LASP: Look-Ahead Spatial Protocol for Vehicular Multi-hop Communication

Rui Meireles[1,2], Peter Steenkiste[1], João Barros[2] and Daniel C. Moura[2]

[1]Department of Computer Science, Carnegie Mellon University, USA
[2]Instituto de Telecomunicações, FEUP DEEC, University of Porto, Portugal

***Abstract*—Vehicular multi-hop protocols typically employ distance-based metrics, which do not capture the complexity of vehicular connectivity. In this work we present LASP, a geographic protocol that uses a more accurate *spatial connectivity*-based metric. Spatial connectivity describes the historical probability of successfully delivering a packet from one geographic area to another. Analysis of data collected from a vehicular testbed showed that, unlike other metrics, spatial connectivity indirectly captures all major factors affecting wireless connectivity. Moreover, it is temporally stable, which makes it useful in estimating the quality of future co-located links. When forwarding, LASP uses spatial connectivity information to pick a well-connected geographic forwarding zone, inside which multiple nodes cooperate in relaying through a distributed prioritization scheme. Compared with other techniques where the sender picks a specific next hop relay, cooperative forwarding improves resilience to losses through *vehicle diversity*. We evaluated LASP on a 30-node testbed, where it achieved a 30% increase in packet delivery ratio over the benchmark GPSR protocol.**

I. Introduction

Vehicular Ad-Hoc Networks (VANETs) enable wireless communication among vehicles and between vehicles and infrastructure through the IEEE 802.11p standard [1]. In open areas, 802.11p can achieve a communication range of over 1 Km. But, in cities, this can drop to less than 100 m [2], due to line of sight obstructions and multipath. This makes multi-hop communication key to better vehicular wireless coverage.

Multi-hop can be divided into routing, the discovery of a path between source and destination; and forwarding, the process of conveying packets along said path. Both are challenging in vehicular networks. First, fast movement combined with a rich scattering environment lead to *unstable topologies* [3]. This precludes global dissemination of connectivity information, and hence the use of traditional topology-based protocols. Ergo, routing must rely on heuristics. Second, connectivity varies with spatial location—*spatial heterogeneity*. Both in terms of relay availability and link quality, depending on line of sight conditions created by buildings and other obstacles. This makes it difficult to define an effective routing metric.

Most vehicular protocols use geographic information to build paths hop-by-hop, often picking the neighbor closest to the destination as a relay (e.g., GPSR [4], MoVe [5]). Building paths one hop at a time increases robustness to topology changes. However, maximizing traveled distance assumes a strong correlation between distance and connectivity, which, due to obstacles, has been shown not to exist in vehicular networks [6]. Also, having the sender choose a specific next hop relay leaves it susceptible to instability on the chosen link.

Other schemes use road maps to find a sequence of roads to forward along, thus avoiding obstacles (e.g., SAR [7], ACAR [8]). Data packets are however not bound to roads like cars are, so requiring them to follow roads excludes forwarding opportunities made possible by the environment's topography.

With this in mind we introduce the concept of *spatial connectivity*. Spatial connectivity describes the measured communication probability between a pair of geographic areas. Since factors affecting spatial connectivity change slowly, it can be used to create a graph abstraction of the network's topology. Our Look-Ahead Spatial Protocol (LASP) creates such a spatial graph and uses it to look ahead at network conditions beyond the local neighborhood. Since it is independent of the actual links between nodes, this representation is more stable than traditional node-level topology graphs. For increased accuracy, historical data is replaced by real-time data in the local neighborhood, where it is available.

Once formed, the spatial graph is used to build, hop by hop, routes that maximize end-to-end delivery probability. Although LASP can be used with traditional node-based forwarding, we maximize its benefits by employing zone-based forwarding. Instead of being forwarded to a specific node, packets are forwarded to a geographic zone [9]. Zone nodes then coordinate to determine which of them will actually forward the packet, based on their individual end-to-end delivery probabilities. This maximizes reliability, as different vehicles experience different losses—*vehicle diversity*.

Our contributions can be summarized as:

- We justify the use of spatial connectivity through an analysis of vehicular connectivity data collected from two testbeds spanning over 100 vehicles (§II).
- We design and implement LASP, a multi-hop protocol that leverages spatial connectivity and vehicle diversity to maximize end-to-end delivery probability (§III and §IV).
- We experimentally evaluate LASP against the GPSR protocol using a 30-node testbed (§V).

II. Spatial connectivity

We introduce the spatial connectivity-based metric used in LASP and evaluate its suitability for routing and forwarding.

A. Concept

Real-time Packet Delivery Ratio (PDR) makes for a good link metric [10]. However, given vehicular networks' unstable topologies, network-wide PDR dissemination is impractical. We could also try to model signal propagation, but doing it accurately would be very complex. We propose a more pragmatic approach. Since link quality is mostly determined by the combination of distance and line of sight conditions, links between co-located but distinct vehicle pairs will be similar. This leads us to the concept of *spatial connectivity*. We define it, for a pair of geographic areas, as the likelihood that a packet sent by a vehicle in one of the areas will be received by one or more vehicles in the other area. In short, the PDR between a pair of geographic areas.

Because line of sight is mostly affected by topography, which changes slowly, we claim spatial connectivity based on historical PDRs can be used to identify future geographic paths. Also, since it is measurement-based, it indirectly captures all factors influencing vehicular connectivity.

In order for it to make sense to use spatial connectivity in this context, a number of requirements must be met. Namely:

1) Many protocols do road map-based routing. Does spatial connectivity uncover connectivity patterns that are not captured by a road map?
2) Most protocols use distance as an heuristic. Is spatial location a better PDR predictor than distance alone?
3) LASP uses historical data to abstract network connectivity. Is spatial connectivity temporally stable?

Following, we use testbed data to answer these questions.

B. Data collection

A data set with spatially-indexed PDR data is required for our analysis. Since no such data existed [11], we collected it from two vehicular testbeds in the city of Porto, Portugal.

The first is HarborNet [12], a 30 node (25 On-Board Units (OBUs)), 5 Road-Side Units (RSUs)) network deployed in the Leixões Harbor — Fig. 1. The OBUs are installed in the trucks that carry the shipping containers. The harbor occupies an area of just 1 Km^2, so node density is high. Stacked containers often obstruct the line of sight, like buildings in a city.

The second testbed, PortoVANET, has around 100 OBUs deployed in taxi cabs and buses operating in the city of Porto. Although it is larger, node density in this testbed is lower.

All units in both testbeds implement the IEEE 802.11p and WAVE standards [13]. Observed communication ranges were similar in both testbeds, which shows that the harbor can be used as a reasonable proxy for an urban setting.

To gather connectivity data for analysis, we had all nodes broadcast Cooperative Awareness Messages (CAMs) [14] at a 10 Hz rate and keep two time-indexed (1 Hz resolution) logs:

1) A location log storing the node's location coordinates and number of sent CAMs for each timestamp.
2) A message reception log storing the number of received CAMs from each sender for each timestamp.

The merged logs yield a complete network topology trace.

Fig. 1. HarborNet topology recorded on June 2nd 2014, 12 PM. Balloons represent nodes and lines represent inter-node links. Imagery ©2014 Google, Cnes/Spot Image, DigitalGlobe and IGP/DGRF (41 11' 25" N, 8 41' 24" W).

C. Lost opportunities in road-based routing

Road-based protocols forward packets towards the destination along roads, relaying packets between vehicles in the same or in adjacent road segments. We now assess how many relaying opportunities are lost due to this limitation. We used two weeks worth of PortoVANET data and a city road map for this analysis (the harbor does not have a road map). The map is a graph where each intersection is a node and each pair of nodes with a road between them is an edge.

For each received CAM, we took the sender's and receiver's positions and mapped them to the closest road segments. We then classified the results into 3 classes: i) *same*, when both vehicles mapped to the same segment, ii) *adjacent*, when the vehicles mapped to segments that are direct neighbors in the road map graph and, iii) *non-adjacent*, for all other cases.

Tab. I summarizes the results. Only 6% of links mapped to a single segment, and 14% to adjacent ones. Most exchanges (80%) occurred between non-adjacent segments, confirming that the road network graph used by vehicles is fundamentally different from the wireless connectivity graph.

	Frequency	
Type	**Absolute**	**Relative**
Within same road segment	6930	6%
Between adjacent segments	15801	14%
Between non-adjacent segments	90404	80%
Total	113135	100%

TABLE I
COMMUNICATION CLASSIFICATION RELATIVE TO THE ROAD NETWORK.

D. Distance as a routing metric

We now look at the relationship between PDR and sender-receiver distance or link length. Many protocols, e.g. GPSR [4] and SAR [7], assume that closer distances equate to better connectivity. However, studies [2], [6] have shown that distance alone can not fully explain PDR. To confirm this we took one month of connectivity data collected from the HarborNet testbed (chosen for its higher node density) and plotted PDR as a function of link length (i.e. distance between endpoints).

The boxplot of Fig. 2a shows the results. As expected, a large variance in PDR can be seen in all distance bins. This makes it hard to infer link performance from distance, which represents a strong case against its use as a forwarding metric.

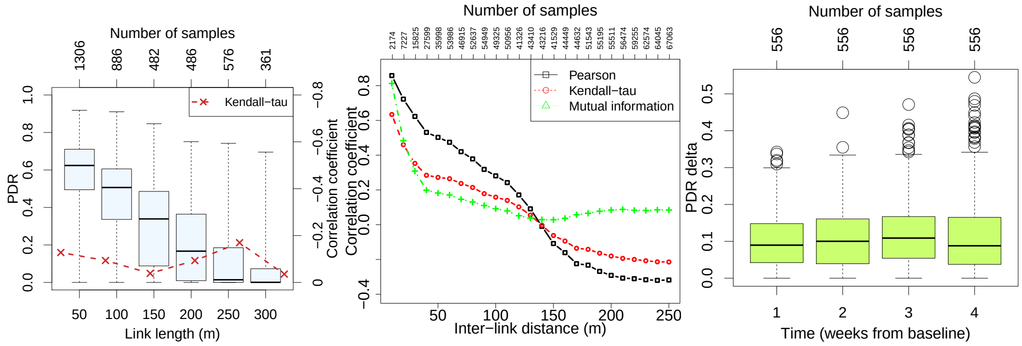

(a) PDR versus link length (b) Spatial correlogram (c) PDR temporal variation

Fig. 2. Results show that: i) distance alone can not predict PDR, ii) PDR exhibits spatial autocorrelation and iii) spatial connectivity is stable in time.

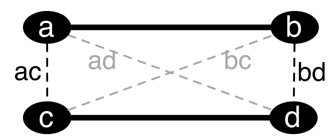

Fig. 3. Inter-link distance is defined as the minimum sum of distances between endpoints. Thus, in the figure, $d(ab, cd) = d(a, c) + d(b, d)$.

E. Spatial location as a predictor for PDR

To quantify how well spatial location can predict PDR we studied the spatial autocorrelation in PDR data, i.e., whether co-located links exhibit similar PDRs. We started by defining an inter-link distance metric as follows. A link is defined by its endpoints, e.g., ab represents a link between a and b. Let the inter-link distance be the minimum sum of distances between endpoints of the two links, as per Eq. 1:

$$d(ab, cd) = \min\left\{[d(a,c)+d(b,d)], [d(a,d)+d(b,c)]\right\}, \quad (1)$$

where $d(a, c)$ is the distance between a and c. Fig. 3 shows an example. Due to symmetry, we need only consider two endpoints combinations for links ab and cd: $d(a, c) + d(b, d)$ and $d(a, d) + d(b, c)$. In this case, $d(a, c) + d(b, d)$ is minimal.

We used one month of HarborNet data. Space was discretized into a 10×10 m grid. We then computed the correlation between PDRs as a function of inter-link distance. We used 3 distinct metrics: Pearson, which captures linear relationships; Kendall-τ, which captures monotonic relationships; and mutual information, which can capture any relationship.

Fig. 2b shows the results. All metrics showed high correlation for short inter-link distances, confirming the relationship between spatial location and PDR. Correlation dropped gradually with distance, hitting zero around 150 m. At larger distances, mutual information stayed close to zero, while the other metrics exhibited negative correlation: long distance, low PDR links correlated negatively with short, high PDR links.

F. Temporal stability of spatial connectivity

If spatial connectivity were to change quickly over time, historical data would not be useful to find future paths. To test the stability of spatial connectivity we took one week of HarborNet data, computed the PDRs for each pair of discretized 10×10 m spatial cells and set it as a baseline. We then repeated the process for the following weeks, and calculated the delta relative to the baseline.

Fig. 2c shows the results. Since containers are constantly being moved, the harbor can be considered a worst-case scenario. Still, the median PDR delta relative to the baseline hovered around 10%. Variance increased with time, but slowly.

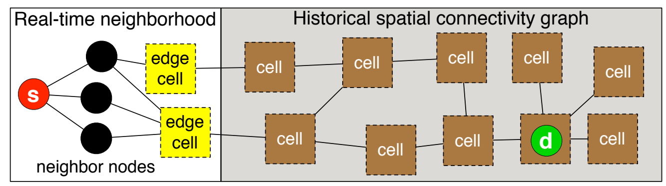

Fig. 4. LASP mixes real-time and historical spatial connectivity information.

III. LASP ROUTING DESIGN

LASP is a geographic unicast multi-hop protocol that tries to maximize end-to-end delivery probability. It is split into routing and forwarding modules. This section focuses on routing, which generates an abstraction of the network topology that is then used to make per-packet forwarding decisions (§IV).

A. Overview

Routing is traditionally based on a topology representing connectivity between network nodes. But, since vehicular links are unstable, network-wide dissemination of node-level connectivity is impractical. LASP's approach is to model the network topology using two graphs, as per Fig. 4: i) a real-time node-level graph for the local neighborhood and, ii) a historical spatial connectivity graph for the rest of the network.

Vertices in the spatial graph represent not network nodes but small geographic regions we call *spatial cells*. Space is discretized into a grid so each physical location maps to a single cell. An edge (u, v) in the spatial graph represents the historical probability that, if a packet is sent by a node on cell u, it will be received by at least one node in cell v. The product of edge probabilities over a path yields the delivery probability between any two spatial cells. We refer to this graph as a *look-ahead*, as it leverages PDR spatial autocorrelation (§II-E) to abstract connectivity patterns beyond the node's immediate vicinity. It provides a stable topology representation (§II-F).

Since historical data may not reflect current conditions, LASP uses real-time information for the local neighborhood (i.e. up to n-hops away), where it is feasible. The two graphs must be melded to yield a complete view of the network's topology. To allow this, the real-time graph has both nodes and spatial cells as vertices (see Fig. 4). Direct neighbors of the node s building the graph become vertices with direct edges to node s. The remaining vertices are spatial cells at the boundary of the neighborhood (i.e., cells that are reachable after n hops in an n-hop neighborhood). The edges between the neighbor nodes and boundary cells represent the spatial connectivity between neighbors and cells (described in §III-C).

This strategy makes the real-time graph and spatial look-ahead graphs compatible, while simultaneously encoding fine-grained information about the available next hop nodes. The combined graph is used to compute the end-to-end packet delivery probability for each possible next hop. This information is then fed to the forwarding module (§IV).

B. Historial spatial connectivity graph

The historical spatial connectivity graph is crowd-sourced from network nodes. Nodes broadcast CAMs periodically and maintain a spatially-indexed log of how many CAMs were sent and received, as in §II-B. A server merges these logs to create the global historical spatial look-ahead graph for the entire network. It is defined as $G_h = (V_h, E_h)$:

1) Space is discretized into cells, each becoming a vertex $v \in V_h$. For simplicity we use 50×50 m cells, a size that ensures good spatial autocorrelation (Fig. 2b).
2) For every cell pair (u, v) that can communicate, let edge $e = (u, v) \in E_h$. Each edge e is weighted with the historical spatial delivery probability P_{hdel}.

P_{hdel} for $e = (u, v)$ is the probability that, if there is a node in cell u, there are also one or more nodes in v, of which at least one can receive packets sent by u. It is calculated as:

$$P_{hdel}(u, v) = P(nodesAt(v) \mid nodesAt(u)) \times P_{u \to v}. \quad (2)$$

The first factor is obtained from the connectivity logs. The second is the ratio of received to sent CAMs in the two cells. Since it is conditioned on the presence of nodes at u, the P_{hdel} of a path is the product of the edge P_{hdel} probabilities.

C. Real-time neighborhood graph

The real-time neighborhood graph is built independently by each node from a neighborhood table with locations and link PDRs for nodes up to n-hops away, which is populated from CAM reception data. The graph $G_r = (V_r, E_r)$ is defined as:

1) Let s, the node building the graph, and all of its immediate neighbors t be vertices in V_r.
2) Let $e = (s, t) \in E_r$ be an edge between node s and each of its neighbors t, annotated with the link PDR.
3) Take the set $\mathcal{Q}$ of nodes that are n-hops away from s (i.e., at the boundary). For each $u \in \mathcal{Q}$, discretize u's position to get a spatial cell c. Let c be a vertex in V_r.
4) For each combination of neighbor t and spatial cell vertex $c \in V_r$, let (t, c) be an edge $e \in E_r$. Each edge is annotated with the real-time P_{rdel} of delivering a packet from neighbor t to cell c (i.e., to any node in c). Since there can be many paths to reach a node in c from t, we use the P_{rdel} of the best-connected path. Formally:

$$P_{rdel}(t, c) = \max_{Path_{t \to c}} \left[\prod_{(u,v) \in Path_{t \to c}} P_{u \to v} \right]. \quad (3)$$

G_r and G_h combined yield a graph with node-level data for the first hop and spatial cell-level data for the remainder.

D. Estimating per-neighbor end-to-end delivery probability

Given a destination, the end-to-end delivery probability P_{edel} for a specific neighbor defines its quality as a forwarder, and is the information input to the forwarding module (§IV).

It is estimated from the real time G_r and historical G_h graphs as follows. For each possible next hop relay t, P_{edel} can be seen as the joint probability of 3 simpler events. First, the packet must reach t from the current node s. Then it must go from t to a cell c_e at the outer edge of the local neighborhood. Finally, the packet must travel from c_e to the destination cell c_d. The joint probability is the product of the 3 event probabilities, where the first two come from G_r and the last from G_h. Formally:

$$P_{edel}(s, t, c_d) = P_{s \to t} \times \max_{c_e \in \mathcal{C}_n} \Bigg\{ P_{rdel}(t, c_e) \times \max_{Path_{c_e \to c_d}} \left[\prod_{(u,v) \in Path_{c_e \to c_d}} P_{hdel}(u, v) \right] \Bigg\}. \quad (4)$$

Note that since there are multiple possible neighborhood edge cells c_e and spatial paths between each c_e and c_d, we consider the instantiation that maximizes P_{edel}.

This computation can be done for all next-hops and destinations using the Floyd-Warshall algorithm in $\mathcal{O}(|V^3|)$ time. When forwarding we want to use the latest real time data, so in our prototype the final probabilities are computed on demand for each packet. This can be done cheaply because, since G_h is constant, the paths' historical portions can be precomputed.

IV. LASP FORWARDING DESIGN

LASP routing can be coupled with a myriad of forwarding schemes. One option is simply to forward the packet to the neighbor that maximizes P_{edel} (we call this design variant LASP-SF, for Single Forwarder, and test it in §V). But that is not our preferred method because individual vehicular links are unstable. We note that, although at a macro level co-located links share similar properties (e.g., PDR), at a micro level individual packet losses are independent. We leverage this *vehicle diversity* by employing zone-based forwarding in LASP, which we describe in this section.

A. Overview

In zone-based forwarding, packets are addressed to a geographic forwarding zone, which we define to be composed of one or more spatial cells. Then, nodes inside the zone run a coordination algorithm to elect the best forwarder among the packet receivers, while others stand by to ensure reliability.

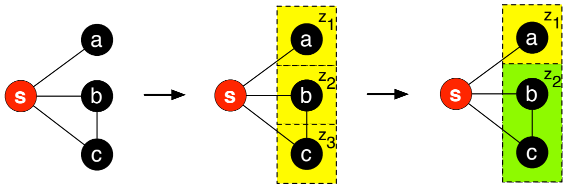

Fig. 5. Candidate forwarding zone formation.

More concretely, in LASP the sender starts by inspecting the packet header to find the destination's location (assumed known to the source, e.g., via a location service [15]). Then it aggregates individual P_{edel} values to find the forwarding zone that maximizes the expected end-to-end delivery probability to the destination's location. Next, the packet is sent and receivers inside the zone prioritize themselves according to their individual P_{edel} estimates. Following, we detail these computations and describe a recovery mechanism that is triggered if a network hole is found.

LASP forwarding is a generalization of DAZL [9] zone-based forwarding, which used distance to prioritize forwarders and could not recover from network holes.

B. Forwarding zone selection

A large forwarding zone is desirable to maximize vehicle diversity gains, but it also increases the likelihood of hidden terminals. LASP uses cells as a routing unit, but it is beneficial to decouple cell and zone sizes so we can balance this tradeoff.

Thus a forwarding zone is defined as a set of one or more spatial cells. LASP starts by defining a set of candidate forwarding zones $\mathcal{Z}$. The candidate zones start as single cells corresponding to each neighbor location and then zones that can communicate with probability $P > P_{min}$ are merged to maximize the zones' sizes. Fig. 5 exemplifies. Here there are three neighbors: a, b and c. Since b and c can communicate, they both share zone z_2, while z_1 remains separate.

Now, for each zone $z \in \mathcal{Z}$, LASP calculates the aggregate zone end-to-end delivery probability P_{zondel}. Let $\mathcal{T}$ be the set of neighbors of the current node s that are inside z. P_{zondel} is the complement of the probability that all neighbors in $\mathcal{T}$ fail to either get the message from s or deliver it to the destination cell c_d. For each $\mathcal{T}_i \in \mathcal{T}$, the probability of reaching c_d is P_{edel} from Eq. 4, with $\mathcal{T}_i$ seen as the starting node s. Formally:

$$P_{zondel}(s, \mathcal{T}, c_d) = 1 - \prod_{i=1}^{|\mathcal{T}|} 1 - \left\{ P_{s \to \mathcal{T}_i} \times \max_{t \in neigh(\mathcal{T}_i)} [P_{edel}(\mathcal{T}_i, t, c_d)] \right\}. \quad (5)$$

The packet is forwarded to the zone that maximizes P_{zondel}.

Let us consider the example scenario in Fig. 6. Node s wants to choose a forwarding zone for a packet destined to d. Assume a single hop real-time graph. First, the neighbors' locations determine the candidate zones. Here, a and b fall into the same zone z_1, while c is in another, z_2. The two zones do not communicate, so they can not be merged. Now, s uses G_r to find the probability of reaching each neighbor ($P_{s \to \text{neigh}}$) and cell (P_{rdel}, Eq. 3), and G_h to find P_{hdel}. The product of all three factors yields the P_{edel}. Assume the results are:

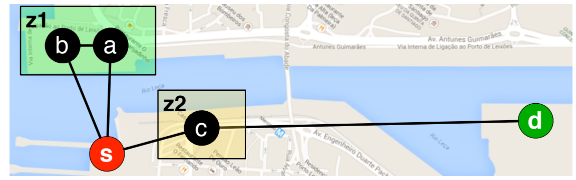

Fig. 6. LASP example operation.

Neighbor	$P_{s\to\text{neigh}}$	P_{rdel}	P_{hdel}	P_{edel}
a	0.6	1	0.5	0.3
b	0.4	1	0.5	0.2
c	0.6	1	0.5	0.3

Now s can compute the end-to-end delivery probability for each zone, P_{zondel}. With a 1-hop look-ahead, the last term of Eq. 5 uses only historical data. Therefore, the probability of failure of a given neighbor becomes the complement of its previously calculated P_{edel}. The results are the following:

Forwarding zone	P_{zondel}
z_1	$1 - [(1 - P_{edel}(a) \times (1 - P_{edel}(b))] = 1 - (0.7 \times 0.8) = 0.44$
z_2	$1 - (1 - P_{edel}(c)) = 1 - 0.7 = 0.3$

Zone z_1 will therefore be the chosen forwarding zone.

C. Forwarder coordination

Having multiple potential forwarders benefits reliability, but having all of them forward would cause unwanted congestion. The goal is to have nodes with higher P_{edel} take precedence. To do this, each node starts by independently ranking itself against other candidates and then waits for a period of time inversely proportional to its rank before forwarding. If, while waiting, the node overhears another's transmission, it cancels its own, as it would now be redundant. This achieves: i) coordination without explicit communication, ii) prioritization of the best candidates and, iii) replication avoidance.

To rank them, each node must estimate the set of candidate forwarders, $\mathcal{F}$, and the quality of its members. This depends on the size of the real-time neighborhood. If it is 2 hops or larger, the previous hop's neighbors are known and the ones in the forwarding zone define $\mathcal{F}$. Otherwise, $\mathcal{F}$ is approximated as the set of the node's neighbors that are inside the zone.

The quality of each candidate in $\mathcal{F}$ is P_{edel} from Eq. 4. The set is then ordered in decreasing P_{edel} order to produce the ranking. The node computes its waiting time as the product of its rank by *ibf*, interval between forwarders, a parameter.

Let us again use Fig. 6 as an example. The zone is z_1, with nodes a and b. a has b in its neighbor table, and vice versa, so they will agree that $\mathcal{F} = \{a, b\}$. Assume that a has the highest P_{edel}. What will happen then is that a will forward immediately, while b waits. When b overhears a's transmission, it will cancel its own, preventing replication.

D. Recovering from dead ends

LASP's heuristic may occasionally fail, so a recovery mechanism is needed. If after sending a packet, no acknowledgment

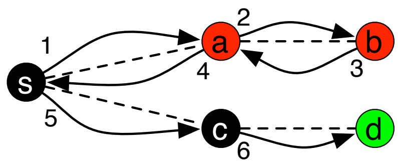

Fig. 7. LASP backtracking example.

is heard from within the chosen forwarding zone, the node triggers a recovery procedure. Recovery starts by sending the packet to the zone with the second-highest P_{zondel}. If this also fails, the packet is sent to the third-best zone, and so on. If none of these attempts is successful, the node sends the packet backtracking to the node it first received it from. The procedure will then be repeated until forwarding is successful.

Backtracking may fail if the packet holder becomes disconnected from the previous hop sender. In this case, the packet is stored until a new neighbor is encountered.

In short, LASP operates as a depth-first search where the order in which children are visited is dictated by P_{zondel}. To avoid loops, packets include a list of previously visited nodes.

Let us use Fig. 6 as an example. Assume that, as aforementioned, node s chooses zone z_1 and inside z_1, a takes precedence over b. Assume also that this was a bad choice, for there is no actual path between z_1 and the destination. Fig. 7 shows what will follow. The packet will go from s to a. Node a, having no neighbors other than b, will use z_1 again as a forwarding zone. Node b will take the packet and, realizing all of its neighbors (s and a) were already visited, will backtrack the packet to a, which will in turn backtrack it to s. Node s now filters out the previously visited zone z_1 and sends the packet to z_2, where c relays it to the destination.

V. Experimental evaluation

A. Setup

We evaluated LASP on the HarborNet testbed (the larger PortoVANET is too sparse for multi-hop communication). For comparison purposes, we ran multiple protocols in parallel. The main benchmark was GPSR [4]. GPSR is a geographic protocol that picks the neighbor closest to the destination as the next hop. If a local minimum is reached, it enters a recovery mode in which it traverses the network graph using a right-hand rule (picking the first node counter-clockwise relative to the line segment defined by the current node and destination locations), until greedy forwarding can resume. GPSR was chosen because: i) it represents a large class of protocols; ii) does not need a road map (unavailable for the harbor) and; iii) has a well-defined implementation provided by the authors.

We also tested LASP-SF (SF for Single Forwarder), the LASP variant that picks a specific next hop node instead of a forwarding zone. This let us isolate the contributions of vehicle diversity and spatial connectivity to overall performance.

Tab. II summarizes the parameters used. The neighborhood table used by all schemes came from CAMs sent at 10 Hz. It kept each neighbor's location and PDR over the last 5 s. The cell size was chosen in accordance with the PDR spatial autocorrelation results in §II-E. The traffic pattern was as follows: each node sent a 100 byte packet once a second to a node picked uniformly at random.

Parameter	**Applicability**	**Value**
Max. number of hops	All	20
Max. retransmissions per hop	All	5
Retransmission delay (ms)	All	100
Backtracking delay (ms)	LASP, LASP-SF	400
Spatial cell size (m)	LASP, LASP-SF	50×50
Real-time neighb. size (hops)	LASP, LASP-SF	1
Max. number of forwarders	LASP	5
Interval between forwarders (*ibf*) (ms)	LASP	40
Min. PDR between forward. zone cells	LASP	0.7

TABLE II
PROTOTYPE CONFIGURATION PARAMETERS.

B. Results

We consider four end-to-end metrics: PDR, a measure of reliability; path length and transmission count, measures of network load; and delay, a measure of temporal overhead.

1) Packet Delivery Ratio (PDR): It is hard to compute absolute PDRs since the ground truth of whether a packet is deliverable is unknown. We tackled this in two ways. First, we considered the PDR of packets delivered by at least one protocol, i.e., proved deliverable by example. GPSR delivered 72% of these. LASP-SF delivered 15% more (83%), and LASP 30% more (94%). This means that GPSR's greedy distance-based strategy was more prone to failure than LASP's. Also, since LASP-SF yielded roughly half of LASP's improvement, it can be said that both vehicle diversity and spatial connectivity contribute to LASP's performance in similar amounts.

Further, we used time-indexed CAM reception data (as per §II-B) to estimate the optimal end-to-end PDR. For each node pair, we defined an erasure channel with a PDR equal to the CAM PDR measured between them. All PDR-maximizing paths were then found from the resulting graph (note that these paths are PDR-optimal for a single forwarder scheme; their PDRs can be surpassed by using a multi-forwarder protocol).

Fig. 8a shows PDR results as a function of the estimated optimum. GPSR delivered less than 80% of the packets with an estimated PDR of 100%. Relative to GPSR, LASP-SF provided an improvement of roughly 15% over the entire PDR range. LASP outperformed LASP-SF, but the delta varied with the estimated PDR. For packets with high estimated PDRs, the improvement was smaller. For example, for packets with 100% estimated PDR the improvement was 6% (93% for LASP versus 88% for LASP-SF). This is natural since when link quality is good, the additional gains from node diversity are limited. When links are poor, the opposite happens. For example, for the packet subset with 40% estimated PDR, LASP doubled LASP-SF's performance (64% versus 32%), exceeding even the estimated optimal single-forwarder PDR.

2) Path length: Fig. 8b depicts the Empirical Cumulative Distribution Function (ECDF) for the hop count of the paths found by all 3 tested protocols. GPSR yielded the shortest paths of all, which is consistent with its greedy distance-based metric. LASP-SF found the longest paths, with roughly double

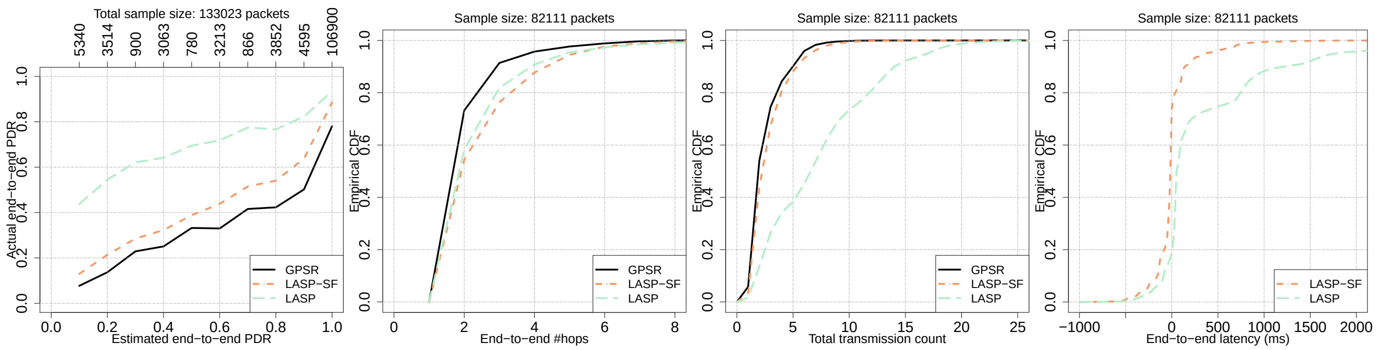

(a) Actual vs estimated PDR (b) Number of hops (c) Transmission count (d) Delay

Fig. 8. Experimental results: LASP improves PDR at a cost in terms of transmissions and delay. LASP-SF improves PDR (less) without compromise.

the number of paths 3 hops or longer relative to GPSR (20% of the ECDF above 3 for LASP-SF vs less than 10% for GPSR). LASP fell in between the other two. This means that relative to LASP-SF, zone-base forwarding uncovered shorter paths that spatial connectivity alone could not predict existed.

3) Transmission count: Fig. 8c presents the ECDF of the per-packet total transmission count. GPSR and LASP-SF behaved similarly in terms of transmissions, despite the longer paths used by LASP-SF. This leads us to conclude that the spatial connectivity-based metric chose better connected neighbors, reducing the need for retransmissions. LASP transmitted more than the other two protocols. GPSR and LASP-SF cross the ECDF's 80$^{\text{th}}$ percentile at around 3, while LASP does so around the 12 transmission mark. This leads us to believe that further tuning of the waiting time before forwarding is necessary to prevent redundant transmissions.

4) Delay: Since that end-to-end acknowledgments were not implemented, Round Trip Time (RTT) could not be used to measure delay. We thus defined the time at which a packet was delivered by GPSR to be zero, and measured LASP and LASP-SF's delay relative to that baseline. Fig. 8d presents the ECDF of this relative delay. LASP-SF's ECDF crosses zero at around the 70$^{\text{th}}$ percentile, meaning most of its packets arrived earlier than GPSR's. This is an indication that the spatial connectivity metric chose better connected nodes, which needed fewer retransmissions and hence, less time. LASP crosses zero at around the 30$^{\text{th}}$ percentile and features a long tail beyond the 80$^{\text{th}}$ percentile. This added delay comes from the forwarder coordination and recovery mechanisms.

VI. Related work

Traditional ad-hoc network routing protocols are based on proactive dissemination of topology information, e.g., OLSR [16], Batman [17] and Babel [18]. Abolhasan *et al.* [19] tested all three and reported convergence times between 15 and 60 s, which render them unsuitable for vehicular networks.

Others are topology-based but reactive, e.g., DSR [20], AODV [21] and TORA [22]: routes are found on demand, by flooding. DSR tested well in a 5 vehicle-test [20]. However, flooding does not scale to larger networks, as Jaap *et al.* [23] confirmed using simulations with up to 600 nodes.

From the need to cope with more dynamic networks came geographic routing. Geographic protocols route packets towards the destination's location using heuristics such as relative positions, trajectories, road maps and statistical traffic data. Distance is the most popular heuristic. Protocols that use it compute paths on-the-fly, hop by hop, by having the packet holder choose the neighbor closest to the destination as the next hop. GPSR [4] was seminal, spanning many variants. GPCR [24], for example, varies only in the recovery mechanism used to get out of local minima, and MoVe [5] uses trajectory data to improve distance estimations. However, this use of relative positions assumes spatial connectivity homogeneity, which does not exist in vehicular networks.

Later came the idea of using road maps to support geographic routing. For example, TrafRoute [25] uses a map to limit the vehicles forwarding packets to those close to intersections, improving scalability. Other schemes select a chain of roads that connect source and destination and then forward packets along them. Road selection strategies vary from shortest distance (e.g., GSR [26], SAR [7], LOUVRE [27]), to most-connected according to traffic information (e.g., GyTAR [28], ACAR [8]). Between intersections, all but ACAR use GPSR-like greedy forwarding based on distance. ACAR uses observed PDRs to minimize packet error rates. Unlike LASP, these protocols assume homogenous signal propagation across space and preclude the use of areas outside roads.

Studies have shown VANET connectivity to be affected not only by distance but also line of sight obstructions [29], [30]. Meireles *et al.* [6] observed an 80% range reduction under such conditions. Mangel *et al.* [2] reported similar results.

The impact of line of sight on routing has not yet been well explored. TVR [31] proposes the use of tall vehicles such as trucks, since they can communicate over the roof of shorter vehicles. But it ignores topography, the major source of line of sight obstructions.

The concept of look-ahead has been used in different contexts. Terminodes [32] and GeO-LANMAR [33] use GPSR-like geographic forwarding combined with a 2-hop look-ahead. However, they use a simple distance-to-destination heuristic.

Using multiple forwarders for reliability was first introduced in CBF [34]. DAZL [9], which we built upon in LASP, refined the concept by employing density-aware forwarder coordination. An alternative would have been to use a single forwarder coupled with Request-To-Send (RTS)/Clear-To-Send (CTS), as in ALBA-R [35]. However, this would have introduced delay and provided less protection from intermittent losses.

Vehicular network protocol evaluation has previously been done mostly through simulation. Experimentation is rare and when used, small in scale. In 1999, Johnson *et al.* [20] tested DSR with 5 vehicles. In 2008 Jerbi and Senouci [29] performed tests with 6 vehicles and fixed routing paths. In 2013 Boban *et al.* [31] ran multi-hop experiments using 802.11p hardware but only 4 vehicles.

VII. Conclusions

We showed how spatial connectivity can help overcome the lack of global topology information in vehicular multi-hop. Relative to other heuristics, such as distance and road maps, absolute location is a better predictor of connectivity, as we demonstrated with data from the HarborNet testbed.

LASP uses spatial connectivity to pick forwarding zones for cooperative relaying, increasing resilience to losses relative to single-forwarder schemes. Experiments showed an increase in end-to-end PDR of 30% relative to the benchmark GPSR.

Our work highlights the impact of spatial location on connectivity and its implications for routing and forwarding.

References

[1] I. C. Society, "IEEE Standard 802.11p," Jun. 2010.

[2] T. Mangel, M. Michl, O. Klemp, and H. Hartenstein, "Real-World Measurements of Non-Line-Of-Sight Reception Quality for 5.9GHz IEEE 802.11p at Intersections," in *Proc. of the 3rd Int. Conference on Communication Technologies for Vehicles*. Springer-Verlag, 2011.

[3] F. Bai, D. D. Stancil, and H. Krishnan, "Toward Understanding Characteristics of Dedicated Short Range Communications (DSRC) from a Perspective of Vehicular Network Engineers," in *Proc. of the 16th Conference on Mobile Computing and Networking (MobiCom)*. ACM, 2010.

[4] B. Karp and H. T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks," in *Proc. of the 6th Conference on Mobile Computing and Networking (MobiCom)*. ACM, 2000.

[5] J. LeBrun, C.-N. Chuah, D. Ghosal, and M. Zhang, "Knowledge-Based Opportunistic Forwarding in Vehicular Wireless Ad Hoc Networks," in *2005 Spring Vehicular Technology Conference (VTC)*, 2005.

[6] R. Meireles, M. Boban, P. Steenkiste, O. Tonguz, and J. Barros, "Experimental Study on the Impact of Vehicular Obstructions in VANETs," in *Proc. of the 2010 Vehicular Networking Conference (VNC)*. IEEE, 2010.

[7] J. Tian, L. Han, and K. Rothermel, "Spatially Aware Packet Routing for Mobile Ad Hoc Inter-vehicle Radio Networks," in *Proc. of the 2003 Conference on Intelligent Transportation Systems (ITS)*. IEEE, 2003.

[8] Q. Yang, A. Lim, S. Li, J. Fang, and P. Agrawal, "ACAR: Adaptive Connectivity Aware Routing for Vehicular Ad Hoc Networks in City Scenarios," *Mobile Networks and Applications*, vol. 15, 2010.

[9] R. Meireles, P. Steenkiste, and J. Barros, "DAZL: Density-Aware Zone-based Packet Forwarding in Vehicular Networks," in *Proc. of the 2012 Vehicular Networking Conference (VNC), IEEE*, 2012.

[10] D. Couto, D. Aguayo, J. Bicket, and R. Morris, "A High-Throughput Path Metric for Multi-Hop Wireless Routing," *Wireless Networks*, vol. 11, no. 4, 2005.

[11] "CRAWDAD - A Community Resource for Archiving Wireless Data At Dartmouth." [Online]. Available: http://www.crawdad.org

[12] C. Ameixieira, A. Cardote, F. Neves, R. Meireles *et al.*, "Harbornet: a real-world testbed for vehicular networks," *Communications Magazine, IEEE*, vol. 52, no. 9, 2014.

[13] C. Ameixieira, J. Matos, R. Moreira, A. Cardote, A. Oliveira, and S. Sargento, "An IEEE 802.11p/WAVE Implementation with Synchronous Channel Switching for Seamless Dual-channel Access (Poster)," in *Proc. of the 2011 Vehicular Networking Conference (VNC)*. IEEE, 2011.

[14] ETSI, "TS 102 637-2: Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Co-operative Awareness Basic Service," 2011.

[15] J. Li, J. Jannotti, D. De Couto, D. Karger, and R. Morris, "A Scalable Location Service for Geographic Ad Hoc Routing," in *Proc. of the 16th Conference on Mobile Computing and Networking (MobiCom)*. ACM, 2000.

[16] T. Clausen and P. Jacquet, "RFC 3626 - Optimized Link State Routing Protocol (OLSR)," 2003.

[17] D. Johnson, N. Ntlatlapa, and C. Aichele, "A simple pragmatic approach to mesh routing using BATMAN," in *2nd Int. Symposium on Wireless Communications and Information Technology in Developing Countries (WCTID)*. IFIP, 2008.

[18] J. Chroboczek, "RFC 6126 - The Babel Routing Protocol," 2011.

[19] M. Abolhasan, B. Hagelstein, and J. C.-P. Wang, "Real-world Performance of Current Proactive Multi-hop Mesh Protocols," in *Proc. of the 15th Asia-Pacific Conference on Communications (APCC 2009)*, 2009.

[20] D. B. Johnson, D. A. Maltz, and J. Broch, *DSR: The Dynamic Source Routing Protocol for Multi-Hop Wireless Ad Hoc Networks in Ad Hoc Networking*. Addison-Wesley, 2001, ch. 5.

[21] C. E. Perkins, "Ad-hoc On-Demand Distance Vector Routing," in *Proc. of the 2nd Workshop on Mobile Computing Systems and Applications (WMCSA)*. IEEE, 1999.

[22] V. Park and M. S. Corson, "Temporally-Ordered Routing Algorithm (TORA) Version 1 Functional Specification," Internet-Draft, Tech. Rep., 2001.

[23] S. Jaap, M. Bechler, and L. Wolf, "Evaluation of Routing Protocols for Vehicular Ad Hoc Networks in City Traffic Scenarios," in *Proc. of the 11th EUNICE Open European Summer School on Networked Applications*, 2005.

[24] C. Lochert, M. Mauve, H. Füßler, and H. Hartenstein, "Geographic Routing in City Scenarios," *ACM SIGMOBILE Mobile Computing and Communications Review*, 2005.

[25] R. Frank, E. Giordano, P. Cataldi, and M. Gerla, "TrafRoute: A different approach to routing in vehicular networks," in *Wireless and Mobile Computing, Networking and Communications (WiMob), 2010 IEEE 6th International Conference on*, 2010.

[26] C. Lochert, H. Hartenstein, J. Tian, H. Fuessler, D. Hermann, and M. Mauve, "A Routing Strategy for Vehicular Ad Hoc Networks in City Environments," in *Proc. of the 2003 Intelligent Vehicles Symposium (IV)*. IEEE, 2003.

[27] K. Lee, M. Le, J. Harri, and M. Gerla, "LOUVRE: Landmark Overlays for Urban Vehicular Routing Environments," in *Vehicular Technology Conference, 2008. VTC 2008-Fall. IEEE 68th*, 2008.

[28] M. Jerbi, S.-M. Senouci, R. Meraihi, and Y. Ghamri-Doudane, "An Improved Vehicular Ad Hoc Routing Protocol for City Environments," *Proc. of the 2007 Int. Conference on Communications (ICC)*, 2007.

[29] M. Jerbi and S. Senouci, "Characterizing Multi-Hop Communication in Vehicular Networks," in *Proc. of the 2008 Wireless Communications and Networking Conference (WCNC)*. IEEE, 2008.

[30] J. S. Otto, F. E. Bustamante, and R. A. Berry, "Down the Block and Around the Corner – The Impact of Radio Propagation on Inter-vehicle Wireless Communication," in *Proc. of the Int. Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2009.

[31] M. Boban, R. Meireles, J. Barros, P. Steenkiste, and O. Tonguz, "TVR - Tall Vehicle Relaying in Vehicular Networks," *IEEE Transactions on Mobile Computing*, vol. 13, no. 5, 2014.

[32] L. Blažević, S. Giordano, and J.-Y. Le Boudec, "Self Organized Terminode Routing," *Cluster Computing*, 2002.

[33] F. De Rango, M. Gerla, and S. Marano, "A scalable routing scheme with group motion support in large and dense wireless ad hoc networks," *Computers and Electrical Engineering*, vol. 32, no. 1-3, 2006.

[34] H. Füßler, J. Widmer, M. Käsemann, M. Mauve, and H. Hartenstein, "Contention-based forwarding for mobile ad hoc networks," *Ad Hoc Networks*, vol. 1, no. 4, 2003.

[35] C. Petrioli, M. Nati, P. Casari, M. Zorzi, and S. Basagni, "ALBA-R: Load-Balancing Geographic Routing Around Connectivity Holes in Wireless Sensor Networks," *Parallel and Distributed Systems, IEEE Transactions on*, 2014.