, 2024.
- [10] J. Cao *et al.*, "Goal-oriented communication, estimation, and control over bidirectional wireless links," *IEEE Trans. Commun.*, vol. 73, no. 5, pp. 3031–3045, 2025.
- [11] P. Agheli *et al.*, "Integrated push-and-pull update model for goal-oriented effective communication," *IEEE Trans. Commun.*, vol. 73, no. 11, pp. 10914–10928, 2025.
- [12] P. Jia *et al.*, "Goal-oriented digital twin for operational loss minimization in 6G-enabled industrial systems: A joint sensing and control approach," *IEEE Trans. Ind. Informat.*, vol. 22, no. 3, pp. 2505–2516, 2026.
- [13] X. Fang *et al.*, "Sensing-communication-computing-control closed-loop optimization for 6G digital twin-empowered robotic systems," *IEEE J. Sel. Areas Commun.*, vol. 43, no. 10, pp. 3330–3346, 2025.
- [14] O. Ayan *et al.*, "Enabling communication and control co-design in 6G networks," in *IEEE CSCN*. IEEE, 2024, pp. 8–13.
- [15] Q. Li *et al.*, "A novel communication and control co-design method for wireless control systems: A communication perspective," *IEEE Internet Things J.*, vol. 12, no. 16, pp. 33858–33869, 2025.