2

Background and Related Work

The primary motivation for VANETs is to enable Cooperative Intelligent Transportation Systems [16] that enhance transportation safety and efficiency. By allowing vehicles and RSUs to share information in real time, VANETs can proactively warn drivers of hazards (e.g. sudden braking by a vehicle ahead, slippery road conditions) and optimize traffic flow (e.g. dynamic route guidance to avoid congestion). This cooperative approach stands in contrast to autonomous onboard sensors alone; vehicles in a VANET effectively become part of a connected, collective sensor network on wheels.

2.1 COOPERATIVE INTELLIGENT TRANSPORT SYSTEMS (C-ITS)

Cooperative ITS, or C-ITS, refer to the next generation of ITS in which vehicles, infrastructure, and other actors (like networked devices) cooperate by exchanging information to improve road safety, traffic efficiency, and comfort. C-ITS extends the concept of isolated intelligent vehicles by adding continuous connectivity and cooperation between all elements of the transport network. Unlike traditional ITS where information flows mostly one-way (for example, from a traffic management center to drivers via variable message signs), C-ITS enables two-way communication in which vehicles share their status with each other and with infrastructure. This two-way, cooperative approach is promoted as a solution to the persistent road safety challenges: by reducing human reaction delays and in blind spots awareness, C-ITS can significantly decrease accidents. In C-ITS, each participant (vehicle or infrastructure) is referred to as an ITS station. The environment is envisioned as a network of ITS stations – including OBUs in vehicles, RSU, traffic signals, and even pedestrians' smartphones – all following a common communication protocol stack. Through this network, a vehicle continuously connects with its surroundings, effectively extending the driver's perception range beyond line-of-sight.

C-ITS safety services have stringent performance requirements: they demand reliable communication with minimal latency, since even small delays can mean the difference between avoiding a crash or not. For example, drivers' reaction times to warnings range roughly between 0.6 and 1.4 seconds, so to be effective, the network must deliver critical warnings within a fraction of this time limit [55]. Common guidelines restrict the maximum end-to-end latency for safety messages to 100 milliseconds or less, and certain pre-crash sensing alerts are even stricter (requiring <50 ms latency) [42]. These strict deadlines highlight the importance of fast message dissemination because the network's ability to rapidly propagate information is essential for C-ITS safety applications. Table 2.1 presents a categorized set of applications

supported by C-ITS, as outlined in the ETSI TR 102 638 [7] standard. It also includes the approximate latency requirements associated with each application class, reflecting their relevance to time-sensitive and safety-critical vehicular scenarios.

Table 2.1: Classification of C-ITS Applications with Latency Requirements

Application Category	Sub-category	Example Use Cases	Latency Requirement
Active Road Safety	Driver Assistance	Emergency vehicle alerts, proximity warnings, intersection collision alerts	<100 ms
Active Road Safety	Road Hazard Alerts	Emergency brake light signals, wrong-way driving notifications	<100 ms
Active Road Safety	Stationary Vehicle Notifications	Incident warnings, breakdown alerts	<100 ms
Active Road Safety	Traffic Condition Warnings	Congestion alerts, signal violation notifications, roadwork updates	100–500 ms
Active Road Safety	Collision Risk Management	General collision warnings	<100 ms
Active Road Safety	Hazardous Condition Detection	Alerts on low visibility, slippery roads, strong winds, and precipitation	100–500 ms
Traffic Efficiency	Speed and Flow Management	Contextual speed limit alerts, traffic light speed advisory	500 ms–1 s
Traffic Efficiency	Cooperative Navigation	Route guidance, alternate routing, restricted access warnings	500 ms–1 s
Traffic Efficiency	In-Vehicle Signage	Digital road sign notifications	500 ms–1 s
Local Cooperative Services	Location-Based Services	Nearby point-of-interest alerts, parking access control	1–2 s
Local Cooperative Services	ITS-Driven Commercial Services	Media downloads, local e-commerce, automated tolling	1–2 s
Local Cooperative Services	Community and Enterprise Services	Insurance and financial tools, fleet coordination, delivery zone optimization	1–2 s
System Maintenance and Support	ITS Station Lifecycle Management	Remote software updates, sensor and data calibration	2–5 s

Beyond safety, C-ITS encompasses traffic efficiency and infotainment services. Traffic efficiency services aim to optimize traffic flow and reduce congestion. For example, a Green Light Optimal Speed Advisory (GLOSA) service can recommend optimal driving speeds to drivers so that they arrive at intersections during green lights [56]. Such services require vehicles to communicate with traffic signals (V2I) and possibly with each other to share queue information. While the latency requirements for efficiency services are less strict than safety (maximum latencies around 500 ms are often acceptable), they still benefit from timely data dissemination to be effective in real-time traffic management. Infotainment and comfort services, on the other hand, provide non-safety information to drivers and passengers for instance, point-of-interest notifications, internet access, or media streaming in vehicles. These prioritize high data throughput over ultra-low latency and can tolerate occasional delays or message loss better than safety messages. C-ITS, therefore, covers a broad spectrum of use-cases, each with different communication requirements. The challenge for the network is to satisfy the most stringent (safety) requirements without undermining the others, often by employing QoS differentiation and message prioritization.

2.1.1 C-ITS Services and Standards

To ensure interoperability in C-ITS, international standards bodies and industry consortia have developed common communication frameworks and message sets. In Europe, the ETSI has defined a comprehensive ITS station reference architecture that all C-ITS devices should follow [8]. This architecture is layered similarly to the traditional Internet (TCP/IP) stack but with adaptations for mobility and direct wireless exchange. It consists of: (i) an Application layer that hosts the C-ITS applications (safety, traffic efficiency, etc.) and assigns priorities to messages, (ii) a Facilities layer responsible for common services like message encoding/decoding, security, and maintaining a local dynamic map of the vehicle's surroundings, (iii) a Networking & Transport layer that handles addressing and routing including specialized geo-routing protocols and transport protocols, and (iv) an Access layer which encompasses the wireless communication technologies used (such as DSRC radio or cellular). In addition, cross-cutting Management and Security entities supervise the overall operation and protect communications, respectively.

One of the key innovations in the ETSI networking layer is the GeoNetworking protocol [57], a location-based ad-hoc routing protocol proposed for vehicular communications. GeoNetworking uses geographic positions to address and forward packets, supporting modes like geo-broadcast (sending a message to all nodes in a geographic area) and geo-unicast (to a specific node at a location). Built on top of this is the BTP, which functions similarly to UDP (connectionless transport) but is designed to multiplex different message types (especially safety messages) efficiently. These protocols are part of the ETSI ITS-G5 standard suite, which operates in the 5.9 GHz band (using DSRC). ITS-G5 enables direct V2V and V2I communication with a typical range of a few hundred meters and low latency, suitable for immediate hazard warnings.

Corresponding efforts have occurred in other regions. In the United States, the IEEE WAVE standards and the SAE J2735 [58] message set define the communication stack and message contents for vehicular communication. The U.S. system similarly uses the 5.9 GHz band with IEEE 802.11p-based radios, and its basic safety message is the BSM (analogous to Europe's CAM) which vehicles broadcast periodically to inform others of their position, speed, orientation, and other state information. Both CAM and BSM are sent frequently (10 Hz or more in high dynamics) to keep neighbors aware of each other. For event-driven alerts (like a sudden hazard on the road), Europe employs the DENM, whereas in the U.S. the BSM itself can carry event flags or specific safety application messages are defined (e.g. Emergency Electronic Brake Light message). The concept is that routine awareness beacons (CAM/BSM) support cooperative awareness among vehicles, and additional emergency messages are disseminated when specific incidents occur.

It is worth noting that C-ITS communications are not limited to one technology. Initially, the focus was on DSRC/802.11p, but more recently C-V2X has emerged as an alternative. C-V2X includes two modes: a direct communication mode (using PC5 interface) similar in concept to DSRC but leveraging cellular physical layers (LTE-V2X, evolving to 5G-V2X), and a network communication mode (using the regular cellular network, the Uu interface) for broader-area connectivity. The advantage of C-V2X is that it can potentially use existing cellular infrastructure and advanced modulation techniques to extend range and reliability, and it is designed to evolve into future 5G systems. China and some other regions (and increasingly Europe) have leaned towards C-V2X for new deployments, while Japan and Europe originally deployed 802.11p-based systems. Standards bodies like 3GPP (for C-V2X) and ETSI/IEEE are working towards harmonization so that these technologies can coexist in C-ITS. For example, many European pilot projects now use a hybrid approach, employing short-range ITS-G5 for fast local exchanges and cellular networks for wide-area services [59].

The standards supporting C-ITS, whether ETSI ITS-G5, IEEE WAVE, or C-V2X all strive to facilitate these services by defining common languages (message formats) and reliable communication mechanisms. Consequently, protocols at various layers have been optimized to minimize latency and maximize delivery reliability for safety-critical data.

2.1.2 From C-ITS to Routing and Data Dissemination

With vehicles, infrastructure, and standards in place, a key question arises: *How does the information actually propagate through the vehicular network quickly and reliably?* This is the question of routing and data dissemination in VANETs. Routing in generic networks usually means finding a path from a source to a destination. In VANETs, especially for safety messages, the goal is often broader, disseminating a message to all vehicles in a certain area (for example, all vehicles approaching a crash site should receive a warning). This is sometimes referred to as message dissemination or broadcast rather than traditional unicast routing.

Unlike Internet routing where end-to-end paths are relatively stable, VANET routing must cope with the extreme dynamics of vehicle movement and the absence of guaranteed connectivity.

Early approaches attempted to adapt MANET routing protocols (like AODV or DSR) to VANETs, but found that route breakages occur too frequently in high-mobility scenarios, leading to excessive overhead in route repairs [60]. Geographic routing emerged as a promising strategy: instead of establishing a fixed route, each packet is forwarded based on the geographical position of the destination (or the desired region for dissemination) [61]. Techniques like geo-broadcast forward a packet to all nodes in a target region by repeatedly rebroadcasting it and using position-based heuristics to decide which vehicle should forward the packet next. One example is the ETSI GeoNetworking mentioned above, which underpins protocols for sending DENM hazard messages to all vehicles in an area without flooding the entire network unnecessarily. These dissemination protocols must balance speed (delivering the warning quickly), reach (ensuring all relevant vehicles get the message), and efficiency (avoiding a broadcast storm of excessive redundant messages that could congest the channel).

The broadcast storm problem is a well-known challenge in VANET data dissemination. If every vehicle rebroadcasts every message it receives, the number of transmissions grows exponentially, leading to collisions and congestion in the wireless channel. Therefore, fast message dissemination is not simply achieved by blind flooding; it requires intelligent suppression of unnecessary rebroadcasts and selection of the most effective relay nodes. Routing and dissemination protocols in VANET literature employ various strategies to tackle this: some use probabilistic or counter-based methods (a node rebroadcasts only with a certain probability or if it hasn't heard the message too many times), others use distance or location-based suppression (only the farthest vehicle from the sender will rebroadcast, since it extends the coverage the most), and others form clusters or trees to structure the dissemination [62]. We will explore these in detail in the Related Work Section, but the common aim is to propagate messages quickly (low latency) and safely (high delivery ratio), without overwhelming the network.

It's also important to recognize the role of infrastructure in data dissemination. While VANETs are often thought of as purely ad-hoc V2V networks, roadside infrastructure (RSUs) and even cellular networks can assist in disseminating messages faster or farther. For example, if an RSU hears an emergency message, it might broadcast it at high power or feed it into a cloud system that then injects the warning to other RSUs ahead, creating a multi-hop vehicle-to-infrastructure-to-vehicle (V2I2V) dissemination pattern. This combination can significantly extend reach and reduce latency to distant vehicles, as infrastructure can act as a high-speed relay. The remainder of this chapter will discuss how real-world C-ITS pilot deployments have implemented these concepts, and how researchers have developed and evaluated various protocols to achieve fast message dissemination in VANETs. By reviewing those deployments and related research, we will highlight successes, limitations, and gaps that motivate further innovation (including the methodology proposed in this thesis).

2.1.3 C-ITS Real-World Pilots

Real-world deployment programs and pilot projects of C-ITS have been launched in various regions worldwide, each with their own emphasis, technologies, and lessons learned. This section reviews major C-ITS pilots in Europe, United States, China, Japan, and India. For each region, we outline the scale of deployment, the technologies and services demonstrated, and an assessment of effectiveness and remaining gaps. Understanding these real-world efforts is crucial, as they provide practical insights into what works in ve-

hicular communication and what challenges remain when moving from theory and simulation into actual highways and city streets.

Europe

Europe has been at the forefront of C-ITS deployment, driven by strong collaboration between governments, automotive industries, and road operators under a unified vision. The European Commission and organizations like the Car-2-Car Communication Consortium [63] have coordinated numerous projects to ensure interoperability across countries. One flagship initiative is the C-Roads Platform [64], a joint effort of EU member states to pilot and harmonize C-ITS services across Europe. Through C-Roads and related Field Operational Tests (FOTs) (e.g., Drive C2X, Testfeld Telematik, SCOOP@F in France, NordicWay in Scandinavia) [65], Europe has deployed C-ITS infrastructure on thousands of kilometers of roads. Approximately 20,000 km of European roads has already been equipped with short-range C-ITS (ITS-G5) communication, and an even larger network (around 100,000 km of roads) had coverage via cellular C-ITS services [66]. These numbers indicate a significant physical footprint: RSUs broadcasting safety messages or traffic information along major corridors, and many vehicles fitted with OBUs capable of receiving them. Further, European Cross-Border Trials (ConVeX [27], CONCORDA [28], 5G CARMEN [29], 5G MOBIX [30]) tested C-V2X along highways and blind intersections, reporting robust connectivity and sub-100 ms latencies under non-line-of-sight conditions, proving viability even in complex cross-border settings [31].

Despite these achievements, gaps and challenges remain. One is the heterogeneity of technology: Europe initially backed IEEE 802.11p (DSRC) [13], but more recently there has been a policy pivot to also allow C-V2X. This led to debates and some uncertainty among stakeholders. Some of the early deployments are DSRC based, while newer pilot sites (e.g., in Germany and Spain) also test C-V2X. Ensuring these technologies can coexist and still deliver cooperative services is an ongoing task. Another challenge is demonstrating clear cost–benefit outcomes. While safety benefits (reduced crashes) are the ultimate goal, they are statistically hard to prove in small pilots. However, a European Commission study found Benefit-Cost Ratios ranging from 2:1 to 8:1 for various C-ITS services when deployed at scale, suggesting long-term positive returns if initial costs can be managed. Lastly, a gap in Europe's deployment is the urban penetration: many pilots have focused on highways; urban C-ITS (with its complex environments of intersections, pedestrians, bicyclists) is just beginning to be addressed through initiatives like the Urban C-ITS pilots and integration with Smart City programs. In summary, Europe's C-ITS pilots have proven the concept and started deployment for real, but the work continues to ensure full coverage, mixed-technology interoperability, and sufficient market penetration of equipped vehicles to realize network effects.

United States

In the United States, the development of connected vehicle technology (the term "Connected Vehicles" is often used interchangeably with C-ITS) has been spearheaded by the U.S. Department of Transportation (USDOT) through a series of research programs and pilot deployments. A significant milestone was the Safety Pilot Model Deployment in Ann Arbor, Michigan, which equipped about 2,800 vehicles and 25 RSUs in a real-world environment. This model deployment, involving everyday drivers, demonstrated that vehicles could exchange BSM reliably and that safety applications (like forward collision warning, red-light violation warning) indeed functioned to alert drivers of potential dangers [4]. Building on that, the USDOT launched the Connected Vehicle Pilot Deployment Program (CV Pilots) in 2015. Three diverse sites were

chosen: New York City (NYC) [23], Tampa (Florida) [3], and Wyoming (I-80 corridor) [67]. Each site targeted different challenges and use-cases:

- **New York City:** Focused on urban intersection safety in a dense city environment, including vehicle-to-pedestrian interactions. New York City Pilot installed around 470 RSUs at traffic signals in Manhattan and Brooklyn, and outfitted approximately 3,000 vehicles (buses, taxis, trucks, and city fleet vehicles) with DSRC units. The aim was to reduce common city crashes (e.g., intersection collisions, pedestrian accidents) by providing alerts such as pedestrian in crosswalk (to turning buses) and red light violation warnings.
- **Tampa (THEA – Tampa Hillsborough Expressway Authority):** Focused on a freeway and downtown driving environment. It equipped approximately 1,000 private cars, 10 buses, and 10 streetcars with OBU, and installed dozens of RSUs along the Selmon Expressway's reversible express lanes and in downtown. The applications tests included wrong-way driving alerts (on expressway ramps), transit signal priority, and traffic congestion detection/management on the expressway. Drivers in equipped cars could receive forward collision warnings, among other alerts.
- **Wyoming (WYDOT):** Focused on a rural interstate (I-80) prone to severe weather and involving heavy freight traffic. This pilot equipped a fleet of about 327 vehicles (mostly snow-plow trucks and fleet trucks) with V2V/V2I devices. RSUs were installed along I-80 at key locations. The primary service was to broadcast road weather alerts (fog, ice, high winds) and road closure information to trucks, aiming to prevent multi-vehicle traffic jams and manage emergency responses during winter storms. By 2019-2020, these pilots had moved from installation to active operation and evaluation. They showed that large-scale deployment is feasible, over 8,000 vehicles across the three sites, and the system integrated with multiple device vendors and vehicles, demonstrating interoperability and the ability to handle real traffic conditions.

The effectiveness of the U.S. pilots is being measured in terms of improvements in safety metrics and mobility. Early results indicated certain positive trends i.e., Tampa saw drivers with connected vehicle warnings maintaining better speed compliance and responding faster to slower vehicles ahead. Wyoming's system provided dozens of weather alerts that helped truck drivers avoid crashes or excessive delays (and even a small fuel savings by better managing stops was noted in evaluation). However, it is acknowledged that the full safety benefits (such as crash reduction) would require much larger penetration of connected vehicles than a pilot can achieve. The pilots also encountered challenges: for example, some DSRC equipment supply issues, and the regulatory rules (the FCC decided to reallocate a majority of the 5.9 GHz band away from DSRC to other uses) leaving only 30 MHz for ITS and suggesting a transition to C-V2X.

China

China has aggressively pursued C-ITS deployment in recent years, leveraging a top-down approach that combines government support, domestic industry mobilization, and integration with the nation's push for Smart Cities and 5G technology. Unlike Europe and the U.S., which began with 802.11p, China started directly into C-V2X (LTE-V2X) as the primary technology for short-range communication. The government, through the Ministry of Industry and Information Technology (MIIT) and other agencies, has designated several National Pilot Zones for Intelligent Connected Vehicles (ICV) [68], effectively large-scale C-ITS testbeds in major cities or regions. Examples include Wuxi (Jiangsu province), Tianjin (Xiqing district),

Changsha (Hunan), and Chongqing. Further, Wuxi, cited as the world's first city-level LTE-V2X pilot and deployed approximately 800 RSUs covering 400 square kilometers of urban area [5]. These RSUs are dual-mode (capable of both direct PC5 communication and network communication via 4G/5G) and are integrated with upgraded traffic signal controllers and various sensors (video cameras, radar, LIDAR) to create a truly instrumented road network. In Wuxi, a broad range of C-ITS applications have been demonstrated: from V2I safety warnings like traffic signal phase information and speed advisory, to V2V collision warnings (intersection collision avoidance, blind spot warnings), to V2P alerts for pedestrian and bicyclist protection. They also trialed 5G for more advanced use cases like platooning and sensor sharing. The presence of many sensors and edge computing (Mobile Edge Computing (MEC) servers) in these pilot zones indicates an architecture where vehicles not only communicate with each other, but also with a smart infrastructure that can process data and assist in decision-making.

Another example is Tianjin's pilot who had set up a test area including both closed-test fields and open urban roads, experimenting with over 100 use cases. Tianjin's focus includes verifying security mechanisms for C-V2X and multi-industry cooperation (like public transport and emergency services integration) [69]. China also announced an ambitious goal for the next phase: to equip a high percentage of new vehicles with C-V2X and to build out smart road infrastructure. Indeed, some Chinese OEMs (Original Equipment Manufacturers) have already started equipping cars with C-V2X units for the domestic market. By leveraging both direct V2V and cellular network communications, Chinese deployments can disseminate safety messages in multiple ways: directly between nearby cars and via the cellular network from backend traffic management centers to cars (V2N). For example, if a hazard is detected by one car or an infrastructure sensor, an alert can be sent through the cloud to other vehicles that are connected via 4G/5G, even if they are not yet in direct radio range. This multi-channel dissemination strategy can potentially increase the speed and reach of message propagation.

In terms of effectiveness, Chinese pilots have demonstrated technical success in dense urban environments, one of the most challenging scenarios due to high interference and complexity. The Wuxi pilot during live demos [1][2] showed latency on the order of milliseconds for vehicle-to-RSU messages and effective warning distances of several hundred meters for intersection collision alerts. Moreover, the integration with traffic signals in these pilots has enabled "green wave" optimizations and priority for emergency vehicles, showing tangible mobility benefits like reduced stop times and improved emergency response.

Japan

Japan was an early pioneer in certain aspects of vehicular communication and continues to implement C-ITS in a somewhat distinct way, reflecting its specific context. One of the earliest large-scale deployments in Japan was the ITS Spot system (deployed around 2010) [70], which is a V2I communication network primarily for highways. ITS Spot uses RSUs operating at 5.8 GHz to broadcast information to vehicles; it was designed to provide drivers with services like dynamic route guidance, safety alerts (such as congestion or obstacle warnings ahead), and even electronic toll collection in one unified system. By mid-2010s, Japan had installed roughly 55,000 ITS Spot units on local roads and 1,600 units on about 11,000 km of expressways. Around 2015, Japan introduced "ITS Connect", which is cited as the world's first production V2V safety system deployed in commercially available cars [71]. Toyota led this initiative, equipping

[1] https://www.huawei.com/en/huaweitech/industry-insights/outlook/mobile-broadband/lte/use-cases/wuxi-internet-of-vehiclec-project-use-cases

[2] Ford Live Demo: https://tinyurl.com/fhx6m946

some models (e.g., Toyota Prius and later other models) with wireless units using the 760 MHz band for direct V2V and V2I communication. The 760 MHz band is unique to Japan (lower frequency than the 5.9 GHz used elsewhere, which provides better propagation in urban roads). ITS Connect allows vehicles to directly communicate with each other and with compatible traffic signals within about 300 meters. The initial services included intersection collision avoidance (vehicles broadcasting their presence to others to warn if they are on a potential collision course at an intersection) and signal phase information (traffic lights broadcasting their current state and time to change, so cars can know if they will likely need to stop). This system has been deployed at numerous intersections in Tokyo and other cities. Traffic signals were equipped to send data on 760 MHz to vehicles, warning of red-light runners or assisting with smoother acceleration and deceleration. By using OBUs in sold vehicles (Toyota launched ITS Connect [3] in 2015 implementing V2V and V2I communication), Japan moved C-ITS from experimental to consumer reality for those drivers.

Another area Japan has explored is platooning and truck automation on selected highways using both V2V communication and precise GPS, part of its highway innovation programs. Another aspect is the ETC 2.0 (Electronic Toll Collection) infrastructure, which also uses 5.8 GHz to deliver safety information to drivers. This means a single in-vehicle unit can handle toll payments and also receive road hazard warnings from the same RSUs installed on highways. This dual-use strategy helped justify infrastructure deployment by providing multiple services (payment + safety). Japan's philosophy has been to focus on specific high-risk scenarios (like intersections) and address them with very tailored solutions, rather than a broad all-encompassing C-ITS roll-out. Moving forward, Japan is exploring connecting its C-ITS systems with automated driving research (e.g., using V2X communication to enhance autonomous vehicle perception).

India

India's journey with C-ITS is at a nascent stage compared to the regions above, reflecting later adoption but also unique local needs. With one of the highest road accident rates in the world, the motivation for C-ITS in India is extremely high: improving road safety is a national priority [72]. However, the deployment of cooperative systems faces challenges such as heterogeneous traffic (mix of cars, bikes, auto-rickshaws, pedestrians), cost sensitivity, and infrastructure limitations.

As of now, India has not seen large-scale C-ITS pilot deployments on public roads at the level of Europe or China, but groundwork is being laid. The government's National Highway Authority and Ministry of Road Transport have signaled interest in connected vehicle technologies as part of the push for smarter infrastructure. In 2018, the national telecom regulator (TEC – Telecommunication Engineering Centre) released a technical report outlining technologies and standards for ITS in India [34]. This report covers DSRC (802.11p) and C-V2X communications, indicating that India is studying both options and keeping aligned with international developments. It also refers to possible timelines for C-V2X deployment and integration with 5G.

There have been a few localized pilot trials. For example, some academic institutions and industry consortia have done testbed deployments. One noted effort is by a collaboration in the city of Bangalore, where a "Connected Vehicles conclave (CAEV Expo [4])" has spurred interest, and limited trials of V2X communication at certain intersections have been reported. One initiative by the Indian government is the plan to implement an emergency vehicle priority system (Green Corridors for ambulances) in some smart city

[3] https://global.toyota/en/detail/9676551

[4] https://www.futuremobilitymedia.events/

projects. This effectively uses communication between ambulances and traffic signals (via central control) to clear the path. Another is the mandate that all commercial vehicles have GPS trackers and panic buttons (for safety of passengers); while not V2V, it shows the direction of policy towards utilizing communication for safety. The evolution could be that those communication devices become two-way and interactive with infrastructure. Meanwhile, automobile manufacturers (i.e., Kia and MG Motors) and telecom operators (i.e., Jio and Airtel) have started preliminary projects (Kia Connect2.0 [5], MG-Jio Connected Car Solutions) to explore V2V communication features in electric and high-end vehicles to test connectivity on Indian roads [37].

The gaps in India are significant: infrastructure (no RSUs deployed at scale yet), regulation (no dedicated ITS spectrum usage yet, though 5.9 GHz is globally harmonized and available), and validation in mixed traffic conditions (the algorithms might need adaptation to a scenario where traffic behavior is less predictable and lanes are not strictly followed by all). On the other hand, starting late can be an advantage – India can directly adopt the most advanced, cost-efficient technology (perhaps jumping to C-V2X or 5G without having invested in older DSRC/RSUs). Moreover, there is potential for high impact: even simple applications like broadcast of fog warnings on highways or cooperative transit signal priority in congested cities could bring noticeable safety and efficiency improvements.

In summary, India's C-ITS pilot stage is just beginning. The focus in the immediate term is on developing indigenous capacity (standards assessment, domestic manufacturing of V2X devices) and conducting pilot projects in a few corridors. The effectiveness and outcome of these initial pilots will inform larger deployment decisions. Given the scale of India's road network and vehicles, any rollout will face challenges but also yield enormous benefits if successful.

2.2 RELATED WORK

This section discusses the state-of-the-art research on routing protocols and data dissemination techniques in VANETs, with a focus on how they achieve fast message dissemination. We categorize the approaches and analyze the strengths and weaknesses of existing literature. Finally, we highlight the research gaps that motivate the need for improved solutions and contrast the existing works with the methodology proposed in this thesis.

2.2.1 Mobility-Adaptive Routing and Connectivity Management

Maintaining connectivity in a rapidly changing VANET topology requires more than just prediction; many protocols take a reactive or mobility-adaptive approach, continuously reconfiguring routes or exploiting infrastructure support to cope with disruptions. Unlike purely predictive schemes, these protocols emphasize robustness through adaptive clustering, store-carry-forward techniques, or integration of RSUs and other networks.

[5] https://www.kia.com/content/dam/kia2/in/en/images/kia-connect/Kia_Connect_2.0.pdf

Adaptive Clustering for Mobility Considerations

Forming clusters of vehicles is a well-known strategy to impose some stability (connectivity between vehicles) on the network. By electing a CH that acts as a local coordinator, intra-cluster communication can be kept fast and reliable, and only CH's handle inter-cluster packet forwarding, reducing overhead. Numerous clustering algorithms have been proposed, each with different criteria for cluster formation and maintenance:

Cross-Layer Stable Clustering: Shafi et al., [73] proposed a cross-layer cluster-based routing protocol, CCBR, that factors in physical-layer signal quality and link lifetimes to decide cluster formation. By measuring mobility (relative speed, link duration) at the MAC layer and feeding that into CH selection, CCBR produced more stable clusters and longer-lasting routes. CCBR assumes that vehicles maintain consistent relative velocities over time, allowing for accurate prediction of link durations through the probability density function. However, real-world vehicular movement exhibits sudden accelerations, decelerations, lane changes, and direction alterations based on traffic conditions, driver decisions, and roadway characteristics. Recognizing this, our thesis draws on the idea of link-quality-based clustering but simplifies it using velocity vectors as a continuous indicator of alignment and relative mobility between vehicles. This single metric captures link stability (vehicles moving in nearly the same direction remain in a cluster) without requiring full cross-layer integration, thus achieving the stability benefits but with far less overhead.

Mobility/Density-Based Clustering: Ren et al., [74] developed a clustering scheme where CHs are selected based on a combination of traffic mobility patterns, vehicle density, and geographic distribution. Ren's algorithm picks cluster-heads that are centrally located among local vehicles and keeps a safe distance between CHs to avoid interference. This method worked well in city grid scenarios, however, throughput gains were modest; the protocol did not significantly improve end-to-end data rates, suggesting that optimizing cluster formation alone isn't enough to maximize network capacity. This highlights that not only forming stable clusters (using direction alignment) is required but also optimizing the routing is also essential between clusters. Therefore, our thesis incorporates a multi-criteria route selection among cluster-heads to maximize throughput and reliability.

Spectral Clustering on Highways: Khan et al., [75] focused on highway scenarios and proposed CVoEG, a routing model using an Eigengap heuristic for cluster formation on bidirectional highways. By applying spectral graph partitioning techniques on the connectivity graph of vehicles, CVoEG identified natural clusters (via the eigengap in the graph's Laplacian) and routed packets through those clusters. CVoEG's approach was innovative, yet computing eigenvalues for clustering is not trivial for an ad-hoc network that changes rapidly. As a result, network delays actually increased under CVoEG, offsetting some of its routing benefits. This illustrates a recurrent theme: advanced math or optimization can find good routing structures, but the timeliness of those computations is critical in VANETs. Our thesis acknowledges this by favoring greedy but fast algorithms over exhaustive or spectral ones. Our MetaLearn framework (Chapter 4) introduces GWO (a heuristic search) for cluster optimization, but crucially combines it with RL to quickly adapt and avoid lengthy computations each time the topology changes. This yields near-optimal clustering with far less computation than spectral methods, addressing the latency issue seen in CVoEG.

Direction-Based Routing: Hadded et al., [76], proposed an angle-based neighbor detection technique, ACA, and a CH election mechanism that groups vehicles only if they move with an acute velocity-vector angle relative to each other. ACA aimed to eliminate unstable clustering caused by naive direction checks and reduce the overhead generated by frequent re-clustering on multi-segment roads and intersections. However,

while this angle-based mechanism greatly improves cluster stability, it confines itself to local geometric relationships without leveraging wider network intelligence, which can limit adaptation in diverse or dense VANET scenarios. Therefore, In Software-Defined Networking-based Routing Protocol (SDCast), we include directional awareness with the global view of topology through SDN. Further SDCast leverages city traffic models (like lane-changing behavior) to adapt routing for intersections, extending direction-based routing to more complex topologies.

Nature Inspired Heuristic and Route Optimization

Alongside clustering, researchers have applied various nature inspired optimization algorithms to VANET routing problems, treating route selection as an optimization of multiple objectives (latency, reliability, hop count, etc.):

Ant Colony Optimization: Gawas et al., [77] merged the swarm intelligence principles of Ant Colony Optimization (ACO) with the reactive route discovery of AODV and proposed ACO-AODV protocol. ACO-AODV applied ant-like agents that deposit and follow pheromone trails to find efficient paths, initiating route searches only when data transmission is required. This approach suits VANETs by reducing the overhead of maintaining continuous routing tables in highly mobile environments. However, the protocol suffers from high latency and convergence issues due to sequential pheromone updates, and it can prematurely converge on suboptimal paths when network conditions change rapidly. ACO-AODV scalability is also problematic: more vehicles increase the complexity of maintaining pheromone-based routes. Our thesis borrows the idea of multi-constraint optimization but avoids the heavy pheromone mechanics. In the MetaLearn protocol (Chapter 4), we use OBQR's multi-constraint routing by formulating it as a scalar optimization problem (using weighted sums) that can be solved in one pass, along with the GWO's heuristic methods. This yields many of the same benefits (balancing delay, link stability, etc.) without flooding the network with control iterations.

Particle Swarm and Genetic Algorithms: Swarm intelligence and genetic algorithms have been applied to VANET clustering and routing as well. GSO-RBVT-R [78] applied Glow-worm Swarm Optimization (GSO) to enhance road-based routing in VANETs. GSO-RBVT-R models glowworm attraction to high luciferin levels, representing superior routes and integrating road traffic patterns for improved alignment with actual vehicular flows. While road-based adaptation stabilizes routing, GSO-RBVT-R was dependent on static parameter settings that reduced its ability to adapt to frequent changes in network topology. GSO-RBVT-R also lacked a mechanism for immediate path updates based on real-time speed and direction, slowing route convergence. Limited exploration depth further restricts the algorithm's effectiveness in finding truly optimal relay nodes. Further, GKO [79], a nature-inspired algorithm used Giraffe kicking behavior to balance exploration and exploitation during CH selection. GKO paired C-means clustering with multi-fitness functions, emphasizing factors like residual energy, node degree, and distance for reliable cluster formation. An explorative kicking phase, akin to a random walk, helps avoid premature convergence and expands the search space. Despite these advantages, GKO's performance depends heavily on fine-tuned fitness functions, making it challenging to adapt to rapidly shifting network conditions. The overhead introduced by coupling C-means with GKO can also slow response times in real-time VANET environments. In our work, we introduce a meta-heuristic framework MetaLearn (Chapter 4) that specifically addresses convergence and local optima issues by combining GWO (another swarm algorithm) with RL (Q-learning). GWO organizes search agents in a hierarchy (alpha, beta, etc.), which MetaLearn adapts for optimal cluster-head selection using a mobility-aware fitness function. We then integrate Temporal-Difference learning (Q(λ))

to fine-tune routing over multiple hops, which accelerates convergence of the route selection by propagating rewards back multiple steps. By combining these methods, MetaLearn effectively avoids the stagnation problem of single heuristics and does so faster.

Hybrid Optimization (Mathematical + Heuristic): MOCRAW [80], integrated the Dragonfly Algorithm (DA) for CH selection and ACO for routing. MOCRAW first identified optimal CHs by evaluating node density using residual energy, and distance, using DA. Later, ACO used pheromone-based reinforcement to establish reliable paths between these clusters and the base station. This hybridized approach delivers strong energy savings but incurs computational overhead, which can impede responsiveness in time-critical networks. Its reliance on pheromone updates and iterative DA cycles affects rapid adaptation to highly dynamic topologies. Moreover, precise parameter tuning is required to maintain a healthy exploration-exploitation balance. Further, [81] combined the Firefly Algorithm (FA) and Cuckoo Search (CS) to meet multiple QoS constraints (delay, reliability, energy) in VANETs. FA simulated fireflies' attraction behavior to refine routes locally based on QoS metrics, while CS employed Levy flights for broad, randomized searches that help discovering new, potentially better paths. Though robust in theory, merging FA and CS raised computational overhead, especially in high-density networks. The algorithms also struggle with real-time adaptability since the update mechanisms can lag behind rapid topology changes in VANETs. Making more efforts, originally developed for Wireless Sensor Networks (WSNs), [82], proposed Whale Moth Flame Optimization (WMFO) for initial clustering and an Improved African Buffalo Optimization (IABO) for routing. WMFO formed energy-conscious clusters, and IABO selects routes that balance residual energy and distance to maximize network lifespan. While highly effective in conserving energy, this design does not cater to latency-sensitive VANET operations, where real-time data exchanges and high throughput are critical.

In this direction, our thesis contributes MRMH, which uses GWO to explore the solution space and then an SQP step to fine-tune the best route found, satisfying nonlinear QoS constraints. MRMH introduces a dynamic nonlinear parameter decay in the heuristic's search to better balance exploration vs. exploitation under VANET conditions. By combining the speed of heuristics with the precision of SQP, MRMH ensures feasible routes are found even in large-scale networks where pure heuristics might falter or violate constraints. This dual approach is novel in the VANET context and directly addresses the limitations seen in works like [83], where either convergence speed or constraint handling was problematic.

Clustering and heuristic optimization techniques in VANETs seek to improve scalability and reliability, yet many suffer from high complexity or slow adaptation. Most heuristics protocols are evaluated through NS2/NS3 or MATLAB simulations using synthetic mobility traces. A notable exception is the concept of clustering itself, which has influenced standardization: for instance, some European ETSI ITS-G5 standards incorporate the idea of a "local dynamic map" and geo-networking that groups nearby vehicles for message dissemination, loosely resembling clustering. This thesis advances the field by proposing clustering and routing frameworks that capture the spirit of these advanced algorithms (stable clustering, multi-metric optimization).

2.2.2 Predictive Routing Protocols

One of the primary goals in VANET routing research is to satisfy QoS requirements, such as low latency, high packet delivery ratio, and reliable connectivity, despite the dynamic network topology. Ensuring robust QoS in VANETs is challenging due to high node mobility and frequent disconnections. Researchers have

addressed this by making routing protocols prediction-aware, leveraging various forms of mobility prediction, learning algorithms, and context-awareness to anticipate link breaks and select more stable paths.

Reinforcement Learning and Machine Learning based Routing: RL has been widely explored to adapt routing decisions. For example, Wu et al., [84] applied Q-learning to VANET routing by having nodes periodically exchange Q-values in HELLO beacons and reward routes that remain stable. This technique improved route selection robustness, yielding higher packet delivery, but it struggled when vehicles moved rapidly and frequent topology changes caused Q-values to become stale, requiring constant updates. In a follow-up, Wu et al., [85] implemented a Q-learning based forwarder selection scheme that treats each packet as an agent and vehicles as states, achieving high data delivery ratios at the cost of very frequent Q-table maintenance to keep up with network changes. A similar protocol in [86], proposed a hierarchical routing framework, IV2XQ, utilizing Q-learning, where RSUs at intersections maintain a multidimensional Q-table for packet forwarding decisions. This approach outperforms traditional GPSR (geo-location based protocol) and other baseline methods in terms of packet delivery ratio and overhead by selectively forwarding through intersections that strike a balance in vehicular density. However, while IV2XQ effectively addresses intersection-induced disruptions and congestion, it relies on a localized definition of optimal density and does not harness a broader global view of mobility patterns, which can hamper adaptability when node velocity changes unpredictably.

Context-Aware and Predictive Techniques: Not all predictive routing relies on heavy ML; some use domain knowledge of vehicular movement. For instance, Ye et al., [87] developed a Target-Driven Mobility Prediction (TDMP) routing protocol that used the driver's intended destination plus real-time signal strength to anticipate link breaks. TDMP includes a "target-driven" concept, where vehicles sense neighboring dynamics and predict vehicle movement to proactively avoid link failures and local maxima. Despite showing higher packet delivery ratios and reduced delays, TDMP relies predominantly on node-centric metrics and localized decisions without integrating centralized or semi-centralized route control. Utilizing context awareness, another hybrid opportunistic and position-based routing protocol, OPBR [88], combined location-based metrics with probabilistic forwarder selection to address VANETs' high mobility and frequent topology changes. OPBR dynamically selects next-hop nodes based on real-time positioning data and opportunistic forwarding decisions, demonstrating improved delivery ratios and lower overhead. However, the reliance on precise location information and stable connectivity introduced potential performance gaps under heterogeneous traffic scenarios, and varying road conditions. It underscores that pure opportunistic use of periodic messages can't guarantee delivery when network load is high; some explicit routing control is still needed. Further, some protocols incorporate mobility models or predictors explicitly. Liu et al. [89] applied a Kalman filter to continuously predict the next position of neighboring vehicles, thereby speeding up neighbor discovery and route maintenance. The Kalman prediction improved route setup time and reduced the frequency of broken links by anticipating topology changes. The drawback was that the protocol prioritized accurate neighbor information so much that it occasionally sacrificed PDR, by focusing on pruning "dead" links quickly, it sometimes underutilized viable links (hence a lower overall delivery ratio). This indicates that mobility-aware protocols must balance proactivity with reliability: an overly aggressive pruning of links or constant route switching can harm QoS. ABOR [90] employed adaptive beacon intervals to reduce redundant transmissions and manage congestion through opportunistic routing. Although this can be effective, it remains a relatively unstructured solution in very dense settings, potentially introducing broadcast collisions or inefficient data dissemination patterns. By incorporating adaptive beaconing and a clear intra and inter-cluster communication hierarchy, our approach tackles these issues more systematically, ensuring robust, predictable forwarding behavior without sacrificing efficiency.

Summary of QoS/Prediction Strategies and Thesis Contributions: The surveyed QoS-driven and predictive protocols highlight the importance of anticipating VANET dynamics. Techniques like Q-learning and mobility prediction improves routing performance, but they introduce complexity and often assume ideal data availability. This thesis builds on proposing routing strategies that incorporate lightweight prediction and context-awareness without relying on heavy black-box models. Our OBQR protocol (Chapter 3) uses vehicles' direction and simple car-following models to predict short-term topology changes, optimizing route selection for stability and low delay in a computationally efficient manner. By combining multi-metric QoS optimization with real-time traffic context, OBQR addresses the limitation of prior works that either ignored QoS or handled it with expensive algorithms. Moreover, in Chapter 3, we further integrate a RL module within an SDN controller (the basis of our SDCast protocol) to adapt routing policies on the fly, but crucially, we do so in a controlled, centralized environment that avoids the instability of fully distributed learning seen in earlier Q-learning approaches.

2.2.3 Software-Defined Networking and Intelligent Routing Frameworks

As VANET research matured, attention turned to more systemic approaches using SDN principles and advanced machine intelligence to manage the network. SDN decouples the control plane from the data plane, meaning a central controller (or a hierarchy of controllers) can gather a global view of the network and make optimized routing decisions, which are then enforced by vehicles (data plane devices).

SDN and Centralized Control: Traditional VANET routing is fully distributed, each vehicle makes its own routing decisions with local information. Software-Defined VANET (SDVN) proposals instead use central controllers (often located at edge servers or RSU) to compute routes with a broader view of network state. Chahal and Harit in [91], introduced an SDN-based two-stage Stackelberg game that selects among multiple wireless access technologies and disseminates data via global network state awareness. Their technique seeks to optimize both the cost of connecting to diverse networks (e.g., LTE, DSRC) and the efficiency of vehicular data sharing. While this method showed promise in handling multi-radio vehicular architectures, it offered limited mobility adaptation: as the protocol did not explicitly incorporate vehicle direction or rapid topological changes, leading to suboptimal performance under high-mobility scenarios. Another framework, MFCAR [92], leveraged SDN to select optimal V2V paths in 5G-enabled vehicular networks. MFCAR applied a graph-based technique to minimize congestion and route multi-flow traffic over short, uncongested links, demonstrating superior performance over existing centralized approaches. However, the solution's reliance on a global controller raises concerns about overhead and real-time adaptability under diverse traffic loads, indicating a need for more robust validation in large-scale, heterogeneous deployments.

Deep Learning-Enhanced Routing: A deep reinforcement learning framework, DROM [93], provided a mechanism for SDN to automatically adapt link weights to enhance network performance (e.g., delay and throughput). Leveraging the DDPG algorithm, DROM used a policy gradient framework to continuously adjust actions based on real-time traffic matrices and achieve lower latency compared to mainstream protocols like OSPF. However, DROM leveraged an M/M/1/K queue-based offline training model to avoid performance degradation during learning phase thus the architecture is explicitly tied to a static or semi-static network topology. Therefore, DROM model must be retrained significantly every time the network topology changes. Going forward, Zhang et al. [94] integrated SDN with deep RL, to create a secure and optimized routing framework known as CFR-RL. CFR-RL treats the SDN controller as a learning agent that observes the entire network and decides routes that avoid congestion and ensure data integrity. This approach yielded optimized data paths and improved network throughput, and it also inherently provided

a way to enforce security (data immutability) via the controller. The main pitfall was the single point of failure, a common concern in SDN. If the central controller crashes or its connection to vehicles is lost, the network falls back to no control (or a default, which might be suboptimal). Cardenas et al. [95] introduced a VANET routing algorithm, GraTree, based on gradient boosting decision trees (a form of ML) to improve route stability. By training a model to predict link quality classes, GraTree's algorithm could choose routes that were statistically likely to remain stable longer. This machine intelligence improved reliability in their experiments, effectively the model learned from past link failures to avoid fragile links. However, GraTree observed slight delays introduced by the model's processing and its coarse class labels. Essentially, using a learning model in the loop added latency, and the discretization of link states could sometimes misclassify a rapidly changing link. This points to a general caution: complex models might optimize one metric but can hurt another (here, added delay). In RL-SDVN [96], an adaptive clustering and Q-learning–based routing mechanism leveraged Gaussian mixture modeling to dynamically form stable clusters and employed an SDN controller to integrate reinforcement learning for optimal route decisions. However, heavy reliance on accurate feature extraction, real-time Q-learning updates, and centralized SDN coordination poses potential scalability and overhead issues in ultra-dense or highly mobile VANET scenarios, indicating a need for further distributed or hybrid adaptations. In our thesis, we incorporate simpler learning agents (e.g., RL with explainable reward metrics) and use well-understood models (car-following, traffic flow models, hypergraph theory) to inject intelligence.

Further, we introduce a hybrid SDN architecture (Chapter 5) that uses both a central controller and distributed edge controllers (at RSUs) working in tandem. Specifically, we propose a master–slave controller hierarchy: a central controller maintains global network state, while local RSU-based controllers handle nearby vehicles to reduce latency. This division's novelty lies in how we offload policies: the central brain computes high-level policies (e.g., weight functions for metrics), and RSUs execute them, adjusting to local conditions. By doing so, we mitigate the single-point-of-failure issue noted in Zhang et al.'s [94] work if one controller fails, others can locally manage. Additionally, we embed a RL module in the SDN controller (central) so that it can continually learn and refine routing policies based on feedback from the network. This design enables real-time adaptation but within a robust, hierarchical control framework that improves fault tolerance.

Interpretable and Secure AI in VANETs: With the infusion of AI into VANET routing, researchers have also pointed out the importance of understanding and trusting AI based models. To address the black-box nature of deep learning models, techniques beyond the networking domain have been exploring the interpretation of the decision-making processes of CNNs and RNNs. Zhang et al. [97] and Feng et al. [98] developed methods for interpreting convolutional and recurrent neural networks, respectively, in their domains (vision and language). The challenge is to transfer this knowledge (interpretation) to network routing decisions, which involve spatio-temporal data and graph topology, quite different from images or text. As a result, only a few works directly tackled interpretability in VANET ML. Meng et. al., [97] proposes Metis, a framework that addresses the critical challenge of interpreting black-box deep neural networks (DNNs) in networking systems. Metis transforms inscrutable DNNs into human-readable interpretations through a two-pronged approach: decision trees for local control systems and hypergraph analysis for global control systems. We take our inspiration from this framework, however, Metis cannot capture temporal dependencies or contextual reasoning within the network, which are often critical in networking environments. Therefore, we include spatio-temporal network characteristics while proposing our hypergraph based framework. These efforts indicate that as networking algorithms become more intelligent, there's a parallel need to ensure they are transparent and can be debugged or verified. Our thesis, while not specifically an interpretability study, keeps this principle in mind. By using algorithms with explainable parameters (e.g.,

number of aligned neighbors, measured delay, queue length) and avoiding obscure features, we make it easier to reason about why a given route was chosen. Additionally, the hierarchical control in our hypergraph based approach (Chapter 5) inherently allows logging and analyzing decisions at the controller, which aids in interpretability (a benefit of SDN in general, decisions can be audited).

Summary of SDN/Intelligent Frameworks and Thesis Positioning: The introduction of SDN to VANETs is a promising direction, offering a path to manage complexity by leveraging infrastructure. However, these approaches are highly dependent on deployment feasibility, they require RSUs, computing at the edge, and reliable V2I links, which only now with 5G are becoming realistic. Indeed, field trials in Europe (ConVeX [27], 5G-CARMEN [29], etc.) are evaluating 5G-based V2X along highways with multi-access edge computing support, and reporting robust connectivity and sub-100ms latencies in complex scenarios. These trials show that the network side (infrastructure + 5G) can handle fast communication, opening the door for centralized routing logic to be applied practically.

2.2.4 Emerging Technologies: C-V2X and Edge/Fog Computing

It is important to place VANET routing in the context of emerging network technologies that will shape next-generation systems. 5G and beyond wireless networks, mobile edge computing (MEC)/fog computing, and networking concepts are beginning to influence VANET research. While these technologies are in various stages of maturity, practical deployment is still largely experimental. Integrating these emerging technologies is part of the forward-looking aspect of this thesis, though our core contributions remain practical for near-term deployment.

5G and C-V2X: The arrival of 5G New Radio (NR) and Cellular-V2X is a game changer for vehicular networks. Compared to 802.11p, C-V2X offers greater range and reliability by leveraging cellular infrastructure and advanced radio techniques, while 5G NR V2X (Release 16 and beyond [19]) promises even lower latencies (down to 10ms-level) and high throughput. Real-world projects have already tested 5G-based vehicular communications: for example, Qualcomm and Huawei demonstrations achieved sub-30ms latency in field tests using 5G for V2V messages [21]. European cross-border trials (5G-MOBIX [30], 5G CROCO [99]) have shown that MEC (edge computing) deployed alongside 5G can maintain stable connectivity and latency under 100ms even as vehicles cross national network boundaries. These results underscore that the wireless and infrastructure platform is becoming robust enough to support advanced routing algorithms: we can now assume vehicles often have an available roadside connection and a cloud or edge server that can assist with computations.

From a routing perspective, 5G and MEC enable two things: (1) more timely and rich global information (via network slicing, vehicles can share their state to edge servers that aggregate a global view), and (2) offloading computation-heavy tasks to edge/cloud. The literature reflects this through proposals like digital twins for traffic (keeping a real-time digital replica of the road network on the cloud) or offloading parts of route calculation to MEC nodes. Nie et al.'s [100] digital twin with DQN is one example of leveraging powerful computation and global data. Another is the work of Duan et al., [101], it offered a workload distribution solution across edge servers; however, the variability in vehicle data processing affects workload distribution. Another framework, MobiPlace [102], utilized Markov predictors for mobility management and simulated annealing for dynamic controller placement. In the same direction, in [103], authors utilized existing urban infrastructure for placing cloudlets, which is cost-effective and leveraged widespread existing assets like street lamps. Although it reduces installation costs and uses familiar urban elements,

the technique faces challenges against dynamicity and is not adaptable to traffic conditions. In practice, initiatives like the one by the Virginia Department of Transportation with Audi (C-V2X RSUs) show that infrastructure is being deployed to send information to cars (e.g., traffic signal status, work-zone warnings) [26]. The logical next step is two-way: vehicles not only receive data but also send their local view to edge servers, which then compute optimized dissemination strategies (e.g., if an accident is reported, the edge could coordinate multi-hop forwarding of that alert across a region). Our SDFC framework is aligned with this vision, it essentially describes how to orchestrate such edge-based coordination in a VANET, with the parallel dynamic clustering algorithm enabling parallel processing of different tasks.

It's worth noting that current pilot deployments mostly use simple routing/dissemination: for instance, the US pilots used DSRC to broadcast hazard messages to nearby vehicles (one-hop), and in some cases, multi-hop via roadside repeaters for V2I, but not much beyond flooding with priority. With 5G MEC, researchers are now starting to test multi-hop forwarding under central coordination. The literature doesn't report fully autonomous, infrastructure-coordinated routing yet, but the building blocks are being laid. Our thesis does not directly depend on 5G, but it assumes a network context where high data rate and low latency links are available (which 5G provides) and that vehicles might periodically connect to an edge cloud. This assumption is supported by global trends: many countries have allocated dedicated spectrum for V2X and are encouraging deployments.

2.2.5 Chapter Summary and Thesis Contributions

For ease of understanding and systematic evaluation of vehicular communication frameworks, Table 2.2 presents a parameter classification that categorizes the essential elements of vehicular communication systems into four fundamental dimensions. The Communication Environment Profile encompasses aspects

Table 2.2: Parameter Taxonomy for the Frameworks Comparison

Dimension	**Parameters**
Communication Environment Profile	E1: Communication Mode
	E2:Global Communication Mode
	E3: Network Softwarization
	E4: Cloud/Edge/Fog
	E5: Access Layer Technology
Network and Topological Properties	NT1: Network Simulator and Traffic
	NT2: Network Topology
	NT3: Mobility Model
Objectives	O1: Clustering
	O2: QoS Routing
	O3: Directional Awareness
	O4: Resource Management
	O5: Load Balancing
Evaluation Type	Ev1: Simulation
	Ev2: Empirical

such as communication modes, network softwarization approaches, and access technologies that define the operational environment. Network and Topological Properties capture the simulation platforms, network structures, and mobility patterns. The Objectives dimension identifies functional goals including clustering

mechanisms, quality-of-service routing, and resource management strategies, while the Evaluation Type distinguishes between simulation-based and empirical validation approaches. Table 2.3 provides a detailed cross-comparison of distinct frameworks against some established parameters. Through this literature re-

Table 2.3: Comparison of Frameworks Based on Communication, Network, Objective, and Evaluation Properties

Framework	Environmental Properties					Network and Topological Properties			Objectives					Evaluation Metric	
	E1	E2	E3	E4	E5	NT1	NT2	NT3	O1	O2	O3	O4	O5	Ev1	Ev2
M. Chahal [91]	V2V, V2I	Centralized	✓	×	DSRC	NS3, SUMO	OpenSteetMap (Chandigarh, India)	Freeway Mobility	×	✓	×	✓	×	✓	×
ACA [76]	V2V	Distributed	×	×	NA	NS-2, SUMO	OpenStreetMap (San Jose, California)	NA	✓	×	✓	NA	NA	✓	×
TDMP [87]	V2X	Local	×	×	DSRC	OMNeT++, SUMO	OpenStreetMap (Loughborough, UK)	Car following model/Krauss	×	✓	✓	NA	NA	✓	×
IV2XQ [86]	V2V, V2I	RSU based	×	×	DSRC	OMNeT++, SUMO	38 two-way segments, 24 intersections	NA	×	×	✓	×	✓	✓	×
MFCAR [92]	V2V, V2I	Centralized	✓	✓	802.11p	Custom, SUMO	Urban	Manhattan Model	×	✓	×	×	×	✓	×
OPBR [88]	V2V	Local	×	×	DSRC	NS2	Urban	Linear	×	×	✓	×	×	✓	×
GSO-RBVT-R [78]	V2V, V2I	Local	×	×	DSRC	MATLAB	NA	NA	×	✓	×	×	×	✓	×
Rajput et al. [81]	V2V	Local	×	×	DSRC	Custom	Highway, City	Graph-based	×	✓	×	✓	×	✓	✓
ACO-AODV [77]	V2V	Local	×	×	DSRC	NS2, VANETMobiSim	Urban, Highway (North Delhi)	Intelligent Driver Model	×	✓	×	×	×	✓	×
MOCRAW [80]	NA	Static	×	×	DSRC	NS2	Random static clustering	NA	✓	×	×	×	×	✓	×
MHCRT [82]	Peer-to-peer	NA	×	×	NA	Custom Simulation	Static	Static	✓	×	×	×	×	✓	×
GKO [79]	V2V	Static Sensors	×	×	DSRC	Custom	Static	Static	✓	×	×	×	✓	✓	×
DROM [93]	NA	Centralized	✓	×	NA	OMNeT++	Gravity Model	Static backbone topology	×	✓	×	✓	✓	✓	×
JDGO [104]	V2V, V2I	Distributed	✓	✓	NA	Python + TensorFlow	Custom	Gaussian Distribution	×	✓	×	✓	✓	✓	×
BBO [105]	V2I, V2V	Distributed	×	✓	C-V2X	Python	Zonal clustering	Gaussian	✓	✓	×	✓	✓	✓	×
MobiPlace [102]	V2V, V2I	Distributed	✓	✓	DSRC	MATLAB	Highway, Urban	Real trace (Gauss-Markov)	×	✓	×	×	✓	✓	×
SDCast	V2V, V2I	Centralized	✓	×	DSRC	NS3, SUMO, OpenStreetMap	Highway, Urban	Microscopic (Manhattan)	✓	✓	✓	×	×	✓	×
MRMH	V2V, V2I	Local, de-centralized	×	×	C-V2X	NS3, SUMO, OpenStreetMap, MATLAB	Highway, Urban	Grid, Random waypoint, car-following	✓	✓	✓	×	✓	✓	×
Hypergraph	V2V, V2I	Decentralized	✓	×	C-V2X	OMNeT++, SUMO, VEINS, Testbed	Urban, Highway, Grid	Hypergraph modeled	×	✓	✓	✓	×	✓	✓
SDFC	V2V, V2I	Decentralized	✓	✓	C-V2X	OMNeT++, SUMO, VEINS, Testbed	Urban, Highway Grid	Microscopic	✓	✓	✓	✓	✓	✓	✓

view, we have identified a trajectory in VANET routing and data dissemination research: from early adaptations of MANET protocols, to clustering and heuristic optimizations for scalability, to QoS-driven and predictive algorithms, and finally to the integration of SDN, edge computing, and other intelligent frameworks. We have also distinguished between approaches that are simulative/theoretical and those closer to real-world implementation:

Implemented or Standardized: Basic beaconing and flooding protocols in DSRC/WAVE, simple geo-forwarding rules in ETSI ITS-G5 [7], and recently, C-V2X one-hop broadcasts with priority (as proven in pilot deployments) are reliable but relatively used unsophisticated strategies due to deployment necessity.

Proposed (Simulated) in Research: Virtually all advanced protocols, ML-based routing, complex clustering, meta-heuristic algorithms, SDN-based control, etc. fall in this category. They are validated in simulators like NS-2, NS-3, or OMNeT++ with tools such as SUMO for mobility, but have not been

trialed on real roads. The gap between simulation and deployment is non-trivial, often due to deployment constraints, traffic modeling, or neglecting driver behavior nuances.

The contributions of this thesis tries to address the gaps identified in the literature:

- We observed that many clustering approaches improved link stability but introduced too much overhead or were not adaptive to all conditions. Our solution (CSBR and extended clustering in later chapters) uses a novel but simple metric (cosine similarity of velocity vectors) to form and maintain clusters, achieving stability with low overhead. It explicitly tackles the broadcast storm problem via weighted probabilistic forwarding, something only few prior works attempted, and it does so in a way that is analytically tunable. This adds to the literature a practical method for direction-based clustering, which we validate in simulation, demonstrating its viability (Chapter 3 results).
- While prior QoS routing proposals often optimized one aspect (say delay) at the expense of others or required heavy (in terms of computations) algorithms (ACO, etc.), our OBQR protocol contributes a multi-constraint routing mechanism that is computationally efficient. By scalarizing multiple QoS metrics into one composite cost, OBQR can run Dijkstra-like computations in near real-time, yet still account for diverse metrics (link reliability, delay, etc.). We enhance this with live traffic awareness (using car-following models), which uses microsimulation models for routing in VANETs. This approach yields more stable and QoS-compliant routes under high mobility than any single-metric or static algorithm could, thus improving upon works which struggled in highly dynamic environments.
- Recognizing the promise of SDN but also its pitfalls, the thesis introduces SDCast, a hybrid SDN-based routing protocol that implements a two-tier control architecture in a VANET context. By doing so, it explicitly addresses the single controller failure and scalability issues noted in SDN literature. We demonstrate that SDCast outperforms purely distributed and purely centralized protocols in terms of adaptiveness and reliability, effectively marrying the best of both worlds (local responsiveness and global optimization).
- In the area of meta-heuristics and optimization, the MetaLearn and MRMH framework combine techniques that were previously used in isolation. GWO + TDL-learning in MetaLearn accelerates convergence and avoids local optima, directly addressing the shortcomings of algorithms like Particle Swarm Optimization (PSO) and ACO which could stagnate or require manual tuning. MRMH's hybridization of GWO and SQP ensures that, unlike prior purely heuristic methods, our routing solutions meet all QoS constraints even in large-scale scenarios. These contributions, detailed in Chapter 4, show improved performance over baseline meta-heuristics in simulation. While MetaLearn/MRMH are advanced algorithms (not yet ready for immediate deployment), they are designed to be parallelizable and offloadable to edge nodes.
- Finally, we contribute an validation methodology (Chapter 5) that uses both large-scale simulation and physical experimentation (including a scaled-down vehicular testbed, duckietown and Edge/Fog based Infrastructure test bed) to evaluate our proposed protocols. This addresses the literature gap where most protocols lack any physical testing. By using tools like NS-3 integrated with SUMO for realistic scenarios and a Duckietown mini-vehicle testbed for empirical tests, we ensure that the proposed solutions are not only theoretically sound but also practically viable. Our work therefore provides confidence that the improvements we propose can translate to real-world benefits.

In conclusion, the literature review reveals that routing and data dissemination in VANETs have seen advancements, but also that many proposals remain confined to academic exploration due to various practical hurdles. By critically analyzing existing work, we identified key requirements for next-generation VANET routing: adaptability to topology changes, scalability to dense networks, multi-metric QoS support, low overhead operation, and integration with emerging network infrastructure, all while keeping an eye on implementability. The contributions of this thesis encompassing direction-aware clustering, multi-constraint QoS routing, SDN-enabled adaptive control, meta-heuristic optimization, hypergraph based dissemination collectively address these requirements. The next chapters will detail each contribution and demonstrate quantitatively how they outperform previous approaches in the context of the existing literature.