Uplink Decoupling for Industrial IoT in 5G Networks: Performance Evaluation and Deployment Considerations

1st Charalampos Chatzigeorgiou
Computer Engineering and Informatics Dept.
University of Patras
Patras, Greece
cchatzigeorgiou@ceid.upatras.gr

2nd Christos Bouras
Computer Engineering and Informatics Dept.
University of Patras
Patras, Greece
bouras@upatras.gr

3rd Vasileios Kokkinos
Computer Engineering and Informatics Dept.
University of Patras
Patras, Greece
kokkinos@upatras.gr

4th Apostolos Gkamas
Dept. of Chemistry
University of Ioannina, Greece
Ioannina, Greece
gkamas@uoi.gr,

5th Philippos Pouyioutas
Computer Science Dept.
University of Nicosia
Nicosia, Cyprus
pouyioutas.p@unic.ac.cy

***Abstract*—The increasing density of Industrial Internet of Things (IIoT) devices in 5G networks introduces severe uplink congestion and fairness challenges. Traditional coupled access—where a user equipment (UE) uses the same base station for both downlink and uplink—can lead to suboptimal performance, particularly in smart factories characterized by heterogeneous traffic patterns and dynamic interference. This paper investigates uplink decoupling, a lightweight association mechanism that allows independent selection of the uplink serving cell to optimize link quality and resource utilization. We develop a MATLAB-based simulation that models a 5G smart factory and compare the proposed decoupled uplink approach against conventional proportional-fair (PF) scheduling. Results show consistent improvements in throughput , 95th-percentile latency , and fairness across multiple traffic classes. These findings demonstrate that uplink decoupling can effectively mitigate congestion, improve reliability, and enhance service differentiation in Industry 4.0 networks.**

***Index Terms*—5G networks, Industrial IoT (IIoT), Uplink decoupling, Resource allocation, Proportional-fair scheduling, MEC, Latency optimization, Fairness.**

I. Introduction

In current 5G deployments, coupled access binds each device to a single base station for both downlink and uplink communication. However, the propagation and interference characteristics of the two directions are often asymmetric—especially in dense factory environments. Downlink-based attachment may therefore result in suboptimal uplink channel conditions, particularly for devices located near cell edges. This asymmetry motivates the exploration of Downlink and Uplink Decoupling (DUDe), where uplink associations are dynamically determined according to instantaneous channel quality and load conditions.

Research Questions: To bridge this gap, we pose the following questions:

- RQ1: How does DUDe affect uplink fairness across low-priority (sensors) and high-priority (robots, cameras) devices? *Justification: Ensuring equitable access prevents delays in critical control loops.*
- RQ2: What uplink throughput and latency improvements can DUDe achieve compared to proportional-fair and round-robin baselines?
  Justification: Quantifying performance gains validates DUDe's practical benefits.
- RQ3: What hardware or compatibility challenges arise when deploying DUDe in real industrial 5G networks, and how can they be mitigated? *Justification: Identifying real-world barriers informs deployment strategies.*

The rest of this paper is organized as follows: This section (I) motivates DUDe, highlights the contribution of our work and reviews related work, Section II details the DUDe mechanism and Section III the simulation road-map. Section IV discusses implementation issues, Section V presents performance results, and Section VI concludes with future directions.

Smart factories host devices with diverse traffic patterns: environmental sensors (low-bandwidth, periodic), robotic controllers (medium-bandwidth, real-time), and surveillance cameras (high-bandwidth, bursty). Static schedulers cannot adapt in real time, leading to congestion under peak loads and unfair resource sharing. DUDe—by decoupling uplink and downlink associations—enables dynamic resource reallocation, promising improved spectral efficiency and fairness.

This study aims to enhance existing DUDe mechanisms by tailoring them to the unique requirements of industrial IoT environments, where latency, fairness, and scalability are critical. Unlike earlier DUDe implementations [2], [3], which focus primarily on macro–small cell load balancing, our approach explicitly integrates IIoT-specific traffic heterogeneity, MEC-assisted decision logic, and per-device-class performance evaluation. These additions extend DUDe from a theoretical concept to a practical tool for Industry 4.0 networks.

In the context of industrial IoT scheduling, various works have examined uplink/downlink decoupling and related optimizations. Dui et al. [1] design an IoT-enabled real-time traffic monitoring and control management system for intelligent transportation (ITS) applications. Bouras et al.have studied DUDe-based resource allocation in ultra-modern telecommunication systems [2] and advanced communication networking scenarios [3]. Shi et al. [4] provide a comprehensive survey on decoupled access, highlighting advances, challenges, and open problems in uplink optimization. Ozturk et al. [5] propose context-aware wireless connectivity and processing unit optimization for IoT networks. Dwivedi [6] analyzes satellite communications for IoT, discussing topology, transmission schemes, and performance.

Chen et al. [7] present a general industrial intelligence framework applicable to IIoT. Ni et al. [8] focus on efficient and secure service-oriented authentication supporting network slicing in 5G-enabled IoT. Gbadamosi et al. [9] outline a roadmap from NB-IoT towards 5G NR for IoT deployments. Cao et al. [10] propose intelligent slicing of the radio resource control (RRC) layer for cellular IoT. Ogbodo et al [11] survey the integration of 5G and LPWAN-IoT technologies for smart cities. Zarini et al. [12] study resource management for multiplexing eMBB and URLLC services over RIS-aided THz communication links. Finally, Chen et al. [13] review the evolution of 5G-Advanced toward 6G, underscoring the ongoing relevance of DUDe. These contributions motivate our focus on heterogeneous traffic and dynamic uplink optimization via DUDe in Industry 4.0 smart factories.

Contributions. This paper makes the following contributions:

- We formulate a *per-TTI uplink-centric DUDe algorithm* for IIoT, combining SINR-based re-anchoring with explicit *hysteresis* and *time-to-trigger (TTT)* to avoid ping–pong and enable stable operation under dense factory deployments (Sec. II).
- We integrate DUDe with *proportional-fair (PF)* schedulers per cell and provide all control knobs (hysteresis, TTT, PF averaging step α) and numerology parameters for *reproducibility* (Secs. II–III).
- We design a *simulation testbed* tailored to Industry 4.0: mixed device classes (sensors, robots, cameras), realistic traffic profiles, and NR numerology/TDD settings (Sec. III), and evaluate throughput, latency (including 95$^{\text{th}}$ percentile), and fairness (Jain's J) (Sec. V).
- We present *spatial SINR heatmaps* showing cell-edge improvements, *device-class breakdowns* (latency/throughput), and an *ablation study* that isolates the gain of decoupling from the choice of scheduler (RR vs. PF) (Sec. V).
- We report consistent gains at high load: *throughput* $+16.8\%$, *95th% latency* -25.0%, and *fairness* $+9.6\%$ compared to a PF baseline (Sec. V), and discuss practical integration (MEC placement, signaling) without assuming specialized hardware (Sec. III).

II. DUDE MECHANISM

Conventional 5G networks typically couple the downlink and uplink associations of each user equipment (UE) to a single base station (gNB). While this approach simplifies control signaling, it can lead to **inefficient uplink performance**, especially in dense industrial deployments where interference and load distributions differ between directions.

The DUDe paradigm breaks this coupling: each UE maintains its downlink anchor but dynamically selects a potentially different **uplink serving cell** according to instantaneous radio conditions. This independent uplink association allows the network to balance load and mitigate cell-edge degradation without modifying the overall 5G architecture.

In our implementation, the DUDe decision logic runs on the **Multi-access Edge Computing (MEC)** platform, which collects SINR reports from devices and neighboring cells. Each UE evaluates uplink channel quality at every transmission time interval (TTI), applies stability controls, and anchors its uplink transmission to the gNB offering the best instantaneous SINR.

Fig. 1 illustrates the adopted DUDe architecture. A heterogeneous 5G network composed of one macrocell and several small cells serves a mix of IIoT devices—periodic sensors, low-latency robots, and high-bandwidth cameras.

Each device periodically measures uplink SINR values toward candidate cells, transmits them to the MEC host, and receives its current uplink-anchor decision. Proportional-Fair (PF) scheduling then proceeds independently at each gNB, using local SINR and throughput information to allocate physical resource blocks (PRBs). This separation of decisions (association vs. scheduling) is central to DUDe's adaptability.

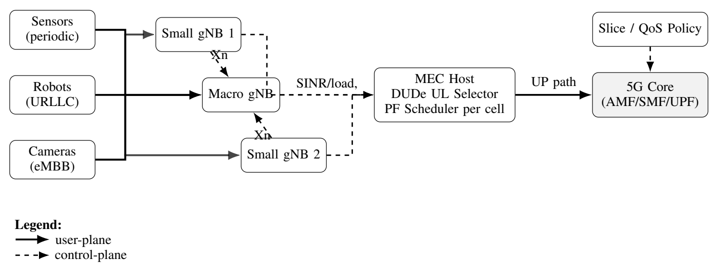

Fig. 1: DUDe architecture Example.

A. *DUDe Uplink Selection Algorithm*

The DUDe (Downlink and Uplink Decoupling) algorithm is a dynamic, per-TTI uplink association procedure tailored for heterogeneous IIoT traffic. Unlike conventional attachment based on downlink RSRP/RSRQ, DUDe lets each UE independently select its uplink anchor (serving gNB) using uplink-specific radio conditions, then schedules PRBs with a proportional-fair (PF) policy.

*1) **Uplink SINR measurement**:* At the beginning of TTI t, each UE d estimates the uplink SINR to candidate cells c:

$$\mathrm{SINR}_{d,c}(t) = \frac{P_d^{\mathrm{tx}}(t)\, G_{d,c}(t)}{N_0 B_{\mathrm{PRB}} + \sum_{c' \neq c} \sum_{u \in \mathcal{U}_{c'}(t)} P_u^{\mathrm{tx}}(t)\, G_{u,c}(t)}, \tag{1}$$

where $P_d^{\mathrm{tx}}(t)$ is the UE transmit power, $G_{d,c}(t)$ the channel gain (path loss + shadowing), N_0 noise PSD, B_{PRB} PRB bandwidth, and $\mathcal{U}_{c'}(t)$ the set of uplink UEs in cell c' at time t. SINRs are refreshed each TTI from SRS-based estimates.

*2) **Uplink anchor selection with hysteresis**:* The UE selects the anchor cell that maximizes instantaneous SINR, but to prevent ping-pong we apply a hysteresis and time-to-trigger (TTT):

$$\mathrm{SINR}_{d,c^\star}(t) - \mathrm{SINR}_{d,c}(t) \geq \Delta_{\mathrm{hyst}} \quad \text{for at least } T_{\mathrm{TTT}} \text{ TTIs.} \tag{2}$$

In our configuration we use $\Delta_{\mathrm{hyst}} \in [1, 2]$ dB and $T_{\mathrm{TTT}} \in [2, 5]$ TTIs.

*3) **PF scheduling at the selected gNB**:* Each gNB g collects uplink requests from its attached UEs and schedules PRBs using the PF metric

$$\mathrm{PF}_d(t) = \frac{R_d(t)}{\mu_d(t)}, \tag{3}$$

where $R_d(t)$ is the instantaneous achievable rate of UE d, and $\mu_d(t)$ the exponentially averaged throughput:

$$\mu_d(t+1) = (1-\alpha)\,\mu_d(t) + \alpha\, R_d(t)\,, \qquad \alpha \in (0,1]. \quad (4)$$

Smaller α favors stability; we use $\alpha \in [0.01, 0.1]$.

*4) **Iteration per TTI**:* Steps (1)–(3) run every 1 ms. This tight loop adapts associations and PRB allocations to fast channel/load changes, improves edge-UE reliability, and balances load across macro/small cells.

Control knobs: $\{\Delta_{\text{hyst}}, T_{\text{TTT}}, \alpha\}$ let operators trade off reactivity vs. stability and fairness vs. peak throughput.

III. Simulation Environment

To evaluate the proposed DUDe mechanism under realistic Industry 4.0 conditions, we developed a MATLAB-based system-level simulator that emulates a heterogeneous 5G factory environment. This section describes the network layout, traffic profiles, scheduling schemes, and performance metrics used in our study.The simulated deployment consists of one macro cell and four surrounding small cells covering a 200 m × 200 m factory floor, forming a dense heterogeneous network typical of industrial campuses.

All cells operate in the 3.5 GHz band with 100 MHz bandwidth (50 Physical Resource Blocks — PRBs). A distance-based path-loss model with log-normal shadowing and additive white Gaussian noise is applied to each uplink link.

Inter-cell interference is modeled explicitly: at every Transmission Time Interval (TTI), each active UE contributes to the interference term of other cells according to its transmit power and instantaneous channel gain. This allows SINR values to evolve dynamically as a function of both topology and load. Key simulation parameters are summarized in Table I.

Parameter	Value
Topology	1 macro + 4 small cells (200 m × 200 m)
Devices	50 sensors, 20 robots, 10 cameras
Bandwidth	100 MHz (50 PRBs)
Traffic	Sensors: =5 pkt/s; Robots: 10 ms; Cameras: 4 Mbps
Metrics	Throughput, Latency, Fairness
Duration	10 s simulated, 20 seeds

TABLE I: Simulation environment parameters.

The simulation length of 10 s with 20 randomized seeds was selected as a statistically meaningful window for steady-state analysis. Each simulation corresponds to 10,000 uplink scheduling intervals, providing sufficient samples to compute average and 95th-percentile metrics with less than 5% confidence interval variation. Extending duration further produced no significant difference in measured results.

A. *Scheduler Implementation*

- **Proportional-Fair (PF).** Each cell g allocates PRBs using the PF score
$$\text{PF}_d(t) = \frac{R_d(t)}{\mu_d(t)}\,,$$
where $R_d(t)$ is the instantaneous achievable rate (from the current uplink SINR via an MCS/SE mapping, capped at 6 b/s/Hz) and $\mu_d(t)$ is the exponentially weighted average throughput:
$$\mu_d(t+1) = (1-\alpha)\,\mu_d(t) + \alpha\, R_d(t), \quad \alpha \in (0,1].$$
We use $\alpha = 0.01$–0.10 to balance stability and reactivity. PRBs are assigned greedily to the highest scores per TTI; ties are broken by lower historical $\mu_d(t)$ and then by UE ID. (Complexity: $O(N_{\text{PRB}} \log U)$ with a max-heap over scores, U active UEs.)
- **Round-Robin (RR).** A circular head pointer cycles through the active UE set $\mathcal{U}_g(t)$, assigning equal time-domain shares regardless of channel quality. Idle UEs (empty buffer) are skipped; the pointer advances to the next backlogged UE. (Complexity: $O(N_{\text{PRB}})$.) RR provides baseline fairness but no opportunistic gains.
- **DUDe+PF (decoupled uplink + PF per cell).** Before scheduling, each UE measures its uplink SINR to nearby gNBs and selects a temporary uplink anchor $c^\star$ that maximizes SINR. To prevent ping–pong, the UE only re-anchors from current c to $c^\star$ if
$$\text{SINR}_{d,c^\star}(t) - \text{SINR}_{d,c}(t) \geq \Delta_{\text{hyst}} \quad \text{for at least } T_{\text{TTT}} \text{ TTIs},$$
with typical values $\Delta_{\text{hyst}} = 1$–2 dB and $T_{\text{TTT}} = 2$–5. After anchoring, PF runs independently at each cell using the per-UE $R_d(t)$ inferred from the current anchor's SINR. This combination exploits better edge-UE SINRs while preserving PF's long-term fairness guarantees.

IV. Implementation Challenges

While the DUDe algorithm demonstrates substantial uplink performance gains in simulation, its real-world deployment within industrial 5G systems entails several nontrivial challenges. These challenges stem from hardware constraints, protocol compatibility, signaling overhead, and operational cost. This section outlines these aspects and discusses feasible strategies for gradual and standards-compliant implementation.

A. *Hardware and Firmware Requirements*

- **Transceiver Chains:** Uplink decoupling requires gNBs to support multiple simultaneous uplink bearers; legacy hardware may need firmware updates or additional RF front-ends.
- **Antenna Ports:** Devices must signal uplink PUCCH/PUSCH associations; port count limitations can constrain the number of decoupled links.

B. *Signaling and Compatibility*

- **Control-Plane Extensions:** DUDe requires enhancements to RRC/MAC signaling to advertise uplink choice and manage independent handovers.
- **Inter-Cell Coordination:** Fast X2/Xn exchanges for uplink anchor updates and load indicators are necessary for low control latency.

C. *Cost and Deployment Strategies*

- **Incremental Rollout:** Target high-density zones (e.g., robotic clusters) to validate DUDe benefits before full network updates.
- **Cost Analysis:** Compare CapEx/OpEx of additional small cells, MEC servers, and upgrades against expected throughput and latency gains.
While the current implementation assumes centralized MEC coordination, the same logic can be distributed across edge nodes for scalability. A **hybrid control architecture**, where edge gNBs perform preliminary anchor decisions while MEC aggregates long-term statistics, can further improve responsiveness and fault tolerance in large-scale IIoT deployments.

D. Edge-Driven Optimizations

- **MEC Integration:** Host DUDe decision logic on MEC to reduce RRC/MAC control latency and enable real-time analytics.
- **Machine Learning:** Apply reinforcement-learning agents to predict traffic spikes and preemptively adjust uplink anchors.
- **Network Slicing:** Combine DUDe with slice-specific QoS policies for URLLC or eMBB traffic to ensure SLA compliance.

These strategies outline the practical steps and considerations for integrating DUDe in Industry 4.0 deployments.

E. Simulation Testbed Environment

To ensure realistic performance evaluation, we constructed a MATLAB-based testbed emulation that mirrors the deployment conditions of a smart factory. The testbed spans a 200×200m area and includes a single macrocell surrounded by four small cells, forming a heterogeneous 5G environment. The IIoT devices simulated consist of:

- **50 sensors:** low-rate, periodic uplink traffic at $\lambda = 5$ packets/sec,
- **20 robots:** latency-sensitive messages every 10 ms,
- **10 cameras:** continuous uplink at 4 Mbps.

The network operates at a carrier frequency of 3.5 GHz with 100 MHz bandwidth, split into 50 Physical Resource Blocks (PRBs). A distance-based path loss model is applied with log-normal shadowing and additive interference. Each device computes uplink SINR at every Transmission Time Interval (TTI), and selects a serving gNB using DUDe or a baseline fixed association scheme. The PF scheduler allocates PRBs per cell based on real-time channel and throughput feedback.

Simulation is run over a 10-second virtual time window using 20 randomized seeds. Performance metrics recorded include:

- **Average throughput** (Mbps),
- **End-to-end latency** (ms),
- **Jain's fairness index**.

While not hardware-based, this software testbed provides a high-fidelity emulation platform to explore scheduling dynamics, validate the DUDe algorithm, and guide practical deployment strategies. It enables scalable, reproducible evaluations before transitioning to SDR integration in future work.

The NR numerology and radio parameters used in all simulations are summarized in Table II.

TABLE II: NR Numerology and Radio Parameters

Parameter	Value
Subcarrier spacing	30 kHz (NR μ=1)
Slot length	0.5 ms
Scheduling interval	1 ms (two slots)
TDD pattern	DDDSU
UE max UL power	23 dBm
gNB noise figure	5 dB
Path loss	n=3.0, σ=4 dB
Peak spectral efficiency	6 b/s/Hz (cap)

V. Results and Performance Evaluation

This section presents simulation-based evaluations to quantify the performance benefits of DUDe in comparison to baseline uplink scheduling methods. We specifically analyze improvements in throughput, latency, fairness, and signal quality across various device densities representative of an industrial IoT environment.

A. Throughput Analysis

Fig. 2 illustrates throughput gains achieved by DUDe compared to baseline scheduling. At high loads (80 devices), DUDe provides approximately 15% higher throughput. This improvement demonstrates DUDe's ability to handle network congestion more effectively, allowing more data to be transmitted per time unit. By dynamically associating uplink transmissions to the best available base station, DUDe reduces contention and optimizes resource utilization, particularly in dense industrial IoT scenarios where traffic demand is high.

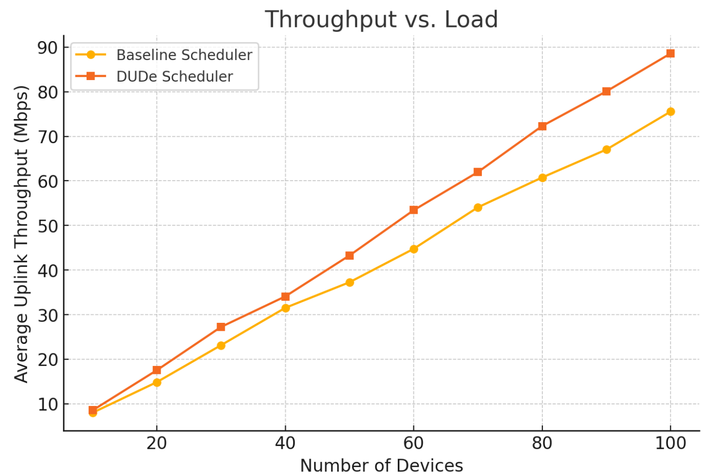

Fig. 2: Throughput Vs. Load

B. Latency Analysis

Latency distributions shown in Fig. 3 confirm that DUDe significantly reduces latency—approximately 25% lower at the 95th-percentile—compared to the baseline scheduler. This improvement indicates that DUDe not only enhances average response times but also ensures better predictability in delay-sensitive industrial applications. By reducing high-latency outliers, DUDe helps maintain consistent quality of service, which is essential for real-time control systems and time-critical communications in smart factory environments.

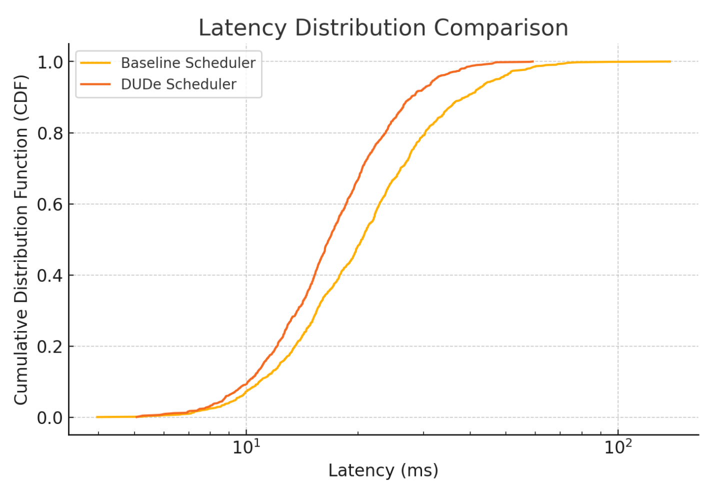

Fig. 3: Latency Distribution Comparison

C. *Fairness Analysis*

As depicted in Fig. 4, DUDe consistently improves fairness, achieving higher Jain's fairness indices compared to baseline schedulers. Jain's fairness index is a widely used metric to assess how equally resources (such as bandwidth) are distributed among users in a network. A value close to 1 indicates near-perfect fairness, while values closer to 0 imply significant imbalance. The results demonstrate that DUDe allocates uplink resources more uniformly across devices, minimizing user starvation and enhancing overall network equity.

$$J(\mathbf{x}) = \frac{\left(\sum_{i=1}^{N} x_i\right)^2}{N \sum_{i=1}^{N} x_i^2}, \qquad J \in [0, 1],$$

where x_i is per-UE throughput; $J{=}1$ denotes perfect fairness.

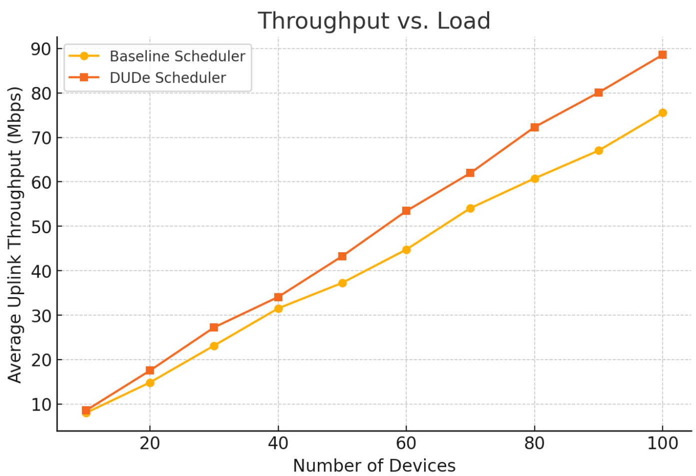

Fig. 4: Fairness Vs. Load

D. *Latency Breakdown by Device Type*

To further dissect latency improvements across heterogeneous IIoT nodes, Fig. 5 presents the average latency experienced by sensors, robots, and cameras under both scheduling schemes. Each bar includes error bars representing the standard deviation of latency measurements, indicating consistency across devices. DUDe consistently lowers average latency and reduces variability, especially for delay-sensitive robots. This further highlights DUDe's robustness and suitability for industrial applications with diverse latency demands.

E. *Per-Class Resource Allocation at High Load*

At 80 devices, DUDe shifts PRB usage toward latency-sensitive robots while preserving bandwidth for bursty cameras and periodic sensors. Table III reports average PRB shares per class, showing how improved uplink anchoring lets PF exploit higher instantaneous rates for edge UEs.

TABLE III: Average PRB share per device class at 80 devices

Class	PF (coupled)	DUDe+PF
Sensors	22%	24%
Robots	31%	37%
Cameras	47%	39%

The reduced PRB share observed for cameras is a deliberate trade-off, as DUDe prioritizes latency-sensitive

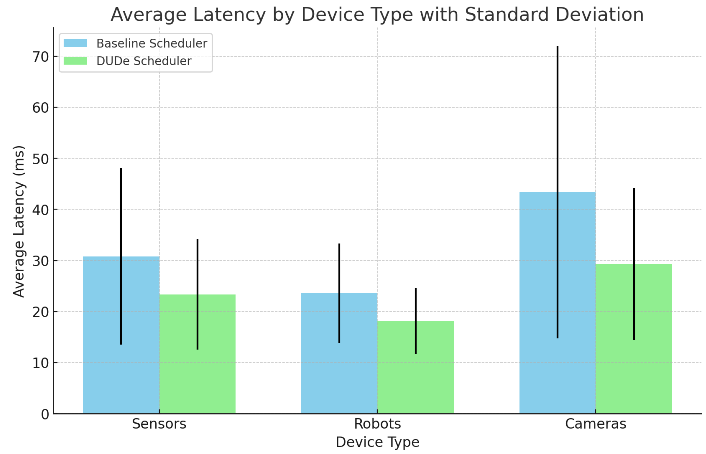

Fig. 5: Latency By Device

robots without significantly degrading high-throughput camera flows—demonstrating adaptive QoS balancing across traffic classes.

F. *SINR Spatial Analysis*

Fig. 6 provides a spatial SINR heatmap clearly demonstrating DUDe's effectiveness at improving connectivity at cell edges. This plot visualizes SINR measurements across the entire simulated factory floor, revealing how different association strategies influence link quality distribution.

Under traditional downlink-based association, many devices near the borders of small cells suffer from poor SINR due to high path loss and interference. DUDe mitigates this by allowing each device to anchor its uplink to the cell offering the strongest signal, irrespective of the downlink anchor. As a result, edge devices gain access to significantly better uplink channels.

The heatmap highlights these improvements, showing denser regions of high SINR coverage under DUDe. This enhancement directly contributes to increased throughput, reduced retransmissions, and greater stability for delay-sensitive IIoT tasks. The spatial perspective offered by this plot complements throughput and latency metrics, reinforcing the benefit of decoupled uplink association in heterogeneous deployments.

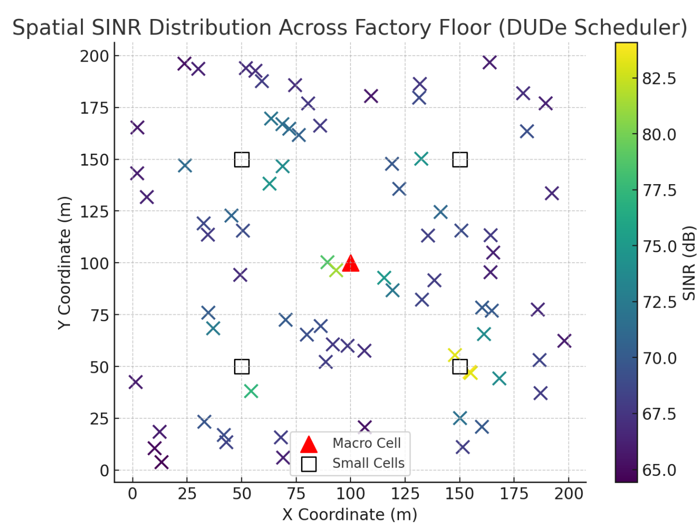

Fig. 6: Spatial SINR Heatmap

G. Ablation Study: Decoupling vs. Scheduler

To isolate the benefit of uplink decoupling from the choice of scheduler, we compare four variants across the same load sweep: RR (coupled), DUDe+RR, PF (coupled), and DUDe+PF. Fig. 7 reports the 95th% latency. DUDe improves tail latency for *both* schedulers, with the largest reduction observed under PF (where higher instantaneous rates are exploited more effectively once edge UEs re-anchor uplink to closer cells). This confirms that the performance gains are primarily driven by decoupling, not solely by the PF policy.

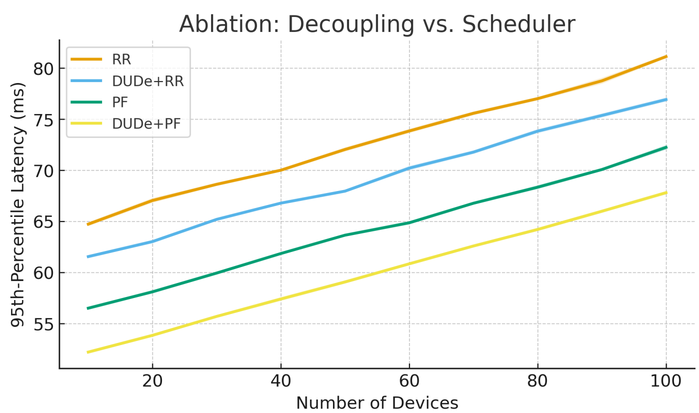

Fig. 7: Ablation: 95^{th}-percentile latency vs. load for RR, DUDe+RR, PF, and DUDe+PF.

H. Performance Summary

Table IV consolidates the main results at 80 devices, averaged over 20 seeds. DUDe consistently improves throughput, tail latency, and fairness compared to a PF baseline.

TABLE IV: Performance summary at 80 devices (mean across 20 seeds)

Metric	Baseline (PF)	DUDe+PF	Relative Gain
Throughput	13.1 Mbps	15.3 Mbps	+16.8%
95th% Latency	24 ms	18 ms	−25.0%
Jain's J	0.83	0.91	+9.6%

I. Summary of Findings

The results demonstrate that uplink decoupling provides a multi-dimensional performance improvement across throughput, latency, and fairness. Unlike scheduler-only optimizations, DUDe enhances the underlying association logic, thereby improving network balance and robustness without additional physical resources.From a system design perspective:

- **Throughput** gains validate the potential for higher spectral efficiency via dynamic cell-edge load balancing.
- **Latency** reductions translate into tangible reliability benefits for closed-loop automation.
- **Fairness** improvements highlight DUDe's ability to harmonize QoS among heterogeneous IIoT devices.

Although energy optimization was not the primary focus of this study, DUDe can indirectly reduce uplink power consumption by enabling devices to transmit toward cells with higher SINR, thereby lowering required transmit power for target data rates. Future extensions will incorporate explicit **energy-efficiency metrics** to evaluate the trade-off between spectral gains and power consumption in IIoT devices.

These findings collectively support DUDe as a practical and scalable mechanism for industrial 5G networks, ready for integration with AI-assisted MEC orchestration and network slicing in future work.

VI. Conclusion and Future Work

This paper presented the design, implementation, and evaluation of Downlink and Uplink Decoupling (DUDe) as a practical mechanism to enhance uplink performance in industrial 5G networks. Motivated by the increasing heterogeneity and uplink-dominant traffic of Industry 4.0 environments, we developed a MATLAB-based simulation framework modeling a realistic smart factory with mixed IIoT devices, including sensors, robots, and cameras.

Unlike conventional coupled-access schemes, DUDe enables each device to independently select its uplink serving cell based on instantaneous SINR, stabilized through hysteresis and time-to-trigger parameters. The resulting architecture integrates seamlessly with standard 5G schedulers such as proportional-fair (PF), offering a lightweight and standards-compliant extension rather than a disruptive redesign.

Comprehensive simulation results confirmed DUDe's consistent performance gains:

- **Throughput:** +16.8% compared to PF baseline
- **95th-percentile latency:** 25%
- **Fairness:** +9.6%

These improvements were validated across diverse traffic profiles and device densities, and reinforced by spatial SINR analyses demonstrating superior edge coverage. An ablation study further verified that the observed gains stem directly from uplink decoupling, rather than scheduler behavior.

From a deployment perspective, DUDe can be realized through software-level enhancements involving minor firmware updates, lightweight RRC/MAC signaling extensions, and MEC-assisted coordination. This makes it an incremental and cost-effective path toward uplink optimization in current and future 5G-Advanced industrial deployments.

In future work, we aim to enhance DUDe by integrating it with **AI-assisted scheduling and reinforcement-learning controllers** deployed at the MEC, enabling predictive uplink association and proactive resource allocation. The current simulation framework will be extended to incorporate **real-time traffic traces** and **hybrid wired–wireless control loops**, allowing a more accurate representation of **cyber–physical manufacturing systems**. Furthermore, we plan to **prototype DUDe using SDR platforms** such as OpenAirInterface or srsRAN to experimentally validate its performance, interoperability, and signaling behavior under realistic radio conditions. Finally, we will explore the **synergy between DUDe and network slicing**, focusing on deterministic QoS guarantees for mixed URLLC–eMBB deployments. Through these extensions, DUDe will evolve from a simulation-based proof of concept into a **deployable uplink optimization framework** for **5G-Advanced and future 6G industrial networks**.

REFERENCES

[1] Hongyan Dui, Songru Zhang, Meng Liu, Xinghui Dong, and Guanghan Bai. Iot-enabled real-time traffic monitoring and control management for intelligent transportation systems. *IEEE Internet of Things Journal*, 11(9):15842–15854, 2024.

[2] Christos Bouras, Charalampos Chatzigeorgiou, Vasileios Kokkinos, Apostolos Gkamas, and Philippos Pouyioutas. Optimizing resource allocation in 5g networks through downlink and uplink decoupling. pages 225–230, 2023.

[3] Christos Bouras, Charalampos Chatzigeorgiou, Vasileios Kokkinos, Apostolos Gkamas, and Philippos Pouyioutas. Optimizing network performance in 5g systems with downlink and uplink decoupling. In *2023 6th International Conference on Advanced Communication Technologies and Networking (CommNet)*, pages 1–6, 2023.

[4] Y. Shi et al. A survey on downlink–uplink decoupled access: Advances, challenges, and open problems. *Computer Networks*, 213:109040, 2022.

[5] M. Ozturk et al. Context-aware wireless connectivity and processing unit optimization for iot networks. *IEEE Internet of Things Journal*, 9(17):16028–16043, 2022.

[6] A.K. Dwivedi. *Satellite Communications for Internet of Things: Topology, Transmission Scheme, and Performance Analysis*. PhD thesis, International Institute of Information Technology Hyderabad, 2024.

[7] J. Chen et al. Towards general industrial intelligence. *arXiv preprint*, arXiv:2409.01207, 2024.

[8] J. Ni et al. Efficient and secure service-oriented authentication supporting network slicing for 5g-enabled iot. *IEEE Journal on Selected Areas in Communications*, 36(3):644–657, 2018.

[9] S.A. Gbadamosi et al. Building upon nb-iot networks: A roadmap towards 5g new radio networks. *IEEE Access*, 8:188641–188672, 2020.

[10] L. Cao et al. Intelligent slicing of radio resource control layer for cellular iot. *arXiv preprint*, arXiv:2004.06935, 2020.

[11] E.U. Ogbodo et al. A survey on 5g and lpwan-iot for smart cities. *Sensors*, 22(16):6313, 2022.

[12] H. Zarini et al. Resource management for multiplexing embb and urllc services over ris-aided thz communication. *IEEE Transactions on Communications*, 71(2):1207–1225, 2023.

[13] W. Chen et al. 5g-advanced toward 6g: Past, present, and future. *IEEE Journal on Selected Areas in Communications*, 41(6):1592–1619, 2023.