Congestion-Aware Geocast Routing in Vehicular Delay-Tolerant Networks

Henrique Nascimento [1], Paulo Rogério Pereira [2],* and Naercio Magaia [3]

[1] INESC-ID, Instituto Superior Técnico, Universidade de Lisboa, 1000-029 Lisboa, Portugal;
    henrique.nascimento@tecnico.ulisboa.pt
[2] INESC-ID/INOV, Instituto Superior Técnico, Universidade de Lisboa, 1000-029 Lisboa, Portugal
[3] LASIGE, Departamento de Informática, Faculdade de Ciências da Universidade de Lisboa, 1749-016 Lisboa,
    Portugal; ndmagaia@fc.ul.pt
*   Correspondence: prbp@inesc.pt; Tel.: +351-213100345

Abstract: Vehicular Delay-Tolerant Networks (VDTNs) are networks of vehicles that communicate wirelessly, where there are no permanent end-to-end connections. VDTNs have a highly variable topology, with frequent partitions, and possibly low node density. Thus, delay-tolerant routing adopts a store-carry-and-forward message transfer paradigm, where messages have a useful Time-To-Live (TTL) and are stored until a good contact opportunity arises. Multiple message replicas can be generated to improve delivery probability at the cost of increasing network congestion. In this paper, we propose the V-GRADIENT geocast routing protocol that monitors node density, buffer occupancy, and interest in geocast groups, to adapt the forwarding techniques used dynamically, to disseminate messages within the geographic region of interest. Simulation results show that the V-GRADIENT is capable of controlling network congestion and efficiently delivering messages resulting in better delivery ratios, lower latencies, and a small increase in overhead when compared with existing protocols.

Keywords: wireless communications; geocast routing; vehicular delay-tolerant networks

1. Introduction

Delay Tolerant Networks (DTNs) are characterized by long or variable delays, intermittent connectivity, asymmetric data rates, and high error rates. In DTNs, routing strategies rely on the store-carry-and-forward paradigm, i.e., on the dissemination of messages to intermediate nodes that can retain, transport, and deliver them either to the final destination or to other intermediate nodes. The complexity of the routing decisions is, for that reason, closely related to the selection of a proper message replication mechanism.

The geocast concept refers to a delivery scheme that filters the eligibility of the network's target nodes using their location as a criterion. With this in mind, the diversity of geocast environments encompassed by DTNs suggests the demarcation of strategies according to generic properties such as the definition of regions of interest (ROI) or the replication mechanisms applied. While unicast routing in DTNs has been extensively studied [1], geocast routing has been studied mainly in Mobile Ad Hoc Network (MANET) scenarios [2], where permanent end-to-end connectivity is assumed, which does not apply to DTNs.

The coupling between the advances in wireless communications technologies and the car industry makes Vehicle-to-Vehicle (V2V) connectivity an enabler for a broad range of applications that can be grouped into two categories: safety-related and commercial-related. For real-time short-range safety communications, several mechanisms have been designed, like using beacons for vehicles to

disseminate information to their direct neighbors [3] with high periodicity (e.g., 10 Hz) or at two hops distance [4]. For non-real-time geocast dissemination of messages through ROIs larger than the wireless communication range, the DTN paradigm of store-carry-and-forward opportunistic dissemination of messages is particularly useful. Therefore, in light of the advances mentioned, the DTN communication paradigm has received a lot of attention in research and has motivated the design of routing strategies for Vehicular Delay Tolerant Networks (VDTNs) [5] that can properly handle the challenges of such distinguishable environments.

In this sense, the design of strategies in VDTNs that can efficiently deliver messages to a group of recipients, which would particularly benefit from the reception of such messages, has become a widespread topic of discussion. To that extent, the implementation of algorithms that exploit information from the vehicles' navigation systems to route messages within specific ROIs motivates the development of the V-GRADIENT geocast routing protocol for VDTN environments.

The proposed routing protocol is the basis for a message dissemination service within an ROI in highly dynamic or sparse vehicular networks. Applications include V2V commercial or safety-related message dissemination such as disseminating maps, software updates, fuel prices, current weather, incident warnings, or others [6]. The proposed protocol provides the following contributions: mechanisms to monitor node density and node's buffer occupancy; mechanisms to control message replication avoiding network congestion and buffer exhaustion; efficient message delivery within the ROI. This paper extends the preliminary version of the V-GRADIENT protocol proposed in [7] by including new dropping policies and replication mechanisms, an overhead analysis, and an extensive scenario analysis.

2. State of the Art

2.1. Unicast Routing

As described in [8], the Direct Delivery protocol is the simplest way to conceive a unicast delivery scheme. The node that carries the message forwards the message only if it establishes a direct contact opportunity with the target node.

The Spray and Wait protocol [9], in turn, is based on the dissemination of a limited set of replicas of a given message to distinct contacted nodes during a first phase called the spray phase. Afterwards, during the wait phase, relay nodes assume the responsibility of delivering the message directly to the final recipient. However, regarding the spray phase, several heuristics can be envisioned to properly distribute messages.

Finally, the Epidemic [10] protocol is classified as a flooding-based scheme, since the idea behind it is based on the provision of unlimited message replication, throughout the network, by exchanging all possible messages whenever contact opportunities arise.

2.2. Geocast Routing

On one hand, protocols such as the Geocasting Spray And Flood (GSAF) [11] and the Geocasting For Opportunistic Networks (Geoopp) [12] focus on distributing messages by all nodes located inside a region set remotely, relative to the position of the nodes that generate such messages. This implies, however, that these algorithms break down the routing process into the two following phases: (1) forwarding (and carrying) the messages to the destination region; (2) delivering the message to all nodes inside the ROI. While both protocols mentioned, regarding the latter phase, follow an epidemic-based behavior, their strategy differs with respect to the first phase. For instance, when using the GSAF protocol, a spraying procedure is applied in the first phase. Messages are assigned a predetermined number of tickets T that denote the number of times a given message can be forwarded to encountered nodes outside the destination region. Thus, each time a message is replicated, the T field is sequentially decreased on both messages until T equals zero.

On the other hand, geocast protocols can be applied to distribute content in ROIs whose center is defined by the position of the nodes responsible for creating the respective messages at the generation time instant. In this sense, the authors of the Floating Content algorithm [13] include in the messages' headers not only the center coordinates and radius of the circularly shaped ROI but also an additional radius parameter to determine the extension of a buffer zone. The buffer zone is defined as an extension of the ROI where carrier nodes may keep the messages stored beyond the boundaries of the ROI. This allows for a more efficient delivery of messages to nodes that approach the ROI.

3. Geocast Strategies Comparison

In this work, geocast focuses on distributing messages during a specified Time-To-Live (TTL) to all nodes that come across an invariant circularly shaped ROI, defined around the node that creates the messages. Moreover, the development of the intended message dissemination algorithm is performed by understanding which features of the considered schemes yield better results under a broad variety of simulations.

In this study, in order to drive such conclusions, the following routing strategies are taken into account, which result from the extension of the definitions regarding the aforementioned unicast routing protocols to a geocast context: Multiple Copy GeoDirectDelivery (MC-GeoDirectDelivery), GeoSpray-and-Wait (GeoS&W), GeoEpidemic, and GeoEpidemic Constrained.

Firstly, when using the MC-GeoDirectDelivery protocol, on one hand, messages are delivered to and kept by interested nodes inside the ROI, at the delivery time instant. On the other hand, messages are retained when nodes leave the ROI, as they may return to the ROI in the future. Because messages are only dropped when their TTL expires or the nodes' buffers are filled up, all nodes that receive a message also become responsible to deliver it to the remaining recipients, enabling the presence of multiple copies of the same message in the network.

Secondly, the GeoS&W scheme couples the principles of the Spray and Wait and the GSAF protocols. In this sense, during the spray phase, messages are replicated, even when nodes are located outside the ROI, using the spraying heuristic adopted by the GSAF protocol. Regarding the wait phase, the GeoS&W behaves the same as the MC-GeoDirectDelivery protocol.

The Epidemic protocol nature may assist message distribution in a geocast context as it already distributes an unbounded number of copies of the same messages by multiple receivers. However, to distinguish the protocol's operation in unicast scenarios from its operation in a geocast routing paradigm, this protocol will be referred as the GeoEpidemic protocol throughout this article. Likewise, the GeoEpidemic Constrained scheme follows the Epidemic protocol's principles as well. However, in this case, the routing scheme under consideration constrains the dissemination of messages within the boundaries of the ROI, since messages are dropped as soon as the nodes carrying them leave this region.

Additionally, it is also crucial to make explicit the following performance assessment metrics utilized: delivery ratio, delivery latency, and overhead. The delivery ratio is computed using the formulation developed in [11]. By monitoring the recipients that reside inside a given message's ROI during its TTL, it is possible to compute a ratio between the number of successful receivers and the total number of nodes belonging to the list of recipients. Obviously, this ratio only translates itself into a per-message result. To that end, to obtain an overall delivery ratio estimation, the per-message metric is extended to all created messages through an average operation. Latency is also monitored, in the first instance, taking into account a single message. However, in this case, the per-message metric is computed by summing the amount of time it takes for the message under consideration to reach each one of its successful receivers and, then, by dividing this outcome by the same number of recipients. Finally, the latency for all messages is averaged. The network overhead, in terms of transmissions, is by definition the ratio between the surplus of transmissions relative to the number of delivered messages, and the number of delivered messages. In spite of being computed similarly to unicast schemes, the

outcome of this metric is expected to be much lower in a geocast context, since the existence of several recipients per message is reflected on a larger number of delivered messages.

4. V-GRADIENT Design

The major points of improvement attained by benchmarking the different geocast strategies relate not only with a moderated replication process applied outside the messages' regions of interest but also with a dropping policy that can dynamically estimate if it is beneficial to keep messages being carried by nodes located outside the ROI. With this in mind, on one hand, with regard to the replication process, the V-GRADIENT algorithm incorporates the principles of the Spray-and-Wait protocol [9] in a geocast context and, during the spray phase, the methodology adopted by the GSAF protocol [11]. On the other hand, the deployed strategy follows the idea behind the Floating Content algorithm [13], i.e., defines a buffer zone whose range varies according to the network conditions and the ROI radius.

4.1. Dynamic Dropping Policy

The dynamic dropping policy implemented uses two mechanisms. The first one considers that if nodes are outside the ROI, they may not be able to return to the ROI within the messages' TTL and thus the corresponding messages can be deleted. The second mechanism controls the ROI radius according to the estimated network conditions.

Assuming that, in urban environments, the vehicle's maximum speed, v_{max}, is usually 50 km/h and the radio interfaces have considerable lower communication ranges relatively to the ROIs' radiuses, the minimum amount of time that would take for a node, located outside the ROI, to reach the ROI's boundaries of a particular message, is approximately given by τ_{min}, as shown in Equation (1). The variables d_{center} and R_{ROI} represent the distance of the vehicle to the ROI's center and the ROI's radius, respectively. Thus, if the remaining TTL of a given message drops below τ_{min}, the bearing node will be out of the ROI's range at least until the TTL expires, as it is advantageous to drop the bundle.

$$\tau_{min} = \frac{d_{center} - R_{ROI}}{v_{max}}. \tag{1}$$

To select dynamically the buffer zone range, the V-GRADIENT algorithm takes advantage of estimations of the network density level and of the message buffer occupation level, updated periodically. Because it is assumed here that nodes do not have access to any knowledge oracles, vehicles rely only on information exchanged during contact opportunities to perceive the surrounding network conditions and, for that reason, both awareness metrics are kept as dictionaries in the nodes' buffers in a decentralized manner. Moreover, the mathematical formulation used to compute both metrics follow an Exponentially Weighted Moving Average (EWMA). For the density level metric, ρ_t, the EWMA is performed every 30 seconds taking into account the number of distinct contacts n_c established in 30 seconds intervals, as shown in Equation (2). The weight α was set to 0.25 to balance old and new measurements. In this way, a 25% weight is used for the number of contacts established during the last 30 seconds interval and a 75% weight is used for the accumulated measurements of the previous intervals. These parameters were set experimentally to values that provided good node density estimations, filtering out very short term variations.

$$\rho_t = \alpha \cdot n_c + (1 - \alpha) \cdot \rho_{t-1}. \tag{2}$$

The message buffer occupation metric β_t is determined by averaging the fractions of the message buffer occupancy, β_{level}, varying from zero to one, retrieved from neighboring nodes in 30 seconds intervals with an EWMA, as shown in Equation (3). Weight γ was set to 0.7 in order to be high for recent buffer occupancy measurements. In this way, the protocol can adapt faster to network

congestion situations. Again, these parameters were set experimentally to values that provided good buffer occupancy estimations.

$$\beta_t = \gamma \frac{\sum^{n_c} \beta_{level}}{n_c} + (1 - \gamma)\beta_{t-1}. \tag{3}$$

In order to carry out a reliable implementation of the mechanism responsible for adjusting the extent of the buffer zone, a wide variety of simulations were executed using the ONE simulator [14] with the Geo ONE package [15], under different density scenarios, considering ROIs with different radius values, and in environments that led to serious buffer congestion issues, as shown in Table 1. In each simulation, an additional dropping policy was included that periodically removed messages from storage if the nodes were located outside of the buffer zone range, defined as the product between the ROI radius value, R_{ROI}, and a threshold R_{+ratio}.

Table 1. Assigned parameters during the simulation procedure.

Map	The Downtown Part of the City of Helsinki [4500 m × 3400 m]
Simulation time	12 h
Movement model	Shortest Path Map-Based Movement
Type of nodes	Vehicles (cars)
Nodes' speed interval	10–50 km/h
Nodes' message buffer size	5 MB
Nodes' wait time	5–30 s
Message size interval	500 kB–1 MB
Message generation interval	30–90 s
Interfaces' data rate	250 kbps
Interfaces' transmission range	50 m
Initial Number of Copy Tickets	3
Message's TTL	150 min
Number of nodes	[20;50;100;200;300;400] nodes
ROI Radius	[200;400;600;800;1000] m
R_{+ratio}	Vector of 100 evenly spaced points between 1 and 10

Figure 1 shows the evolution of the average density level estimated with Equation (2) in simulations encompassing a different number of nodes, proving that it is possible for each node to estimate accurately the average network density level, just by monitoring the number of contacts over time.
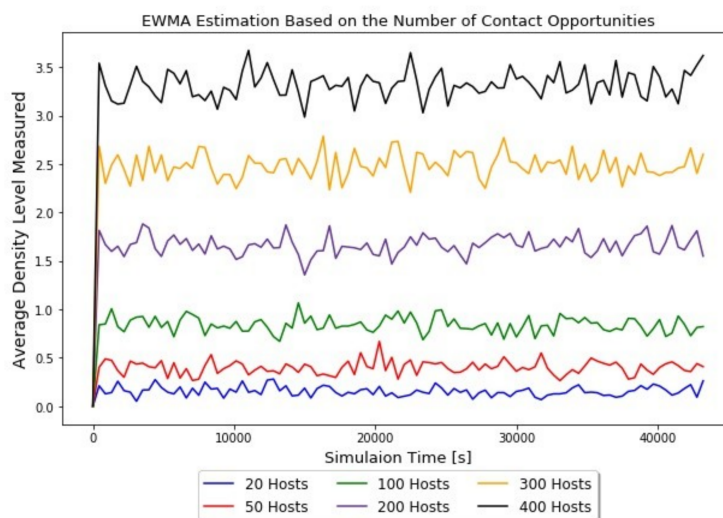

Figure 1. Evolution of the average density level measured considering distinct density conditions.

To exhibit the impact of the network density conditions and the ROIs' radius values on the additional dropping policy formulation, Table 2 includes the optimal R_{+ratio} values, bearing in mind the delivery ratio results obtained. It can also be seen in Table 2, that for a fixed number of nodes deployed in the network, the optimal R_{+ratio} expression can be modeled by an equation obtained through a power regression. The profile of such equations, following a trend depending on the number of nodes and consequently on the density level, also suggest a mathematical formulation of a variable threshold value associated with the extension of the buffer zone as shown in Equation (4).

$$R_{+ratio}(R_{ROI}, \rho_t) = g(\rho_t) \cdot R_{ROI}^{h(\rho_t)}. \tag{4}$$

Table 2. Optimal R_{+ratio} values obtained for simulations considering different density conditions and distinct regions of interest (ROI) radius values.

	50 Nodes $\rho_{avg} = 0.3952$	100 Nodes $\rho_{avg} = 0.8329$	200 Nodes $\rho_{avg} = 1.6470$	400 Nodes $\rho_{avg} = 3.2884$
$R_{ROI} = 200$ m	9.5	8.6	8.1	7.9
$R_{ROI} = 400$ m	3.4	3	2.8	2.5
$R_{ROI} = 600$ m	2	1.7	1.5	1.4
$R_{ROI} = 800$ m	1.5	1.2	1.1	1.1
$R_{ROI} = 1000$ m	1	1	1	1
$\hat{R}_{+ratio}$	$12,872 \cdot R_{ROI}^{-1.366}$	$10,970 \cdot R_{ROI}^{-1.360}$	$9460.3 \cdot R_{ROI}^{-1.348}$	$7649.4 \cdot R_{ROI}^{-1.320}$
R^2 score	0.9964	0.9932	0.9827	0.9706

With respect to the formulation of the $g(\rho_t)$ and the $h(\rho_t)$ expressions, by selecting their corresponding target parameters presented in Table 2, it can be observed, as depicted in Figure 2, that both describe a well-modeled decay relative to the average density level, ρ_{avg}, measured in simulations deploying a different number of nodes.
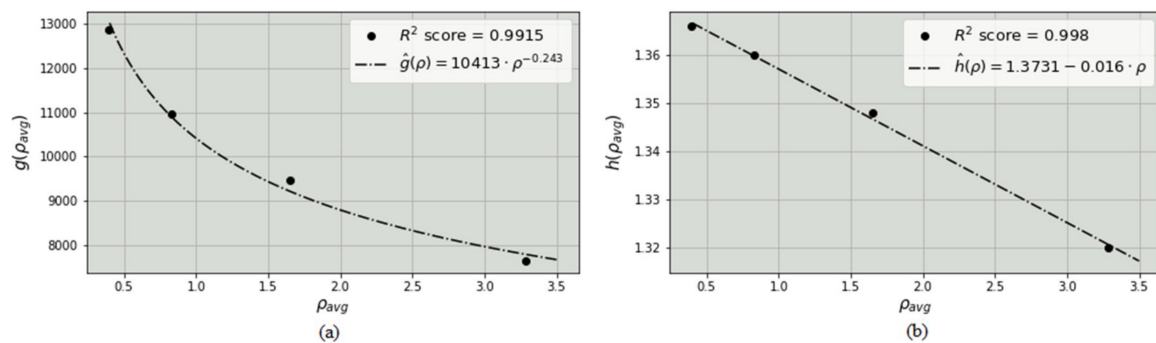

Figure 2. Regression-based model for the: (**a**) $g(\rho_t)$; and (**b**) $h(\rho_t)$ expressions.

In this sense, in scenarios characterized by extreme buffer congestion issues, the V-GRADIENT algorithm allows a dynamic adjustment of the range of the buffer zone for all messages carried by a particular node according to the estimated threshold expression presented in Equation (5), as a result of the best fit from Figure 2. Additionally, the buffer occupation level metric plays an important role in detecting the occurrence of congestion, providing nimbleness to the dropping policy operation. If we let ourselves picture a scenario where only a few messages are generated per hour, it can be concluded that, because no congestion issues arise, the drop policy should be less restrictive. For this reason, a

logistic function $f(\beta_t)$ responsible for regulating the stringency of the policy, with $\beta_t = 0.5$ defining the Sigmoid's midpoint, and $k = 20$, as shown in Equation (6), is included in the final definition of R_{+ratio}:

$$R_{+ratio}(R_{ROI}, \rho_t, \beta_t) = \frac{1}{f(\beta_t)} 10,413 \cdot \rho_t^{-0.243} \cdot R_{ROI}^{1.3731-0.016 \cdot \rho_t}. \tag{5}$$

$$f(\beta_t) = \frac{e^{k(\beta_t-0.5)}}{1 + e^{k(\beta_t-0.5)}}. \tag{6}$$

4.2. Dynamic Replication Mechanism

One of the main conclusions drawn from the analysis of the previous subsection relates to the benefits of the application of a moderate replication process outside the regions of interest, particularly if the vehicle's density is low and ROIs' dimensions are small. This is shown in Figure 3, where the size of the buffer zone R_{MAX} is determined through Equation (5).
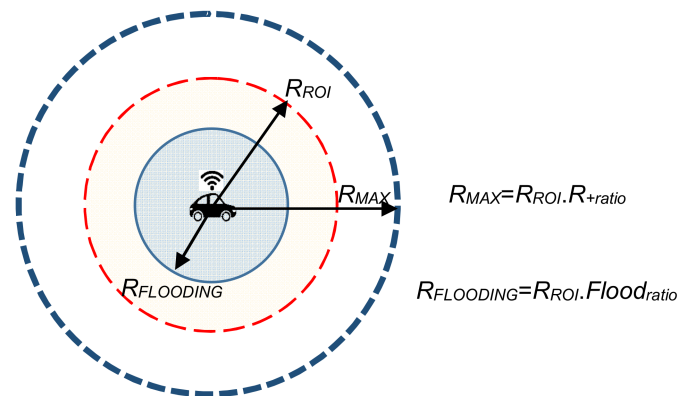

Figure 3. Message flooding zone (light blue) and buffer zone size (dark blue) as compared to the size of the ROI (red).

The replication mechanism inside the ROI should depend upon several parameters. For instance, if all nodes are interested in the messages, corresponding to a penetration rate of $\sigma_{rate} = 100\%$, the messages should be distributed to all nodes within the ROI and all nodes should contribute to spreading the message as fast as possible within the ROI. However, if only a small fraction of the nodes are interested in the messages, corresponding to a low penetration rate, the messages' TTL is small and the network is sparse, and the interested nodes may not be enough for efficiently spreading the messages, resulting in a low delivery ratio. Thus, the V-GRADIENT adopts two mechanisms for improving message dissemination in these situations. First, a flooding zone is defined, smaller than the ROI, where messages are flooded to every node, as in the epidemic protocol. Second, a ticket-based spraying heuristic, as proposed in [11] is used, where some nodes not directly interested in the message are selected to help to disseminate messages within the ROI or the buffer zone. If T tickets are assigned, the number of message replicas transmitted during the spray phase is 2^T, as in [11]. Naturally, if the network is sufficiently dense and/or the ROIs' radiuses are considerably large, messages will run out of tickets in locations closer to the ROI's center and no replication will occur outside of the ROIs. Notably, it was considered enough for the V-GRADIENT protocol to have $T = 3$ tickets assigned to all messages.

Nodes using the V-GRADIENT have to estimate the penetration rate (σ_{rate}) of all distinct geocast groups in the network. When a contact opportunity arises between two nodes, they exchange geocast group identifiers, which they are interested in. This allows them to keep a dictionary with a count of the number of nodes contacted that were interested in each geocast group. By keeping a count of all established contacts also, the penetration rate (σ_{rate}) is simply the fraction of nodes contacted that were interested in each geocast group, as kept in the dictionary, over the total number of nodes contacted.

Figure 4 shows the evolution of the average penetration rate measured by nodes for three different geocast groups with penetration rates of 25%, 50%, and 75%, and for different density conditions of 40, 100, and 200 nodes in the Helsinki scenario. As can be seen, the nodes can have a fairly accurate idea of the real penetration rate for each geocast group, and the convergence speed is faster for denser networks and higher penetration rates, as this results in more frequent contacts between nodes.
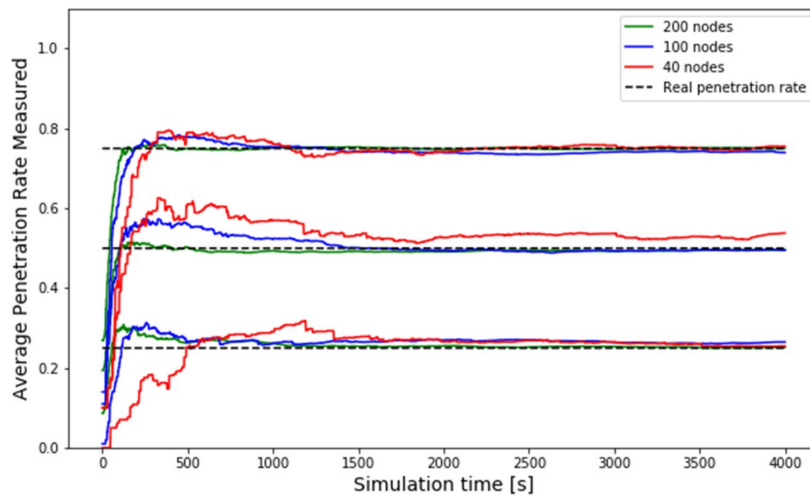

Figure 4. Average geocast group penetration rate measured by nodes.

The V-GRADIENT uses the penetration rate metric to control the number of uninterested vehicles that should cooperate in forwarding messages of geocast groups of low penetration rates. The term control stems from the limited amount of shared resources available, which implies the management of a trade-off. The more forwarding load assigned to nodes that are not interested in messages results in an improvement in the delivery performance of such messages. In contrast, the resources spent on carrying and forwarding messages which the nodes have no interest in end up, inevitably, compromising the delivery of those in which they are interested. Thus, to allow a fair distribution of the forwarding load, the proposed algorithm for the V-GRADIENT defines a flooding region within the ROI whose area covers a portion of the ROIs directly linked to the penetration rate measured. In this sense, the radius of these regions is obtained as the product between the corresponding ROIs' radiuses and a variable, $Flood_{ratio}$, computed as a function of the penetration rate and the messages' initial TTL as shown in Equation (7). It should be noted that messages are disseminated within the flooding region near the ROI center even when they run out of tickets.

$$Flood_{ratio} = 1 - \frac{e^{k(TTL-40[min])}}{1 + e^{k(TTL-40[min])}}\left(1 - \sqrt{\sigma_{rate}}\right). \tag{7}$$

The inclusion of the TTL parameter in the expression of the $Flood_{ratio}$ relates to the volatility of messages with low values of TTL, meaning that such messages are dropped shortly after they are created and, for that reason, should be disseminated quickly as they only constrain the available resources in the network for brief periods of time. A threshold of 40 minutes for the TTL was considered, as messages with shorter TTL values expire very soon. Figure 5 shows the resulting graph regarding the $Flood_{ratio}$ outcome for distinct penetration rates and different initial TTL values (with a logistic function similar to Equation (6), but with $k = 0.25$).
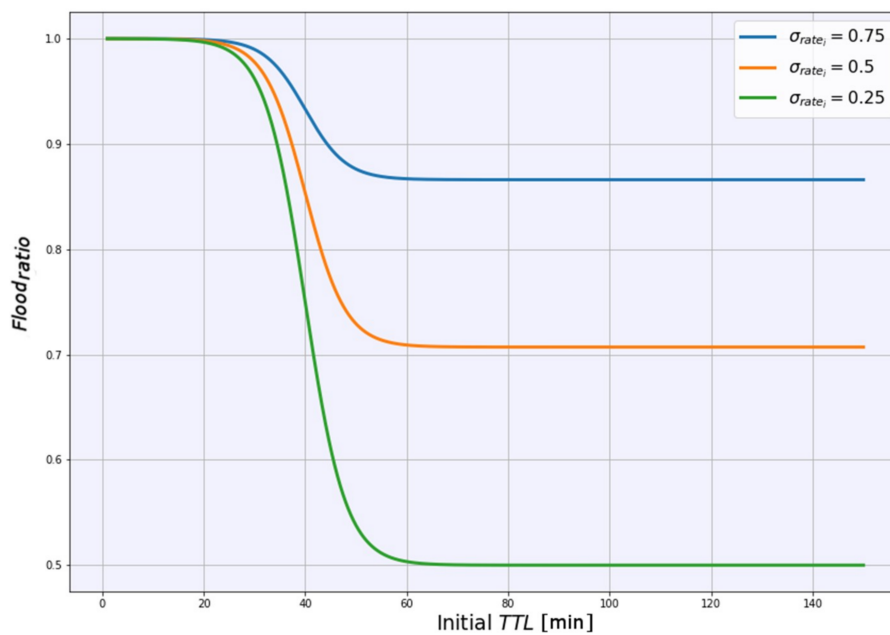
Figure 5. Graph of *Flood$_{ratio}$*, from Equation (7), as a function of the initial Time-To-Live (TTL) and distinct penetration rates.

For very dense networks, there are many contact opportunities, which suggests that nodes can select the relays nodes more carefully, leading to a more even geographic dissemination of messages. Thus, the ticket-based spraying heuristic is modified for very dense networks, handling a message copy to another node and spending a ticket, if the following conditions apply: (i) the network density level measured by the sender node exceeds a minimum threshold ρ_{min}; (ii) the sender node resides within the flooding region of the ROI; (iii) the receiver node is not an eligible receiver of that particular bundle; (iv) the absolute difference between the moving direction of the sender and receiver is larger than 45 degrees, as proposed in [16]; and (v) there are unspent tickets. A threshold value ρ_{min} = 2.8 was used, which only applies to more than 300 nodes in the Helsinki scenario (Figure 1).

5. Performance Evaluation and Results

The V-GRADIENT protocol was compared with different strategies described in Section 3. Most simulation parameters are as shown in Table 1, except the buffer sizes, which were increased to 20 MB on all nodes, so that network congestion was smaller. Three geocast groups were configured as shown in Table 3. Different message generators were deployed in each simulation simultaneously with the configuration shown in Table 4. Four different seed values are used for the movement model's pseudo-random number generator to reinforce the statistical significance of the results obtained. The corresponding 95% confidence intervals are presented in the graphs with the simulation results. In most cases, the confidence intervals are very small, which shows high confidence in the simulation results. A picture of the simulation scenario is presented in Figure 6, showing the Helsinki downtown area with 80 vehicle nodes and corresponding identifier numbers, small green circles depicting the nodes' wireless communication ranges and red, blue, and black circles depicting the ROI of active messages for geocast groups one, two, and three, respectively.

Table 3. Groups of vehicles interested in the distinct geocast services for the Helsinki scenario.

Group of Nodes	Group Size	Interested in Messages with Geocast Group IDs:		
		ID = 1	ID = 2	ID = 3
Group 1	25% of vehicles	✓	✗	✗
Group 2	50% of vehicles	✓	✓	✗
Group 3	25% of vehicles	✓	✗	✓

Table 4. Parameters for the distinct message event generators used in the Helsinki scenario.

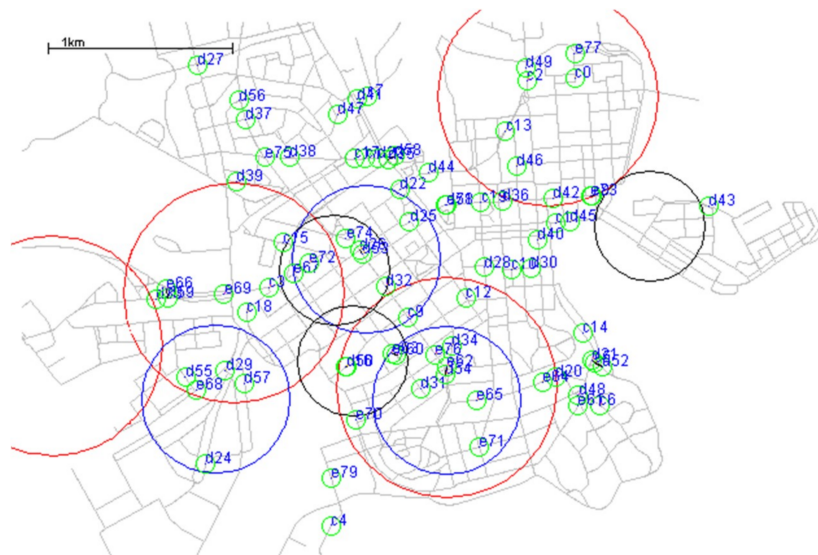| Message Event Generator Parameters	Generator 1	Generator 2	Generator 3
Message size interval	500 kB–1 MB	500 kB–1 MB	500 kB–1 MB
Message generation interval	60–120 s	80–140 s	100–160 s
Messages' initial TTL	100 min	60 min	80 min
ROI radius	600 m	400 m	300 m
Geocast group identifier	ID = 1	ID = 2	ID = 3
Geocast group penetration rate	100%	50%	25%

Figure 6. Simulation map of downtown Helsinki. The streets are marked gray. The map scale is shown at the top. Nodes are numbered in blue. Regions of interest marked with red, blue, and black for geocast groups one, two, and three. Green circles depict the nodes' communication ranges.

5.1. Delivery Rate

Figure 7 depicts the delivery ratio simulation results of the various protocols as a function of the number of nodes deployed in the network.

The delivery rate for most protocols increases as the network density (i.e., the number of nodes) increases, as more contacts between nodes help message dissemination. One exception is the GeoEpidemic protocol that disseminates messages without any geographical constraints, thus wasting resources outside the ROI. This means that nodes can be congested with messages that are not useful, because they are outside their ROI, preventing other useful messages from being disseminated. This is particularly significant for geocast groups with small ROI. Therefore, the GeoEpidemic delivery rate does not increase for cases with more than 160 nodes. The other protocols limit message dissemination outside the ROI, so they are less affected by congestion. The V-GRADIENT always achieves the highest delivery ratio compared with the other protocols. using the case of the 160 nodes as an example, the V-GRADIENT shows a 19.5% increase in the overall delivery rate in relation to the GeoS&W protocol,

40.4% in relation to the MC-GeoDirectDelivery, 56.4% in relation to the GeoEpidemic Constrained, and 84.1% in relation to the GeoEpidemic protocol.
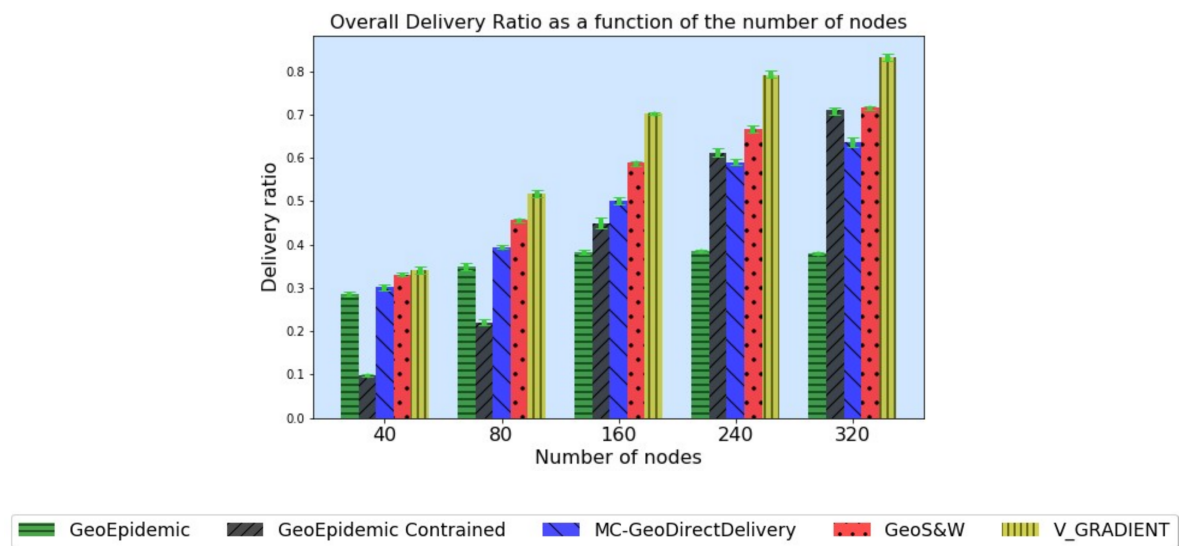

Figure 7. Delivery ratio as a function of node density and protocol.

The GeoEpidemic Constrained protocol has very poor performance for sparse networks (40 nodes), as it drops messages as soon as a node carrying them leaves the ROI. If there are very few nodes in the ROI and they leave it, messages will be permanently deleted, resulting in low delivery rates for sparse networks. For dense networks, flooding messages within the ROI is efficient and deleting them outside the ROI reduces network congestion, resulting in an improvement in the delivery rate.

MC-GeoDirectDelivery retains messages as nodes leave the ROI. This results in better performance for sparse networks, as more message copies circulate in the network, but degrades the performance for dense networks, as the resulting congestion reduces the delivery efficiency outside the ROI.

GeoS&W creates a limited number of message copies in nodes that are not direct destinations. This limits the congestion and also helps to spread messages in sparse scenarios, resulting in the second-best performance.

The V-GRADIENT has the best delivery ratio results proving the benefit of using a dynamic dropping policy that is stricter in denser networks to avoid congestion and softer in sparse networks to improve message dissemination. In addition, the more intelligent spraying of messages is more efficient than in GeoS&W.

5.2. Delivery Latency

Figure 8 shows the average message delivery latency for different protocols and geocast groups shown in Table 3.

Latency is only averaged for delivered messages. This is why the GeoEpidemic Constrained protocol has the lowest (i.e., best) latencies. As messages are dropped as soon as nodes leave the ROI, longer paths through nodes that leave and return to the ROI do not result in a successful delivery and, for sparse networks, the delivery rate is so low that all messages are delivered rapidly near the center of the ROI or not delivered at all.

GeoEpidemic disseminates messages through all nodes, even outside the ROI. Hence, it has the highest (i.e., worse) latencies, as some paths are very long. This is more significant for geocast groups two and three that have a smaller ROI, and for denser networks. The waste of resources for spreading messages outside the ROI causes congestion, degrading the protocol efficiency.

The V-GRADIENT provides a moderate and adequate solution, putting a limit on latency, by means of a dynamic dropping policy, while still achieving good performance regarding the delivery

ratio. For sparse networks, the V-GRADIENT has better delays, but they are similar to GeoS&W, since nodes are allowed to retain messages beyond their respective ROIs to ease further deliveries in the future. On the other hand, for denser networks, the latency slowly converges towards the GeoEpidemic Constrained results, since the buffer zone size is adapted dynamically according to network density and congestion level.
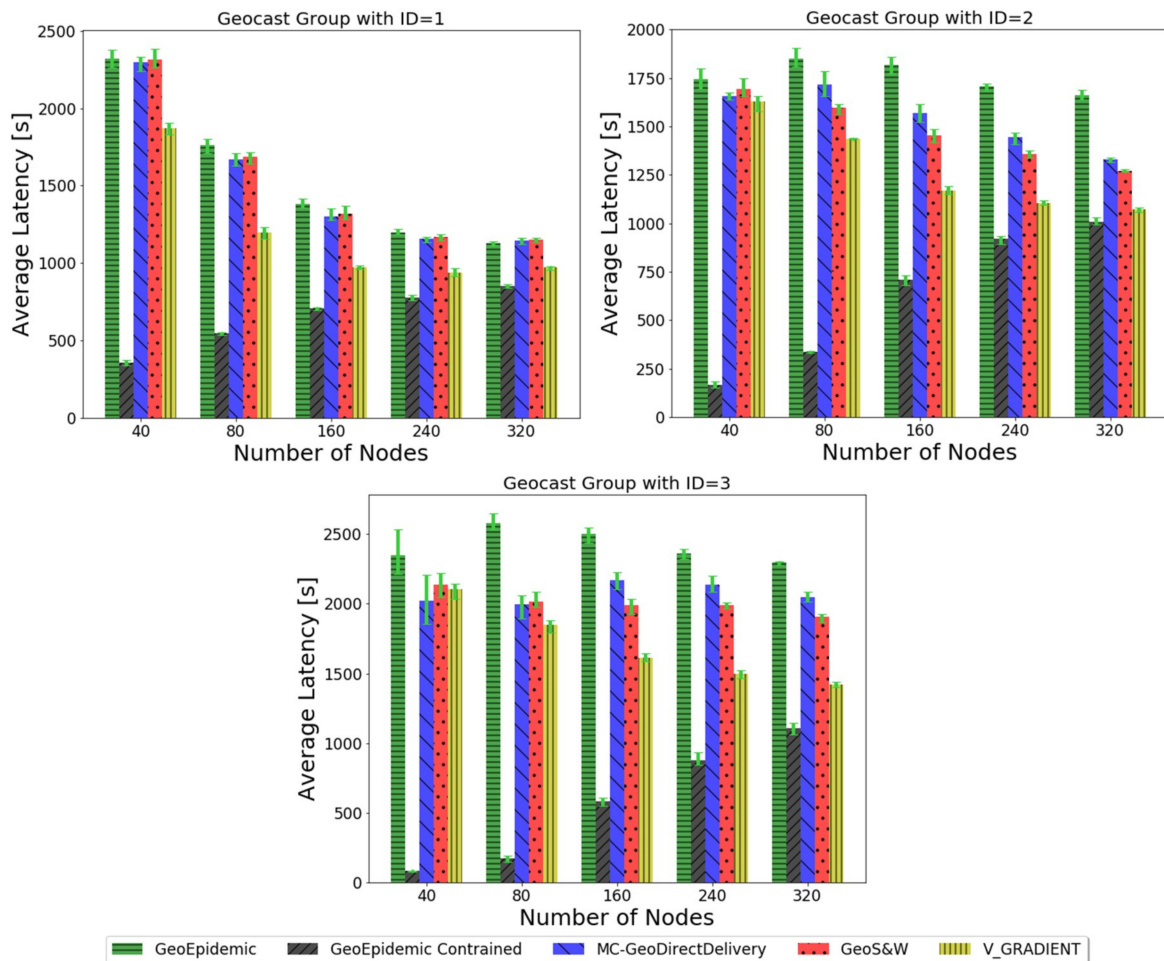

Figure 8. Delivery latency as a function of node density, protocol, and geocast group.

Considering the average latency results obtained for messages pertaining to the geocast group with ID = 3, which encompasses 25% of the total number of nodes, while the V-GRADIENT performance shows improvements as the number of vehicles deployed in the network grows, the results associated with the MC-GeoDirectDelivery and GeoS&W protocols show minor variations. This effect is related to the low degree of cooperation shown among nodes for the latter protocols, which limits the exploitation of an increasing number of contact opportunities. Conversely, the V-GRADIENT protocol carefully selects regions where the replication is not restricted, allowing nodes to seize the surrounding contact opportunities with the aim of enabling a quick propagation of messages locally, even though the selected carrier nodes may not show interest in such relayed messages.

The V-GRADIENT algorithm not only efficiently uses the density level and the buffer congestion metrics to infer the benefits of preserving messages being carried near their respective regions of interest, but also establishes location-based criteria to assess if vehicles that are not interested parties of a given message should cooperate in its dissemination process. Consequently, by making use of dynamic network information intelligently, the V-GRADIENT algorithm estimates which contact opportunities can improve the local spread of messages without compromising the delivery probability. Broadly

speaking, in this simulation scenario, the V-GRADIENT achieves considerably better delay-related results than the other protocols, apart from the GeoEpidemic Constrained protocol.

5.3. Protocol Overhead

Figure 9 shows the average overhead for different protocols and geocast groups shown in Table 3. Since results can vary significantly, a logarithmic scale is used for the overhead. The results for MC-GeoDirectDelivery are not presented, since this protocol, similarly to Direct Delivery, only delivers messages to interested nodes, resulting in an overhead of zero.
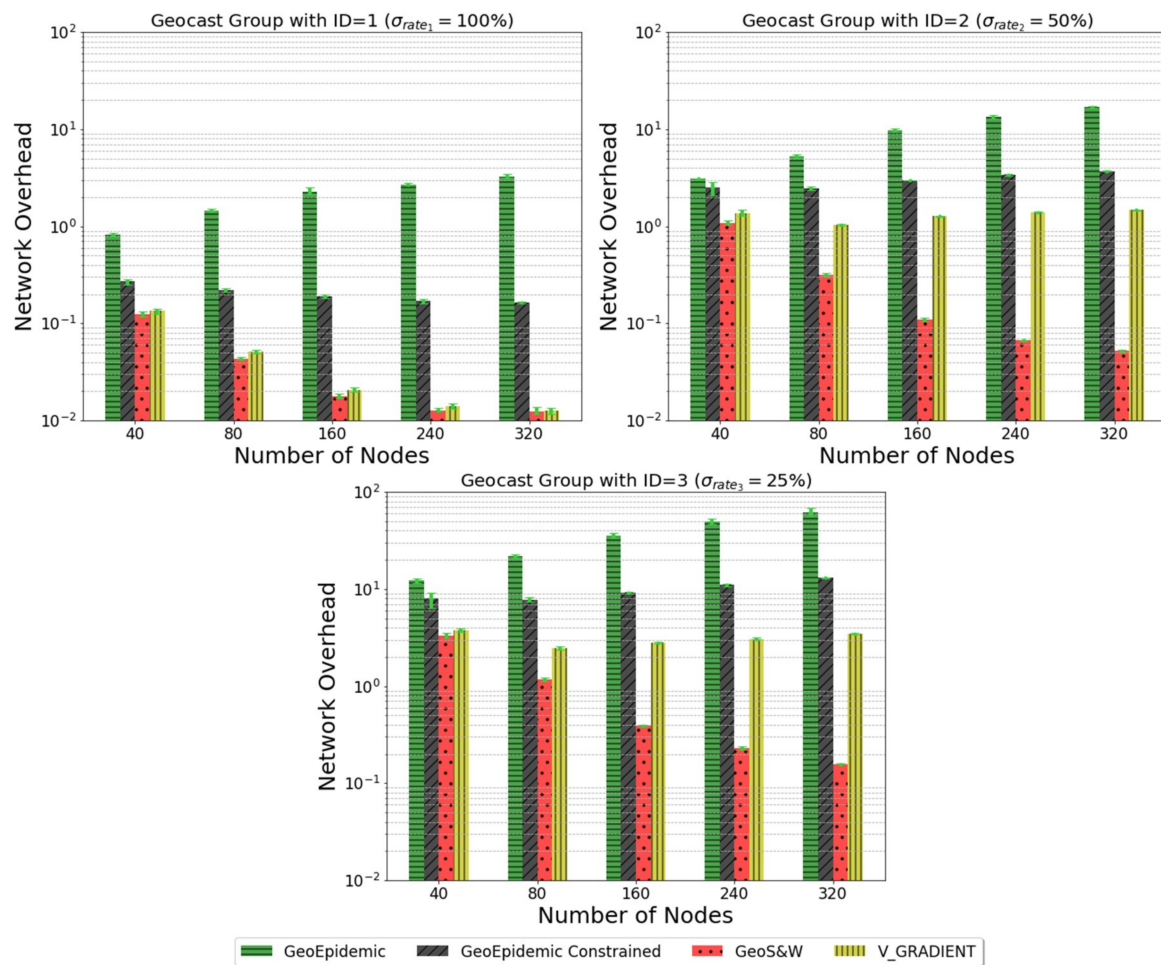

Figure 9. Overhead as a function of node density, protocol, and geocast group.

The network overhead tends to increase for all protocols as the geocast groups' penetration rates drop. For group one, which has 100% interested nodes, the overhead is the lowest, as most message replications are delivered to nodes interested in receiving them. For group three, which only has 25% interested nodes, the overhead is the highest, as nodes will also be replicating messages for many other nodes that do not want to receive those messages.

Looking at the results for geocast group with ID = 1, which includes all nodes deployed in the network, it is possible to see that the performance of the GeoS&W and V-GRADIENT protocols is similar and tends to significantly improve as the density level increases. Naturally, this enhancement is justifiable not only by the use of the same spraying mechanism on both strategies; on the other side, for the remaining geocast groups, the GeoS&W protocol yields better results since it imposes a limit on the number of copies to be transmitted, unlike the V-GRADIENT protocol. In turn, at the cost of jeopardizing the performance in terms of overhead, the V-GRADIENT algorithm exploits node

resources located inside the flooding range, $R_{flooding}$, to increase the delivery probability of geocast groups with low penetration rates, which was considered more important than the corresponding overhead cost.

6. Conclusions and Future Work

We proposed the new V-GRADIENT geocast routing protocol for VDTNs that efficiently disseminates messages within an ROI centered in the source node. The V-GRADIENT monitors node density, buffer occupation, and group penetration rate. Messages are disseminated within a buffer zone with sizes proportional to the size of the ROI and a function of node density and buffer occupation. The V-GRADIENT also uses a spraying mechanism that spreads a limited number of message copies to nodes that are not directly interested in the messages, and helps to relay messages towards their destinations in sparse networks and/or when the geocast groups have low penetration rates. The V-GRADIENT uses both the messages' initial TTL values and an estimation of the penetration rates of distinct geocast groups in order to lay out a flooding region comprised within the ROI.

After a comprehensive analysis of the simulation results was obtained, one can conclude that the V-GRADIENT is not only able to adjust its operation mode dynamically, according to the network density conditions, but can also cope with a wide variety of geocast application requirements. Simulation results show that the V-GRADIENT protocol outperforms the remaining geocast strategies in terms of the overall delivery ratio, which is the most relevant performance metric. The V-GRADIENT protocol achieved the second-lowest delivery latency outcome, surpassed only by the GeoEpidemic constrained, which has a very low delivery ratio as it only delivers messages quickly, near the ROI center. The protocol overhead performance of the V-GRADIENT is a compromise when compared to the GeoS&W and GeoDirectDelivery protocols, particularly for geocast groups with low penetration rates and/or under high-density conditions.

In our future work, we intend to test the V-GRADIENT in other scenarios, implementing message prioritization schemes and evaluating scenarios with misbehaving nodes that do not cooperate with message relaying.

References

1. Cao, Y.; Sun, Z. Routing in delay/disruption tolerant networks: A taxonomy, survey and challenges. *IEEE Commun. Surv. Tutor.* **2013**, *15*, 654–677. [CrossRef]
2. Maihöfer, C. A survey of geocast routing protocols. *IEEE Commun. Surv. Tutor.* **2004**, *6*, 32–42. [CrossRef]
3. Bazzi, A.; Masini, B.M.; Zanella, A.; Ilaria Thibault, I. Beaconing from Connected Vehicles: IEEE 802.11p vs. LTE-V2V. In Proceedings of the IEEE 27th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC) IEEE (2016), Valencia, Spain, 4–8 September 2016. [CrossRef]
4. Zanella, A.; Bazzi, A.; Masini, B.M. Relay Selection Analysis for an Opportunistic Two-Hop Multi-User System in a Poisson Field of Nodes. *IEEE Trans. Wirel. Commun.* **2017**, *16*, 1281–1293. [CrossRef]
5. Pereira, P.R.; Casaca, A.; Rodrigues, J.J.P.C.; Soares, V.N.G.J.; Triay, J.; Cervelló-Pastor, C. From Delay-Tolerant Networks to Vehicular Delay-Tolerant Networks. *IEEE Commun. Surv. Tutor.* **2012**, *14*, 1166–1182. [CrossRef]
6. Abali, A.M.; Ithnin, N.B.; Ebibio, T.A.; Dawood, M.; Gadzama, W.A. A Survey of Geocast Routing Protocols in Opportunistic Networks. In Proceedings of the International Conference of Reliable Information and Communication Technology, Johor, Malaysia, 22–23 September 2019; Springer: Cham, Germany, 2019; pp. 683–694. [CrossRef]

7. Nascimento, H.; Pereira, P.R. V-GRADIENT: A Density-Aware Geocast Routing Protocol for Vehicular Delay-Tolerant Networks. In Proceedings of the 10th Advanced Doctoral Conference on Computing, Electrical and Industrial Systems (DOCEIS'2019), Caparica, Portugal, 8–10 May 2019; IFIP Advances in Information and Communication Technology. Springer: Berlin, Germany, 2019; Volume 553, pp. 296–304. [CrossRef]

8. Spyropoulos, T.; Psounis, K.; Raghavendra, C.S. Single-copy routing in intermittently connected mobile networks. In Proceedings of the First Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks (SECON), Santa Clara, CA, USA, 4–7 October 2004; pp. 235–244. [CrossRef]

9. Spyropoulos, T.; Psounis, K.; Raghavendra, C.S. Spray and wait: An efficient routing scheme for intermittently connected mobile networks. In Proceedings of the ACM SIGCOMM Workshop on Delay-tolerant Networking (WDTN), New York, NY, USA, 22–26 August 2005; ACM: New York, NY, USA, 2005; pp. 252–259. [CrossRef]

10. Vahdat, A.; Becker, D. *Epidemic Routing for Partially-Connected Ad Hoc Networks*; Tech. rep., CS-200006; Department of Computer Science, Duke University: Durham, NC, USA, 2000.

11. Rajaei, A.; Chalmers, D.; Wakeman, I.; Parisis, G. GSAF: Efficient and flexible geocasting for opportunistic networks. In Proceedings of the IEEE 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), Coimbra, Portugal, 21–24 June 2016; pp. 1–9. [CrossRef]

12. Lu, S.; Liu, Y. Geoopp. Geocasting for opportunistic networks. In Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC), Istanbul, Turkey, 6–9 April 2014; pp. 2582–2587. [CrossRef]

13. Ott, J.; Hyytiä, E.; Lassila, P.E.; Vaegs, T.; Kangasharju, J. Floating content: Information sharing in urban areas. In Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom), Seattle, WA, USA, 21–25 March 2011; pp. 136–146. [CrossRef]

14. Keränen, A.; Ott, J.; Kärkkäinen, T. The ONE Simulator for DTN Protocol Evaluation. In Proceedings of the 2nd International Conference on Simulation Tools and Techniques ICST, Rome, Italy, 2–6 March 2009. [CrossRef]

15. Rajaei, A. The G-ONE Package for the ONE Simulator. Available online: https://github.com/aydinrajaei/g-one (accessed on 7 March 2020).

16. Palma, A.; Pereira, P.R.; Casaca, A. Multicast Routing Protocol for Vehicular Delay-Tolerant Networks. In Proceedings of the 5th International Workshop on Selected Topics in Wireless and Mobile Computing STWiMob'2012, 8th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob), Barcelona, Spain, 8–10 October 2012; pp. 761–768. [CrossRef]