Adversarial Marker Rejection

by Nick Clark | Published April 25, 2026

What Adversarial Rejection Specifies

The architecture rejects fake or compromised markers at multiple structural layers. Layer one: credential verification. A marker that fails to verify against the credentialed authority chain is rejected at protocol level. Layer two: sequence consistency. A marker whose payload doesn't match the credentialed segment sequence (e.g., a marker claiming to be on segment X when the actual segment X is elsewhere) fails consistency check. Layer three: spatial-temporal verification. A marker whose declared position doesn't match the vehicle's confidence-bounded position estimate fails consistency.

Adversaries face all three layers simultaneously. Forging a credential requires compromising the credentialing authority. Producing a sequence-consistent marker requires knowing the segment's credentialed sequence. Producing a spatial-temporally-consistent marker requires knowing the vehicle's actual position estimate. The architecture forces adversaries to compromise multiple structural elements rather than just placing physical fake markers.

Why Per-Vehicle Detection Has Structural Limits

Adversarial-marker rejection at the per-vehicle level has known limits. A vehicle's sensor stack may not recognize a fake marker as fake if the fake is sufficiently

realistic. Heuristic detection is brittle; ML-based detection requires training data adversaries can study.

The architectural rejection moves detection above the per-vehicle level. The credentialing chain, the sequence-consistency check, and the spatial-temporal verification all operate structurally. Adversaries cannot defeat all three through physical attack alone.

How Multi-Layer Rejection Composes

The vehicle's marker reader passes credentials through verification first. Markers that pass credential verification then pass sequence-consistency check against the segment's credentialed sequence. Markers that pass both then contribute to the route manifest after spatial-temporal verification confirms the marker's claimed position is consistent with the vehicle's confidence-bounded estimate.

Failures at any layer trigger structural events. A failed credential triggers an authority-compromise alert that propagates through the mesh. A failed sequence consistency triggers a segment-integrity alert. A failed spatial-temporal verification triggers a positioning anomaly alert.

What This Enables for Attack-Resistant Operation

The architecture supports operation in environments where adversarial marker placement is a real concern. Border-crossing routes, defense-relevant corridors, critical-infrastructure access roads, and emerging autonomous-fleet operations all benefit from the structural rejection.

The patent positions the primitive at the layer where adversarial-marker concerns have been handled through per-vehicle heuristics that adversaries can study and defeat. Structural rejection raises the cost of adversarial action substantially.