

Geo-Opportunistic Routing for Vehicular Networks

Kevin C. Lee, *University of California, Los Angeles*

Uichin Lee, *Alcatel-Lucent Bell Labs*

Mario Gerla, *University of California, Los Angeles*

ABSTRACT

Road topology information has recently been used to assist geographic routing in urban vehicular environments to improve overall routing performance. However, the unreliable nature of wireless channels due to motion and obstructions still makes road topology assisted geographic routing challenging. In this article we begin by reviewing conventional road topology assisted geographic routing protocols, and investigate the robust routing protocols that address and help overcome the unreliable wireless channels. We then present topology-assisted geo-opportunistic routing that incorporates topology assisted geographic routing with opportunistic forwarding. That is, the routing protocol exploits the simultaneous packet receptions induced by the broadcast nature of the wireless medium and performs opportunistic forwarding via a subset of neighbors that have received the packet correctly. Our simulation results confirm TO-GO's superior robustness to channel errors and collisions compared to conventional topology-assisted geographic routing protocols.

INTRODUCTION

The sharp increase of vehicles in recent years has made driving more challenging and dangerous. For safe driving, leading car manufacturers have been jointly working with national government agencies to develop solutions to help drivers anticipate hazardous events and avoid traffic jams. One of the recent outcomes is a novel wireless architecture called wireless access for the vehicular environment (WAVE) that provides short-range intervehicular communications to enable fast dissemination of emergency related messages.

While the major objective has clearly been to improve the overall safety of vehicular traffic, industry laboratories and academia have been exploring novel vehicular applications such as traffic management and onboard entertainment. Emerging vehicular applications often necessitate wide-area coverage using multihop routing protocols, which is a major departure from safety applications that require only local coverage.

However, efficient multihop routing in a vehicular ad hoc network (VANET) is challenging for the following reasons. First, it is a highly distributed self-organizing network formed by moving vehicles that are characterized by very high mobility yet constrained by roads. Second, its size can scale up to hundreds of thousands of nodes. Third, nodes could suffer from severe wireless channel fading due to motion and obstructions in urban environments (e.g., building, trees, and vehicles). Finally, the vehicle density changes over time (rush hours), and the distribution of vehicles is non-uniform due to various road widths and skewed popularity of roads. Under this circumstance, most ad hoc routing protocols that discover and maintain end-to-end paths (e.g., Ad Hoc On Demand Vector [AODV], Dynamic Source Routing [DSR]) are less preferable due to high protocol overheads. Therefore, we cannot directly use those protocols to support such emerging vehicular applications.

One of the popular routing protocols in a VANET is geographic routing where the forwarding decision by a node is primarily made based on the position of a packet's destination. A packet is greedily forwarded to a neighboring node whose distance toward the packet's destination is closer than that of the current node (called the greedy mode). If there is no such a node (i.e., a packet has reached a local maximum where it has made the maximum progress toward the destination locally), the protocol then reverts to the recovery mode. Face routing (or perimeter routing) [1], a widely used stateless recovery strategy, planarizes a network graph such that its edges intersect only at their endpoints, and then forwards a packet along one or possibly a sequence of adjacent faces (or edges), thus providing progress toward the destination node.

Geographic routing is preferable in a VANET for the following reasons. First, geographic routing is stateless; it neither exchanges link state information nor maintains established routes as in conventional mobile ad hoc routing protocols. The exchange and route maintenance are very costly in highly mobile vehicular environments. Second, it is becoming easier to support geographic routing as GPS-based navigation systems

are getting cheaper and becoming a common add-on.

In urban vehicular environments, however, it is known that conventional geographic routing protocols such as Geographic Perimeter Stateless Routing (GPSR) [1] may not work well because vehicles have constrained mobility patterns due to the road structure and tend to show heterogeneous density distribution — a mixture of heavily populated and sparse road segments. In particular, face routing could be very costly, because a packet has to travel along a sequence of adjacent faces where each step could make only small progress (as opposed to a nominal radio range) toward the destination when vehicle density is relatively high. Given that road topology is typically planar, Lochert *et al.* incorporated the road topology into geographic routing and proposed Geographic Perimeter Coordinator Routing (GPCR) [2], where packets can *always* be forwarded along the road segments greedily until they reach nodes at junctions/intersections (called junction nodes). Junction nodes then decide to which road a packet must be forwarded based on the packet's current mode.

However, existing topology-assisted geographic routing protocols do not consider error-prone urban wireless channels due to multipath fading and shadowing where the assumption of unit disc propagation does not hold. Geographic routing attempts to greedily forward a packet to the furthest neighboring node that is closest to the packet's destination. The problem is that the further the distance, the higher the attenuation, and the greater the likelihood of packet loss. Therefore, we want to improve the performance of topology-assisted geographic routing protocols by effectively handling unreliable wireless channels.

In this article we first review existing geographic routing protocols such as Geographic Random Forwarding (GeRaF) [3] and Contention Based Forwarding (CBF) [4, 5] that address the unreliable channels using opportunistic forwarding where a sender takes advantage of random packet receptions in its neighboring nodes due to the error-prone wireless channel, and performs opportunistic forwarding via a subset of the neighbors (called a forwarding set) that have received the packet correctly. We find that these protocols often fail to exploit the full benefit of opportunistic forwarding, because they do not take the road topology into account when choosing a forwarding set. To remedy this problem, we then propose TOPOLOGY-assisted Geo-Opportunistic routing (TO-GO), that incorporates road topology information into the forwarding set selection to better exploit the benefit of opportunistic forwarding. Unlike previous approaches [3–5], TO-GO does not rely on the unit-disk propagation assumption, but uses the actual *intersection* of neighbors made available by two-hop neighbor information. Simulation results confirm that TO-GO can effectively avoid poor wireless links and is thus robust to channel impairments. TO-GO can achieve up to a 98 percent packet delivery ratio, which is 40 percent higher than conventional protocols in the error-prone wireless channel scenario under consideration.

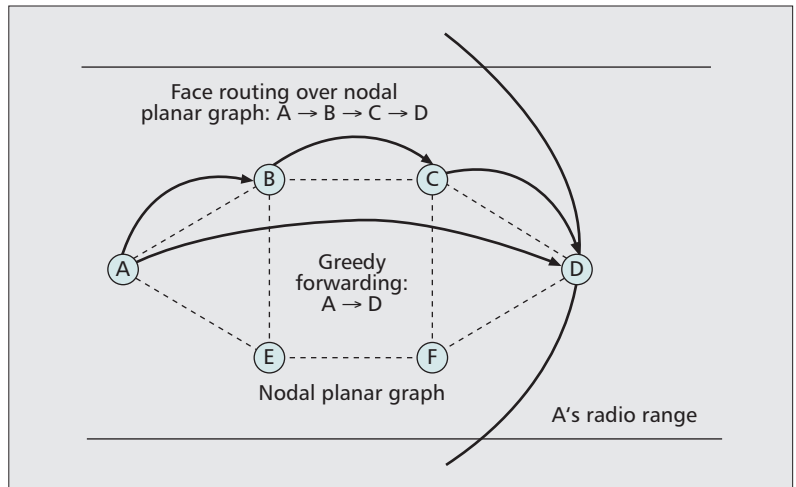


Figure 1. Baby step problem in the recovery mode.

BACKGROUND

In this section we review topology-assisted geographic routing and opportunistic routing protocols, and identify limitations of existing opportunistic routing techniques when used in urban vehicular environments. Readers can find a survey of VANET routing protocols in [6].

TOPOLOGY-ASSISTED GEOGRAPHIC ROUTING

Lochert *et al.* [2] found that a planarized connectivity graph for vehicles along a street could lead to a graph where a vehicle no longer sends packets to the neighboring node with the largest forward progress, which is called a *baby step* problem in the recovery mode. Recall that planarization is to transfer a local connectivity graph into a planar graph by eliminating redundant edges such that its edges intersect only at their endpoints. This problem is illustrated in Fig. 1 where we can greedily forward a packet along a road segment in a single hop (from A to D), but the recovery mode that uses face routing over the nodal planar graph requires three hops. For this reason, instead of relying on planarization of nodes, Lochert *et al.* [2] proposed GPCR, which takes advantage of the fact that an urban map naturally forms a planar graph where a junction (or intersection) is a node, and a road segment is an edge in the graph. In GPCR junctions are the only places a routing decision takes place. Packets are always greedily forwarded along the street from one junction to the other (even in the recovery mode), which solves the baby step problem. Moreover, GPSR-like face routing (using a right hand rule) is performed over the road topology graph in the recovery mode.

GpsrJ+ [7] enhances GPCR by noting that nodes do *not* necessarily need to stop at each junction node (Fig. 2). The key idea is that not every packet must be stored and forwarded by a junction node; in other words, the junction is not a necessary stop. More precisely, a packet must be stored and forwarded by a junction node only when it needs to make a left or right turn at that junction. This greatly reduces the dependence on junction nodes. In GpsrJ+ a

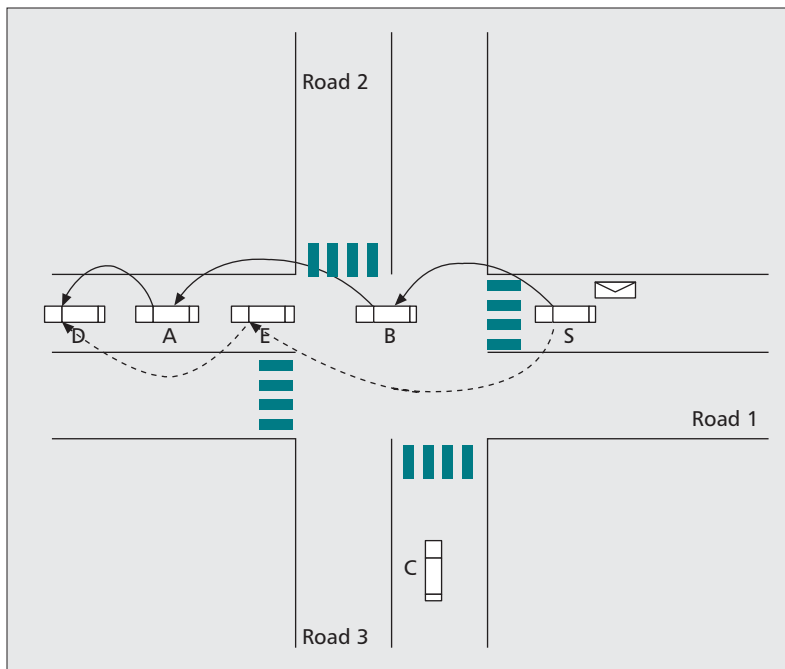


Figure 2. Dashed arrows are GpsrJ+ and solid arrows are GPCR.

forwarding node uses two-hop neighbor information to detect advantageous junction turns and also to better estimate a routing path. Upon learning that there are no advantageous turns, GpsrJ+ simply bypasses the junction. This two-hop prediction reduces hop counts, increases the packet delivery ratio, and obviates the need to distinguish junction nodes from ordinary nodes.

Topology-assisted geographic routing protocols can be further enhanced by checking connectivity of road segments to avoid forwarding packets along disconnected road segments [8]. Note that besides stateless geographic routings where a forwarding decision is made in each junction (e.g., GPCR and GpsrJ+), it is also possible to compute a shortest path using an urban map and then embed a set of junctions in the packet to perform source-based routing, as in Geographic Source Routing (GSR) [9]. This approach may fail to provide end-to-end connectivity due to disconnected road segments; thus, we need to proactively collect connectivity information of road segments to prune disconnected road segments as in Landmark Overlays for Urban Vehicular Routing Environments (LOUVRE) [10]. In this article we focus on stateless approaches such as GPCR and GpsrJ+ that do not require network-wide information exchanges, and our goal is to improve their performance by taking error-prone wireless channels into account. Note that the protocol proposed in this article can also exploit the aforementioned techniques to further enhance its performance.

OPPORTUNISTIC ROUTING

Geographic routing tries to greedily forward a packet to the furthest neighboring node that is closest to the packet's destination, but the problem is that the further the distance, the higher

the attenuation, and the greater the likelihood of packet loss. This fact brought forth the concept of opportunistic routing [11, 12], where a sender takes advantage of random packet receptions in its neighboring nodes due to the error-prone wireless channel and opportunistic forwarding by a subset of the neighbors that have received the packet correctly. The key challenge is to select a subset of neighbors that can make the best progress toward the destination, yet without the hidden terminal problem. When a higher-priority node transmits a packet, other low-priority nodes should be able to suppress forwarding to prevent redundant packet transmissions and collisions. Most opportunistic routing protocols (also called anypath routing) such as ExOR [11] and Least Cost Opportunistic Routing (LCOR) [12] that do not use geographic information, require global topology and link quality information (like link state routing) to find a set of forwarding groups toward the destination; thus, they are more suitable for static wireless mesh or sensor networks.

In practice, geographic routing can also benefit from opportunistic forwarding as in GeRaF [3] and CBF [4, 5], although it is not optimal due to the lack of global knowledge. For forwarding set selection, researchers typically used a geometric shape faced toward the destination (e.g., triangle or lens shape) [4, 13] where nodes can hear one another. For instance, Fig. 3a shows a lens shape forwarding set that contains nodes *A* and *B*. Nodes in this forwarding region contend for packet forwarding based on a distance-based timer (i.e., the further the distance from the sender, the shorter the packet expiration timer) [3, 5, 13]. In the figure node *A* has higher priority than node *B* because node *A* is closer to the destination. Lower-priority nodes will cancel their impending transmissions when they hear a higher priority transmission.

In urban vehicular environments, however, choosing a direction toward the destination often yields a suboptimal set in terms of its size and progress because the destination may not lie on the same road segment as the current forwarding node. For example, Fig. 3 shows that the forwarding region toward the destination contains two nodes, whereas the forwarding region toward the furthest node on the current road has four nodes in the forwarding region; the latter is more robust than the former. That is exactly what TO-GO does: TO-GO focuses on a more effective forwarding set between the sender and the *target node* that is the furthest node on the current road segment. By incorporating the road topology information, it can better exploit opportunistic forwarding.

TO-GO DESIGN

In this section we present the Next-hop Prediction Algorithm (NPA), which determines a packet's target node; the Forwarding Set Selection (FSS) algorithm, which finds a set of candidate forwarding nodes; and the priority scheduling method, which suppresses redundant packet transmissions based on a distance-based timer.

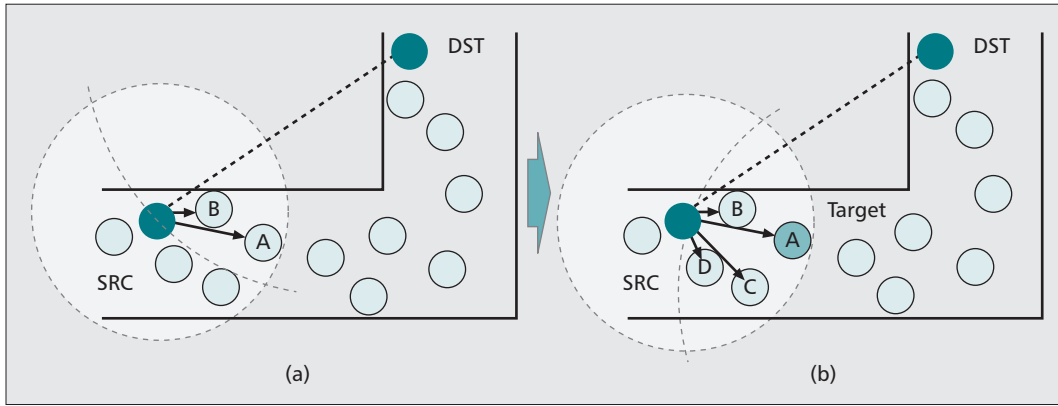


Figure 3. The lens shaped area is the forwarding region established between source and destination nodes in existing schemes, and between the source and the furthest node on the current road segment (called target node) in TO-GO: a) existing schemes; b) TO-GO.

NPA

As in GpsrJ+ [7], the conventional hello beacon of a node E is augmented to include the *furthest neighbors (and their locations)* in each direction on the urban map (typically, only two neighbors except for intersection nodes). This is required to support junction forwarding prediction in both greedy and recovery modes. The beacon also contains the Bloom filter representing a set of E 's neighbors, and the size of this set. Since a Bloom filter is a space-efficient membership checking data structure, it enables the construction of a forwarding set while keeping the broadcast overhead at a minimum. For instance, a filter size of 150 bits (19B) can represent 15 items at a false positive rate smaller than 1 percent. Upon receiving a beacon, a node would have a neighbor list that contains its neighbor, every neighbor's furthest neighbors, and a Bloom filter of their neighbors and its size.

TO-GO uses this enhanced beacon to predict the target node, which is either the furthest node or the *junction node*. Here, a junction node is a node that is located at the junction and can forward packets in any direction. In the greedy mode, the best forwarding node is the furthest node when its neighboring junction node's neighbor closer to the destination lies on the same road segment as the furthest node (i.e., a packet will not make left/right turns at the junction). Otherwise, the best forwarding node is the junction node. The two-hop information in enhanced beacons enables TO-GO to make an advanced decision on whether to bypass the junction node.

FSS

After finding the target node, the current forwarding node C must determine which nodes will be in a forwarding set. In principle, the forwarding set should be selected such that nodes in the set can hear each other to prevent hidden terminal collisions. A brute force algorithm to find a forwarding set in which nodes hear one another is analogous to finding a maximal *clique* in which every node has a connection to every other node. Such a problem is NP-complete. We propose a simplified scheme to obtain an approx-




Figure 4. Forwarding set selection approximation: a) shaded region contains neighbors of C that can hear both C and T ; b) shaded region contains neighbors of C that can hear both M and T , and can also hear each other.

imate forwarding set by first eliminating C 's neighbors that cannot hear the target node. Out of the neighbors that remain, we then pick the neighbor that has the largest number of neighbors. Denote this neighbor as M . For each neighbor N of the current forwarding node, test its membership in M 's Bloom filter. If N is in the Bloom filter and N 's Bloom filter contains M , test N 's membership using the Bloom filters of existing elements in the forwarding set. If N is in the Bloom filters of all these elements, add N to the set. Continue adding such N until all the neighbors of C have been checked. The algorithm takes $O(n^2)$ where n is the number of C 's neighbors.

The intuition behind the approximate algorithm is that the neighbor M that has the most neighbors is in the most dense area. Despite irregular and different radio ranges, nodes selected from that region are more likely to have one another as neighbors. The forwarding set produced thus should be close to a maximal set that provides the largest number of nodes as potential next hop forwarders. Note that the resulting forwarding set represented in a Bloom filter is embedded into the data packet for distributed priority scheduling.

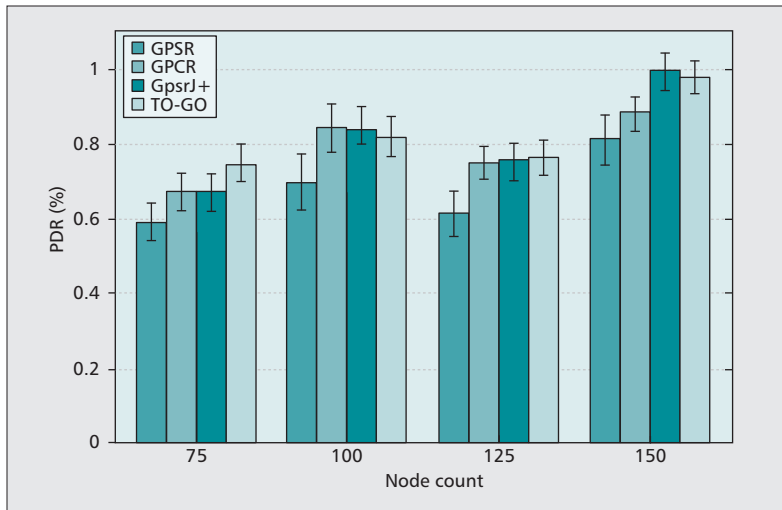


Figure 5. PDR vs. node count.

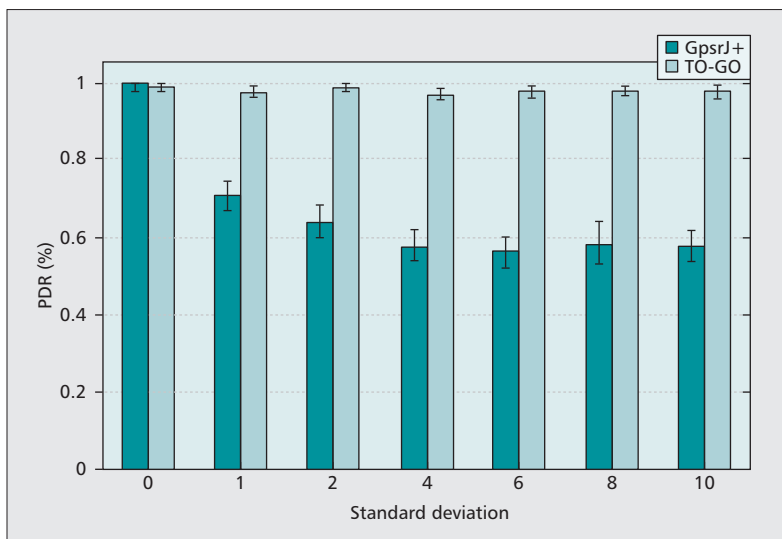


Figure 6. PDR vs. different degrees of shadow fading (σ).

The shaded region in Fig. 4a contains a set of C 's neighbors (denoted \mathcal{S}) that can hear both current node C and target T . From the set \mathcal{S} , node C then picks the neighbor M that has the largest number of neighbors. In Fig. 4b the resulting shaded region represents a subset of \mathcal{S} that contains neighbors of C that can hear both M and T , and can also hear each other.

PRIORITY SCHEDULING

Having found the forwarding set, we want a node closer to the target node to become the next forwarder, because the shorter the distance between the receiving node and the target node, the greater the progress, and therefore the shorter the timer. Unlike the timer formula in [5] where the authors assume that there is a fixed radio range R , and this range is used for normalization, we use this distance between the sending node and the target node for normalization, by noting the fact that radio range differs from vehicle to vehicle in reality. Hence, we set the timer T as follows:

$$T = C \times \frac{\text{dist}(\text{receiving node}, \text{target node})}{\text{dist}(\text{sending node}, \text{target node})},$$

where C is the maximum forwarding delay that varies with the transmission rate and processing time.

PERFORMANCE EVALUATION

SIMULATION SETUP

The evaluation was conducted on a QualNet simulator 3.95 with IEEE 802.11b DCF as the medium access control (MAC) with a transmission rate of 2 Mb/s and transmission range of 250 m. We assume that nodes on different roads cannot talk to each other because of obstacles (trees, buildings, etc.). The mobility traces are generated using VanetMobiSim [14] that produces realistic urban mobility traces using macro- and micro-mobility features of the vehicular environment. Intersections are controlled by stop signs, and road segments contain speed limitations. All roads have a single lane in each direction and a speed limit of 15 m/s (54 km/h). We use a grid topology in an urban area of size 1800 m \times 300 m where the side length of a single grid is 300 m.

We use a simple log-normal shadow fading model where we can vary the degree of shadow fading using a single parameter [15],

$$PL(d)[dB] = PL(d_0) + 10n \log\left(\frac{d}{d_0}\right) + X_\sigma,$$

where n is the path loss exponent which indicates the rate at which the path loss increases with distance, d_0 is the close-in reference distance determined from measurements close to the transmitter, d is the transmitter-receiver distance, and X_σ is a zero-mean Gaussian distributed random variable with standard deviation σ to account for random and distributed log-normal shadow fading. We use $n = 2$ for the path loss exponent and $d_0 = 0.025$ for the reference distance, which is a default setting in the QualNet simulator. We vary the standard deviation σ of the zero-mean Gaussian distributed random variable X to simulate different magnitudes of shadowing effects and thereby different probabilities of packet loss.

We compare the performance of GPSR, GPCR, GpsrJ+, and TO-GO. GpsrJ+ is enhanced by enabling junction prediction in both greedy and recovery modes. The number of nodes in the network ranges from 75 to 150, with 25-node increments. We configure the constant in the timer equation as $C = 0.1$. This value maximizes throughput under channel fading conditions when the number of nodes is 150. For each node trace, we run 20 simulations and report the average value with 95 percent confidence interval. The duration of each run is 180 s. In each simulation we select 10 random source-destination pairs for every 10 s where each pair transfers a stream of 1460-byte packets at a constant rate (1 packet/s).

SIMULATION RESULTS

Figure 5 shows the packet delivery rate (PDR)

of GPSR, GPCR, GpsrJ+, and TO-GO with respect to node density in an error-free wireless channel. We set the σ value as zero to model 0 percent dropping probability. A superficial observation indicates that while GPCR, GpsrJ+, and TO-GO are almost similar to one another in PDR, GPSR always lags behind. The performance hit is due to making baby steps in the recovery mode; that is, due to nodal planarization, each hop makes only a small progress toward the destination. As node density increases, the frequency of falling into the recovery mode decreases; thus, GPSR's PDR gradually increases to about 82 percent. Moreover, when there are more nodes in the network, TO-GO gains because there are more opportunities for packets to be delivered to nodes closer to the target.

We now introduce errors into the channel by varying the standard deviation σ of the Gaussian distributed random variable X ranging from 0 to 10 (in a 150-node scenario). Recall that the larger the deviation, the greater the channel error. Here, we only compare the performance of GpsrJ+ and CBF because GpsrJ+ is an enhancement of GPCR, and GpsrJ+ outperforms GPCR. We plot the average PDR and latency in Figs. 6 and 7, respectively. When the error increases, TO-GO maintains the PDR above 96 percent, but GpsrJ+ keeps on dropping. At $\sigma = 10$, TO-GO's PDR remains at 98 percent, while GpsrJ+'s PDR drops to 58 percent. The relatively higher latency of TO-GO from $\sigma = 1$ to $\sigma = 10$ is due to averaging these values, which are not accounted for in GpsrJ+ because packets are dropped. In general, those protocols with high PDR tend to show high hop count and longer latency, because a packet has to travel more hops in order to discover a path to the destination. Note that in TO-GO, additional delay can be incurred for retransmission due to packet collision as it always broadcasts packets, and priority scheduling in each hop also contributes to the delay.

CONCLUSION

In this article we review road topology assisted geographic routing that uses road topology information to enhance geographic routing, and illustrate that the unreliable wireless channels in urban environments make this goal challenging. For this reason, we investigate existing geographic opportunistic routing protocols that address the unreliable channels by opportunistic forwarding. We find that these protocols fail to exploit the full benefit of opportunistic forwarding, because they do not take the road topology into account when choosing a forwarding set. To overcome this limitation, we propose TO-GO, a geographic opportunistic routing protocol that exploits road topology information in opportunistic packet reception to improve packet delivery. As the goal in vehicular routing is to maximize the expected packet advancement to the destination, TO-GO defines a candidate forwarding set between the current sender and the target node. This set is selected using a simple junction prediction algorithm with topology information and enhanced beaconing. The for-

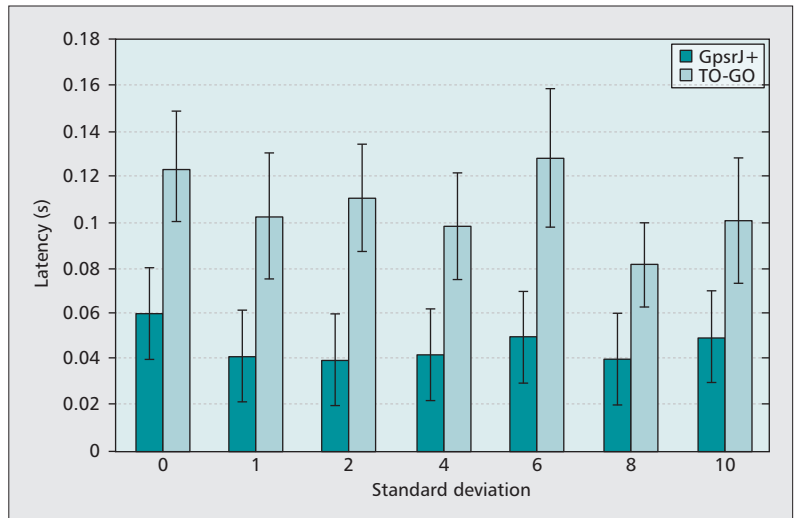


Figure 7. Latency vs. different degrees of shadow fading (σ).

warding set is then adjusted to reduce packet duplication and collision. We have validated the robustness of TO-GO under wireless channel errors via extensive simulations.

REFERENCES

- [1] B. Karp and H. T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks," *MobiCom '00*, Boston, MA, Aug. 2000.
- [2] C. Lochert et al., "Geographic Routing in City Scenarios," *ACM SIGMOBILE Mobile Comp. Commun. Rev.*, vol. 9, no. 1, 2005, pp. 69–72.
- [3] M. Zorzi and R. Rao, "Geographic Random Forwarding (GeRaF) for Ad Hoc and Sensor Networks: Multihop Performance," *IEEE Trans. Mobile Comp.*, vol. 2, no. 4, Oct.–Dec. 2003, pp. 337–48.
- [4] H. Fühler et al., "Contention-Based Forwarding for Street Scenarios," *WIT '04*, Hamburg, Germany, Mar. 2004.
- [5] H. Fühler et al., "Contention-Based Forwarding for Mobile Ad Hoc Networks," *Ad Hoc Net.*, vol. 1, no. 4, 2003, pp. 351–69.
- [6] K. C. Lee, U. Lee, and M. Gerla, "Survey of Routing Protocols in Vehicular Ad Hoc Networks," chapter in *Advances in Vehicular Ad-Hoc Networks: Developments and Challenges*, M. Watfa, Ed., IGI Global, Oct. 2009.
- [7] K. C. Lee et al., "Enhanced Perimeter Routing for Geographic Forwarding Protocols in Urban Vehicular Scenarios," *AutoNet '07*, Washington, DC, Nov. 2007.
- [8] D. Forderer, *Street-Topology Based Routing*, Master's thesis, Univ. of Mannheim, 2005.
- [9] C. Lochert et al., "A Routing Strategy for Vehicular Ad Hoc Networks in City Environments," *IEEE IVS '03*, June 2003.
- [10] K. C. Lee et al., "LOUVRE: Landmark Overlays for Urban Vehicular Routing Environments," *IEEE WiVeC*, 2008.
- [11] S. Biswas and R. Morris, "ExOR: Opportunistic Multi-Hop Routing for Wireless Networks," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 35, no. 4, 2005, pp. 133–44.
- [12] H. Dubois-Ferriere, M. Grossglauser, and M. Vetterli, "Least-Cost Opportunistic Routing," *Allerton '07*, Urbana, IL, Sept. 2007.
- [13] R. C. Shah et al., "When Does Opportunistic Routing Make Sense?" *PERCOMW '05*, Kauai, HI, Mar. 2005.
- [14] J. Häerri et al., "VanetMobiSim: Generating Realistic Mobility Patterns for VANETS," *VANET '06*, Los Angeles, CA, Sept. 2006.
- [15] T. Rappaport, *Wireless Communications: Principles and Practice*, Prentice Hall PTR, 2001.

BIOGRAPHIES

KEVIN C. LEE (kclee@cs.ucla.edu) is currently a Ph.D. student in the Computer Science Department at the University of California, Los Angeles (UCLA). He received his B.S. degree

in computer science engineering and his B.A. degree in mathematics from the University of Pennsylvania in 2002. He also received his M.S. degree in computer science from Carnegie Mellon University in 2004. He has published several papers in the fields of vehicular ad hoc networks ranging from developing and running a P2P application on a real vehicular testbed, designing an optimized geographic routing protocol in urban scenarios, developing a lightweight loop-free geographic routing protocol, and proposing link-state routing based on density. His current research interests include vehicular ad hoc network routing, theoretical analysis of computer networks, worldwide interoperability for microwave access networks, QoS, and mobile applications for vehicular ad hoc networks.

UICHIN LEE [M'09] (uichin.lee@bell-labs.com) received his B.S. in computer engineering from Chonbuk National University in 2001, his M.S. degree in computer science from Korea Advanced Institute of Science and Technology (KAIST) in 2003, and his Ph.D. degree in computer science from UCLA in 2008. He is currently a member of technical staff in Bell Labs, Alcatel-Lucent. His research interests include distributed systems, mobile wireless networking systems, and performance modeling/evaluation.

MARIO GERLA [F'02] (gerla@cs.ucla.edu) is a professor in the Computer Science Department at UCLA. He holds an Engineering degree from Politecnico di Milano, Italy and a Ph.D. degree from UCLA. At UCLA, he was part of the team that developed the early ARPANET protocols under the guidance of Prof. Leonard Kleinrock. At Network Analysis Corporation, New York, from 1973 to 1976, he helped transfer ARPANET technology to government and commercial networks. He joined the UCLA faculty in 1976. At UCLA he has designed and implemented network protocols including ad hoc wireless clustering, multicast (ODMRP and CodeCast), and Internet transport (TCP Westwood). He led the \$12 million, six-year ONR MINUTEMAN project, designing the next-generation scalable airborne Internet for tactical and homeland defense scenarios. He is now leading two advanced wireless network projects under Army and IBM funding. His team is developing a vehicular testbed for safe navigation, urban sensing, and intelligent transport. A parallel research activity explores personal communications for cooperative, networked medical monitoring (see <http://www.cs.ucla.edu/NRL> for recent publications).