# TECHBOOST

# Multi-Source Data Fusion in Intelligent Transportation system

## Technology demo report

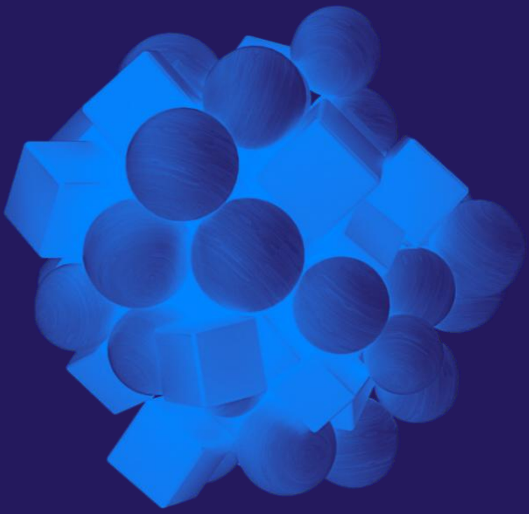

Inspired by Conveqs Oy

Elinkeino-, liikenne- ja ympäristökeskus
Työ- ja elinkeinoministeriö Arbets- och näringsministeriet
Tampereen ammattikorkeakoulu
Aalto-yliopisto Aalto-universitetet Aalto University
Metropolia
HAMK Hämeen ammattikorkeakoulu
OAMK Oulun ammattikorkeakoulu
Tampereen yliopisto

# Multi-Source Data Fusion in Intelligent Transportation System

Aalto University
18.12.2025
TECHBOOST Project Report
Aleksi Pippuri

# Contents

# 1. Introduction

Intelligent transportation systems (ITS) are increasingly reliant on accurate, timely, and robust situational awareness to improve road safety, traffic efficiency, and environmental sustainability. While roadside sensing infrastructure such as radar plays a central role in current adaptive traffic management systems, its performance can be limited by occlusions, sensing geometry, and environmental conditions. At the same time, connected vehicles (CVs) equipped with advanced onboard sensors provide high-quality local perception but only for a limited subset of traffic participants.

This project addresses these limitations by developing and validating a lightweight, scalable multi-source data fusion method that combines vehicle detections from a connected vehicle and roadside radar into a unified representation of traffic state. By operating at the decision level and utilizing only essential object-level information, the proposed approach enables collaborative perception without imposing excessive communication or computational requirements.

The work was carried out using real-world data collected in the Jätkäsaari Mobility Lab, leveraging research vehicles from Aalto University and existing roadside infrastructure seen in figure 1. The resulting system demonstrates how current traffic management assets can be augmented with connected vehicle sensing to enhance situational awareness. The outcomes of this project demonstrate the feasibility of collaborative localization in live urban traffic conditions and establish a foundation for further quantitative evaluation and future deployment in intelligent transportation systems.

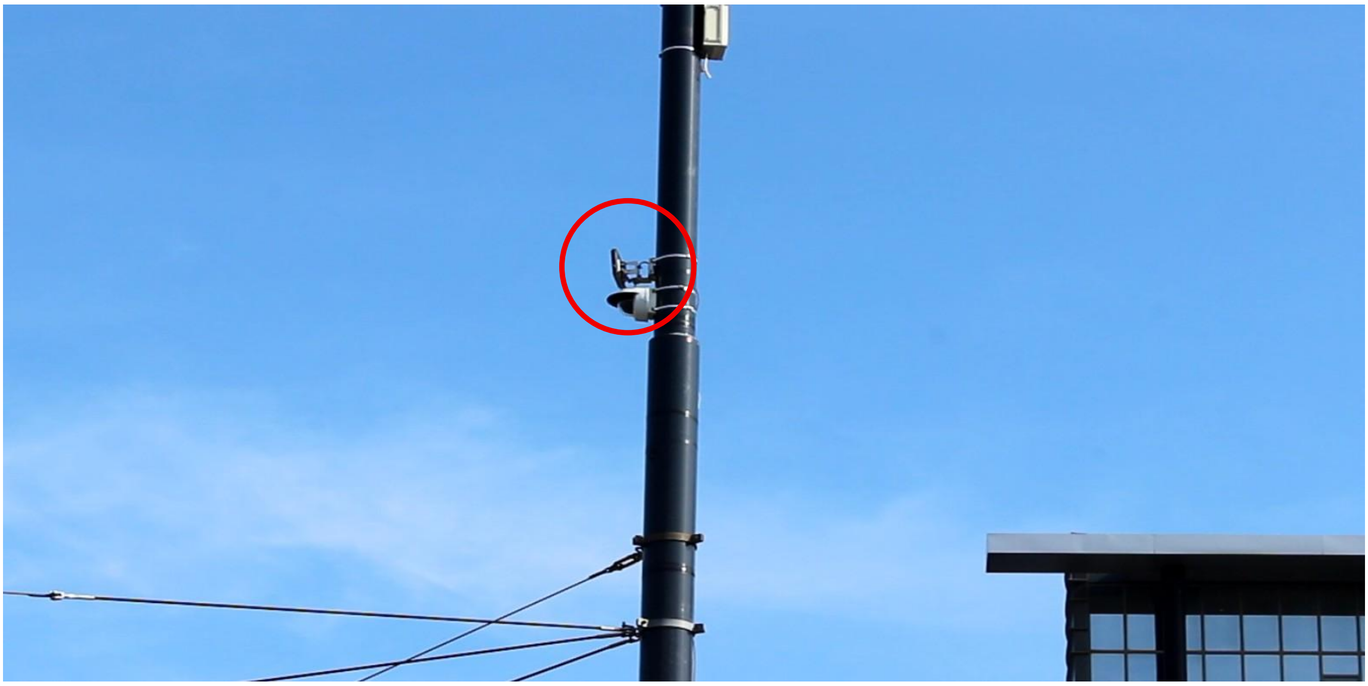

*Figure 1: Roadside radar above a camera in the intersection between Mechelininkatu and Jätkäsaarenlaituri.*

# 2. Project Overview

## 2.1 Background

Intelligent transportation systems (ITS) and closely related connected vehicles (CV) aim to enhance the future of traffic management by increasing road-safety and traffic flow as well as reducing emissions. These three goals are to be achieved by multiple technological advancements in the future. In this project we developed and tested a lightweight and scalable solution for data fusion between a CV and the traffic management system. This collaboration perception method can fuse together vehicle detections from a CV and similar detections from a road-side radar. By enabling multi-source vehicle detections, the issues that can affect stationary sensors as a part of adaptive traffic management systems can be mitigated. These include but are not limited to occlusion, inaccuracies and drift.

The resulting data can then be used in conjunction with other data sources such as public transport GPS data and traditional loop detectors to provide traffic management systems with high-precision situational awareness compared to the systems in wide use today. More precise understanding of the state of traffic allows the traffic management system to control traffic more optimally and thus affect the three underlying goals positively.

## 2.2 Objectives & Scope

The goal of this project is to develop and test together with the collaboration company an application layer method for fusing vehicle detections made by a CV and road-side radar. The method was tested with real-world data gathered at the Jätkäsaari Mobility Lab using research vehicles provided by Aalto University Autonomous Mobility Lab and a road-side radar located in Jätkäsaari by Conveqs.

The system can be divided into data gathering and data fusion. The former comprises of the CV which utilizes a lidar operating in ROS 2 (Robot Operating System) and open-source based Python programming for further data processing, as well as the radar side which includes proprietary programming and open-source based programming in data processing. The data fusion method was developed in house using open-source tools and is published as open-source.

The system is designed to be lightweight and scalable, but in this project the aim was to collect real-world data along a single road in Jätkäsaari. The collected test data comprised of 15 runs, each consisting of CV, radar and target vehicle data. This provided enough data to deem the data fusion method as viable for further development and testing.

## 2.3 Key Technologies & Innovations

Core technologies:

- Road-side radar: Smartmicro UMRR-11 TYPE 44 STOP+ADVANCE 4D/HD
- CV lidar and sensing seen in figure 2: Velodyne VLP32C, Novatel PwrPak7D-E2 GNSS unit, Pointpillars detection model (open source), ROS2
- Data fusion: Python based programming utilizing in-house developed sequential confidence aware fading memory Kalman filtering (SCAFM-KF)

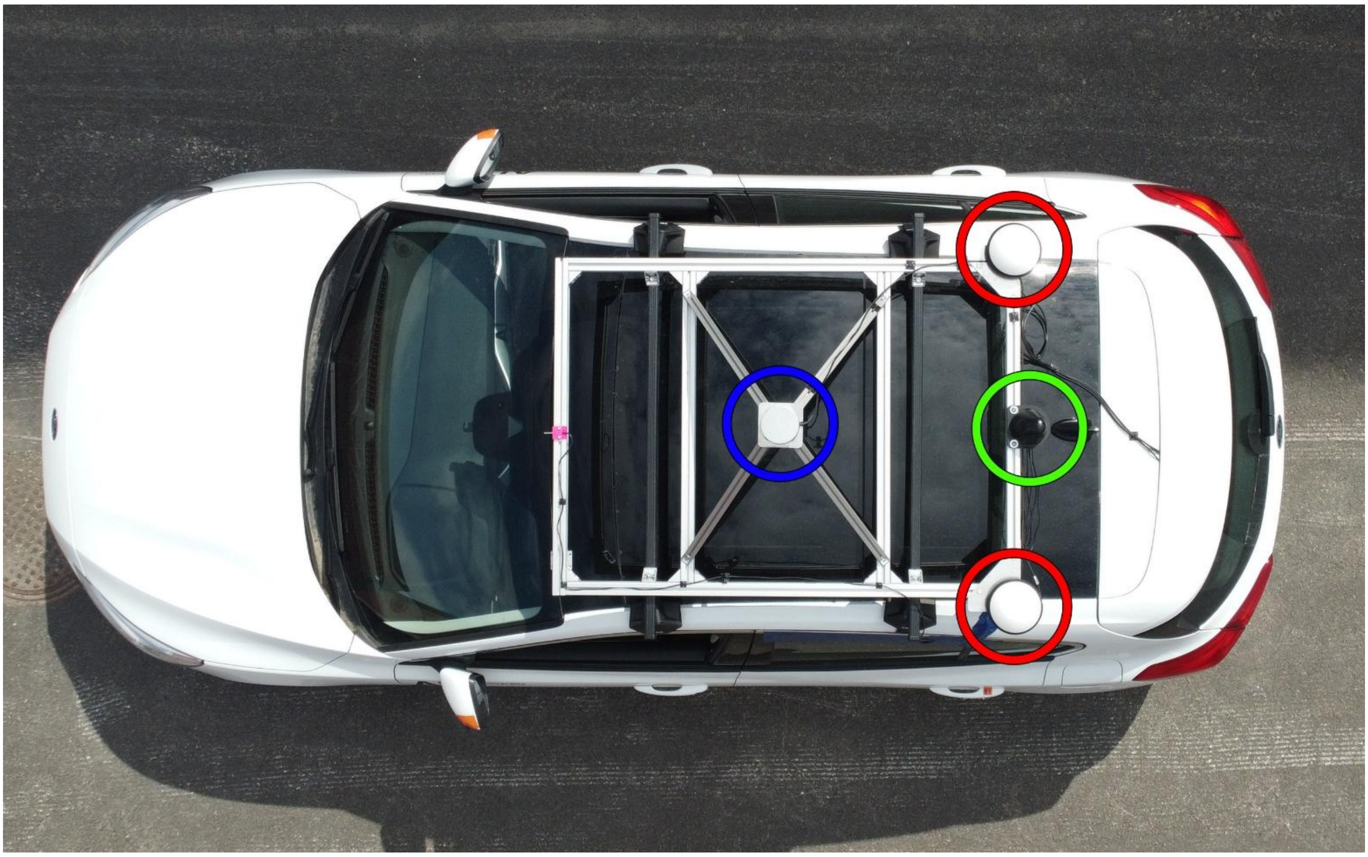

*Figure 2: Blue - lidar, red - GNSS antennas and green - V2X antenna*

Methodology highlights:

- Multi-source: combating sensor occlusion and reflection using only static sensors can prove difficult in highly dynamic scenarios such as traffic. Combining multiple sensors, static and dynamic, these challenges can be mitigated.
- Lightweight: by only utilizing simple location and velocity data of detected vehicles, the system cycle time can be reduced and thus allow real-time operations without expensive processing units or added latency from cloud computing.
- Modular: the system is developed to incorporate multiple non-sustained observers and thus is suitable for regular vehicle traffic.
- Real-world data: most of other research done in the context of data-fusion is done by using simulated which cannot consider all of the factors that affect the accuracy of such systems.
- Easily appliable: the data formats are widely used and thus provide a simple way to incorporate the method into already existing systems as an additional data source.

User value:

- Companies developing ITS: With connected capabilities becoming more widely available the fufure, this system allows users to enhance existing adaptive traffic management and data gathering systems with additional situational awareness.
- Infrastructure operators and municipalities: The method allows current roadside radar installations to be leveraged beyond their original purpose by augmenting them with collaborative perception and tracking capabilities.
- Project owners and funders: The system is built on open, widely adopted technologies and an extensible software architecture, reducing vendor lock-in and enabling long-term reuse in research, pilots, and scaled deployments.
- Researchers and system integrators: The use of standardized data formats and decision-level fusion lowers the barrier to experimentation and deployment, allowing the method to be easily incorporated as an additional data source into existing ITS, V2X, or traffic analytics pipelines.

# 3. Implementation Strategy (Technical)

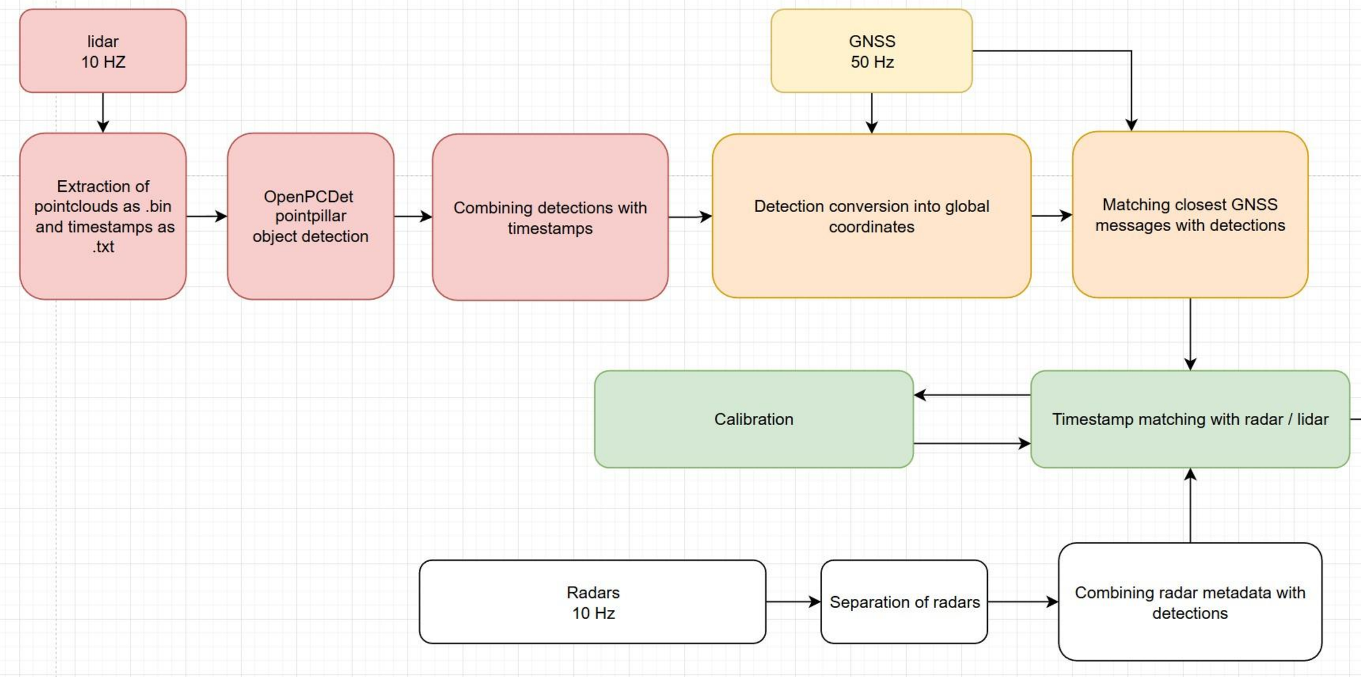

*Figure 3: Flowchart showing the data gathering and preprocessing of the system.*

## 3.1 Design & Development

The system was designed as a modular, multi-sensor experimental platform for evaluating collaborative localization and decision-level data fusion between connected vehicles and roadside infrastructure. The overall design integrates mechanical sensing hardware, on-board and edge computing resources, and a software pipeline that enables synchronized data collection, preprocessing, and fusion across heterogeneous sensors.

The experimental setup consisted of three primary sensing entities:

1. a research vehicle acting as a mobile observer,
2. a roadside radar installation, and
3. a target vehicle providing ground-truth localization.

The observer vehicle was equipped with a Velodyne VLP-32C lidar and a Novatel PwrPak7D-E2 GNSS unit, with all sensor data recorded using ROS 2 into replayable ros2bag files. This configuration enabled precise timestamping and offline reproducibility of LiDAR point clouds, ego-vehicle pose, and orientation information. In parallel, the roadside radar of interest streamed detection data via a REST API to a separate logging computer, ensuring that infrastructure-based observations were captured independently of the vehicle platform. The target vehicle recorded its own GNSS position to a dedicated rosbag, allowing its trajectory to serve both as a calibration reference and as ground truth for quantitative validation in separate steps. This configuration produced a dataset in which both observers—mobile and static—monitored the same traffic participants within overlapping spatial and temporal windows, enabling controlled evaluation of collaborative perception and fusion.

To reduce computational complexity and avoid numerical issues related to Earth curvature, all processing was performed in a local East-North-Up (ENU) coordinate frame centered at a fixed reference location of the radar. Radar detections natively expressed in radar-centric ENU coordinates were used directly, while LiDAR-based detections and GNSS measurements expressed in WGS-84 latitude–longitude coordinates were transformed into the same local frame prior to fusion. This design choice enabled efficient linear motion modeling and simplified Kalman filter updates, while remaining sufficiently accurate for the spatial scales relevant to urban intersection and corridor scenarios. By enforcing a common local coordinate system, the system avoided repeated global-to-local conversions during runtime and ensured consistent interpretation of spatial uncertainty.

All detections were processed locally either on the vehicle platform or on an edge computing unit associated with the traffic management system. This design supports low-latency operation and avoids the need to transmit raw sensor data over the network, which would be infeasible due to bandwidth constraints. Importantly, the fusion method itself remains agnostic to the upstream detection algorithm; any localization source capable of providing position estimates and an associated confidence metric could be incorporated without modification to the fusion logic.

The fusion and tracking process shown in figure 4 is implemented using a Sequential Extended Kalman Filter (EKF) augmented with a confidence-aware fading memory mechanism. Each tracked object is represented by a planar state vector consisting of position and velocity components in the local ENU frame. The filter assumes a constant turn-rate and velocity motion model, which is appropriate for urban driving scenarios over short time horizons and allows reliable interpolation between asynchronous sensor updates. Process noise accounts for unmodeled accelerations and maneuvering behavior.

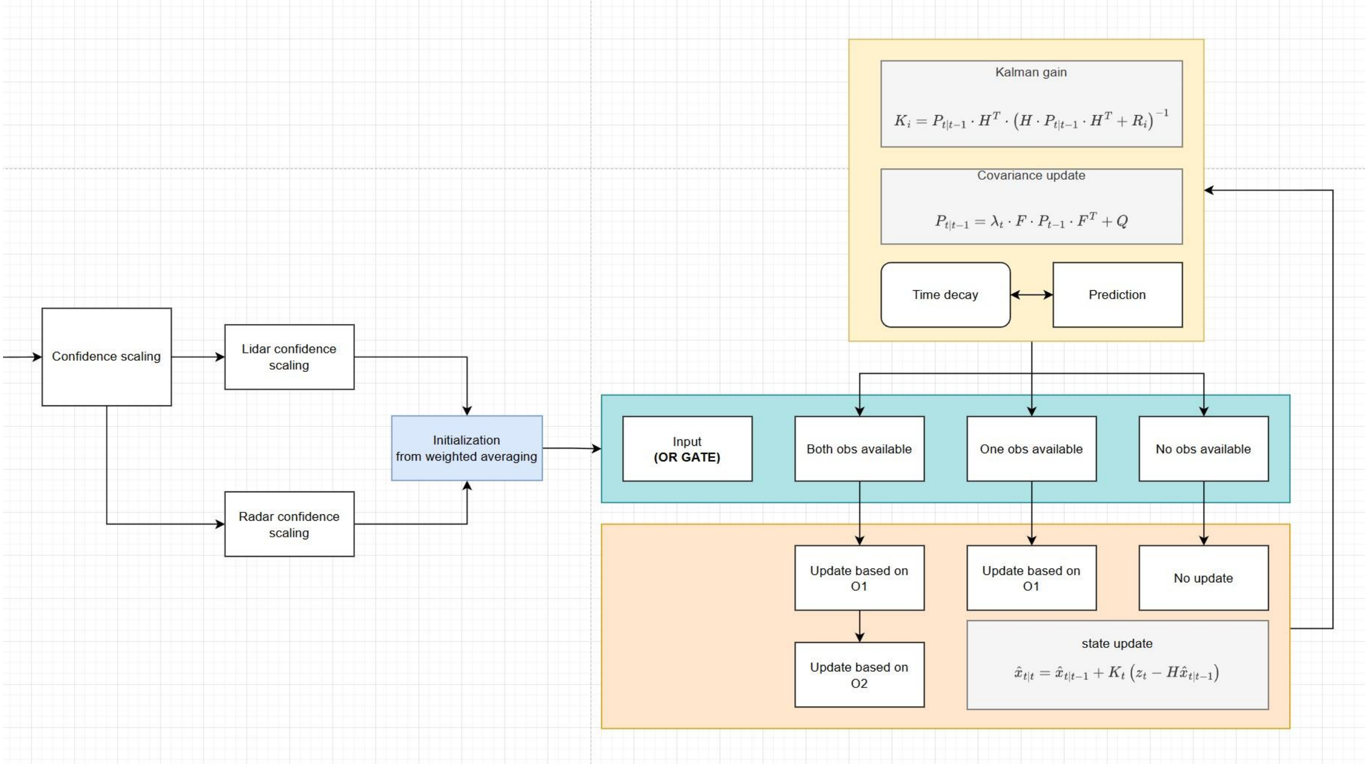

*Figure 4: High-level flowchart showing the the EKF based datafusion loop.*

Sensor observations are processed sequentially. At a given timestep, a track may receive:

- a LiDAR-based observation from the connected vehicle,
- a radar-based observation from the roadside unit,
- both observations, or
- no observation at all.

When both observations are available, the filter applies two consecutive update steps—one per source—using each sensor's measurement model and uncertainty. This preserves the individual noise characteristics of each sensor and avoids the need to artificially combine heterogeneous measurements into a single pseudo-observation.

Each detection is accompanied by a confidence score reflecting the reliability of the underlying sensor output. These confidences are used to adaptively scale the measurement noise covariance during the EKF update.

Observations with higher confidence are assigned lower effective measurement noise and thus exert greater influence on the state estimate, while low-confidence detections are naturally down-weighted. This mechanism allows the filter to remain robust in situations where one sensing modality temporarily degrades (e.g., LiDAR sparsity at long range or radar multipath effects), without requiring hard thresholding or manual sensor prioritization.

## 3.2 Testing & Validation

Testing and validation were conducted to verify the functional correctness, operational robustness, and practical feasibility of the proposed collaborative localization and fusion system under real-world urban traffic conditions. All experiments were carried out in the Jätkäsaari Mobility Lab during daytime hours in order to reflect typical traffic flow, signal timing, and environmental conditions. The experimental setup consisted of the observer vehicle (Car A), the target vehicle (Car B), and a statically mounted roadside radar 269.1 installed on a lamp post facing Mechelininkatu at the intersection with Jätkäsaarenlaituri seen in figure 5. This configuration enabled simultaneous observation of the same traffic participants from both mobile and infrastructure-based sensing perspectives.

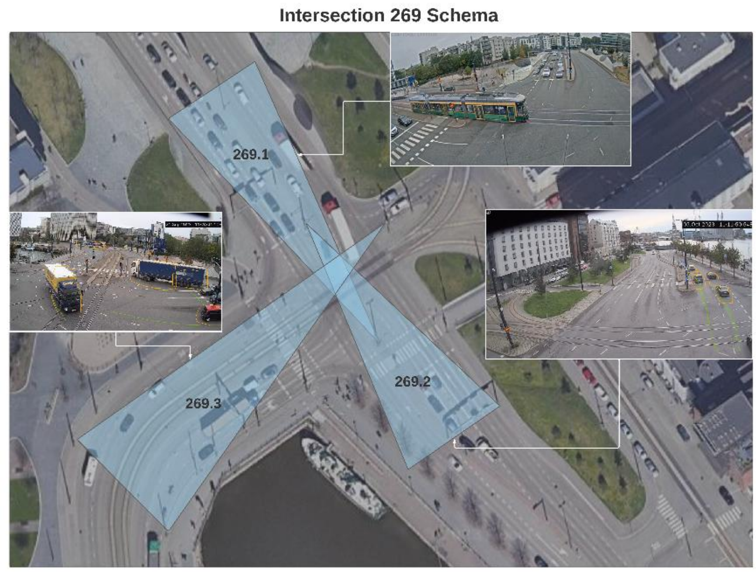

*Figure 5: Intersection 269 schema.*

A predefined route was driven repeatedly by the vehicles shown in figure 6. Each traversal of the route naturally resulted in two distinct sensing scenarios: one in which the vehicles moved away from the radar and another in which they approached it. These two scenarios were intentionally included, as they introduce differing radar viewing geometries, relative velocities, and detection characteristics, all of which are relevant for evaluating collaborative perception systems. In total, 15 experimental runs were conducted, yielding 30 unique test measurements when accounting for both approach and departure segments.

*Figure 6: Connected vehicle (left) and target vehicle during measurement run.*

All sensor data were recorded synchronously throughout the experiments. LiDAR point clouds and GNSS/INS data from the observer vehicle were logged into ROS 2 bag files, while radar detections were logged independently via the infrastructure interface. The target vehicle recorded its own GNSS trajectory, which was used as a reference for validating the fused localization estimates. Because the experiments were conducted in live traffic, each run exhibited natural variability in surrounding vehicles, spacing, and occlusions. This variability was preserved intentionally, as it provides a realistic test of system behavior under operational conditions rather than idealized laboratory settings.

The validation strategy focused on confirming correct system operation across the full processing chain. This included verifying that heterogeneous sensor data were ingested and temporally synchronized as intended, that spatial calibration and coordinate transformations were applied correctly, and that object detections from different sources were successfully associated over time. Attention was paid to the stability of the tracking output, ensuring that the fusion filter maintained consistent track identities and produced physically plausible trajectories even when individual sensor observations were temporarily unavailable.

Qualitative validation was performed through trajectory visualization and consistency checks against the GNSS reference trajectory of the target vehicle. The fused estimates evolved smoothly over time, responded appropriately to new measurements, and remained stable during short interruptions in sensor coverage. The system produced valid track estimates in both approach and departure scenarios, despite variations in traffic density and the presence of other road users.

Overall, the experiments demonstrated end-to-end system functionality in a real urban environment. The results confirm that the proposed fusion architecture, calibration procedures, and tracking logic operate as intended and are suitable for further quantitative evaluation. The conducted tests establish the feasibility and robustness of the system and provide a solid foundation for subsequent accuracy and performance analysis.

### 3.3 Safety & Compliance

All experimental activities were conducted in accordance with Finnish traffic laws and regulations governing the operation of research vehicles in public road environments. The tests were carried out under normal traffic conditions without any modification to existing infrastructure or traffic control systems, ensuring that the experiments did not introduce additional risk to other road users.

No personal data was stored, or published as part of the experiments. The system processed only sensor-derived measurements related to vehicle motion and object localization, without capturing identifiable information. As a result, the data collection and processing activities complied with applicable data protection and privacy requirements.

# 4. Impact & Performance Evaluation

## 4.1 Key Success Metrics

The success of the proposed system is evaluated using metrics that reflect both its operational feasibility and its technical suitability for collaborative localization and tracking in real traffic environments. Given that the primary objective at this stage is to demonstrate correct system behavior and readiness for quantitative evaluation, the selected metrics focus on consistency, plausibility, and robustness.

A central success criterion is the system's ability to produce stable and temporally consistent vehicle tracks when combining heterogeneous sensor inputs. This includes maintaining persistent object identities over time, avoiding track fragmentation, and ensuring that estimated trajectories evolve smoothly in accordance with physical vehicle motion. Stability of the fused output is considered a prerequisite for downstream applications such as traffic state estimation, safety analysis, or adaptive signal control.

Another key metric is spatial consistency with reference GNSS data from the target vehicle. While raw GNSS measurements themselves contain noise and bias, they provide a practical baseline against which fused estimates can be compared qualitatively and statistically. Agreement in overall trajectory shape, direction of travel, and relative positioning along the driven route serves as an indicator that the fusion method correctly integrates observations from mobile and infrastructure-based sensors.

The system is also evaluated based on its robustness to sensor availability and variability. Successful operation is defined by the ability to continue producing reasonable state estimates when one sensing modality is temporarily degraded or unavailable, such as during occlusions, reduced LiDAR point density, or intermittent radar detections. This robustness is assessed by observing filter behavior during such intervals, including uncertainty growth and recovery once reliable measurements resume.

Together, these metrics establish that the system not only functions correctly but also meets the practical requirements necessary for subsequent quantitative performance evaluation and real-world ITS deployment.

## 4.2 Results & Efficiency Gains

The experiments demonstrate that the proposed collaborative localization and fusion system operates in real-world urban traffic conditions and fulfills its intended functional objectives. The system successfully processed heterogeneous sensor inputs, performed spatial and temporal alignment, associated detections across sources, and produced vehicle tracks.

Qualitative comparison against the GNSS reference trajectory from the target vehicle showed strong agreement in overall trajectory shape, direction of travel, and relative positioning along the route. While minor deviations were observed—consistent with expected GNSS noise and sensor-specific measurement uncertainty—the fused estimates followed the reference motion closely and remained bounded throughout each run. This indicates that the fusion method correctly integrates observations from both mobile and infrastructure-based sensors into a coherent state estimate.

The system also demonstrated robustness to intermittent sensor degradation. During periods of reduced detection, the fusion filter continued to propagate vehicle tracks using prediction, with uncertainty increasing in a controlled manner. Once reliable measurements resumed, the filter recovered smoothly without introducing discontinuities or abrupt corrections in the estimated trajectory. This behavior confirms the effectiveness of the confidence-aware weighting and fading memory mechanisms in handling real-world sensing variability.

## 4.3 Challenges & Solutions

While the experiments demonstrate that the proposed collaborative localization and fusion system functions reliably under real-world conditions, several challenges and limitations were identified that influence performance and should be considered when interpreting the results.

A primary challenge arises from the heterogeneous nature of the sensing modalities. Roadside radar and vehicle-mounted LiDAR differ significantly in measurement characteristics, update rates, and uncertainty profiles. Radar detections are sensitive to viewing geometry and may exhibit angular uncertainty or occasional false targets due to multipath reflections, particularly in dense urban environments. Conversely, LiDAR-based detections provide higher spatial resolution but are susceptible to occlusions, reduced point density at longer ranges, and temporary loss of detection in complex traffic scenes. Although the fusion method mitigates these effects through confidence-aware weighting, residual inconsistencies remain inherent to the sensors themselves.

Another limitation relates to calibration and timing accuracy. The fusion approach assumes that spatial calibration between sensors and temporal alignment across logging systems are sufficiently accurate. Small errors in sensor mounting parameters, clock offsets, or GNSS reference alignment can manifest as systematic biases in the fused output. While such effects were not observed to destabilize the filter, they can influence absolute positioning accuracy and should be carefully addressed in larger-scale or long-term deployments.

From a methodological perspective, the fusion strategy assumes conditional independence between sensor observations and relies on confidence values provided by upstream detection modules. Miscalibration or inconsistent interpretation of confidence scores across sensing modalities can affect the relative weighting of observations and, in turn, the behavior of the fusion filter. Ensuring consistent confidence semantics across sensors remains an important consideration for future system refinement.

# 5. Conclusion & Future Work

## 5.1 Key Takeaways

This work demonstrates the feasibility of collaborative localization using decision-level fusion between connected vehicles and roadside infrastructure in real urban traffic conditions. The proposed system successfully integrates heterogeneous sensor observations from vehicle-mounted LiDAR and a statically mounted radar into stable, temporally consistent vehicle tracks.

Overall, the results show that the proposed fusion architecture forms a solid foundation for quantitative performance evaluation and future ITS deployments that seek to enhance situational awareness using existing sensing infrastructure and connected vehicles.

## 5.2 Scalability & Potential for Expansion

The proposed system is inherently scalable due to its decision-level fusion architecture and use of lightweight, standardized data representations. Additional observers, such as further connected vehicles or roadside sensors, can be incorporated without modifying the core fusion logic, enabling gradual expansion as sensing and communication capabilities increase.

The modular design allows the method to be deployed alongside existing traffic management and perception systems, either as a complementary data source or as part of a broader collaborative perception framework. With further validation, the approach can be extended to more complex traffic scenarios, higher object densities, and additional sensor modalities, supporting long-term adoption in intelligent transportation systems.

## 5.3 Recommendations for Further Development

Further development should focus on expanding the experimental dataset to include a wider range of traffic densities, environmental conditions, and sensor geometries in order to enable comprehensive quantitative performance evaluation. Improving calibration procedures and systematically validating confidence estimates across sensing modalities would further enhance fusion reliability. The scalability of the system could also be validated by integrating more observers into the pipeline. These could include additional CVs or a larger network of static sensors.

In addition, integrating the system with existing traffic management platforms and evaluating its impact on downstream applications, such as traffic state estimation or adaptive signal control, would help demonstrate practical value. These efforts would support the transition from experimental validation toward scalable deployment in real-world ITS environments.