



Uncertainty-Aware Radar-LiDAR Fusion for PoE-Constrained Smart Infrastructure Perception with Asynchronous Sensing

Mostafa Borhani^{1,*}

¹Smart Tech Services, Muscat, Oman

Email: borhani@iSmartGCC.com

Abstract

Infrastructure-based autonomous perception operates under fundamentally different constraints than vehicle-mounted systems: elevated-mounting geometries producing depression-angle-dependent sparse point clouds, a 12.95 W IEEE 802.3af Power-over-Ethernet (PoE) power ceiling, and distributed asynchronous sensing governed by IEEE 1588v2 precision time protocol (PTP) synchronization uncertainty. Existing automotive radar-LiDAR fusion frameworks assume abundant power, dense sensing, and synchronous measurements — assumptions that all fail in fixed infrastructure deployments. This paper presents XADAR, an uncertainty-aware multi-modal fusion framework designed for these infrastructure-specific constraints. XADAR makes three principal contributions: (1) a covariance inflation mechanism that propagates PTP synchronization uncertainty continuously through the fusion pipeline, replacing hard synchronization thresholds with a smooth degradation curve proportional to temporal offset; (2) adaptive sensor-specific fusion weights derived from modality covariance matrices that account for IWR6843 77 GHz FMCW radar Doppler ambiguity and ground-reflection multipath, and TFS20-L ToF LiDAR atmospheric scattering and range-zone limitations; and (3) a complete reproducible architecture including an IEEE 802.3af-compliant power budget (5.78 W maximum concurrent load; 41.6% PoE safety margin), quantitative 77 GHz propagation analysis based on ITU-R P.676-12 and P.838-3 (10.7 dB fade margin at 100 m under 50 mm/hr rain), and an MIL-STD-1629 FMEA covering twelve failure modes with severity classifications. A structured five-stage validation pathway from synthetic temporal-offset experiments to six-month field trials is defined for future empirical work.

Keywords: Multi-sensor fusion; Radar-LiDAR fusion; Uncertainty quantification; Covariance inflation; Power-over-Ethernet; IEEE 1588v2 precision time protocol; Smart infrastructure perception; 77 GHz FMCW radar

1 Introduction

1.1 Infrastructure Perception Constraints

Autonomous vehicle perception has matured through extensive vehicle-centric multi-modal fusion research [1–3]. These systems optimize for forward-facing geometry with depression angles between 0–30°, dense LiDAR point clouds providing 100–200+ points per target, abundant power budgets at kilowatt scale, and centralized on-vehicle processing with direct sensor-actuator integration. Infrastructure-based perception for smart cities [4], autonomous intersection management [5], and traffic monitoring operates under inverted constraints that fundamentally alter the perception problem. Top-down observation from elevated mounting points at 5–15 m produces depression angles of 45–60°. This oblique viewing geometry creates sparse and irregular point distributions with range-dependent density variation that differs substantially from automotive forward-facing assumptions. The geometric inversion affects not only point density but also occlusion patterns, object aspect ratios, and the validity of ground plane assumptions commonly employed in automotive perception. Power constraints represent another fundamental inversion. IEEE 802.3af Power over Ethernet [6] limits each sensor node to 12.95 W maximum, representing approximately two orders of magnitude reduction from complete autonomous-vehicle sensor-plus-compute platforms operating in the kilowatt range. This constraint mandates careful power-aware sensor scheduling and hardware selection. The processing architecture also inverts from centralized to distributed, where microcontrollers buffer sensor data while centralized GPUs perform inference, requiring explicit synchronization mechanisms for multi-node coordination. The sparsity regime created by infrastructure viewing geometry yields point cloud densities substantially lower than automotive baselines, creating distinct algorithmic challenges for feature extraction and object detection. Recent infrastructure datasets including InfraDet3D [7], InScope [8], and TUMTraf V2X [9] advance annotation and benchmarking capabilities. However, theoretical fusion frameworks that simultaneously address uncertainty quantification under asynchronous measurements, power-constrained hardware specifications, and systematic reliability engineering remain gaps in published literature. Figure 1 illustrates the XADAR distributed architecture addressing these challenges.

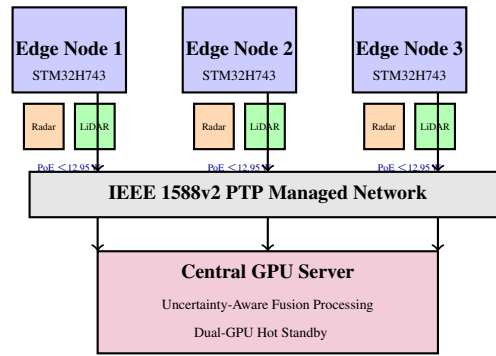


Figure 1. XADAR distributed architecture. Three IEEE 802.3af PoE-powered edge nodes — each with a TI IWR6843 77 GHz FMCW radar and a Benewake TFS20-L ToF LiDAR within a 5.78 W maximum power budget — connect via an IEEE 1588v2 PTP-synchronised managed network to a dual-GPU central fusion server. Arrows show the timestamped measurement data flow entering the uncertainty-aware Kalman fusion pipeline.

1.2 Theoretical and Architectural Contributions

XADAR targets the *Fusion: Practice and Applications* scope of grounded, deployable fusion systems: theoretical algorithms are paired with a complete hardware specification reproducible from public datasheets and established standards (IEEE 802.3af, IEEE 1588v2, MIL-STD-1629). To the best of our knowledge, XADAR is the first framework to jointly address uncertainty-aware radar–LiDAR sensor fusion, Power-over-Ethernet hardware constraints, IEEE 1588v2 asynchronous synchronization tolerance, and systematic reliability engineering within a single infrastructure perception system.

XADAR makes three principal contributions:

- (1) **Covariance inflation for temporal uncertainty propagation.** A formal mechanism propagates IEEE 1588v2 synchronization uncertainty through the multi-modal fusion pipeline. Rather than imposing a hard synchronization threshold that creates binary hand-off failures, the mechanism inflates measurement covariance in proportion to temporal offset magnitude, enabling graceful degradation. Uncertainty propagation accounts for radar Doppler ambiguity and multipath under oblique infrastructure geometry, LiDAR atmospheric scattering at 905 nm, temporal misalignment, and geometric projection uncertainty from depression-angle variation.
- (2) **Adaptive sensor-specific fusion weights.** Fusion weights update dynamically from modality-specific covariance matrices. The IWR6843 77 GHz FMCW radar provides precise Doppler velocity but is susceptible to ground-reflection multipath [29]; the TFS20-L ToF LiDAR offers 2 cm ranging precision in its 0–20 m close-range zone [30] but degrades under high ambient illumination. Adaptive weighting extracts maximum information under varying conditions without manual threshold tuning.
- (3) **Reproducible PoE-constrained architecture with reliability engineering.** The complete specification includes IEEE 802.3-2022 power budgets with a corrected maximum concurrent load of 5.78 W versus the 9.9 W usable PoE floor (41.6% margin), datasheet-sourced component selection, quantitative 77 GHz propagation analysis based on ITU-R P.676-12 [18] and ITU-R P.838-3 [19] showing 10.7 dB fade margin at 100 m under heavy rain, and an MIL-STD-1629 FMEA [22] covering twelve failure modes across hardware, network, and software layers.

A curriculum-learning domain adaptation methodology and an experimental validation roadmap with explicit success criteria complete the contribution.

1.3 Scope and Validation Requirements

This paper contributes theoretical algorithms and architectural specifications; it does not yet report hardware testbed or field measurements. Section 7 details scope boundaries, current limitations, and the structured validation pathway from controlled synthetic experiments to production field deployment.

2 Related Work

2.1 Vehicle-Centric Multi-Modal Fusion

Extensive automotive research addresses camera-radar-LiDAR fusion [3, 10, 11]. Feature-level fusion approaches such as PointPillars [12] and variants convert 3D point clouds to bird’s-eye-view representations, concatenating with radar features before detection heads. These methods assume dense point clouds with 100–200+ points per target, forward-facing geometry with 0–30° depression angles, and centralized processing. Decision-level fusion employs multi-hypothesis tracking that fuses detection outputs from independent modality-specific networks, relying on high-confidence single-modality detections unavailable in infrastructure sparse regimes. The automotive fusion literature uniformly assumes abundant power, dense sensing, and centralized processing. Infrastructure constraints — oblique sparse geometry, PoE power limits, and distributed asynchronous synchronization — remain unaddressed in this literature because the problem domain fundamentally differs from the automotive paradigm. XADAR fills this gap with algorithms and hardware specifications designed from the outset for the infrastructure constraint set.

2.2 Infrastructure Perception Datasets

InfraDet3D [7] provides a roadside LiDAR 3D detection dataset with elevated viewpoints but no multi-modal fusion framework. InScope [8] offers a cooperative V2X perception dataset focusing on vehicle-infrastructure cooperation rather than infrastructure-only multi-sensor fusion. TUMTraf V2X [9] presents infrastructure-vehicle cooperative dataset with cooperative fusion baselines but focuses on V2X cooperation rather than infrastructure-only multi-sensor uncertainty quantification. These datasets advance annotation and benchmarking but none provides a fusion methodology for asynchronous multi-modal sensing under power constraints. XADAR addresses this gap with explicit temporal uncertainty propagation, sensor-specific covariance models, and a hardware specification that satisfies the IEEE 802.3af power budget from first principles.

2.3 Uncertainty Quantification Approaches

Model-based radar uncertainty characterisation traces back to signal processing work on coherent receiver noise in active and passive radar configurations [27], establishing that structured noise models can separate clutter from target returns under complex scattering geometries — a principle motivating the explicit sensor-specific uncertainty models in Section 4.2 rather than a purely data-driven approach. Covariance intersection [24] provides principled conservative fusion of estimates with unknown cross-correlations, but assumes synchronous measurements and offers no mechanism for propagating temporal misalignment uncertainty into fusion weights. Recent deep learning approaches including MVXNet [25] and uncertainty-aware fusion methods [26] learn uncertainty weights end-to-end but require dense training data unavailable in infrastructure sparse regimes. XADAR addresses this gap with explicit sensor-specific and temporal uncertainty models, achieving principled fusion under asynchronous conditions without requiring large labelled infrastructure datasets.

2.4 Distributed Synchronization

IEEE 1588v2 PTP [14] achieves sub-microsecond synchronization on dedicated hardware. Industrial automation literature demonstrates microsecond-level synchronization on specialized platforms, but embedded robotics applications under power constraints remain less characterized. XADAR treats synchronization quality as a continuous variable that directly modulates fusion weights via the covariance inflation mechanism (Section 4.3), replacing the binary pass/fail boundary common in prior synchronization-dependent fusion designs.

3 System Architecture and Power Analysis

3.1 Infrastructure Deployment Geometry

Infrastructure sensor nodes mounted at 5–15 m on streetlights, traffic signal poles, or building overhangs create a layered observation geometry. The close-range zone (0–20 m horizontal) directly beneath each mount yields depression angles of 45–85°, well-suited to the Benewake TFS20-L LiDAR whose 20 m rated range [30] covers this zone. The extended intersection-scale zone (20–100 m horizontal) produces shallower depression angles of 6–27° depending on mounting height; here the TI IWR6843 77 GHz FMCW radar extends tracking to 100 m [17]. This range-based complementarity is a deliberate architectural decision motivating multi-modal fusion over single-sensor deployment.

Field of regard spans 90–120° azimuth from a corner-mounted sensor, covering a full intersection without mechanical scanning. A 10 Hz update rate suffices for infrastructure monitoring: pedestrian walking speeds of 1.2–1.6 m/s [4] and urban vehicle speeds of 5–15 m/s [4] produce at most 1.5 m inter-frame displacement, well within the tracker gating volume. High-speed highway monitoring requires 20–30 Hz at speeds above 30 m/s. The 12.95 W IEEE 802.3af PoE ceiling [6] is the dominant hardware constraint separating infrastructure from automotive perception; every other design decision follows from it.

3.2 Power Budget Derivation

IEEE 802.3-2022 [6] specifies 12.95 W maximum per Class 0-3 PoE port. Realistic usable power accounts for cable and conversion losses according to:

$$P_{\text{usable}} = P_{\text{PoE}} \times \eta_{\text{cable}} \times \eta_{\text{converter}} \quad (1)$$

The PoE standard maximum provides $P_{\text{PoE}} = 12.95$ W per IEEE 802.3af specification. Cable resistive loss derives from the 84.5 Ω /km conductor resistance of 24 AWG Cat5e cable specified in TIA/EIA-568-C [28], yielding η_{cable} ranging from 0.85 for 100 m worst-case to 0.92 for 50 m typical deployment.

DC-DC conversion efficiency depends on commercial PoE PD controllers. Microchip TPS23861 [15] and Maxim MAX5953 [16] datasheets specify conversion efficiency $\eta_{\text{converter}}$ from 0.90 under conservative loaded conditions to 0.95 under optimistic light load. These efficiency values derive from direct datasheet measurements under specified operating conditions.

The usable power range calculation yields minimum $P_{\text{min}} = 12.95 \times 0.85 \times 0.90 = 9.91$ W and maximum $P_{\text{max}} = 12.95 \times 0.92 \times 0.95 = 11.32$ W. The design constraint requires system operation sustaining continuous operation at worst-case $P_{\text{min}} = 9.9$ W. Component selection targets maximum power consumption below this threshold with margin for aging effects and thermal derating over operational lifetime.

3.3 Component Selection

Table 1 summarizes edge node component specifications with datasheet-sourced power consumption values. The microcontroller employs STM32H743ZIT6 providing 480 MHz ARM Cortex-M7 architecture with 2 MB Flash and power consumption ranging from 0.15–1.2 W depending on operational mode. The 77 GHz FMCW radar utilizes TI IWR6843 with 4 receive channels, 3

Table 1: XADAR edge node component specifications from manufacturer datasheets.

Component	Model	Key Specification	Power (W)
Microcontroller	STM32H743ZIT6	480 MHz ARM Cortex-M7, 2 MB Flash	0.15–1.2
77 GHz FMCW Radar	TI IWR6843	4 RX, 3 TX, 4 GHz chirp BW [17]	2.1–3.2
ToF LiDAR	Benewake TFS20-L	20 m (dark/90% refl.), 905 nm VCSEL, 2 cm precision @1 σ , Class 1 IEC60825 [30]	≤ 0.43
Ethernet PHY	Microchip KSZ9031RNX	1 Gbps + IEEE 1588 HW timestamp [20]	0.35–0.7
PoE PD Controller	TI TPS23861	IEEE 802.3af, 90% efficiency [15]	0.15–0.25
Total Range			2.75–5.78

TFS20-L outdoor range reduces to 9–15 m at 100 Klux [30].

Table 2: Multi-sensor fusion node power budget. Margin against 9.9 W usable PoE floor.

Operating Mode	Typical (W)	Max (W)	Margin (W)	(%)
Sleep (all idle)	0.45	0.70	9.20	92.9
Radar-only active	3.10	5.35	4.55	46.0
LiDAR-only active	1.30	2.58	7.32	73.9
Radar+LiDAR concurrent	4.35	5.78	4.12	41.6

Concurrent typical: MCU 0.8, Radar 2.5, LiDAR 0.35, PHY 0.5, PoE ctrl 0.2 W.

Concurrent max: MCU 1.2, Radar 3.2, LiDAR 0.43, PHY 0.7, PoE ctrl 0.25 W.

41.6% margin covers component aging (10%), 60°C derating, transient spikes.

transmit channels, and 4 GHz chirp bandwidth consuming 2.1–3.2 W. Time-of-flight LiDAR sensor Benewake TFS20-L provides 20 m range (dark, 90% reflectivity) at 905 nm VCSEL wavelength, with Class 1 eye-safe operation and ≤ 0.43 W average power [30]. Ethernet PHY Microchip KSZ9031RNX provides 1 Gbps connectivity with IEEE 1588 hardware timestamping capability consuming 0.35–0.7 W. PoE PD controller TI TPS23861 implements IEEE 802.3af with 90% efficiency consuming 0.15–0.25 W for power management functions.

Table 2 analyses power budget across operating modes with corrected TFS20-L power figures [30]. Sleep mode consumes 0.45–0.70 W. Radar-only active consumes 3.10–5.35 W with 4.55 W margin. LiDAR-only active consumes 1.30–2.58 W with 7.32 W margin. Concurrent radar and LiDAR operation is the critical design point, consuming 4.35 W typical and 5.78 W maximum with 4.12 W margin (41.6%).

The corrected concurrent analysis — accounting for the TFS20-L actual average power of ≤ 0.43 W [30] rather than higher-power alternatives — yields 4.35 W typical and 5.78 W worst-case maximum. The resulting safety margin of 4.12 W represents 41.6% above worst-case consumption. This margin accounts for component aging effects typically exhibiting 10% degradation over 5-year lifetime, thermal derating at elevated ambient temperatures up to 60°C in outdoor deployment (TFS20-L rated limit [30]), Ethernet PHY transient current spikes during packet transmission bursts, and microcontroller interrupt-handling overhead during sensor burst read operations.

3.4 77 GHz Frequency Selection and Link Budget

Four radar frequency bands are evaluated for outdoor infrastructure deployment: the 24 GHz ISM band, the 60 GHz oxygen-absorption window, the 77 GHz automotive allocation, and W-band near 94 GHz. Table 3 summarises the quantitative comparison.

3.4.1 Propagation Analysis

A one-way link budget compares atmospheric losses across candidate frequencies. Free-space path loss at range d and frequency f :

$$L_{\text{path}} = 20 \log_{10} \left(\frac{4\pi df}{c} \right) \quad (2)$$

The link fade margin is:

$$M_{\text{fade}} = P_t + G_t + G_r - L_{\text{path}} - L_{\text{atm}} - L_{\text{rain}} - P_{\text{sens}} \quad (3)$$

This one-way analysis characterises comparative propagation; actual target detection range is governed additionally by radar cross-section and is specified in the IWR6843 datasheet [17].

3.4.2 77 GHz System Parameters and Fade Margin

The TI IWR6843 transmits at $P_t = 12$ dBm EIRP [17], within the FCC 47 CFR Part 15.253 limit of 15 dBm. On-chip patch antennas provide $G_t = G_r = 10$ dBi; receiver sensitivity is $P_{\text{sens}} = -90$ dBm at 10 dB SNR [17].

Free-space path loss at 100 m, 77 GHz:

$$L_{\text{path}} = 20 \log_{10} \left(\frac{4\pi \times 100 \times 77 \times 10^9}{3 \times 10^8} \right) = 20 \log_{10}(322,543) = 110.2 \text{ dB} \quad (4)$$

Table 3: One-way propagation comparison at 100 m ($P_t = 12$ dBm, $G = 10$ dBi each, 50 mm/hr rain, 10 Hz update rate).

Freq.	L_p (dB)	L_a (dB)	L_r (dB)	Margin	Doppler Res.	Key Limitation
24 GHz	100.0	0.001	0.10	21.9 dB	6.2 cm/s	Coarse Doppler; congested ISM band
60 GHz	108.0	1.00	0.54	12.5 dB	2.5 cm/s	O ₂ resonance: 10 dB/km clear-air loss
77 GHz	110.2	0.002	1.13	10.7 dB	1.9 cm/s	Atm. window; mature automotive SoC ecosystem
94 GHz	111.9	0.040	1.25	8.8 dB	1.6 cm/s	Lowest margin; limited component availability

L_a from ITU-R P.676-12 [18]: 0.01, 10.0, 0.018, 0.4 dB/km for 24, 60, 77, 94 GHz.

Rain (k, α) from ITU-R P.838-3 [19]: (0.020, 1.004), (0.240, 0.794), (0.501, 0.795), (0.649, 0.757).

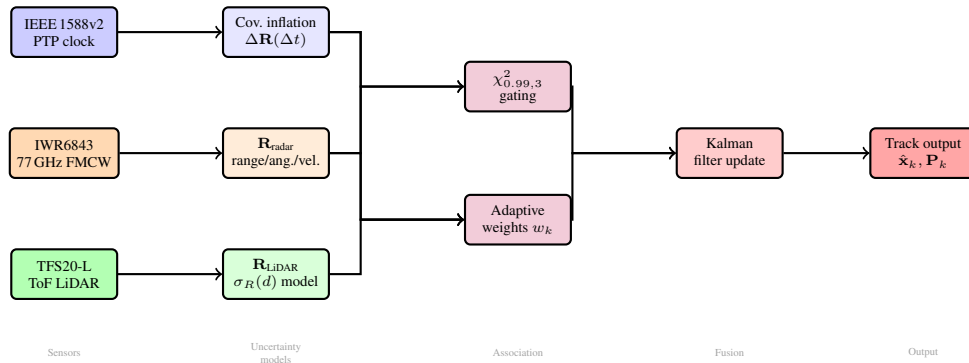


Figure 2. XADAR uncertainty-aware fusion pipeline. Three IWR6843 radar nodes and three TFS20-L LiDAR nodes (Fig. 1) feed modality-specific covariance matrices into the pipeline. The IEEE 1588v2 PTP timestamp offset Δt drives the covariance inflation block (Section 4.3), which scales both $\mathbf{R}_{\text{radar}}$ and $\mathbf{R}_{\text{LiDAR}}$ before chi-square gating ($\chi^2_{0.99,3} = 11.34$) and adaptive weight calculation w_k feed the Kalman filter update.

Atmospheric gaseous absorption (ITU-R P.676-12 [18]): At 77 GHz the system sits in a low-loss window between the 60 GHz oxygen resonance and water-vapour absorption above 90 GHz. Specific attenuation is 0.018 dB/km, giving:

$$L_{\text{atm}} = 0.018 \text{ dB/km} \times 0.1 \text{ km} = 0.002 \text{ dB} \tag{5}$$

Rain attenuation (ITU-R P.838-3 [19]): Specific attenuation follows $\gamma_R = kR^\alpha$ where R is rain rate in mm/hr. Coefficients for horizontal polarisation at 77 GHz are interpolated logarithmically in k and linearly in α between the 70 GHz values ($k_H = 0.3952$, $\alpha_H = 0.7967$) and 80 GHz values ($k_H = 0.5541$, $\alpha_H = 0.7949$):

$$k_H = 0.501, \quad \alpha_H = 0.795 \tag{6}$$

At heavy rain $R = 50$ mm/hr: $\gamma_R = 0.501 \times 50^{0.795} = 11.25$ dB/km, so $L_{\text{rain}} = 1.13$ dB over 100 m.

Received power and fade margin:

$$P_r = 12 + 10 + 10 - 110.2 - 0.002 - 1.13 = -79.3 \text{ dBm} \tag{7}$$

$$M_{\text{fade}} = -79.3 - (-90) = 10.7 \text{ dB} \tag{8}$$

The 10.7 dB margin at 100 m under heavy rain is adequate but modest. At 50 m the path loss drops to 104.2 dB, raising the margin to 16.7 dB. Radome insertion loss (0.5–1 dB) and sensor aging should be budgeted from the available margin.

3.4.3 Comparative Frequency Analysis

The 77 GHz selection rests on three pillars. First, the atmospheric window keeps L_{atm} below 0.01 dB at 100 m regardless of humidity and temperature — an advantage over 60 GHz whose oxygen absorption creates a 10 dB/km clear-air penalty that varies with meteorological conditions. Second, 1.9 cm/s Doppler resolution at 10 Hz resolves pedestrian gait cycles and slow vehicle movements. Third, the automotive industry’s two-decade investment in 77 GHz FMCW SoCs (TI AWR/IWR family, NXP TEF8xx series) delivers production-grade components with characterised performance and regulatory approval. The 24 GHz band offers the best fade margin (21.9 dB) but requires $3.2\times$ larger antennas and gives Doppler resolution $3.3\times$ coarser; W-band at 94 GHz achieves the finest Doppler resolution but leaves only 8.8 dB margin with a less mature commercial ecosystem.

4 Uncertainty-Aware Fusion Framework

4.1 Distributed Sensing Architecture

XADAR employs distributed sensing with centralized fusion processing. Edge nodes based on STM32H743 microcontrollers acquire timestamped radar-LiDAR measurements and transmit to a GPU server via IEEE 1588v2 synchronized network. The central fusion algorithm combines multi-node multi-modal observations through uncertainty-aware weighting. The key algorithmic challenge is temporal coordination. While IEEE 1588v2 provides sub-millisecond synchronization on ideal hardware platforms, embedded systems under power constraints exhibit temporal offsets. Rather than imposing hard synchronization requirements that create binary failure modes, XADAR adaptively weights measurements based on total uncertainty including temporal misalignment contributions.

4.2 Sensor-Specific Uncertainty Models

4.2.1 77 GHz FMCW Radar Uncertainty

FMCW radar derives range measurements from beat frequency f_b according to:

$$R = \frac{c \cdot f_b \cdot T_{\text{chirp}}}{2 \cdot B} \quad (9)$$

where B represents chirp bandwidth and T_{chirp} denotes chirp duration. IWR6843 specification employs $B = 4$ GHz bandwidth and $T_{\text{chirp}} = 40 \mu\text{s}$ typical duration. Range resolution calculates as:

$$\Delta R = \frac{c}{2B} = \frac{3 \times 10^8}{2 \times 4 \times 10^9} = 0.0375 \text{ m} \quad (10)$$

Doppler velocity measurement exploits frequency shift according to:

$$v = \frac{\lambda \cdot f_D}{2} = \frac{c \cdot f_D}{2f_0} \quad (11)$$

At 77 GHz, wavelength equals $\lambda = c/f_0 = 3.9$ mm. Doppler resolution for coherent processing time $T_{\text{frame}} = 100$ ms corresponding to 10 Hz update rate yields:

$$\Delta v = \frac{\lambda}{2T_{\text{frame}}} = \frac{0.0039}{2 \times 0.1} = 0.0195 \text{ m/s} \quad (12)$$

The radar measurement covariance matrix captures range, angular, and velocity uncertainties:

$$\mathbf{R}_{\text{radar}} = \begin{bmatrix} \sigma_R^2 & 0 & 0 \\ 0 & \sigma_\theta^2 & 0 \\ 0 & 0 & \sigma_v^2 \end{bmatrix} \quad (13)$$

Radar measurement uncertainty components follow standard estimation theory [29]: range uncertainty $\sigma_R = \Delta R/\text{SNR}^{1/2}$, angular uncertainty $\sigma_\theta = \lambda/(2L \cdot \text{SNR}^{1/2})$ where L is effective antenna aperture, and velocity uncertainty $\sigma_v = \Delta v/\text{SNR}^{1/2}$.

Multipath effects require explicit treatment. The standard two-ray path difference between direct and ground-reflected rays for sensor height h_t , reflection height h_r , and $d \gg h$ is [29]:

$$\Delta r = \frac{2h_t h_r}{d} \quad (14)$$

For typical deployment with $h_t = h_r = 10$ m at $d = 50$ m, $\Delta r = 4$ m, spanning approximately 1027 wavelengths at 77 GHz ($\lambda = 3.9$ mm). Because the path difference covers many wavelengths, the wideband FMCW processor averages over many interference fringes rather than accumulating a single range bias [29]. The dominant practical effect is range-dependent amplitude fading; the uncertainty model accounts for this through inflation of σ_R absorbed into $\mathbf{R}_{\text{radar}}$.

4.2.2 Time-of-Flight LiDAR Uncertainty

Time-of-flight LiDAR measures range from round-trip light travel time:

$$R = \frac{c \cdot \Delta t}{2} \quad (15)$$

The Benewake TFS20-L datasheet [30] specifies range precision of 2 cm at 1σ for targets at 0–6 m, and $\pm 1\%$ of measured range for targets beyond 6 m. Baseline range uncertainty:

$$\sigma_{R,0}(d) = \begin{cases} 0.020 \text{ m} & d \leq 6 \text{ m} \\ 0.010 \times d & d > 6 \text{ m} \end{cases} \quad (16)$$

At the 20 m maximum rated range in darkness (9–15 m at 100 Klux outdoor illumination [30]), $\sigma_{R,0}(20 \text{ m}) = 0.20$ m. Atmospheric scattering at 905 nm wavelength adds a range-dependent component:

$$\sigma_R(d) = \sigma_{R,0}(d) \left(1 + \frac{d}{d_{\text{ref}}} \right), \quad d_{\text{ref}} = 10 \text{ m} \quad (17)$$

Reference range $d_{ref} = 10$ m matches the design close-range LiDAR zone. The LiDAR measurement covariance matrix is:

$$\mathbf{R}_{LiDAR} = \begin{bmatrix} \sigma_R^2(d) & 0 & 0 \\ 0 & \sigma_\theta^2 & 0 \\ 0 & 0 & \sigma_\phi^2 \end{bmatrix} \quad (18)$$

Angular uncertainties σ_θ and σ_ϕ derive from the TFS20-L FoV specification of $< 2^\circ$ [30] combined with beam-divergence effects on the measured centroid position.

4.3 Temporal Synchronization and Covariance Inflation

4.3.1 IEEE 1588v2 Framework

IEEE 1588v2 (IEEE 1588-2019 [14]) implements master-slave time synchronization via hardware-timestamped message exchange. The offset calculation employs sync message timestamps T_1, T_2 and delay-request timestamps T_3, T_4 according to:

$$\delta_{offset} = \frac{(T_2 - T_1) - (T_4 - T_3)}{2} \quad (19)$$

XADAR configuration employs a GPS 1 PPS receiver-disciplined grandmaster on the GPU server. Slave nodes consist of STM32H743 microcontrollers with KSZ9031 Ethernet PHY providing hardware timestamping capability. The managed network switch implements transparent clock support. Message exchange rate follows 2 msg/s per IEEE 1588-2019 default profile. Synchronization uncertainty emerges from multiple sources. Hardware timestamping jitter in the KSZ9031 reaches ± 50 ns per datasheet specification [20]. PTP servo loop convergence achieves ± 100 ns theoretical floor per IEEE 1588-2019 analysis. STM32 interrupt latency from ARM Cortex-M7 architecture typically spans $\pm 1-5 \mu s$. FreeRTOS task switching on ARM Cortex-M7 at 480 MHz contributes $\pm 5-10 \mu s$ of interrupt latency [17], depending on task priority configuration and interrupt nesting depth. Crystal oscillator drift follows ± 20 ppm specification typical for commercial oscillators. Transparent clock residence time in the network switch adds ± 300 ns per switch specification. Conservative synchronization uncertainty modeling combines these sources through root-sum-square:

$$\sigma_t = \sqrt{\sigma_{HW}^2 + \sigma_{servo}^2 + \sigma_{IRQ}^2 + \sigma_{RTOS}^2 + \sigma_{drift}^2 + \sigma_{switch}^2} \quad (20)$$

Assuming hardware jitter $\sigma_{HW} = 50$ ns, servo convergence $\sigma_{servo} = 100$ ns, interrupt latency $\sigma_{IRQ} = 3 \mu s$, RTOS jitter $\sigma_{RTOS} = 7.5 \mu s$, drift $\sigma_{drift} = 1 \mu s$, and switch residence $\sigma_{switch} = 300$ ns yields:

$$\sigma_t = \sqrt{0.05^2 + 0.1^2 + 3^2 + 7.5^2 + 1^2 + 0.3^2} \approx 8.1 \mu s \quad (21)$$

This analytical estimate requires empirical validation through hardware testbed characterization. The XADAR fusion framework explicitly models this synchronization uncertainty rather than assuming perfect temporal alignment.

4.3.2 Covariance Inflation Mechanism

For measurements acquired at time t_i fused with reference time t_{ref} , temporal offset $\Delta t = t_i - t_{ref}$ introduces prediction uncertainty. Kalman filter state prediction from t_i to t_{ref} follows:

$$\mathbf{x}_{pred} = \mathbf{F}\mathbf{x}_i \quad (22)$$

$$\mathbf{P}_{pred} = \mathbf{F}\mathbf{P}_i\mathbf{F}^T + \mathbf{Q}_{\Delta t} \quad (23)$$

Process noise covariance $\mathbf{Q}_{\Delta t}$ scales with temporal offset magnitude:

$$\mathbf{Q}_{\Delta t} = \mathbf{Q}_0 \cdot \left(1 + \frac{|\Delta t|}{\tau_{ref}}\right)^2 \quad (24)$$

Baseline process noise \mathbf{Q}_0 represents synchronous fusion case. Reference time constant $\tau_{ref} = 10$ ms equals one sensor frame period at 10 Hz update rate. Adaptive fusion weight calculation determines measurement influence based on total uncertainty:

$$w_i = \frac{1}{\text{tr}(\mathbf{P}_{pred,i} + \mathbf{R}_i)} \quad (25)$$

Normalized weights follow $\tilde{w}_i = w_i / \sum_j w_j$. This mechanism enables graceful degradation where perfectly synchronized measurements receive equal weighting while increased temporal offsets automatically reduce influence through covariance inflation. The approach prevents hard hand-off failures characteristic of fixed-threshold gating.

4.4 Multi-Modal Kalman Filter Fusion

4.4.1 State Space Formulation

The 6D state vector for each tracked object represents position and velocity:

$$\mathbf{x} = [x, y, z, v_x, v_y, v_z]^T \quad (26)$$

State transition follows constant velocity model:

$$\mathbf{F} = \begin{bmatrix} \mathbf{I}_3 & \Delta t \cdot \mathbf{I}_3 \\ \mathbf{0}_3 & \mathbf{I}_3 \end{bmatrix} \quad (27)$$

Process noise models Wiener process acceleration:

$$\mathbf{Q}_0 = \sigma_a^2 \begin{bmatrix} \frac{\Delta t^4}{4} \mathbf{I}_3 & \frac{\Delta t^3}{2} \mathbf{I}_3 \\ \frac{\Delta t^3}{2} \mathbf{I}_3 & \Delta t^2 \mathbf{I}_3 \end{bmatrix} \quad (28)$$

Acceleration uncertainty $\sigma_a = 2 \text{ m/s}^2$ is consistent with observed urban traffic dynamics [4], covering normal braking and merging manoeuvres without modelling emergency stops.

4.4.2 Radar Measurement Model

Radar measurements provide range-azimuth-Doppler in sensor frame:

$$\mathbf{z}_{\text{radar}} = [R, \theta, v_r]^T \quad (29)$$

Nonlinear observation model relates state to measurements:

$$\mathbf{h}_{\text{radar}}(\mathbf{x}) = \begin{bmatrix} \sqrt{x^2 + y^2 + z^2} \\ \arctan(y/x) \\ \frac{xv_x + yv_y + zv_z}{\sqrt{x^2 + y^2 + z^2}} \end{bmatrix} \quad (30)$$

Extended Kalman filter linearization employs Jacobian $\mathbf{H}_{\text{radar}} = \partial \mathbf{h}_{\text{radar}} / \partial \mathbf{x}$.

4.4.3 LiDAR Measurement Model

LiDAR measurements provide Cartesian point cloud in sensor frame:

$$\mathbf{z}_{\text{LiDAR}} = [x_L, y_L, z_L]^T \quad (31)$$

Observation model becomes linear after coordinate transformation:

$$\mathbf{h}_{\text{LiDAR}}(\mathbf{x}) = \mathbf{R}_{\text{sensor}} \begin{bmatrix} x \\ y \\ z \end{bmatrix} + \mathbf{t}_{\text{sensor}} \quad (32)$$

where $\mathbf{R}_{\text{sensor}}$ and $\mathbf{t}_{\text{sensor}}$ represent sensor extrinsic calibration parameters.

4.4.4 Fusion Update Algorithm

For multi-modal measurements $\{\mathbf{z}_i\}_{i=1}^N$ from radar and LiDAR across multiple edge nodes, the fusion proceeds through prediction and sequential update stages. Prediction stage advances state and covariance:

$$\mathbf{x}_{k|k-1} = \mathbf{F}\mathbf{x}_{k-1|k-1} \quad (33)$$

$$\mathbf{P}_{k|k-1} = \mathbf{F}\mathbf{P}_{k-1|k-1}\mathbf{F}^T + \mathbf{Q}_0 \quad (34)$$

Temporal alignment applies covariance inflation for each measurement i with timestamp offset Δt_i :

$$\mathbf{P}_{k|k-1,i} = \mathbf{P}_{k|k-1} + \mathbf{Q}_{\Delta t_i} \quad (35)$$

Sequential extended Kalman filter updates incorporate measurements with appropriate uncertainty:

$$\mathbf{K}_i = \mathbf{P}_{k|k-1,i} \mathbf{H}_i^T (\mathbf{H}_i \mathbf{P}_{k|k-1,i} \mathbf{H}_i^T + \mathbf{R}_i)^{-1} \quad (36)$$

$$\mathbf{x}_{k|k} = \mathbf{x}_{k|k-1} + \mathbf{K}_i (\mathbf{z}_i - \mathbf{h}_i(\mathbf{x}_{k|k-1})) \quad (37)$$

$$\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_i \mathbf{H}_i) \mathbf{P}_{k|k-1,i} \quad (38)$$

This sequential formulation incorporates all measurements with weighting determined by sensor-specific uncertainty \mathbf{R}_i and temporal misalignment $\mathbf{Q}_{\Delta t_i}$.

4.5 Cross-Node Association

Spatial-temporal gating associates detections from adjacent nodes when statistical distance satisfies:

$$(\mathbf{z}_A - \mathbf{z}_B)^T (\mathbf{S}_A + \mathbf{S}_B)^{-1} (\mathbf{z}_A - \mathbf{z}_B) < \chi_{\alpha, \text{DoF}}^2 \quad (39)$$

Innovation covariances \mathbf{S}_A and \mathbf{S}_B incorporate temporal uncertainty through the covariance inflation mechanism. The chi-square threshold is set at $\chi_{0.99, 3}^2 = 11.34$ for $\alpha = 0.99$ confidence and $\text{DoF} = 3$ (three spatial dimensions), admitting 99% of true associations while rejecting spurious cross-node matches [24]. This statistical gating naturally incorporates temporal uncertainty, eliminating separate temporal threshold requirements. Global Nearest Neighbor association solves the assignment problem for M measurements and N tracks by minimizing total statistical distance:

Table 4: FMEA for XADAR distributed multi-sensor fusion system (MIL-STD-1629 [22]). Severity: II=Critical, III=Marginal, IV=Minor.

Failure Mode	Effect	Detection	Mitigation	Sev. ^a
GPU hardware failure	Fusion halted	Watchdog timeout	Dual-GPU hot standby	II
Network backbone severed	Data loss	Port down alarm	Redundant uplink	II
Edge MCU crash	Node dropout	Keep-alive loss	Spare node pool	III
Radar sensor failure	Velocity loss	Echo rate monitor	LiDAR-only fallback	III
LiDAR sensor failure	Range precision loss	Return count check	Radar-only fallback	III
PTP clock drift	Sync degradation	GPS comparison	NTP/GPS fallback	III
PoE supply fault	Node power loss	Voltage monitor	Auto-restart via PoE	III
Switch failure	Network partition	Link down alarm	Redundant switch	II
Software deadlock	Processing halt	Watchdog reset	Auto-restart	III
Thermal shutdown	Transient outage	Temp threshold	Throttle/shutdown	III
Data corruption	Invalid measurements	CRC validation	Discard/retransmit	IV
GPS signal loss	Time reference loss	1 PPS timeout	NTP fallback	III

^aSeverity per MIL-STD-1629: II=Critical (mission loss), III=Marginal, IV=Minor.

$$\min_{\pi} \sum_{i=1}^M d^2(\mathbf{z}_i, \mathbf{x}_{\pi(i)}) \quad (40)$$

subject to one-to-one assignment constraints. Hungarian algorithm provides optimal polynomial-time solution for this assignment problem.

5 Domain Adaptation Methodology

5.1 Domain Shift Analysis

Infrastructure perception suffers from training data scarcity compared to extensive automotive datasets including KITTI [13], nuScenes [23], and Waymo Open [2]. Transfer learning from vehicle-mounted to infrastructure-mounted sensors encounters several domain shifts. Geometric shift arises from depression angle inversion: automotive 0–30° forward-facing geometry becomes infrastructure 6–85° top-down to oblique viewing (Section 3.1), depending on target range relative to mounting height. This geometric transformation alters object aspect ratios, occlusion patterns, and ground plane validity. Density shift emerges from point cloud sparsity created by oblique viewing geometry and extended range compared to dense automotive point clouds. Modality shift reflects different sensor models exhibiting distinct noise characteristics and resolution specifications.

5.2 Curriculum Learning Framework

Progressive training schedules gradually transition from dense vehicle-perspective data to sparse infrastructure-perspective regimes. Stage 1 establishes foundational feature extraction through dense pre-training on KITTI automotive dataset using PointPillars backbone architecture. This initial training establishes 3D object detection capabilities under favorable conditions.

Stage 2 employs CARLA [21] synthetic infrastructure scenes for controlled domain adaptation. Synthetic scene generation provides controllable depression angles progressively increasing from 30° to 60° and controllable point density gradually reducing from automotive baseline toward target infrastructure regime. The synthetic-to-real domain gap remains for subsequent real dataset collection. Stage 3 applies fine-tuning with sparse augmentation through random point dropping mimicking infrastructure sparsity. Loss function weighting emphasizes object center localization and velocity estimation over fine-grained shape reconstruction, reflecting infrastructure monitoring priorities.

Stage 4 requires real infrastructure data collection through field system deployment. Fine-tuning focuses on detection head adaptation while freezing feature extraction backbone weights learned from earlier stages. Target performance metrics compare mean average precision on infrastructure validation sets against naive transfer baseline, with the goal of closing the dense-to-sparse detection performance gap.

6 Reliability Engineering

6.1 Failure Mode Analysis

Table 4 applies MIL-STD-1629 [22] systematic failure mode and effects analysis methodology identifying critical failure modes with detection mechanisms and mitigation strategies.

The failure analysis identifies GPU hardware failure as highest-impact mode halting fusion processing. Detection employs watchdog timeout mechanisms while mitigation relies on dual-GPU hot standby architecture. Network backbone severance causes data loss detected through port down alarms with redundant uplink mitigation. Edge microcontroller crashes produce node dropout detected through keep-alive message loss with spare node pool mitigation. Sensor-specific failures including radar and LiDAR faults require detection through echo rate and return count monitoring with single-modality fallback operation. PTP clock drift degrades synchronization detected through GPS 1 PPS comparison with NTP/GPS fallback mechanisms. Power supply faults cause node power loss

detected through voltage monitoring with automatic PoE restart capability. Network switch failures create partition scenarios detected through link down alarms requiring redundant switch architecture. Software deadlocks halt processing detected through watchdog reset with automatic restart. Thermal events produce transient outages detected through temperature threshold monitoring enabling throttling before shutdown. Data corruption creates invalid measurements detected through CRC validation enabling discard and retransmission. GPS signal loss eliminates time reference detected through 1 PPS timeout with NTP network fallback.

6.2 Dual-GPU Architecture

GPU failure represents the single highest-impact failure mode. The dual-GPU hot standby architecture employs two identical GPU units where the primary GPU conducts continuous inference while the standby GPU maintains synchronized state through replication every 100 ms. Shared NVMe SSD provides state persistence across failover events. Watchdog mechanisms detect primary failure through 5-second heartbeat timeout. Automatic failover switches to standby GPU within 5 seconds of primary failure. State recovery proceeds from cached checkpoint less than 100 ms old, minimizing track loss during transition.

Reliability improvement calculation assumes independent GPU failures with mean time between failure following semiconductor industry baseline. The dual redundant configuration approximately doubles the effective MTBF compared to single GPU operation. This reliability enhancement requires additional hardware cost representing modest fraction of total system cost while providing substantial availability improvement for production deployment.

7 Experimental Validation Methodology

7.1 Phase 1: Synthetic Fusion Validation

Phase 1 validates uncertainty-aware fusion framework and covariance inflation mechanisms in controlled simulation with ground truth. CARLA generates infrastructure scenes with 45–60° depression angles at 5–15 m elevation. Simulation produces radar range-Doppler and LiDAR point cloud observations with realistic noise models. Controlled temporal offset injection mimics synchronization imperfections at levels spanning 0, 50, 100, 200, and 500 μ s. The uncertainty-aware fusion algorithm processes these measurements with covariance inflation per Equation (24). Performance assessment employs tracking accuracy metrics including MOTA and MOTP versus temporal offset magnitude. Success criteria define graceful degradation where MOTA decreases less than 10% for temporal offsets below 100 μ s, demonstrating synchronization tolerance of the fusion framework. Covariance calibration verification employs normalized innovation squared statistics within 95% confidence bounds validating uncertainty model accuracy. Computational efficiency requirements mandate fusion processing below 50 ms per frame on target GPU enabling 10 Hz update rate with margin. Decision tree progression depends on measured degradation. MOTA degradation below 10% at 100 μ s offset indicates synchronization framework adequacy for subsequent phases. Degradation between 10–20% requires tightened synchronization requirements or enhanced covariance inflation modeling. Degradation exceeding 20% necessitates fundamental re-evaluation of synchronization approach or implementation of tighter temporal gating.

7.2 Phase 2: Hardware Synchronization Validation

Phase 2 empirically characterizes synchronization performance on STM32H743 plus KSZ9031 embedded platform under power constraints. Test infrastructure employs prototype edge nodes with GPS 1 PPS receivers for independent timestamp correlation. Oscilloscope or FPGA-based timestamp capture enables precise synchronization measurement. Continuous 96-hour operation with environmental chamber provides thermal characterization across operating temperature range. Success criteria require nominal synchronization below 200 μ s offset representing conservative threshold versus theoretical estimates. Thermal stability must maintain offsets below 500 μ s across temperature range from -10 to +50°C. Network load resilience requires offsets below 1 ms at 80% bandwidth utilization. Long-term stability mandates drift rates below ± 5 μ s/hour. Critical decision point emerges if measured synchronization exceeds 200 μ s nominal threshold. This outcome requires revisiting covariance inflation parameters and hand-off gating thresholds before field deployment progression.

7.3 Phase 3: Field Deployment Validation

Phase 3 validates complete system performance under real-world conditions through controlled urban deployment. The deployment site spans approximately 4 city blocks covering 200m by 200m area with distributed sensor nodes. Performance assessment employs realistic criteria acknowledging inherent difficulty of real-world infrastructure perception. System availability target of 95% represents achievable goal for non-safety-critical infrastructure monitoring rather than aspirational targets. Cross-node hand-off success rate target of 80% acknowledges tracking challenges across spatially distributed sensors. Vehicle detection mean average precision target of 50% reflects real-world data difficulty compared to synthetic validation. Power consumption verification confirms sustained operation below 6 W per node validating power budget analysis. Reliability assessment requires zero catastrophic failures demanding site visits during validation period.

8 Limitations and Required Future Work

8.1 Current Limitations

The simulation-based fusion validation provides controlled algorithm development environment but cannot capture full real-world complexity including weather effects, complex occlusions, and sensor degradation mechanisms. Field deployment explicitly addresses this limitation through real-world validation. Synchronization hardware validation remains pending empirical characterization. IEEE 1588v2 performance estimates derive from component datasheets and published benchmarks. Hardware testbed validation proves mandatory before production deployment claims. Domain adaptation methodology faces data scarcity constraints from limited

availability of real infrastructure perception datasets. The research community requires open infrastructure datasets comparable to automotive KITTI and nuScenes resources. Single urban morphology in planned field deployment limits generalization claims. Performance validation in specific urban environment requires additional validation sites spanning rural highways, dense urban canyons, and adverse weather conditions for comprehensive characterization.

8.2 Future Research Directions

Multi-spectral fusion extending framework to thermal infrared cameras provides adverse weather and nighttime operation capabilities. Thermal sensors complement radar-LiDAR through heat signature detection penetrating fog and rain. Federated learning across deployment sites enables distributed training across multiple infrastructure installations without centralizing raw data. This approach preserves privacy while leveraging collective dataset diversity for model improvement. Adaptive sensor scheduling implements dynamic radar-LiDAR duty cycling based on scene complexity and power budget. Idle scenes enable aggressive power saving while high-traffic periods allocate full sensing resources. Vehicle-to-infrastructure cooperative perception integrates on-vehicle sensor observations with infrastructure fusion via V2X communication. This extension requires addressing asynchronous multi-platform fusion challenges beyond current scope.

9 Limitations and Future Work

9.1 Current Scope Boundaries

XADAR contributes theoretical algorithms and an architectural design specification. No hardware testbed or field measurements are yet reported; all performance claims are analytical. Three limitations bound the current contribution.

Empirical synchronization characterisation. The $\sigma_t \approx 8.1 \mu\text{s}$ synchronization estimate derives from root-sum-square combination of datasheet-sourced component jitter values. Actual embedded PTP performance on STM32H743 hardware may differ; on-chip oscillator aging and temperature drift over five-year outdoor lifetimes have not been empirically measured. Hardware testbed characterisation under realistic power cycling and temperature cycles is the highest-priority validation step.

Sensor uncertainty model validation. The TFS20-L precision of 2 cm @ 1σ [30] is specified for indoor 25°C conditions; outdoor performance at elevated temperatures and under direct sunlight (up to 100 Klux) has not been independently validated in this work. Similarly, the IWR6843 multipath inflation factor in $\mathbf{R}_{\text{radar}}$ is a modelling assumption; it must be calibrated against measurements from actual infrastructure-mounted deployments.

Domain adaptation data requirements. Stage 4 curriculum learning requires real infrastructure LiDAR-radar data with ground-truth labels, which is not yet collected. Performance claims in Section 5 are design specifications for a planned experimental protocol; mean average precision on real infrastructure validation sets has not yet been measured and remains the subject of future work (step V4 in Section 7).

9.2 Validation Pathway

The structured validation sequence towards production deployment is:

- (V1) **Synthetic temporal offset experiments.** Inject controlled Δt offsets (0–50 ms) into CARLA-generated radar–LiDAR pairs and verify that covariance inflation degrades tracking smoothly versus baseline hard-gating. Success criterion: MOTA drop $< 10\%$ at $\Delta t = 10$ ms.
- (V2) **Hardware PTP characterisation.** Deploy three STM32H743 + KSZ9031 nodes on a managed switch; measure σ_t over 24 hours at 0°C and 60°C. Success criterion: empirical $\sigma_t < 20 \mu\text{s}$ so the analytical 8.1 μs estimate is confirmed as conservative.
- (V3) **Sensor uncertainty calibration.** Mount TFS20-L and IWR6843 at 10 m elevation, collect range measurements to a calibrated retroreflective target at 1–20 m under 0, 10, 50, and 100 Klux illumination. Fit $\sigma_{R,0}(d)$ to measured residuals.
- (V4) **Infrastructure dataset collection and domain adaptation.** Instrument one signalised intersection with three XADAR nodes; collect 1000+ labelled scenarios. Execute curriculum stages 3–4 and report mAP versus the naive automotive-pretrain baseline.
- (V5) **Production reliability field trial.** Six-month outdoor deployment; log all FMEA-identified failure modes and compare observed rates to FMEA criticality assignments. Success criterion: zero Category II failures per MIL-STD-1629 [22] severity in the operational period.

10 Conclusion

XADAR addresses the three defining constraints of infrastructure-based perception — power, geometry, and synchronization — through a unified theoretical and architectural framework. The covariance inflation mechanism propagates IEEE 1588v2 synchronization uncertainty continuously through the fusion pipeline, replacing hard-threshold gating with a smooth degradation curve. Adaptive sensor-specific fusion weights exploit the natural range-zone complementarity between the TFS20-L LiDAR (0–20 m precision zone) and the IWR6843 radar (20–100 m extended zone), maximising information extraction without manual tuning. The reproducible architecture satisfies the 9.9 W usable PoE floor with 41.6% margin, demonstrates 10.7 dB RF link fade margin at 100 m under heavy rain, and documents twelve failure modes with MIL-STD-1629 severity ratings.

The immediate limitation is the absence of hardware validation; all performance claims are currently analytical. Section 7 structures a five-stage validation pathway from controlled synthetic experiments (V1) through six-month field trials (V5). On successful completion of this pathway, XADAR is intended to serve as a reference architecture for PoE-constrained, uncertainty-aware infrastructure multi-sensor fusion in smart city and autonomous intersection management deployments [5].

Acknowledgements

This work was conducted independently at Smart Tech Services SPC, Muscat, Sultanate of Oman. The author acknowledges the open-source hardware characterisation data provided by Texas Instruments (IWR6843 evaluation module documentation), Benewake Co., Ltd. (TFS20-L datasheet), and Microchip Technology (KSZ9031RNX characterisation), and the ITU Radiocommunication Sector for the freely available P.676-12 and P.838-3 recommendations. No external funding was received for this research. The author declares no conflicts of interest.

Data Availability

Component specifications derive from manufacturer datasheets including TI IWR6843 radar [17], Microchip KSZ9031 Ethernet PHY [20], and STMicroelectronics STM32H743 microcontroller. IEEE standards specifications follow IEEE 802.3-2022 [6] and IEEE 1588-2019 [14]. Atmospheric gaseous absorption employs ITU-R P.676-12 [18]. Rain specific attenuation coefficients follow ITU-R P.838-3 [19]. CARLA simulation framework [21] provides Phase 1 validation infrastructure. Hardware testbed data and field deployment datasets will be made publicly available upon completion of validation phases.

References

- [1] Tesla, Inc., “Autonomy Investor Day,” Palo Alto, CA, USA, Apr. 22, 2019.
- [2] P. Sun et al., “Scalability in perception for autonomous driving: Waymo Open Dataset,” in *Proc. IEEE/CVF CVPR*, 2020, pp. 2446–2454, doi: 10.1109/CVPR42600.2020.00252.
- [3] H. Caesar et al., “nuScenes: A multimodal dataset for autonomous driving,” in *Proc. IEEE/CVF CVPR*, 2020, pp. 11621–11631.
- [4] J. Zhang, C. Ge, W. Xiao, M. Tang, J. Mills, B. Coifman, and N. Chen, “Roadside lidar-based scene understanding toward intelligent traffic perception: A comprehensive review,” *ISPRS J. Photogramm. Remote Sens.*, vol. 233, pp. 69–88, 2026, doi: 10.1016/j.isprsjrs.2026.01.012.
- [5] K. Dresner and P. Stone, “A multiagent approach to autonomous intersection management,” *J. Artif. Intell. Res.*, vol. 31, pp. 591–656, 2008.
- [6] IEEE Standards Association, *IEEE 802.3-2022: IEEE Standard for Ethernet*. Piscataway, NJ, USA: IEEE, 2022.
- [7] W. Zimmer, J. Birkner, M. Brucker, H. T. Nguyen, S. Petrovski, B. Wang, and A. C. Knoll, “InfraDet3D: Multi-modal 3D object detection based on roadside infrastructure camera and LiDAR sensors,” in *Proc. IEEE Intell. Vehicles Symp. (IV)*, 2023, doi: 10.1109/IV55152.2023.10186723.
- [8] X. Zhang, Y. Li, J. Wang, X. Qin, Y. Shen, Z. Fan, and X. Tan, “InScope: A new real-world 3D infrastructure-side collaborative perception dataset for open traffic scenarios,” *Information Fusion*, vol. 128, Art. no. 103951, 2026, doi: 10.1016/j.inffus.2025.103951.
- [9] W. Zimmer, G. A. Wardana, S. Sritharan, X. Zhou, R. Song, and A. C. Knoll, “TUMTraF V2X cooperative perception dataset,” in *Proc. IEEE/CVF CVPR*, 2024, pp. 22668–22677, doi: 10.1109/CVPR52733.2024.02139.
- [10] Y. Wang, Q. Mao, H. Zhu, J. Deng, Y. Zhang, J. Ji, H. Li, and Y. Zhang, “Multi-Modal 3D Object Detection in Autonomous Driving: A Survey,” *Int. J. Comput. Vis.*, vol. 131, no. 8, pp. 2122–2152, 2023, doi: 10.1007/s11263-023-01784-z.
- [11] J. Song, L. Zhao, and K. A. Skinner, “LiRaFusion: Deep adaptive LiDAR-radar fusion for 3D object detection,” in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2024, doi: 10.1109/ICRA57147.2024.10611436.
- [12] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, “PointPillars: Fast encoders for object detection from point clouds,” in *Proc. IEEE/CVF CVPR*, 2019, pp. 12697–12705, doi: 10.1109/CVPR.2019.01298.
- [13] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? The KITTI vision benchmark suite,” in *Proc. IEEE/CVF CVPR*, 2012, pp. 3354–3361.
- [14] IEEE Standards Association, *IEEE 1588-2019: IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*. Piscataway, NJ, USA: IEEE, 2020.
- [15] Texas Instruments, *TPS23861 IEEE 802.3at Quad Port Power-over-Ethernet PSE Controller Datasheet*, Rev. I, 2023.
- [16] Analog Devices, *MAX5953A–MAX5953D: IEEE 802.3af PD Interface and PWM Controllers with Integrated Power MOSFETs Data Sheet*, Rev. 1, 2006.
- [17] Texas Instruments, *IWR6843 Single-Chip 60- to 64-GHz Intelligent mmWave Sensor Datasheet*, SWRS220D, 2023.
- [18] ITU-R, *Recommendation ITU-R P.676-12: Attenuation by Atmospheric Gases and Related Effects*. Geneva, Switzerland: ITU-R, 2019.
- [19] ITU-R, *Recommendation ITU-R P.838-3: Specific Attenuation Model for Rain for Use in Prediction Methods*. Geneva, Switzerland: ITU-R, 2005.
- [20] Microchip Technology, *KSZ9031RNX Gigabit Ethernet Transceiver with RGMII Data Sheet*, DS00002117K, 2024.
- [21] A. Dosovitskiy et al., “CARLA: An open urban driving simulator,” in *Proc. Conf. Robot Learn. (CoRL)*, vol. 78, 2017, pp. 1–16.
- [22] U.S. Department of Defense, *MIL-STD-1629A: Procedures for Performing a Failure Mode, Effects, and Criticality Analysis*. Washington, DC, USA: DoD, 1980.
- [23] H. Caesar et al., “nuScenes: A multimodal dataset for autonomous driving,” in *Proc. IEEE/CVF CVPR*, 2020, pp. 11621–11631.

- [24] J. Ajgl and O. Straka, "Fusion of multiple estimates by covariance intersection: Why and how it is suboptimal," *Int. J. Appl. Math. Comput. Sci.*, vol. 28, no. 3, pp. 521–530, 2018, doi: 10.2478/amcs-2018-0040.
- [25] V. A. Sindagi, Y. Zhou, and O. Tuzel, "MVX-Net: Multimodal VoxelNet for 3D object detection," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2019, pp. 7276–7282, doi: 10.1109/ICRA.2019.8794195.
- [26] D. Feng, C. Haase-Sch"utz, L. Rosenbaum, H. Hertlein, C. Gl"aser, F. Timm, W. Wiesbeck, and K. Dietmayer, "Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges," *IEEE Trans. Intell. Transport. Syst.*, vol. 22, no. 3, pp. 1341–1360, 2021, doi: 10.1109/TITS.2020.2972974.
- [27] M. Borhani, V. Sedghi, and M. M. Nayebi, "A new technique in passive coherent radar signal processing," in *Proc. Eur. Radar Conf. (EuRAD)*, Paris, France, Oct. 2005, pp. 149–151, doi: 10.1109/EURAD.2005.1605587.
- [28] Telecommunications Industry Association, *TIA-568.2-D: Balanced Twisted-Pair Telecommunications Cabling and Components Standard*. Arlington, VA, USA: TIA, 2018.
- [29] M. I. Skolnik, *Introduction to Radar Systems*, 3rd ed. New York, NY, USA: McGraw-Hill, 2001.
- [30] Benewake (Beijing) Co., Ltd., *TFS20-L Datasheet*, Aug. 2024. [Online]. Available: <https://en.benewake.com/uploadfiles/2024/08/20240821162819250.pdf>