Performance Analysis of MANET Routing Protocols in Urban VANETs

Antonio Di Maio[1(✉)], Maria Rita Palattella[2], and Thomas Engel[1]

[1] SnT - University of Luxembourg, Luxembourg City, Luxembourg
antonio.dimaio@uni.lu

[2] Luxembourg Institute of Science and Technology, Esch-sur-Alzette, Luxembourg

Abstract. Infrastructure-less communications between moving vehicles present emblematic challenges because of high node mobility and link volatility, which may harm the performances of different categories of emerging vehicular applications. In order to move data between vehicles that are not in direct communication range, several distributed routing protocols have been proposed and tested in vehicular networks, highlighting their strengths and weaknesses. Some previous works report disagreeing claims about routing protocol performances in similar vehicular scenarios. Therefore, in this work, we evaluate the performances in terms of Packet Delivery Ratio (PDR), packet delay, frame collision rate, and signaling rate of three well-known routing protocols (AODV, DSDV, and GPSR), simulating them in a realistic Manhattan scenario. Furthermore, we evaluate the impact of typical urban obstacles (e.g. buildings) on the considered performance metrics. We observed that, in the proposed urban scenario, AODV provided the best PDR, GPSR the best packet delay, and DSDV failed to provide satisfactory performances due to signaling-induced congestion. Simulations showed that considering the shadowing effects induced by the buildings in an urban scenario drastically changes the observed performances, i.e. reduces the frame collisions, decreases the PDR, and increases the packet delay.

Keywords: Mobile ad-hoc networking · Vehicular networks · Routing · AODV · DSDV · GPSR · Simulation · Performance evaluation

1 Introduction

Connected vehicles of the future will provide users with a wide range of different applications that will need to exchange information with high data rate, high reliability, and low communication delay. For example, some modern vehicular applications that aim at improving the users' safety (e.g. emergency remote control [10]) require a video data throughput of up to 4 Mbit/s and a control packet

This work was supported by the FNR CONTACT project, CORE/SWISS/15/IS/10487418.

M. R. Palattella et al. (Eds.): ADHOC-NOW 2019, LNCS 11803, pp. 432–451, 2019.
https://doi.org/10.1007/978-3-030-31831-4_30

delay in the order of milliseconds. The incoming revolution in vehicular networks brought by 5G will heavily rely on heterogeneous and Device-to-Device (D2D) communications, which will have the potential to offload a substantial part of vehicular data traffic from the core network infrastructure to the Vehicular Ad-Hoc Network (VANET). VANETs are a particular class of Mobile Ad-Hoc Networks (MANETs), in which the network nodes are extremely mobile and the communication links are very volatile. Vehicles in a VANET will be able to exchange data over multi-hop routes when they are not in direct communication range, and several protocols to compute these routes have been proposed so far. Due to the specific characteristics of dynamicity and instability of VANETs, we would expect that applying traditional MANET routing protocols to VANETs, without special adjustments, would lead to suboptimal network performances. This issue has already been addressed in literature, yet without reaching a unanimous understanding about the achievable performances in these scenarios. For this reason, we conducted further simulations and analyses to determine whether MANET routing protocols are able to satisfy vehicular applications' requirements.

Among all the proposed routing protocols, we chose to compare Ad-hoc Online Distance Vector (AODV) [18], Destination-Sequenced Distance-Vector Routing (DSDV) [19], and Greedy Perimeter Stateless Routing (GPSR) [11] because they are the best-known protocols belonging to their own categories: reactive, proactive table-driven, and proactive position-based protocols, respectively [4]. Since they have been the most studied routing protocols, they are also the ones about which the highest number of contrasting claims has been reported (see Sect. 3).

The rest of the article is organized as follows. In Sect. 2, we will describe the operation and characteristics of the selected MANET routing protocols. In Sect. 3, we will present some previous studies about their performances and discuss their similarities, strengths, and limitations. In Sect. 4, we explain the rationales behind the choice of the simulation parameters and their expected impact on the performance. In Sect. 5, we analyze and present the findings originated from the data collected from the simulations. Section 6 concludes the article and highlights future research questions.

2 Routing Protocols for MANETs

There exist several distributed routing protocols for MANETs, whose main aim is to provide *next-hop* information to the intermediate nodes along the path between source and destination. Routing protocols can be classified into *reactive* and *proactive* protocols, according to the events that trigger the exchange of signaling traffic [4,16]. Reactive protocols exchange signaling packets to establish a valid route between source and destination only when the source needs to send information. Conversely, proactive protocols exchange signaling packets regardless of the presence of active data traffic between nodes. In the category of proactive protocols, some algorithms exchange periodic signaling packets to update the local network view, whereas others rely on external mechanisms (e.g. Link-Layer Acknowledgements) to trigger network-wide network status updates.

Hereafter, we will describe the operation and the characteristics of the selected routing protocols: AODV, DSDV, and GPSR.

2.1 AODV

Operation—AODV [18] is a reactive routing protocol for MANETs, thus the routes from source to destination are created upon request. When a source node needs to send data to a destination node and does not have a valid route entry in its routing table, it initiates a route discovery by broadcasting a Route Request (RREQ) message to its neighbors. Upon reception of the RREQ, each neighbor checks whether it has already received a RREQ for the same route discovery and, if not, it checks if it has a route entry in its routing table to reach the destination. If the node cannot find a valid entry, it rebroadcasts the RREQ to its neighbors. Otherwise, if it has a valid entry to reach the destination (or the node is the final destination itself), it unicasts a Route Reply (RREP) message to the route discovery originator. The route discovery process is represented in Fig. 1b. The unicast route gets finally established when the RREP reaches the source node, because each intermediate node creates reverse and forward path entries in its routing table as the RREQ and RREP messages cross the network. If any of the intermediate nodes along the active route diappears, the node upstream of the link break will detect the topology change (e.g. with missing Link-Layer Acknowledgements) and unicast a Route Error (RERR) message back to the source. This RERR message informs every intermediate node about the topology change, and they will accordingly modify their routing tables. At this point, the source node can reinitiate a route discovery process.

Features—The routes generated by AODV are guaranteed to be loop-free because of the node sequence numbers that are associated to each signaling packet. The protocol offers many techniques for optimizing its operations, such as a local repair for broken routes or gratuitous RREPs for efficient bidirectional route instantiation. For sake of simplicity, they have not been considered and studied in the present work.

2.2 DSDV

Operation—DSDV [19] is a proactive table-driven routing algorithm for MANETs. Each mobile node periodically broadcasts information about viable routes to reach every other destination node in the network to their one-hop neighbors (Fig. 1c). In particular, the distributed information is a set of route entries, each of them associated with the distance in number of hops (or any other metric) between the sender and the route destination, accompanied by a *sequence number*. This data structure is referred to as the *distance vector* (DV).

The *sequence number* is needed to maintain only the freshest route entry received by a node, and to guarantee that the computed routes are loop-free.

Its value is determined exclusively by the destination node, and it is an even number when the associated route is viable. When a link break occurs, all the route entries towards destinations that were routed through the unreachable hop are modified by setting their metrics to $+\infty$ and setting the relative sequence number to the next odd number. After having done this, the intermediate node must immediately advertise the event to the neighbors by broadcasting an update that contains the new route entries.

In order to reduce network signaling traffic, the routing-protocol updates can be incremental, rather than full dumps. Full dumps happen less regularly than incremental updates, and convey all of the information stored in the routing table of a node to its neighbors. Incremental updates happen more frequently than full dumps and can be either periodic or triggered by significant events (e.g. immediately after a link break).

Features—One of the biggest drawbacks of applying traditional DV-based routing protocols to highly dynamic MANETs is that the routing tables in the mobile nodes could contain stale network information status. Furthermore, small inaccuracies of the network state contained in each router can lead to routing loops. Nevertheless, DSDV is immune to routing loops because of the embedded sequence numbers in its signaling packets. DSDV suffers from low scalability: each node is required to maintain a routing table entry for every destination in the network, determining a linear space complexity $\mathcal{O}(n)$ of the protocol.

2.3 GPSR

Operation—GPSR [11] is a proactive position-based routing algorithm for MANETs. Each vehicle encapsulates its address and geographical position in a *beacon* and broadcasts it to its one-hop neighbors, which use it to build a *neighbors list*. When a node does not receive a beacon from a neighbor after an expiration time, it deletes that neighbor from its neighbor list. When a node needs to send a packet to a destination, it forwards it to the neighboring node that is closest to the destination. In case there is no node closer to the destination than the sender, the packet goes in *perimeter mode*. When a packet enters the perimeter mode, it temporarily gets routed further from the destination with the hope of finding a route that goes around the void area (Fig. 1d).

In specific cases, the perimeter mode might fail to find an existing route to the destination. This might happen especially when the graph is *non-planar*, i.e. with crossing edges. For this reason, GPSR must employ some *planarization* techniques. The first naive approach is applying the *no-crossing* heuristic, in which a random crossing edge is removed from the graph. The disadvantage of this technique is the possibility of partitioning the graph. More sophisticated approaches are the *Relative Neighborhood Graph* (RNG) and the *Gabriel Graph* (GG) planarizations. In short, they work by checking, per each couple of nodes in the network, if a third node is present in an area between them. If this is the case, the edge between the two considered nodes is removed from the graph. For RNG, the area is shaped as the intersection of two circles centrered on the selected nodes. For GG, the area is a circle centered at the median point of the segment connecting the two selected edges, with a diameter equal to the distance between the two nodes.

Features—GPSR is more scalable than other table-driven routing protocols because it needs and stores only local information regarding the network topology. The protocol's signaling rate is (i) constant per each node in the network, (ii) depends uniquely on the beaconing frequency, and (iii) is unrelated to the network traffic load, vehicular mobility, and vehicular density. The reaction speed to topology changes is influenced by the beaconing frequency and the vehicular density and mobility. When the beaconing frequency is low, the signaling rate is low but the network registers topology changes slowly. When the beaconing frequency is high, the signaling rate is higher but the neighbors lists are updated more promptly. In case the local product of beaconing frequency and vehicular density is very high, the contention-based access to the channel can influence the *freshness* [12] of the neighbors' positions stored in the local neighbors list.

One further foreseeable problem of stateless routing protocols like GPSR is that they do not remember poor routing decisions. If a packet enters a local maximum and travels around a void in perimeter mode, so will all the subsequent packets. Each packet will run through the same suboptimal route discovery, leading to longer paths and lower PDR.

2.4 Signaling Traffic Analysis

The signaling rate of proactive protocols like GPSR and DSDV depends exclusively on the beaconing frequency or on how dynamic and large the network topology is, and never on the rate of creation and duration of new network flows. On the contrary, reactive protocols like AODV have a signaling rate that depends on the mobility of the relay nodes along active routes, but also on the rate of creation and duration of new network flows. This is because every time a route is distrupted, the route discovery process is restarted or the route repair procedure is initiated. For reactive protocols, the neighborhood of the active routes is flooded with information on how to update routing tables at each new flow instantiation or disconnection of active relay nodes. For proactive protocols, the whole network is flooded with information on how to update the routing tables of every node for each topology change in the network.

The overall signaling traffic is much higher for DSDV because local topology-change information must be iteratively spread throughout the whole network, whereas for reactive protocols the topology changes trigger signaling traffic in the vicinity of the route (i.e. route repair mechanisms or route discovery repetition). However, we mostly care about the signaling traffic that is generated in the vicinity of the active routes, as it might induce congestion and application packet loss. Under this point of view, the signaling rates of DSDV and of reactive protocols become comparable, even though we still expect reactive protocols to produce lower signaling traffic.

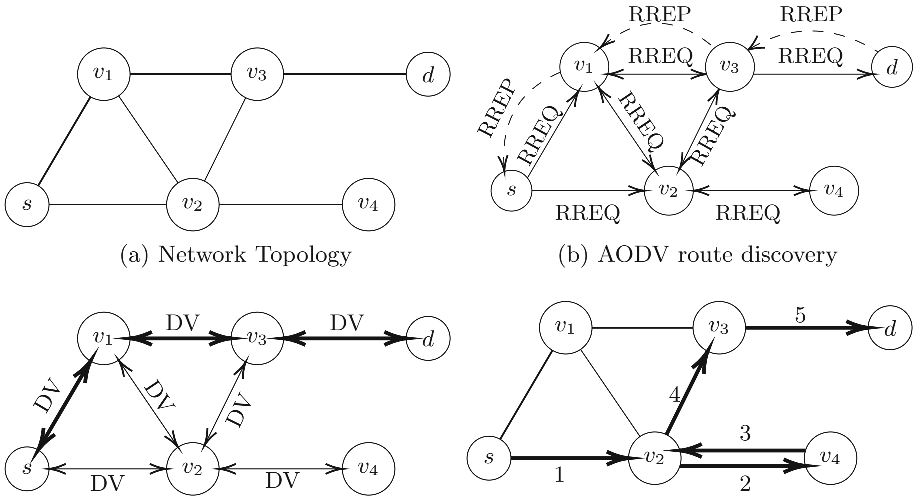

(a) Network Topology

(b) AODV route discovery

(c) DSDV route discovery

(d) GPSR perimeter mode

Fig. 1. Features of AODV, DSDV, and GPSR. On (a) an example of VANET connectivity graph, with a source vehicle s and a destination vehicle d, (b) illustrates the AODV's route discovery process, (c) shows the DSDV's distance vector broadcasting, with the computed route in bold, and (d) shows the route of a packet in GPSR perimeter mode to bypass a void area, where the numbers are the sequence of links on which the data packet is forwarded.

3 Related Works

Broch et al. [5] simulate AODV and DSDV in a MANET scenario, and evaluate their performances in terms of PDR and signaling traffic. This work does not consider the presence of obstacles (e.g. buildings) that might shadow wireless transmissions and that would have a dramatic impact on networking performances (as we show in Sect. 5). The mobility model used in this study is the Random Waypoint (RWP) model [3], in which nodes pick random destinations, move towards them at a constant speed and, after reaching the target, wait for a constant number of seconds defined as the *pause time*. This work studies the influence of pause time on the different aforementioned metrics, per each evaluated routing protocol. The results obtained by this pioneering piece of work are not directly applicable to VANETs because the RWP mobility model is not suitable to approximate realistic vehicular mobility patterns [24]. It is therefore indispensable to perform simulations based on realistic vehicular mobility models, such as those offered by Simulation of Urban MObility (SUMO) [14].

Karp and Kung [11] propose, describe, and evaluate GPSR in a simulated 1500 m × 300 m scenario without obstacles, in which nodes move according to a RWP mobility model and have a 250 m transmission range. Several other works (e.g. [5]) also use this scenario for their evaluations. The study compares GPSR and Dynamic Source Routing (DSR) [9] in terms of PDR and signaling traffic, showing the superior scalability and performances of their protocol for both metrics. As claimed by the authors, an elongated, rectangular scenario *forces the use of longer routes between nodes, compared to a square scenario with identical node density*. However, the lack of physical obstacles, combined with the high ratio between the fixed transmission range and the length of the smaller edge of the scenario, generates mono-dimensional routes and hardly allows GPSR to go into *perimeter mode*. This means that, in this scenario, the source can almost always reach the destination with greedy forwarding alone. This justifies why in both evaluations [5,11], the routes generated by the evaluated protocols are almost always as short as physically possible. From these two articles it is not evident how the authors evaluated the behavior of MANET routing algorithms when the data packets must route around an obstacle or a void area. For this reason, in our work, the scenario's aspect ratio and the vehicles' transmission ranges are set so that the routes have a higher degree of freedom and are more likely to be forced to route around a void area.

Naumov et al. [17] compare the performances of AODV and GPSR for the RWP mobility model and for a realistic vehicular mobility model. The study shows that the performances of both routing protocols are dramatically lower when considering realistic vehicular mobility models. GPSR outperforms AODV in the RWP mobility scenario and the opposite happens in the realistic vehicular mobility scenario. In both scenarios, 550 vehicles move at urban speeds within an area of 6 km^2. Shadowing effects are taken into account through analytical models that ignore the specific location of the obstacles. GPSR's beaconing interval is set to 3 s (a relatively high value) because the authors claim that higher values would significantly increase frame collisions. Unsurprisingly, they also report that nodes' neighbor tables are often stale, causing up to 80% of next-hop forwarding decisions to be incorrect. The majority of the currently-standardized protocol stacks for vehicular networks (i.e. IEEE WAVE and ETSI ITS-G5) do not recommend any specific beaconing frequency but, for several delay-sensitive applications in the literature, the beaconing frequency typically ranges from 10 Hz [21]. For this reason, for our performed simulations, GPSR's beaconing interval is set to 100 ms without observing any detrimental congestion effects.

In [20], the authors compare the packet delay performances of AODV and DSDV in the context of vehicular safety applications in VANETs. The article proposes a cooperative collision avoidance application on highways: the overtaking of a vehicle at the end of a line of vehicles is canceled if another vehicle occupies the lane in the opposite direction. The simulation lasts 10 seconds and the highway scenario contains from 3 to 12 vehicles, each provided with a 1 Mbit/s IEEE 802.11b interface and a 250 m nominal transmission range. The authors claim that DSDV is the only protocol able to support safety applications in VANETs and they discourage the use of AODV.

Ali et al. [1] measure PDR, packet delay, and packet burst loss for AODV and GPSR both in a Manhattan grid and in a section of London's map. The simulated fading model (Nakagami Two Ray), MAC Layer (IEEE 802.11p), and vehicular mobility (SUMO) are realistic. The scenarios contain 100 vehicles that move at urban speeds and share between 200 and 1000 data flows containing 100 packets each, with a duration of 10 s/flow. The study claims that GPSR provides the most stable PDR and lowest packet delay under different traffic load scenarios.

In [2], the authors evaluate PDR, packet delay, throughput, and signaling traffic of AODV, DSDV, and GPSR in a realistic urban map of Oujda, Morocco. The vehicular density ranges between 20 and 90 vehicles over an area of 1.7 km $\times$ 1.5 km, with realistic mobility and moderate urban speeds. The data flow is modeled as a UDP CBR stream of 5 packets/s and 512 B/packet. The study claims that DSDV provides a high PDR and a high throughput, whereas GPSR provides the lowest signaling traffic and packet delay. AODV shows the highest signaling traffic among all the evaluated protocols.

Inconsistent Claims About Routing Performances—Many of the cited performance studies [1,2,11,17,20] unanimously claim that proactive protocols provide the best packet delay in a variety of scenarios. However, we also notice that several works claim discordant findings regarding PDR, throughput, and signaling rate for the selected routing protocols, even in similar vehicular scenarios and network conditions. For example, [5,17] claim that reactive routing protocols provide a better PDR and a better throughput than proactive protocols, whereas [1,2,11] claim the opposite. Regarding signaling rate, [5] claims that reactive protocols outperform proactive protocols, whereas [2,11,17] claim the opposite. In order to provide new evidence to solve the disagreement, we hereby investigate the performances of some of the most widely studied MANET routing protocols to verify the claims reported in previous articles, and check whether traditional MANET routing protocols can satisfy the performance requirements of future VANET applications.

4 Simulation Setup

The present work aims at comparing MANET routing algorithms in urban scenarios. Therefore, we designed a set of suitable simulation scenarios that emphasized strenghts and weaknesses of the evaluated protocols. To provide statistical relevance to the results, we repeated the simulation for each configuration 20 times, each time setting a different seed for the random number generators that control the randomness of the network and mobility simulators.

Road Network, Obstacles, and Vehicular Mobility—We perform our simulations in Manhattan grids, with characteristics defined in Table 1. We fixed the road topology and the maximum vehicular speed in an urban area, and we generated different vehicular densities to test the performances of the selected routing protocols in different vehicular traffic conditions. The vehicular mobility is highly realistic: each vehicle is modeled with its own physical characteristics such as its unique acceleration, deceleration, size, and category. Drivers are modeled with different driving skills and respect for the road rules. Each vehicle enters the scenario at a random location and plans a trip to reach a random arrival point

Table 1. Road network and vehicular mobility parameters

Manhattan grid	1 km × 1 km, 6 roads × 6 roads
Inter-road distance	200 m
Road width	6 m (Two 3 m lanes)
Road features	No traffic lights, 2 opposite-directional lanes
Buildings' base dimensions	180 m × 180 m
Inter-building distance	20 m
Vehicular mobility generator	SUMO Netgen, Discretization 1 s
Vehicular density	20 to 200 veh/km^2, increments of 20 veh/km^2
Vehicular maximum speed	Uniformly distributed from 30 to 50 km/h
Car-following model	Krauss [15]
Starting and arrival points	Random trips, minimum distance 500 m
Road path choice	Shortest path (Dijkstra)

on the map, provided that the starting and arrival points are at least 500 m apart. The road path to reach the arrival point from the starting point is computed using the Dijkstra shortest-path algorithm, where the edge weight is the road segment's length. The vehicles chosen as data source and destination do not move and are geographically fixed at the opposite corners of the simulation scenario. This is to ensure that the vehicles keep the same distance from one another and remain in the simulation for the whole duration of the data flow. This also prevents the length of the computed routes from varying according to the geographical distance between source and destination vehicles, consequently altering the performance of the routing protocols.

Inter-vehicular Communications—Our scenario features a single data flow from the data-source vehicle to the data-destination vehicle, respectively located at the top-left and bottom-right corners of the simulation scenario (Fig. 2). The transmission power of their wireless network interfaces, and the presence of shadowing objects (i.e. buildings), prevent them from communicating directly. Therefore, a multi-hop path must be established between them by one of the selected routing protocols. The characteristics of the application data flows and the parameters of the MAC/PHY layers are reported in the respective sections of Table 2.

We chose to evaluate the performances of the selected routing algorithms in scenarios with a single data flow. When the network is not congested, simulating multiple simultaneous flows instead of a single flow would not affect the signaling rate for the proactive routing algorithms, but would increase AODV's signaling rate proportionally with the number of simultaneous flows. This is because, unlike proactive protocols, AODV's signaling rate depends also on the number of active flows. The increased signaling traffic, generated by AODV for routing multiple flows, would decrease the PDR due to a high amount of frame collisions.

Table 2. Networking parameters

Application Layer	
Application packet size	1200 B = 9600 bit
Application interpacket interval	512, 128, 32 ms
Application data rates	18.75, 75, 300 kbit/s
Number of flows	1
Flow duration	180 s
MAC and PHY Layers	
Protocol	IEEE 802.11p single channel, no priority
Transmission power	15 mW
Receiver sensitivity	−89 dBm
Transmission capacity	6 Mbit/s
MAC queue capacity	100 frames

Furthermore, none of the compared routing protocols implement congestion-avoidance, load-balancing, or flow-distribution techniques. Therefore, simulating multiple flows instead of a single flow in different scenarios would not provide a meaningful discrimination element to compare the performances of the evaluated protocols. In addition, the performances of the routing algorithms would be influenced by the spatial allocation and the uncontrolled overlapping of the allocated flows, which could direct multiple flows over the same path and congest the network.

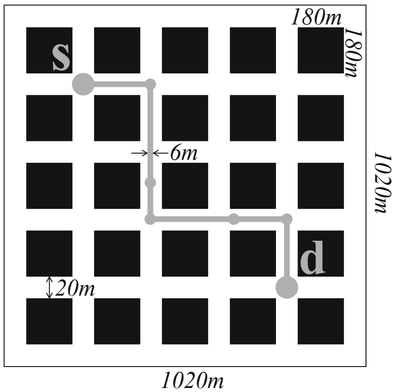

Fig. 2. Representation of the simulated Mahattan scenario. The larger gray nodes are the fixed source and destination vehicles, while the smaller nodes are the intermediate relays selected by the routing protocol. We observe an example of an established unicast route over 6 hops.

Simulating multiple flows could aggravate the congestion induced by application-layer messages and consequently show worse PDR and packet delay. Considering that the selected routing protocols are not designed for multi-flow management, the congestion induced by application-layer messages is unrelated to the protocols' intrinsic performances. Therefore, we focus our analyses on the congestion generated by the signaling packets at the network layer and leave application-induced congestion out of the scope of this work.

We fix source and destination nodes' positions at the opposite corners of the simulation scenario and we simulate realistic urban mobility between them, with different vehicular densities per each scenario (Table 1). The source node sends a Constant Bit Rate (CBR) stream of fixed-size UDP packets to the destination, with a bitrate never higher than a twentieth of the transmission capacity of the wireless channel (Table 2). This is because we do not want the application-layer packets to generate congestion, as it is not the focus of the present study. We are interested in comparing the signaling rates of the different routing protocols and measuring the data packet loss due to route reestablishments and due to the congestion induced by network-layer signaling packets.

Choice of MANET Routing Parameters—For each investigated routing protocol, a parameter study has been performed in order to detemine the configurations that would lead to the highest PDR for each scenario. For sake of brevity, we omit the details of this preliminary parameter study and we report, in Table 3, the best parameter configurations that we chose to perform the comparative simulations. The parameters that have not been reported in Table 3 are set to the default values as indicated in the relative implementations.

Table 3. Parameters of routing protocols

AODV	
`ActiveRouteTimeout`	3 s
`AskGratuitousRREP`	false
`UseLocalRepair`	false

DSDV	
`HelloInterval`	2048 ms
`RouteLifetime`	4096 ms
`UseFullDumps`	false

GPSR	
`BeaconingInterval`	100 ms
`NeighborValidity`	450 ms
`Planarization`	GG

Notably, the parameters that have been optimized are the `ActiveRouteTimeout` for AODV, the `HelloInterval` for DSDV, and the `BeaconingInterval` for GPSR. AODV's `ActiveRouteTimeout` is the maximum time interval during which a route entry can remain unused. After its expiration, the route becomes inactive and subsequently deleted. An overly short `ActiveRouteTimeout` would increase signaling traffic and packet loss, as routes that are still valid would be rediscovered during temporary periods in which the source does not send any data to the destination. An overly long `ActiveRouteTimeout` could make nodes store a route to an unreachable destination, due to the mobility of the network nodes. DSDV's `HelloInterval` is the time that elapses between two consecutive one-hop broadcasts of a *distance vector* containing local routing-table information. Decreasing the `HelloInterval` would increase the signaling rate but also increase the reliability of the local network-connectivity-status information, and vice versa. GPSR's `BeaconingInterval` is the time that elapses between two consecutive one-hop broadcasts of a beacon containing the node's address and geographical position. Decreasing the `BeaconingInterval` would increase the signaling rate but reduce the error between the list of neighbors maintained by the node and the neighbors that are actually reachable. In order to avoid synchronization effects [7], GPSR beacons are randomly delayed with a maximum jitter of 50 ms.

DSDV's `RouteLifetime` is the maximum time interval a route entry can be inactive before being deleted, and it has been set to double the `HelloInterval`. Due to implementation limitations, DSDV does not provide periodic full dumps as indicated by the protocol specifications [19], therefore each variation of a route entry will trigger a network-wide signaling broadcast. For GPSR, a node v removes a neighbor w from its neighbor list if v has not received a beacon from w for at least 4.5 times the duration of the `BeaconingInterval`, as suggested in [11]. The planarization mode is Gabriel Graph (GG) [8] and we assume that the precision of the vehicle's geographical position can be represented with 8 B. Assuming that a vehicle's L3 address is 4 B, the GPSR beacon will be 4 B+8 B = 96 bit long.

Implementation Details—The implementations of the evaluated routing protocols are based on the INET[1] library. In particular, AODV's implementation does not provide route repair, so every time a route is interrupted, the route discovery procedure is restarted from the source. The implementation of DSDV does not distinguish between full dumps and incremental updates of routing tables. Therefore, each node that receives a *useful* distance vector (called *hello message* in the implementation) from a neighbor, propagates it to its neighbors with a uniformly randomized delay between 10 ms and 500 ms without aggregation, causing a substantial signaling traffic increase.

The vehicular mobility was simulated with SUMO[2] [14], and the network protocols were simulated with Objective Modular Network Testbed in C++ (OMNeT++)[3] [23]. These two tools communicate through a TraCI interface wrapped by Veins INET[4] [22].

5 Performance Evaluation

Performance Metrics—In this work, we use *Packet Delivery Ratio (PDR)*, *packet delay*, *frame collision rate*, and *signaling rate* to compare the performances of the evaluated routing protocols in each scenario. The PDR is defined as the ratio between the number of packets that have been correctly received by the destination over the number of packets sent by the source. The frame collision rate is defined as the total number of frames incorrectly decoded by the network interfaces of every vehicle in the simulation, divided by the duration of the data flow. The signaling rate is defined as the total volume of transmitted signaling messages, specific to each routing protocol, divided by the duration of the data flow. For the three abovementioned metrics, each plotted point shows the average and the standard deviation of the metric computed for each simulation repetition. The packet delay is defined as the total time needed for a packet to travel from source to destination across the network. For this metric, each

[1] INET-v4.1.1 (hash *ce69d08*), https://github.com/inet-framework/inet.
[2] SUMO-v1.2.0 (hash *1d09773*), https://github.com/eclipse/sumo.
[3] OMNeT++-v5.5 (hash *6942b44*), https://github.com/omnetpp/omnetpp.
[4] Veins-v4.7.1 (hash *550e246*), https://github.com/sommer/veins.

plotted point shows the average and the standard deviation of the packet delay of all the received packets in all the simulation repetitions. The *goodput* provided by the different routing protocols is proportional to the PDR by a factor equal to the transmission data bitrate at the source (e.g. 300 kbit/s for the PDR in Figs. 3 and 4).

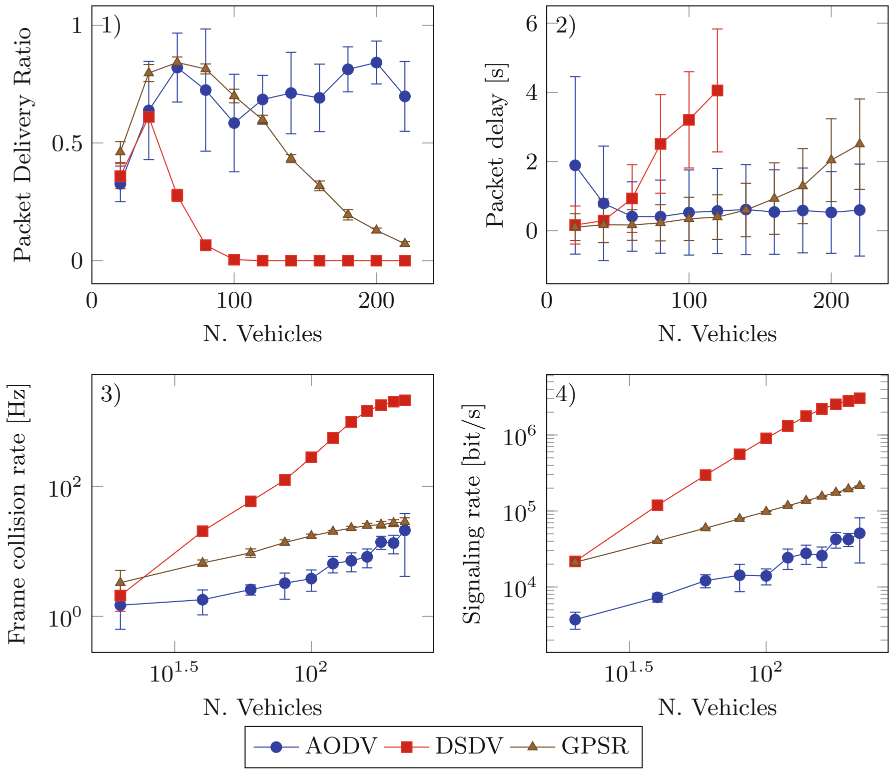

Fig. 3. Performances of the selected routing algorithms (AODV, DSDV, and GPSR) for a 300 kbit/s data flow, ignoring the shadowing effects caused by the buildings.

Figures 3 and 4 show the performance of the evaluated routing algorithms in terms of the aforementioned metrics for a data bitrate of 300 kbit/s, respectively ignoring and considering the shadowing effects induced by the presence of the buildings in the Manhattan grid. More simulations were performed for other data bitrates (75 kbit/s and 18.75 kbit/s), but the obtained results were similar to those displayed in Figs. 3 and 4, and are omitted for brevity.

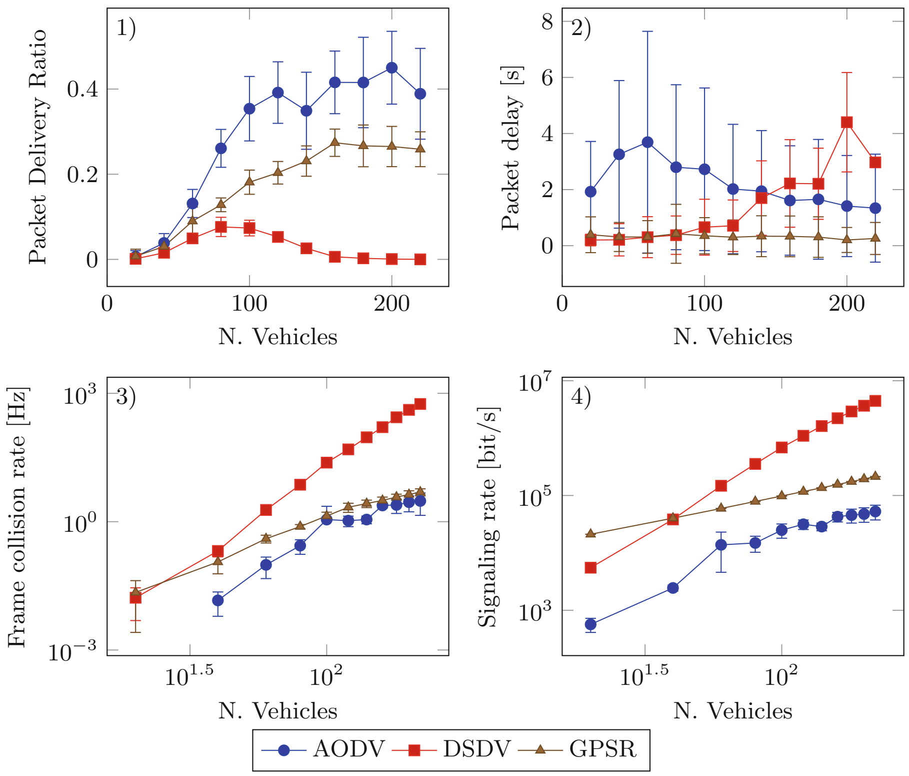

Fig. 4. Performances of the selected routing algorithms (AODV, DSDV, and GPSR) for a 300 kbit/s data flow, considering the shadowing effects caused by the buildings.

5.1 Signaling and Frame Collision Rate

From Figs. 3.4 and 4.4, we notice that the signaling rate of all the considered routing protocols is positively correlated to the vehicular density, although with different coefficients. AODV's and GPSR's signaling rates have similar sensitivity to vehicular density, i.e. comparable ratios of signaling rate increment to vehicular density increment. This can be justified because GPSR's signaling rate is linearly bound to vehicular density and beaconing frequency (which is fixed in this work), whereas AODV's signaling rate for a single flow depends on the route break frequency and on the vehicular density. For AODV, a higher vehicular density means a higher average number of neighbors that must forward a RREQ during a route discovery phase. However, a higher vehicular density leads to lower average vehicular speeds, and therefore to more stable routes.

DSDV displays a much higher sensitivity to vehicular density compared to the other two protocols. In fact, DSDV's signaling rate depends on two factors. The first factor is the network nodes' mobility, because each node must immediately inform its neighbors about link breaks. The second factor is the vehicular density, because each vehicle broadcasts a distance vector to its one-hop neighbors and each topology change must reach every node in the network. DSDV's signaling rate can be up to two orders of magnitude greater than the signaling rates of the other protocols.

We can observe that the signaling rate of each routing protocol does not dramatically vary between scenarios in which we consider shadowing effects and scenarios in which we ignore them. This suggests that the degree of connectivity of the network graph has a moderate impact on the amount of signaling traffic *transmitted* by each node. For GPSR, it has no effect at all because the beaconing frequency is fixed. Both AODV's and DSDV's signaling rate curves are slightly shifted upwards for the scenario without building shadowing because the higher nodes' average degree increases the number of neighbors that must forward an AODV's RREQ or a DSDV link-break notification.

We notice that the signaling rate of the three protocols is correlated to the relative frame collision rate (Figs. 3.3 and 4.3) which, in turn, influences the PDR.

5.2 Packet Delivery Ratio

From Figs. 3.1 and 4.1, we can observe that the PDR is generally lower for scenarios in which the shadowing effects are considered, as the routes become longer and more difficult to establish and maintain (even though there are relatively fewer frame collisions). For particularly low vehicular densities (e.g. 20 vehicles/km^2), PDRs are low for all protocols and configurations. This is because, with such low vehicular densities and having fixed the transmission power for all the network interfaces, there is a relatively low probability of having enough vehicles that can form a route from source to destination.

DSDV's PDR peaks when the frame collision rate is slightly below 10 Hz, which corresponds to a vehicular density of 20 vehicles/km^2 when we ignore shadowing effects, and a vehicular density of 80 vehicles/km^2 when we consider them. When the vehicular density increases, so does the frame collision rate, and the PDR of DSDV quickly converges to zero. Nodes running DSDV must broadcast their distance vectors to one-hop neighbors and propagate important network-update information across the whole network. This generates an enormous amount of signaling packets that congest the wireless medium, increasing the chances of frame collisions involving a data packet, and therefore decreasing the PDR. As an additional side effect, in congested-medium scenarios, signaling packets get queued and therefore delayed, slowing down the convergence of the protocol. Due to the slow convergence of the protocol, network nodes are likely to keep stale routing table entries or delete expired but valid entries. This would make nodes forward data packets to incorrect intermediate nodes or drop data packets when the correct intermediate node is still in range. As a consequence, we notice that the PDR drops for high vehicular densities and high signaling rates.

AODV's PDR increases proportionally to the vehicular density up to a moderate density and then converges to a steady-state value for higher vehicular densities. This is because, for lower vehicular densities, packets get dispersed due to unstable links and sporadic route establishments. For higher densities, routes become easier to set, but the increased signaling traffic hinders the data packet delivery.

As the vehicular density increases, GPSR's PDR first increases, then converges to a steady-state value, and then decreases. For low vehicular densities it is hard to find a set of vehicles that can offer a route between source and destination, whereas for high vehicular densities the beacons congest the wireless medium and collide with data packets, consequently decreasing the PDR. This phenomenon is clearly visible for the scenario in which the shadowing effects are ignored, and not present at all in the scenarios in which shadowing effects are taken into account. This happens because the presence of buildings in the Manhattan grid shadows the nodes from hearing excessive signaling traffic from their neighbors, therefore preventing the data packets from colliding with signaling packets.

5.3 Packet Delay

The curves in Figs. 3.2 and 4.2 report the average and standard deviation of the correctly-delivered packets' delay. Considering that the number of delivered packets varies per each simulation, the stability of the reported averages and standard deviations is not uniform for all the vehicular densities. For some simulations in which no packet was correctly delivered, no packet delay data could be displayed.

The average delay of packets routed by DSDV and GPSR increases as the vehicular density, signaling rate, and frame collision rate increase. This is because when the wireless medium is congested and very busy, the transmission of the frames containing the data packets are deferred. The packet delay curves increase faster in the scenarios without building shadowing than in the scenarios with shadowing, due to the missing protective effect of the physical obstacles against interference.

For AODV, we notice that the average packet delay decreases as the vehicular density increases. This is due to easier route establishment in scenarios with higher vehicular densities. In scenarios with building shadowing, the packet delay decays slower than in scenarios without building shadowing. This phenomenon can be justified because, under the same vehicular density, routes are longer and more difficult to set.

We notice that some of the packet delay curves are highly heteroscedastic. This means that the standard deviation of the packet delay varies across scenarios with different vehicular densities. A high packet delay variance means that the length of the packet queues in the network nodes are very variable over time, which also means that the packets' transmission is bursty. In the designed scenario, this may happen (i) because the routes are established and interrupted frequently, (ii) because the packets are routed over paths of very variable length, or (iii) because of dishomogeneous zones of congestion that migrate over time according to vehicular mobility.

AODV establishes a single route per data flow. In the case of a highly-dynamic network topology, such as in VANETs, this path is frequently broken. Each time a path is broken, the route discovery must be restarted from the source. When a route is interrupted, the packets received by the intermediate node without a valid next hop are queued and accumulate delay until a route is re-established. Upon route re-establishment, the queued packets are transmitted and leave the queue. If the network is not congested, the packets are cleared from the queue in bursts and the newly arrived packets are forwarded with low delay until the next path break happens. Proactive protocols do not show this behavior: for GPSR, the concept of a route does not exist at all, and for DSDV, the local information of the global connectivity status is constantly refreshed and therefore the routes are locally repaired. The packet delay variance of GPSR is particularly low because there is no route establishment phase and no routing table buildup; all the routing decisions are made only according to the node's and the node's neighbors' current locations.

6 Conclusion and Future Work

In this work, we investigated the applicability of some well-known MANET routing protocols to VANETs in realistic urban scenarios. We found out that using DSDV as a routing protocol in VANETs is unfeasible because, even for modest vehicular densities, the signaling traffic generated for updating the routing tables is unsustainably high. When using DSDV as a routing protocol, the PDR peaks for a medium-low vehicular density and decreases sharply after that due to intense signaling-induced congestion. AODV provided the best PDR, especially for higher vehicular densities. We also noticed that AODV's packet delay decreased as the vehicular density increased. GPSR offered intermediate PDR and signaling rate, with an overall low packet delay that showed a moderate sensitivity to vehicular density. This is because GPSR's beaconing congests the network only when network nodes have a very high average number of neighbors. In terms of signaling rate, DSDV has turned out to be up to two orders of magniture more demanding than the others. None of the evaluated routing protocols could satisfy modern vehicular application requirements in all the generated scenarios, therefore further development of vehicular routing protocols is desirable.

With this study, we observed that reactive protocols provide a higher PDR and throughput compared to proactive protocols, as claimed in some of the previous works [5,17]. We also observed that proactive protocols provide the best packet delay, as stated by the majority of the cited previous works. In the proposed single-flow scenario, AODV outperformed the proactive protocols in terms of signaling rate. For reactive protocols, the signaling traffic is highly dependent on the data traffic load, therefore we cannot extend this claim to multi-flow scenarios.

Future Work—For each of the evaluated routing protocols, the packet delays are too high to satisfy the strict requirements of modern vehicular applications.

One reason for this is the complexity of the traditional TCP/IP stack, which features some functions that are non-essential or even detrimental [13] for ad-hoc networks. These functions increase packet delay because of unnecessary MAC address resolutions, unicast packet acknowledgments, and various other operations. Therefore, a new architecture to provide a high PDR and a low packet delay in ad-hoc networks is needed.

One way to achieve such a challenging goal is to delegate some routing intelligence from the ad-hoc network to the network infrastructure, which has a global view of the ad-hoc network state and can make globally-optimal routing decisions. One paradigm that has been proposed to improve performances in VANETs is Software-Defined Vehicular Networking (SDVN) [6]: in this view, each vehicle becomes a simple L2 forwarding device, delegating all the routing decisions to a centralized controller.

The literature about SDVN is still young and its benefits not yet fully explored. For example, there is a lack of studies regarding the minimum requirements for the signaling traffic between vehicles and the SDVN controller to ensure such improved performances. In ultra-dense SDVNs (e.g. in road congestion scenarios), controllers could quickly become overloaded in terms of computational and network traffic loads, as every managed node reports its information to the controller at a fixed frequency.

In modern 5G-assisted ultra-dense SDVNs, vehicles communicate with the SDVN controller via valuable cellular links. The scarcity of bandwidth between vehicles and the controller makes it even more important to study the signaling traffic features of modern SDVNs architectures, and to devise novel techniques that can reduce the signaling rate requirements between vehicles and the controller while maintaining high network performances.

From an architectural standpoint, not enough studies consider scenarios in which the controller's knowledge domain covers only part of the total amount of connected vehicles. Therefore, it is important to study a hybrid approach that can benefit from the controller's coordination when it is available and fall back on a distributed routing scheme when it is not.

References

1. Ali, A.K., Phillips, I., Yang, H.: Evaluating VANET routing in urban environments. In: 2016 39th International Conference on Telecommunications and Signal Processing (TSP), pp. 60–63, June 2016. https://doi.org/10.1109/TSP.2016.7760829
2. Bengag, A., Elboukhari, M.: Performance evaluation of VANETs routing protocols using SUMO and NS3. In: 2018 IEEE 5th International Congress on Information Science and Technology (CiSt), pp. 525–530. IEEE (2018)
3. Bettstetter, C., Resta, G., Santi, P.: The node distribution of the random waypoint mobility model for wireless ad hoc networks. IEEE Trans. Mob. Comput. **2**(3), 257–269 (2003)
4. Boukerche, A., Ahmad, M.Z., Turgut, D., Turgut, B.: A taxonomy of routing protocols for mobile ad hoc networks. Algorithms and Protocols for Wireless and Mobile Ad Hoc Networks, p. 129 (2009)

5. Broch, J., Maltz, D.A., Johnson, D.B., Hu, Y.C., Jetcheva, J.G.: A performance comparison of multi-hop wireless ad hoc network routing protocols. In: MobiCom., vol. 98, pp. 85–97 (1998)
6. Di Maio, A., Palattella, M.R., Engel, T.: Multi-flow congestion-aware routing in software-defined vehicular networks. In: 2019 IEEE 90th Vehicular Technology Conference (VTC Fall). IEEE (2019)
7. Floyd, S., Jacobson, V.: The synchronization of periodic routing messages. IEEE/ACM Trans. Netw. **2**(2), 122–136 (1994). https://doi.org/10.1109/90.298431
8. Gabriel, K.R., Sokal, R.R.: A new statistical approach to geographic variation analysis. Syst. Zool. **18**(3), 259–278 (1969)
9. Johnson, D.B., Maltz, D.A.: Dynamic source routing in ad hoc wireless networks. In: Mobile computing, pp. 153–181. Springer, Boston (1996). https://doi.org/10.1007/978-0-585-29603-6_5
10. Kang, L., Zhao, W., Qi, B., Banerjee, S.: Augmenting self-driving with remote control: challenges and directions. In: Proceedings of the 19th International Workshop on Mobile Computing Systems & Applications, pp. 19–24. ACM (2018)
11. Karp, B., Kung, H.T.: GPSR: greedy perimeter stateless routing for wireless networks. In: Proceedings of the 6th Annual International Conference on Mobile Computing and Networking, MobiCom 2000, pp. 243–254. ACM, New York (2000). https://doi.org/10.1145/345910.345953, http://doi.acm.org/10.1145/345910.345953
12. Kaul, S., Yates, R., Gruteser, M.: Real-time status: how often should one update? In: 2012 Proceedings IEEE INFOCOM, pp. 2731–2735, March 2012. https://doi.org/10.1109/INFCOM.2012.6195689
13. Klingler, F., Dressler, F., Sommer, C.: IEEE 802.11p unicast considered harmful. In: 2015 IEEE Vehicular Networking Conference (VNC), pp. 76–83. IEEE (2015)
14. Krajzewicz, D., Erdmann, J., Behrisch, M., Bieker, L.: Recent development and applications of SUMO - Simulation of Urban MObility. Int. J. Adv. Syst. Meas. **5**(3&4), 128–138 (2012)
15. Krauß, S., Wagner, P., Gawron, C.: Metastable states in a microscopic model of traffic flow. Phys. Rev. E **55**(5), 5597 (1997)
16. Lee, K.C., Lee, U., Gerla, M.: Survey of routing protocols in vehicular ad hoc networks. In: Advances in vehicular ad-hoc networks: Developments and challenges, pp. 149–170. IGI Global (2010)
17. Naumov, V., Baumann, R., Gross, T.: An evaluation of inter-vehicle ad hoc networks based on realistic vehicular traces. In: Proceedings of the 7th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc 2006, pp. 108–119. ACM, New York (2006). https://doi.org/10.1145/1132905.1132918
18. Perkins, C.E., Royer, E.M.: Ad-hoc on-demand distance vector routing. In: Proceedings WMCSA 1999 Second IEEE Workshop on Mobile Computing Systems and Applications, pp. 90–100, February 1999. https://doi.org/10.1109/MCSA.1999.749281
19. Perkins, C.E., Bhagwat, P.: Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In: Proceedings of the Conference on Communications Architectures, Protocols and Applications, SIGCOMM 1994, pp. 234–244. ACM, New York (1994). https://doi.org/10.1145/190314.190336
20. Santoso, G.Z., Kang, M.: Performance analysis of AODV, DSDV and OLSR in a VANETs safety application scenario. In: 2012 14th International Conference on Advanced Communication Technology (ICACT), pp. 57–60, February 2012
21. Sommer, C., Dressler, F.: Vehicular Networking. Cambridge University Press, Cambridge (2015)

22. Sommer, C., German, R., Dressler, F.: Bidirectionally coupled network and road traffic simulation for improved IVC analysis. IEEE Trans. Mob. Comput. **10**(1), 3–15 (2011). https://doi.org/10.1109/TMC.2010.133
23. Varga, A., Hornig, R.: An overview of the OMNeT++ simulation environment. In: Proceedings of the 1st International Conference on Simulation Tools and Techniques for Communications, Networks and Systems & Workshops, Simutools 2008, pp. 60:1–60:10. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), Brussels, (2008). http://dl.acm.org/citation.cfm?id=1416222.1416290
24. Yoon, J., Liu, M., Noble, B.: Random waypoint considered harmful. In: IEEE INFOCOM 2003 Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies (IEEE Cat. No. 03CH37428), vol. 2, pp. 1312–1321. IEEE (2003)