Abstract

Consider a driver on a busy highway where their vehicle immediately receives alerts about a collision occurring three cars ahead, well before it comes into their line of sight. Or imagine a navigation system dynamically rerouting the driver to avoid a newly formed traffic jam just a few kilometers away. This is the capability offered by Vehicular Ad-hoc Networks (VANETs), which enable direct communication between vehicles and roadside infrastructure, establishing a real-time digital network that enhances road safety and optimizes traffic flow. However, deploying VANETs is complex. On actual roads, particularly in developing countries, traffic is highly heterogeneous ranging from cars and buses to motorcycles and auto-rickshaws, all traveling at varying speeds, frequently changing lanes unpredictably, and constantly joining or leaving the network. Traditional communication protocols, originally designed for relatively stable and homogeneous networks, struggle under these conditions: communication links frequently break, message flooding occurs in dense traffic scenarios (known as "broadcast storms"), and critical safety messages that require delivery within 100 milliseconds often fail to meet the stringent latency requirements. This thesis addresses these fundamental networking challenges (packet routing) by developing and validating next-generation routing and data dissemination techniques that maintain reliability despite high node mobility and diverse traffic conditions. Our contributions encompass three different directions: direction-aware forwarding mechanisms, metaheuristic-based clustering frameworks, and hypergraph-based communication models.

We first propose three orientation-informed routing protocols, Cosine Similarity-Based Routing (CSBR), Orientation-Based QoS Routing (OBQR), and SDCast leveraging vehicles' movement direction and road context to improve data dissemination. Unlike conventional broadcast or shortest-path schemes, these protocols dynamically bias message forwarding along the direction of traffic flow, reducing redundant transmissions and avoiding relays on vehicles that are likely to move out of range. In CSBR, a cosine similarity metric between vehicle velocity vectors is used to select relay candidates, ensuring that only vehicles with aligned directions participate in rebroadcasting. OBQR builds on this by incorporating multi-constraint QoS metrics (link stability, transit delay, etc.) into routing decisions, using a weighted optimization to find routes that honor safety-critical latency and reliability requirements. Finally, SDCast introduces a hybrid Software-Defined Networking (SDN) architecture into the VANET: a two-tier controller system (a central controller working with local Roadside Units) that orchestrates cluster-based forwarding policies.

To improve communication stability and QoS in dynamic conditions, we next develop advanced clustering and routing optimization techniques using metaheuristics. Two frameworks, i.e., MetaLearn and Multi-constraint Routing using Hybrid Metaheuristics (MRMH) are introduced to intelligently organize vehicles into semi-stable clusters and optimize multi-hop routes within and between these clusters. MetaLearn employs a hybrid learning approach: it uses meta-heuristic algorithms (Grey Wolf Optimization (GWO)) to bootstrap efficient clustering, and then applies reinforcement-learning principles (fast adaptation based on prior outcomes) to continually refine routing policies as conditions change. This enables the routing strategy to "learn" from the network's behavior, quickly adapting to recurring traffic patterns (e.g., rush hour flows) and thereby improving long-term performance. MRMH, on the other hand, hybridizes multiple optimization techniques (GWO and Sequential Quadratic Programming (SQP)) to solve the routing problem under multiple constraints (such as latency, link durability, and bandwidth) simultaneously. By hybridizing metaheuristics methods, MRMH avoids the pitfalls of single-metaheuristic approaches (like premature convergence or high computational cost) and finds high-quality routes that satisfy all QoS requirements even as the network scales.

Finally, we present an approach using Spatio-Temporal Information-Aware Hypergraph formulation that generalizes the traditional network graph model to a hypergraph structure. In a hypergraph, an edge (now called a hyperedge) can connect any number of vertices, which in our context means a communication event can directly involve multiple vehicles. This representation is paired with a deep learning-driven routing strategy that uses spatial (geographic/positional) and temporal (time-dependent) dynamics of vehicles and network conditions to make optimized decisions. By capturing higher-order relationships (beyond simple pairwise links) and feeding them into a deep learning algorithm, the network can better anticipate and adapt to changes. Further, we introduce a Software-Defined Fog Computing (SDFC) framework for VANETs, which pushes computational intelligence and control closer to the network edge (the vehicles and roadside units). This enables data processing and decision-making to occur in proximity to where data is generated. By doing so, fog computing can drastically reduce end-to-end communication delays and offload traffic from the core network. In our SDFC framework, VANET management functions (such as cluster formation, routing control, and load balancing) are distributed across a hierarchy of cloud, fog, and edge layers. This design improves scalability and reliability by avoiding single points of failure and by adapting to local conditions.

To ensure experimental validations, all proposed techniques were implemented using standard VANET simulation tools and real hardware. Simulations leveraged frameworks like NS-2/NS-3 and OMNeT++ (with the Vehicle in Network Simulations (VEINS) open source library) for network-layer behavior, standard Vehicle-to-Everything (V2X) communication technology (IEEE 802.11p, Cellular-V2X) and Simulation of Urban Mobility (SUMO) for generating realistic vehicle mobility on road layouts imported from OpenStreetMap (openly-licensed data from national mapping agencies and other sources). We seeded simulations with actual city road maps and traffic patterns (including heterogeneous vehicle types, intersections and traffic lights) to closely mirror real-world conditions. Key performance metrics, end-to-end latency, packet delivery ratio, routing overhead, cluster membership time, throughput, and route discovery time were measured across a range of scenarios (urban environments, highways, varying vehicle densities from sparse to congested). Furthermore, the algorithms were tested on a physical testbed: our Duckietown setup (miniature autonomy testbed) and an edge computing platform with Raspberry Pi and JetsonNano devices (working as Onboard Units and Roadside Units) allowed us to verify that the protocols run within real-time constraints on resource-constrained hardware.