

Doppler-Induced Orthogonality Loss in Semantic-ISAC: Provable Gap, Digital-Twin Validation, and the Missing O-RAN Observable

Azharul Islam, *Member, IEEE*, Abdulahi Abiodun Badrudeen, *Member, IEEE*,
Mohsin Ali, *Member, IEEE*, and Sunwoo Kim, *Senior Member, IEEE*

Abstract—Semantic communication, integrated sensing and communication (ISAC), and O-RAN control have each matured around the same OFDM waveform, but under assumptions that do not hold together. Importance-weighted semantic subcarrier mapping presumes that subcarriers remain orthogonal; ISAC deliberately operates at high mobility, where Doppler-induced inter-carrier interference (ICI) breaks that orthogonality; and the standardized O-RAN measurement set reports neither Doppler nor orthogonality loss, so the controller that should arbitrate between the two has no interoperable, multi-vendor view of either. In this paper, we formalize this conflict. We prove that the ICI leakage power grows as $(\pi\varepsilon)^2/3$ in the normalized Doppler ε , that an orthogonality-assuming importance map becomes strictly suboptimal whenever ICI reorders the per-subcarrier SINRs, and that the optimal Doppler-compensation fraction depends on mobility and sensing priority, hence every static design incurs positive regret. Monte-Carlo simulations over 3GPP TDL channels show that the effect is mobility-gated, appearing in 1% of realizations at 3 km/h but in 93–97% at 120–400 km/h, and a ray-traced digital-twin study on NVIDIA Aerial Omniverse reproduces this prevalence curve (0.909 at a measured 119.8 km/h) on physical urban scenes. Finally, we propose a Doppler-aware importance mapper, a per-UE compensation policy, and three candidate E2 measurement extensions that would allow the RAN Intelligent Controller to close this loop.

Index Terms—semantic communication, integrated sensing and communication, OFDM, Doppler, inter-carrier interference, O-RAN, digital twin

I. INTRODUCTION

Three research communities currently optimize the same OFDM waveform on assumptions that are individually reasonable but jointly inconsistent. Semantic communication improves end-to-end task performance through importance-weighted resource allocation, in which high-value semantic units are mapped to high-quality subcarriers under unequal error protection [1]–[3]. Such designs implicitly assume that subcarriers remain orthogonal, so that per-subcarrier quality is well defined. ISAC, in contrast, targets vehicular, UAV, and high-speed-rail scenarios where the Doppler shift is large, because Doppler is the very observable from which radial

velocity is estimated [4]. Meanwhile, O-RAN closes its control loops on the standardized key-performance-measurement (KPM) set exposed over the E2 interface [5], drawn from the 3GPP measurement catalogue [6].

These assumptions collide in ISAC’s own operating regime: Doppler-induced inter-carrier interference (ICI) destroys the orthogonality that the semantic importance map relies on, while the same Doppler is signal to the sensor and interference to the semantic link. The controller that should arbitrate this trade-off cannot see it, since the standardized measurement set carries no per-UE Doppler, orthogonality-loss, or semantic-fidelity observable. To the best of our knowledge, no prior work jointly proves the suboptimality of orthogonality-assuming importance mapping under Doppler ICI and ties it to a missing O-RAN observable: existing studies benchmark semantic links on quasi-static channels, treat Doppler purely as an estimand, or standardize SINR-type KPIs.

Our contributions are threefold. First, we prove a quadratic orthogonality-loss lemma, a strict-suboptimality theorem for orthogonality-assuming importance maps under ICI reordering, a sensing–semantics regret theorem with a closed-form context-optimal compensation fraction, and a corollary identifying the three observables missing from the standardized O-RAN measurement set (Sec. III). Second, we validate the theory by link-level Monte-Carlo over 3GPP TDL channels in Sionna and by a ray-traced digital-twin study on NVIDIA Aerial Omniverse Digital Twin (AODT), which reproduces the predicted mobility-gated prevalence within a few percentage points (Sec. V). Third, we propose a Doppler-aware importance mapper, a per-UE compensation split, and three candidate E2 service-model extensions that close the loop on the RAN Intelligent Controller (RIC) (Sec. IV).

II. SYSTEM MODEL AND PROBLEM FORMULATION

Consider one OFDM symbol with N subcarriers and subcarrier spacing Δf . A propagation path carrying Doppler shift ν defines the normalized Doppler

$$\varepsilon = \nu/\Delta f, \quad (1)$$

where ν is the Doppler shift in Hz. After cyclic-prefix removal and the N -point DFT, a Doppler shift common to the symbol

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2024-00409492). (*Corresponding author: Sunwoo Kim.*)

Sunwoo Kim, Azharul Islam, Abdulahi Abiodun Badrudeen, and Mohsin Ali are with the Department of Electronic Engineering, Hanyang University, Seoul, South Korea (e-mail: {remero, azharul001, aabadrudeen, mohsin0987}@hanyang.ac.kr).




Fig. 1. System scenario at $v = 120$ km/h ($f_c = 3.6$ GHz, $\Delta f = 30$ kHz, $\varepsilon \approx 0.013$). The per-path Doppler $\nu_l = (\mathbf{v} \cdot \mathbf{k}_l)f_c/c$ is signal to ISAC but interference to the semantic link, and the standardized KPM set of the Near-RT RIC reports neither.

produces the inter-carrier leakage

$$I_{k,m}(\varepsilon) = S(m-k+\varepsilon), \quad S(\ell) = \frac{\sin(\pi\ell)}{N \sin(\pi\ell/N)} e^{j\pi\ell(N-1)/N}, \quad (2)$$

where $I_{k,m}$ is the leakage coefficient from transmit subcarrier m to demodulated subcarrier k and $S(\cdot)$ is the Dirichlet kernel; the algebra is identical to classical carrier-frequency-offset analysis [7]. Since a frequency shift followed by a unitary DFT is unitary, $\sum_m |I_{k,m}(\varepsilon)|^2 = 1$ for every k . The post-DFT per-subcarrier SINR is then

$$\text{SINR}_k(\varepsilon) = \frac{|I_{k,k}(\varepsilon)|^2 S_k}{\sigma^2 + \sum_{m \neq k} |I_{k,m}(\varepsilon)|^2 S_m}, \quad (3)$$

where $S_m \geq 0$ is the transmit power on subcarrier m and σ^2 is the noise power.

The semantic layer assigns an importance-weight vector to subcarriers through an assignment π , giving the total semantic distortion

$$D(\pi; \varepsilon) = \sum_k w_{\pi(k)} g(\text{SINR}_k(\varepsilon)), \quad (4)$$

where $w_1 \geq \dots \geq w_R$ are the importance weights (one per semantic unit) and g is a strictly decreasing per-unit distortion function of the SINR. The model also covers learned joint source-channel codecs [2]: a channel-robust decoder flattens g and shrinks the absolute distortion, but Sec. III needs only a strictly decreasing g , since ICI changes *which* subcarriers are best, not merely how noisy they are. In practice, the designer computes π from the orthogonality-assumed qualities $q_k = |H_{kk}|^2/\sigma^2$, where H_{kk} is the static channel gain of subcarrier k , i.e., from (3) with the ICI term deleted. Fig. 1 summarizes the scenario.

III. THEORETICAL RESULTS

A. Orthogonality Loss Is Quadratic in Doppler

Lemma 1: As $\varepsilon \rightarrow 0$,

$$\sum_{m \neq k} |I_{k,m}(\varepsilon)|^2 = 1 - |I_{k,k}(\varepsilon)|^2 = \frac{(\pi\varepsilon)^2}{3} + O(\varepsilon^4), \quad (5)$$

and for equal powers $S_m \equiv S$, $\text{SINR}_k(\varepsilon)$ is strictly decreasing in $|\varepsilon|$ near 0.

Proof sketch. Taylor-expanding (2) at $\ell = \varepsilon$ gives the diagonal gain; the leakage follows from energy conservation. The $1/3$ coefficient agrees with the classical average-ICI bound [7] and the value 0.3333 measured in Sec. V.

Monotonicity holds for equal powers only; under a non-uniform power profile, ICI can reorder the per-subcarrier SINRs—the mechanism exploited next.

B. The Importance Map Breaks Under Doppler

To leading order in ε , (3) can be written as

$$\text{SINR}_k(\varepsilon) \approx \frac{(1 - c\varepsilon^2) \gamma_k}{1 + c\varepsilon^2 \eta_k}, \quad (6)$$

where $c = \pi^2/3$, $\gamma_k = S_k/\sigma^2$ is the static SNR of subcarrier k , and η_k is its leakage-weighted neighbor-interference-to-noise ratio. Consider two semantic units with weights $w_{\text{hi}} > w_{\text{lo}}$ and two subcarriers a, b with $\gamma_a > \gamma_b$ but $\eta_a/\eta_b > \gamma_a/\gamma_b$, i.e., a is better statically but sits in a noisier ICI neighborhood. The static optimum assigns w_{hi} to a , yet from (6) the SINR order flips for all

$$\varepsilon^2 > \varepsilon_1^2 := \frac{\gamma_a - \gamma_b}{c(\gamma_b \eta_a - \gamma_a \eta_b)} > 0, \quad (7)$$

where ε_1 is the flip threshold, which is positive under the stated condition. The static assignment is then strictly suboptimal on an open interval above ε_1 . The general case is the following.

Theorem 1 (Doppler-blind importance mapping is strictly suboptimal): Let the weights $\{w_i\}$ and the ICI-aware SINRs $\{\text{SINR}_k(\varepsilon)\}$ be distinct. Let π_0 pair weights (descending) with subcarriers ordered by the orthogonality-assumed quality q_k , and let π^* pair them with subcarriers ordered by the true $\text{SINR}_k(\varepsilon)$. If the two orderings differ in at least one position, then

$$\Delta(\varepsilon) := D(\pi_0; \varepsilon) - D(\pi^*; \varepsilon) > 0, \quad (8)$$

where $\Delta(\varepsilon)$ is the excess distortion of the Doppler-blind map. *Proof.* Since g is strictly decreasing, $\{g(\text{SINR}_k)\}$ reverses the SINR order, and by the strict rearrangement inequality [8] the sum (4) is uniquely minimized by pairing the largest weight with the largest SINR, that is, by π^* . If the q -order differs from the $\text{SINR}(\varepsilon)$ -order, then $\pi_0 \neq \pi^*$, so π_0 is not the unique minimizer.

At $\varepsilon = 0$ the two orders coincide and $\Delta = 0$, so the theorem predicts no effect on static links. For $\varepsilon > 0$, the fraction of realizations in which the ranking changes is exactly the prevalence $\text{frac}(\Delta > 0)$ measured in Sec. V; since g is arbitrary, each such instant degrades any SINR-monotone task metric (BLER, throughput, semantic accuracy).

C. Sensing–Semantics Regret of Static Doppler Handling

Let $\rho \in [0, 1]$ be the fraction of Doppler compensated before detection, leaving the residual Doppler $(1-\rho)\varepsilon$. By Lemma 1, the semantic distortion behaves as $D_{\text{sem}}(\rho) = a((1-\rho)\varepsilon)^2 + d_0$, where $a > 0$ and d_0 are model constants, and it is strictly decreasing in ρ . Compensation, however, removes the very observable used for velocity estimation, so the velocity Cramér–Rao bound scales as $\text{CRB}_v(\rho) = b/(1-\rho)^p$, where

$b > 0$ depends on the link and $p > 0$ on the estimator; coherent single-tone frequency estimation gives $p \approx 3$ [9], and our measured fit is $p \approx 2.98$ (Sec. V).

Theorem 2 (Static compensation incurs positive regret): Every $\rho \in [0, 1)$ is Pareto-efficient. For the scalarization $J_\lambda(\rho) = \lambda D_{\text{sem}}(\rho) + (1 - \lambda) \text{CRB}_v(\rho)$, where $\lambda \in (0, 1)$ weights semantics against sensing, the minimizer $\rho^*(\varepsilon, \lambda)$ is unique and non-constant in ε . Hence, for any static ρ_{fix} and any non-degenerate mobility/priority distribution $\mathcal{D}$,

$$\text{Regret}(\rho_{\text{fix}}) = \mathbb{E}_{(\varepsilon, \lambda) \sim \mathcal{D}}[J_\lambda(\rho_{\text{fix}}) - J_\lambda(\rho^*(\varepsilon, \lambda))] > 0. \quad (9)$$

For $p = 2$,

$$u^*(\varepsilon, \lambda) = \text{clip}_{(0,1]} \left(\frac{1}{\varepsilon} \sqrt{\frac{(1-\lambda)b}{\lambda a}} \right), \quad \rho^* = 1 - \sqrt{u^*}, \quad (10)$$

where $u = (1 - \rho)^2$ is the squared residual fraction.

Proof sketch. J_λ is strictly convex in u for any $p > 0$, which gives uniqueness. The interior optimum balances an ε^2 -scaled term against an ε -free term, so ρ^* varies with ε , and the expectation in (9) vanishes only on a measure-zero set.

At low mobility $\rho^* = 0$, preserving the sensing observable; at high mobility $\rho^* > 0$, protecting the semantic link. The right ρ is therefore a control decision rather than a design constant. This also delimits the deployment scope: low-mobility ISAC (indoor sensing, industrial monitoring) has $\rho^* = 0$ and a near-optimal static map, so the framework targets vehicular, UAV, and high-speed-rail ISAC, where Sec. V shows the effect is near-universal.

D. The Missing O-RAN Observable

Corollary 1: Realizing $\rho^*(\varepsilon, \lambda)$ and the ICI-aware map of Theorem 1 requires three per-UE observables on the RIC control interface: (i) the Doppler spread or normalized ε , (ii) an orthogonality-loss (residual-ICI power) indicator, and (iii) a semantic-fidelity measurement. Based on the O-RAN E2SM-KPM catalogue [5] and the 3GPP measurement set it draws on [6], none of the three is standardized, so the optimal control law is not implementable on today's RIC through standardized measurements.

This is a standardization gap, not a fundamental limit of O-RAN observability: the Near-RT RIC is extensible via xApps and custom E2 service models, so a single vendor could expose these quantities today; what is missing is an interoperable, standardized information element a multi-vendor xApp can rely on—precisely what we propose next.

IV. DOPPLER-AWARE SEMANTIC-ISAC CONTROL

The theory converts directly into the control architecture of Fig. 2. **C1 (Doppler-aware importance mapping)** replaces the orthogonality-assuming assignment π_0 with $\pi^*(\varepsilon)$ computed against the ICI-coupled SINR (3); by Theorem 1 this closes the gap $\Delta(\varepsilon)$ exactly and is a drop-in change to existing importance-based OFDM pipelines. **C2 (context-adaptive compensation split)** applies $\rho^*(\varepsilon, \lambda)$ from (10) per UE, trading semantic distortion against velocity CRB on the Pareto frontier instead of using a fixed equalizer. **C3**




Fig. 2. Proposed closed loop: a Near-RT RIC xApp consumes three new per-UE measurements $\{\hat{\varepsilon}, \hat{\sigma}_{\text{ICI}}^2, \Phi_{\text{sem}}\}$ (C3) and applies, per coherence interval, the ICI-aware map of Theorem 1 (C1) and the split ρ^* of Theorem 2 (C2).

TABLE I
SIONNA MONTE-CARLO (TDL-C, JAKES, 400 REALIZATIONS/SPEED)

UE speed	ε	mean ICI/signal	frac($\Delta > 0$)
3 km/h	0.0003	4.8×10^{-7}	0.01
30 km/h	0.0032	5.3×10^{-5}	0.44
60 km/h	0.0065	2.0×10^{-4}	0.76
120 km/h	0.0130	7.9×10^{-4}	0.93
250 km/h	0.0270	3.1×10^{-3}	0.96
400 km/h	0.0432	9.7×10^{-3}	0.97

(the missing observables and the RIC loop) defines three candidate E2SM-KPM measurement extensions, namely per-UE Doppler spread, residual-ICI power, and semantic fidelity, together with an xApp that consumes them to set (π^*, ρ^*) per UE per coherence interval, making Corollary 1 implementable.

V. PERFORMANCE EVALUATION

A. Link-Level Monte-Carlo over 3GPP Channels

We first verify the closed forms with the exact single- ε ICI operator ($N = 64$): the measured leakage coefficient is 0.3333 (Lemma 1 predicts $1/3$), and $\Delta(\varepsilon) \geq 0$ throughout. We then replace the idealization with the exact doubly-selective channel matrix built from time-varying 3GPP TDL-C path gains with Jakes Doppler spectra in Sionna [10], [11]: $f_c = 3.5$ GHz, $\Delta f = 30$ kHz, 20 dB SNR, and 400 realizations per speed; a static channel yields an exactly diagonal effective matrix (self-test). Table I and Fig. 3 report the results. Mean ICI grows quadratically in ε as predicted, and the Theorem 1 gap is mobility-gated, negligible for pedestrians yet near-universal in ISAC's high-mobility regime.

At FR1 with 30 kHz spacing the absolute gaps are small ($\varepsilon \lesssim 0.043$), so prevalence is the main story; sweeping FR2 (28 GHz, 60/120 kHz) and TDL-A/B/C raises ε to 0.216 with roughly tenfold larger gaps and prevalence of 0.57–0.99 across all delay profiles, confirming that the effect is not an FR1 or TDL-C artifact.

For Theorem 2, we measure $D_{\text{sem}}(\rho)$ on the same channels with ρ -fractional compensation and estimate velocity with a periodogram estimator; the measured error-variance exponent is $p \approx 2.98$, matching the coherent CRB model [9]. Fig. 4 shows the measured distortion family and the resulting $\rho^*(\varepsilon)$, which rises from 0.17 to 0.71 over 30–400 km/h.




Fig. 3. Monte-Carlo over 3GPP TDL-C with Jakes Doppler spread. Left: measured mean ICI/signal vs. a quadratic reference (Lemma 1). Right: mean and median Theorem 1 distortion gap vs. normalized Doppler.




Fig. 4. Theorem 2 on Sionna. Left: measured semantic distortion $D_{\text{sem}}(\rho)$ per mobility class. Right: context-optimal $\rho^*(\varepsilon)$ (measured D_{sem} and empirical CRB, $p \approx 3$) vs. the best static ρ .

The best static $\rho = 0.61$ —itself a tuned, Doppler-aware but context-blind baseline—carries positive regret (0.002, and 0.012/0.006/0.003 for $p = 1/2/3$ in the closed-form model), confirming (9).

B. Ray-Traced Digital-Twin Validation

Statistical TDL channels assume a Doppler spectrum, whereas a real scene produces per-path Doppler $\nu_l = (\mathbf{v} \cdot \hat{\mathbf{k}}_l) f_c / c$ from physical geometry, where $\mathbf{v}$ is the UE velocity, $\hat{\mathbf{k}}_l$ the direction of path l , and c the speed of light. We therefore validate on NVIDIA AODT 1.4.0 [12], which ray-traces time-evolving channels on geo-referenced urban scenes (Fig. 5). Each electromagnetic simulation (FR1 3.6 GHz, $\Delta f = 30$ kHz, one RU at 43 dBm, one UE on a straight 67 m track, 11 mobility samples) exports per-path channel impulse responses; we estimate per-path Doppler from the ray geometry, synthesize the intra-symbol time variation, and reuse the Sionna effective-matrix builder.

Two native runs bracket the theory. The pedestrian null control (measured 2.11 m/s, $\varepsilon = 0.0008$) gives $\text{frac}(\Delta > 0) = 0.000$, the low-mobility prediction. The positive test at a measured 33.28 m/s (119.8 km/h, $\varepsilon = 0.0126$) gives $\text{frac}(\Delta > 0) = 0.909$, against the Sionna reference 0.93 at 120 km/h, with $\rho^* \approx 0.52$ –0.54. Because the instantaneous ray geometry is velocity-independent and $\nu_l \propto |\mathbf{v}|$ exactly, scaling the UE speed along the real heading yields a rigorous ray-

TABLE II
THEOREM 1 PREVALENCE $\text{frac}(\Delta > 0)$: SIONNA VS. AODT SCENES

UE speed	Sionna TDL-C (statistical)	AODT Hanyang (≤ 25 paths)	AODT Tokyo (≈ 500 paths)
3 km/h	0.01	0.00	0.00
30 km/h	0.44	0.55	0.28
60 km/h	0.76	0.73	0.51
120 km/h	0.93	0.91	0.66
250 km/h	0.96	1.00	0.70

traced speed sweep (Fig. 6): prevalence rises monotonically and tracks the TDL Monte-Carlo within a few percentage points.

C. Cross-Scene Effect of Multipath Richness

Repeating the sweep on the much richer Tokyo scene (≈ 500 paths/frame) yields the comparison of Table II. Both scenes confirm Theorem 1, with mobility-gated and monotonically rising prevalence. The measured nuance is that heavy diffuse multipath partially averages out the SINR-ranking flip, lowering the 120 km/h prevalence from 0.909 (dominant-path) to 0.66; the theorem is strongest in dominant-path channels and narrows, but does not vanish, under rich scattering. All native runs are single-UE with isotropic panels, so no cross-population regret is claimed from AODT.




Fig. 5. AODT ray-traced digital-twin scenes used for validation: a Tokyo city scene with rich multipath (≈ 500 ray paths per frame, top) and the Hanyang University campus scene (dominant-path, bottom).




Fig. 6. AODT velocity sweep on real ray geometry (Hanyang scene). Left: mean normalized Doppler vs. UE speed. Center: orthogonality loss (Lemma 1). Right: Theorem 1 prevalence, from 0 (pedestrian) to 1.0 (250 km/h).

VI. CONCLUSIONS

We formalized a structural conflict between semantic communication, ISAC, and O-RAN control. Doppler-induced ICI invalidates the orthogonality assumption of semantic importance mapping precisely in ISAC's operating regime; the optimal compensation fraction is context-dependent, so any static design carries positive regret; and the standardized O-RAN measurement set exposes none of the observables needed to control either effect. The theory was validated on statistical 3GPP channels and ray-traced digital-twin scenes, which reproduce the mobility-gated prevalence within a few percentage points. The proposed remedy, an ICI-aware importance mapper, a per-UE compensation policy, and three E2 measurement extensions, is implementable on the RIC; multi-UE Doppler spread, native mmWave scenes, an xApp prototype on a live RIC, and end-to-end evaluation with trained semantic codecs on task-level metrics are left for future work.

REFERENCES

- [1] A. Islam and K. Chang, "Navigating the future of wireless networks: A multidimensional survey on semantic communications," *ICT Express*, vol. 10, no. 4, pp. 747–773, 2024.
- [2] E. Boursoulatzé, D. B. Kurka, and D. Gündüz, "Deep joint source-channel coding for wireless image transmission," *IEEE Trans. Cogn. Commun. Netw.*, vol. 5, no. 3, pp. 567–579, Sep. 2019.
- [3] A. Islam and K. Chang, "Beyond bit fidelity: Dual-mode domain-aware knowledge-enhanced semantic text communication framework," *IEEE Trans. Netw. Sci. Eng.*, vol. 13, pp. 7964–7986, 2026.
- [4] F. Liu *et al.*, "Integrated sensing and communications: Toward dual-functional wireless networks for 6G and beyond," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 6, pp. 1728–1767, Jun. 2022.
- [5] O-RAN Alliance, "O-RAN Work Group 3, Near-Real-Time RAN Intelligent Controller, E2 Service Model (E2SM) KPM," O-RAN.WG3.E2SM-KPM-R003, Tech. Spec., 2024.
- [6] 3GPP, "Management and orchestration; 5G performance measurements," TS 28.552, v18, 2024.
- [7] Y. Li and L. J. Cimini, "Bounds on the interchannel interference of OFDM in time-varying impairments," *IEEE Trans. Commun.*, vol. 49, no. 3, pp. 401–404, Mar. 2001.
- [8] G. H. Hardy, J. E. Littlewood, and G. Pólya, *Inequalities*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1952.
- [9] D. C. Rife and R. R. Boorstyn, "Single-tone parameter estimation from discrete-time observations," *IEEE Trans. Inf. Theory*, vol. 20, no. 5, pp. 591–598, Sep. 1974.
- [10] J. Hoydis *et al.*, "Sionna: An open-source library for next-generation physical layer research," *arXiv preprint arXiv:2203.11854*, 2022.
- [11] 3GPP, "Study on channel model for frequencies from 0.5 to 100 GHz," TR 38.901, v17.0.0, Mar. 2022.
- [12] NVIDIA, "Aerial Omniverse Digital Twin," 2025. [Online]. Available: <https://docs.nvidia.com/aerial/aerial-dt/>