Front. Comput. Sci., 2012, 6(6): 713–724
DOI 10.1007/s11704-012-1154-7

RESEARCH ARTICLE

Publish/subscribe based information dissemination over VANET utilizing DHT

Tulika PANDEY (✉)[1], Deepak GARG[2], Manoj Madhav GORE[3]

1 Department of Computer Science and IT, Sam Higginbotom Institute of Agriculture, Technology and Sciences, Allahabad 211007, India

2 Department of Computer Science and Engineering, Thapar University, Patiala 147004, India

3 Department of Computer Science and Engineering, Motilal Nehru National Institute of Technology, Allahabad 211004, India

Abstract A vehicular ad-hoc network (VANET) can be visualized as a network of moving vehicles communicating in an asynchronous and autonomous fashion. Efficient and scalable information dissemination in VANET applications is a major challenge due to the movement of vehicles which causes unpredictable changes in network topology. The publish/subscribe communication paradigm provides decoupling in time, space, and synchronization between communicating entities, and presents itself as an elegant solution for information dissemination for VANET like environments. In this paper, we propose our approach for information dissemination which utilizes publish/subscribe and distributed hash table (DHT) based overlay networks. In our approach, we assume a hybrid VANET consisting of stationary info-stations and moving vehicles. These info-stations are installed at every major intersection of the city and vehicles can take the role of publisher, subscriber, or broker depending upon the context. The info-stations form a DHT based broker overlay among themselves and act as rendezvous points for related publications and subscriptions. Further, info-stations also assist in locating vehicles that have subscribed to information items. We consider different possible deployments of this hybrid VANET with respect to the number of info-stations and their physical connectivity with each other. We perform simulations to assess the performance of our approach in these different deployment scenarios and discuss their applicability in urban and semi-urban areas.

Keywords vehicular ad-hoc network (VANET), peer-to-peer overlay network, publish/subscribe communication paradigm, distributed hash table (DHT)

1 Introduction

A vehicular ad-hoc network (VANET) can be defined as a distributed, self-organizing communication network of moving vehicles and stationary roadside info-stations. VANET can be utilized to disseminate important information like dangerous locations (construction site, post crash obstacles, road conditions) to vehicles and can provide comfort of driving, better managed traffic, and other benefits. Vehicles can get information about the current and expected traffic conditions (expected delays, better routes, etc.). Information regarding the location of the next intersection and signal timing can also be disseminated so that vehicles can notify drivers of the optimal speed. If a vehicle travels at the optimal speed, the traffic signal is likely to be green and the driver will not have to slow down or stop the vehicle resulting in increased traffic flow and fuel economy for vehicles.

Cellular networks such as 2 GB, 3 GB, 4 GB can serve as alternatives for information dissemination among vehi-

Received August 22, 2011; accepted January 21, 2012

E-mail: tulika_mishra@rediffmail.com

cles. However, scarcity of bandwidth is being experienced more frequently as the number of subscribers grows rapidly. Taking this issue into account, the 802.11 working group of IEEE is standardizing 802.11p also known as dedicated short range communication (DSRC) [1] for vehicular communication. Major players in the automotive industry such as BMW, Toyota, and Nissan have already begun producing vehicles equipped with short range communication technologies. Further, experiments conducted in [2] suggest that a vehicle traveling past a roadside access point with 802.11 a/b/g wireless interface at 80 km/h is able to transfer up to 50 MB of data, which indicates that even if constant connectivity is not present and frequent changes in topology are possible, the data rate achieved is enough for common applications. The previous arguments all indicate that in near future short range communications between vehicles can be realized and network applications will be built over them.

Vehicles in a VANET move along roads with different speeds, stopping times, and directions. Developing applications over such dynamic and distributed environment is challenging due to the unpredictable and continuous changes to the underlying network topology. Moreover, it is not possible to establish any central administrative authority in such mobile, autonomous and ad-hoc environments.

The publish/subscribe communication paradigm [3] presents itself as the most elegant solution for building applications where the underlying interaction mechanisms are flexible, asynchronous and highly dynamic in nature. The main strength of this paradigm lies in decoupling in time, space, and flow among event producers (publishers) and event consumers (subscribers).

In this paper, we present our publish/subscribe information dissemination framework for VANET which utilizes a DHT based peer to peer overlay of stationary info-stations as rendezvous points for matching and notifying publications to interested subscribers. The initial idea of the framework was proposed earlier in [4]. Our framework has following salient features.

- It enables vehicles to publish, or subscribe to, relevant information and route it towards the nearest info-stations through other vehicles using multi-hop connections.
- The DHT based overlay enables info-stations to act as designated placeholders for subscriptions and publications having similar content based attributes.
- Info-stations also assist in locating vehicles. Using the DHT methodology, every info-station takes responsibility to store location information of subscriber vehicles still waiting for notifications.
- Notifications are routed towards only those vehicles that have previously subscribed. This is done by utilizing the overlay links of the info-station and vehicle to vehicle multi-hop connections.

To locate vehicles, we utilize message exchanges between vehicles passing through major intersections and info-stations installed there. Vehicles can inform the info-stations that which direction they are coming from and going to. This information, along with timestamp, is stored at a responsible info-station according to vehicle to info-station mapping provided by DHTs information storage policy. The last info-station passed and direction information is utilized to locate vehicles.

We consider two specific deployment scenarios. In the first scenario, we assume that info-stations are connected to the internet and DHTs virtual links between them utilizing the underlying IP-routing mechanism of the internet. In the second scenario, we consider that info-stations are not physically connected to each other and DHT among them utilizes multi-hop routes formed by vehicles moving across them as information carriers. We simulate both these approaches with different numbers of info-stations, speeds of vehicles, and densities of vehicles. The density of vehicles is dependent on the time of day in realistic scenarios. Simulated results allow us to argue about the applicability of the above deployment scenarios in urban or semi-urban areas.

Section 2 provides a technical description of the publish/subscribe notification service over DHT based structured peer to peer overlays. In Section 3 we provide our system model. Section 4 presents algorithms of our approach and possible deployment scenarios. Section 5 describes the simulation environment and results. In Section 6 we survey related work and Section 7 concludes the paper.

2 Publish/subscribe over DHT

Information dissemination in dynamic distributed networks such as VANETs is challenging due to the unpredictable and restricted mobility of vehicles which results in dynamic topology changes. The design challenges for information dissemination applications can be summarized as:

Timeliness Vehicles should be notified of obstacles or hazards, such as accidents, or traffic, with sufficient time for a new route to be chosen.

High mobility with dynamic topology should be able to tolerate the high mobility of vehicles and consequent rapid topology changes. The mobility is constrained, as vehicles usually move along controlled trajectories.

Sparse and uneven distribution of vehicles should be able to cope up with network partitions which are very common in VANETs where vehicles are often distributed unevenly. The distribution of vehicles is dependent upon the deployment area (urban or rural) and time of the day (before office hours, after office hours, night hours, etc.).

Minimal infrastructure should be able to perform effectively in scenarios where communication infrastructure is minimal. Should be able to utilize moving vehicles for disseminating information by using multi-hop communication links in ad-hoc manner.

Scalability should be able to scale so that more vehicles can be added anywhere and at any time without affecting the performance.

Considering these design challenges we require decoupling and asynchrony in interaction among communication partners. Furthermore, a distributed, decentralized and autonomous network substrate is needed to efficiently store, lookup and route the information to intended users.

2.1 Publish/subscribe communication paradigm

VANET presents itself as an unpredictable environment where many-to-many communication takes place and it is not clear in advance who needs what information and at which moment. The publish/subscribe communication paradigm is most suitable for such environments due to its inherent decoupled nature. This paradigm provides decoupling from time, flow, and space. Decoupling from time means that the event subscriber and publisher need not be up at the same time whereas decoupling from flow implies that sending and receiving does not block participants. Further, space decoupling enables subscribers to move from one location to another without informing the publisher.

This communication paradigm consists of three basic components: publisher, subscriber, and broker. Publishers are the sources of information whereas subscribers are known as consumers. The system may consist of multiple publishers, subscribers and brokers. A user shows interest in receiving certain types of events by generating subscriptions in the system. A broker acts as a mediator between the publisher and the subscriber. It is used to store and match the publications and subscriptions and to notify matched publications to interested subscribers.

Publish/subscribe system can be broadly classified into either topic based or content based [3] depending on the manner in which a publication and subscription can be described. For our application we utilize a mix of both topic based and content based methods. We statically define a list of attributes beforehand to be utilized for describing attribute-value pair based publications and subscriptions.

2.2 Structured P2P overlay based on DHT

Structured peer to peer overlays networks [5–7] create a virtual structured topology above the basic transport layer. These overlays are based on DHTs and offer various features like robust wide-area routing architectures, efficient search of data items, storage, and self-organization.

In our approach, we utilize chord [7], a simple structured peer to peer overlay based on DHT. Chord utilizes SHA-1, a consistent hashing mechanism to assign m-bit identifier to each node and content object. This m-bit identifier space is represented as ring which wraps up at 2^{m-1}. In chord, each node maintains a link to its successor-node in the identifier ring. Lookup requests are passed around the ring via these successor links. This process terminates when a node is found which is responsible for the desired identifier getting looked up. In order to improve the routing performance, each node maintains the additional routing information in form of *finger-table*. This finger-table is a routing table of at most m entries (for m-bit identifier space). The ith entry in this finger table at any node n contains identifier of a node that succeeds n by at least 2^{i-1} on the identifier ring, where $1 \leqslant i \leqslant m$ and all operations are modulo 2^m.

We make use of the content-to-node mapping provided by chord. The content-to-node mapping maps the set of subscriptions and publications to the set of nodes. This mapping has to satisfy the intersection rule that an event and its matched subscription have to be mapped to the same node. These nodes are termed rendezvous nodes. When any node wants to publish or subscribe, the attributes of publication, or subscription are hashed to generate content ID. The node uses its successor list and finger-table to find the storage point of this publication or subscription. Publications and subscriptions are stored at the node whose ID is the immediate successor of the publication or subscription ID.

3 System model

We assume a city based scenario where the info-stations are installed at major positions (e.g., important intersections) of the city. These info-stations act as rendezvous points for publications and subscriptions by forming a DHT structure among them. These info-stations and the vehicles are equipped with omni-directional antenna and they have fixed transmission range. Each vehicle and info-station has its unique identification number. We assume that the info-stations have a transmission range of 1 000 m whereas vehicles have a transmission range of 500 m. The minimum distance between two info-stations is assumed to be 3 000 m. Info-stations broadcast their IDs periodically for the vehicles moving around them. Vehicles can send publications/subscriptions to info-stations directly if they are in their range. Otherwise, publications and subscriptions are forwarded hop by hop by utilizing other vehicles moving on the road.

Info-stations act as the ultimate meeting points or brokers for publications and subscriptions and they forward matching publications to interested subscribers. Each vehicle and info-station has its own unique identification number. Info-stations form a chord like DHT structure.

Chord is utilized to create hashed IDs of info-stations to form a logical overlay ring. The published and subscribed messages get hashed content IDs. These content IDs are termed keys. Both content and node IDs are chosen from same identifier space which is taken big enough to avoid node and content IDs mapping to the same hash value. This overlay of info-stations is used to form rendezvous points for publication and subscriptions. We assume following two deployment scenarios:

- Infrastructure deployment: Info-stations are connected to the Internet. The underlying communication among them is through IP based Internet.
- Infrastructure less deployment: Info-stations are not connected to the Internet. The underlying communication among them is multi-hop via vehicles.

Further, we assume that vehicles are not equipped with global positioning system (GPS) or navigation systems in our approach.

3.1 Skewed vehicle distribution

It is believed that distribution of vehicles running around any area is not uniform. It is dependent on time of day, with some popular hot spots in particular area. For example, during morning and evening hours, traffic density is maximal due to crowds of office goers. Similarly, in night hours traffic density is minimal. Further, some areas can be considered as hot spots during overall low traffic density times. We observe that around movie theatres, airports, railway stations, night clubs, etc., traffic density is more during off hours.

This skewed distribution of vehicles has an interesting effect on the DHT of info-stations when we consider infrastructure-less deployment. In high density situations, all the info-stations are able to run the DHT maintenance algorithm and, as a result single DHT of info-stations is formed. On the other hand, under skewed distribution and low density this single DHT is broken up into smaller DHTs as only those info-stations which are surrounded by vehicles can maintain their neighborhood information. This split of DHTs is depicted in Fig. 1 and Fig. 2.

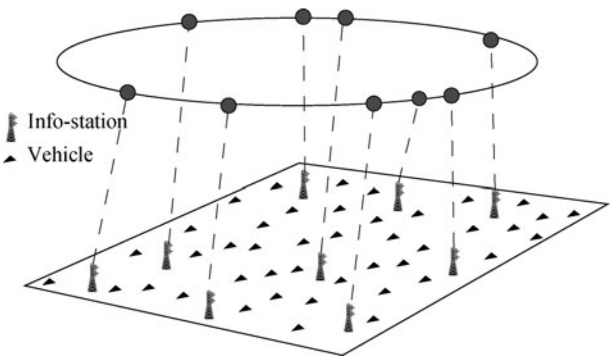

Fig. 1 DHT of info-stations

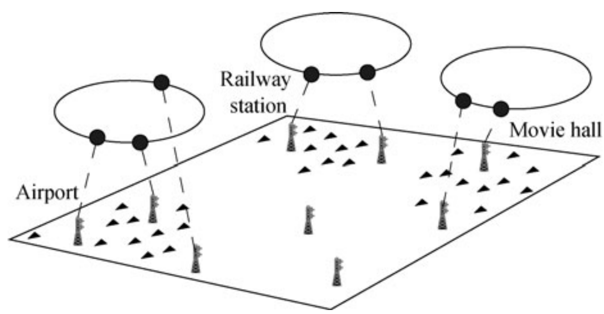

Fig. 2 Splitting of DHT

4 Description of algorithms

In this section we present the algorithms for publication, subscription, location determination, and notification used in our approach. Each vehicle has the following data structures to manage publications and subscriptions and to provide location services:

- Subscription table: to store active subscriptions.
- Publication table: to store active publications.
- Forwarding table: to store others publications and sub-

scriptions for further forwarding.

- Vehicle ID: identification number of the vehicle.
- Last info-station time stamp: identifier of the last info-station passed and the time stamp of this event.

Info-stations also have subscription table, publication table, info-station ID, and forwarding table. Additionally, they maintain a location table which stores the location information of vehicles they are responsible for.

4.1 Publishing, subscribing and matching

In our approach, the publish primitive is defined as publish (publication specification, TTL) where TTL is the time for which a publication is considered to be active. Subscribe primitive is defined as subscribe (subscription specification, subscriber_id, TTL) where subscriber_id is the identification of the vehicle which subscribes and TTL is the time for which a subscription is valid. Figure 3 provides the procedures for publishing and subscribing.

```
Procedure publish (publication, TTL)
1. Store publication in publication table
2. if (publication is active) //check TTL
3. if (any vehicle in 1-hop range)
4. Send publication
5. else
6. Keep on moving for some time
7. goto Step 2.
8. else (discard the publication)

Procedure subscribe (vehicle ID, subscription, TTL)
1. Store subscription in subscription table
2. if (subscription is active) // check TTL
3. if (any vehicle in 1-hop range)
4. Send subscription
5. else
6. Keep on moving for some time
7. goto Step 2.
8. else (discard the subscription)
```

Fig. 3 Publication and subscription installation

Vehicles willing to publish or subscribe send the publication/subscription to the nearest info-station through other vehicles running between itself and the info station. Vehicles forward the publication/subscription to one hop neighbors in range. Vehicles which are running ahead of the vehicle (that is publishing or subscribing) in the same direction and the vehicle running in opposite direction carry the publication or subscription and send it again to their one hop neighbor in their ranges. This process continues till the publication or subscription reaches an info-station. If traffic is dense, then publication/subscription is transferred to the info-station with less delay. On the other hand, if there are very few vehicles on road, then there may be substantial delay in sending the publication/subscription to info-station.

In this process of hop-by-hop transfer of publication and subscription towards info-stations, it may be possible that some of the vehicles (through which the publication/subscription is forwarded) receive matching publications or subscriptions. In this case, in addition to forwarding the publication/subscription, they can also act as rendezvous point for publications and subscriptions. Once the publication/subscription reaches an info-station, it is sent to the rendezvous node by utilizing the DHT structure connecting the info-stations. Figure 4 provides algorithms which outline forwarding and matching processes for moving vehicles.

```
Procedure forwarder (pub/sub/notification, vehicle ID)
1. Store pub/sub/notification in forwarding table
2. Case: (notification)
3. if (notification for itself)
4. Accept
5. else goto Step 17
6. Case: (subscription)
// match with publication table and forwarding table
7. if (match=True) then build notification
8. goto Step 17 to notify concerned subscriber
9. else goto Step 17 to forward subscription further
10. Case: (publication)
// match with subscription table and forwarding table
11. if (match=True) in subscription table then
12. Accept as notification
13. else if (match =True) in forwarding table then
14. Build notification
15. goto Step 17 to notify concerned subscriber
16. else goto Step 17 to forward publication further.
17. if (pub/sub/notification still active) // check TTL
18. if (any vehicle in 1-hop range)
19. Forward pub/sub/notification
20. else keep on moving for some time
21. goto Step 17.
22. else (discard pub/sub/notification)
```

Fig. 4 Forwarding and matching

Each vehicle, after receiving a publication or subscription from other vehicles first performs the matching operation with pre-stored publications and subscriptions. At this instant, each vehicle assumes the role of broker. For example, every received publication is first checked with the subscriptions in the subscription table and forwarding table. If the received publication matches with any subscription of subscription table then it is accepted as a notification by the vehicle receiving the publication. If the publication matches any subscription in forwarding table then the corresponding subscriber is notified by the vehicle itself. If none of the cases mentioned above occurs, then the publication is forwarded to other vehicles in the one hop range. This process is repeated until a matching subscription is found at vehicles on the way or the publication reaches at the nearest info-station.

When a publication or subscription reaches any info-station it is routed to a rendezvous info-station by utilizing the DHT routing substrate. These info-stations are connected to each other through virtual links provided by the underlying transport layer. We predefine a specific attribute schema for describing publications and subscriptions. Each publication and subscription specification has some attributes associated with it. These attribute names are hashed to find the rendezvous info-station of publications and subscriptions. Consequently, publications and subscriptions related to the same attributes are routed towards the same info-station. The publications and subscriptions are matched with each other at rendezvous info-stations and corresponding subscribers are notified with the help of vehicles passing info-stations.

4.2 Tracking subscribers location

As vehicles move and subscribe simultaneously, the location of subscribed vehicles will very likely be different when notifications are ready to be delivered. Vehicles may issue subscription in one region of a city and at the time of notification they might be in any other region. In our approach, we have not assumed GPS enabled vehicles. Instead, subscribed vehicles and info-stations work in a cooperative manner to locate vehicles.

Location information of subscribed vehicles is maintained in the DHT of info-stations in a distributed fashion. Every info-station is responsible for storing the location information for a set of vehicles. The vehicle ID is hashed to discover the info-station which stores the location details of a vehicle. Each vehicle broadcasts its ID and direction at the time of passing an info-station. That info-station hashes the vehicle ID to find out which info-station is responsible for the location base of the vehicle and updates the location information at that info-station. In this way, the location of every subscribed vehicle is up to date at info-stations.

It should be noted that the system is not required to maintain the location details of all the vehicles. It would be very costly overhead to regularly update the locations of all vehicles all the time. The location information is maintained only for vehicles that have subscribed to some events and are thus expecting notifications. The information consumers are not interested in knowing the ID or location of information providers. They are only interested in getting the desired information. On the other hand, subscribed vehicles must be located to forward the notifications to them at a given time.

The info-stations are installed at major intersections of city in our approach. As mentioned earlier, info-stations periodically broadcast their IDs to vehicles moving past them. This assists vehicles as they know that they are nearing info-stations and certain actions should be taken. Roads leading away from info-stations are marked with direction tags (North, South, East, and West) as depicted in Fig. 5. Vehicles can inform the info-stations that which direction they are coming from and going to.

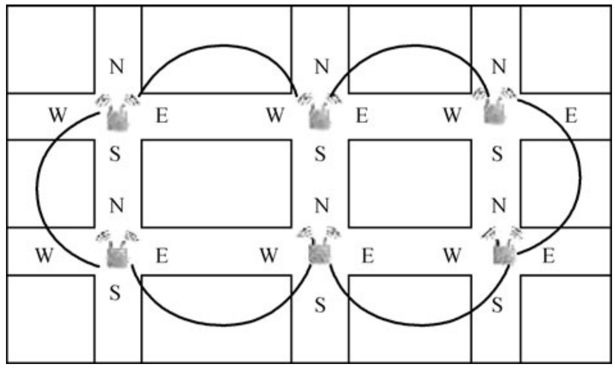

Fig. 5 Tracking vehicle location

How this direction information is generated can be vehicle or application specific. At an info-station, vehicle can either turn left, right or go straight. Smart vehicles can record their steering movement while they are passing info-stations to know what direction is taken.

The direction information with timestamp is stored at the info-station responsible for the location of the vehicle. This last info-station crossed and direction information is also stored in the vehicles. For example, suppose a notification is ready at an info-station for any vehicle. The info-station enquires about the location of vehicle by contacting the info-station maintaining the location of a vehicle via DHT lookup. The returned info-station reports the last info-station passed data for that vehicle. The notification is forwarded to that last passed info-station using the underlying DHT routing substrate which uses the direction information to forward the notification by using other vehicles as carriers that are passing the info-station and moving in that direction. If a notification is ready at any vehicle (matching publications and subscriptions on vehicles met on the way), then the notification is first transferred to nearest info-station which forwards it to the target info-station.

In our design, subscribed vehicles are not required to be located in a fine grained manner. Our vehicle tracking approach is rather a coarse-grained one. It provides details of the info-stations between which a vehicle can be found at any given time. Once vehicle is located between two info-stations, then other vehicles moving in the direction of the target vehicle can be utilized to forward notifications.

4.3 Notification delivery

Notifications are delivered using the location information of subscriber. They are forwarded towards target info-stations which in turn forward them towards the target vehicle. This process is explained by the following example. Suppose we learn that a target subscriber vehicle is in between info-stations *A* and *B* at a given time (has passed *A* and is moving towards *B*). The notification is routed towards *A* which forwards it through vehicles moving towards info-station *B*. There may be a situation that more than one vehicle acts as a carrier for this notification. There can be two cases: 1) the notification is delivered to the target vehicle but it is still being forwarded towards *B* (because more than one vehicle was carrying it); 2) notification is not delivered to target vehicle (perhaps because the target vehicle has already passed *B* at the time the notification arrived). These two cases are handled by our approach in following manner:

- In the first case, when the subscriber vehicle passes *B*, it informs *B* that it has received notifications (with their ID) while it was in between *A* and *B*. Thus when duplicate notifications reach at info-station *A*, they are discarded and not forwarded further.
- In the second case, *B* knows which direction the target vehicle has moved in. Consequently, when a notification reaches *B* (and the subscribed vehicle has not acknowledged the receipt of the notification), it is forwarded towards the target subscribed vehicle.

5 Simulation environment and results

We have simulated our approach by using Oversim [8] over objective modular network testbed in C++ (OMNeT++) [9] and traffic control interface (TraCI) client for OMNeT++/MiXiM [10] framework. Our approach for simulating the peer to peer overlay network of publishers, subscribers, brokers and info station is depicted in Fig. 6.

OverSim [8] is an open-source simulation framework for peer to peer overlay networks to be used over OMNeT++ [9] simulation environment. MiXiM [10] is a communication network simulation package for OMNeT++ with a focus on wireless networks. TraCI client [11] uses simulation of urban

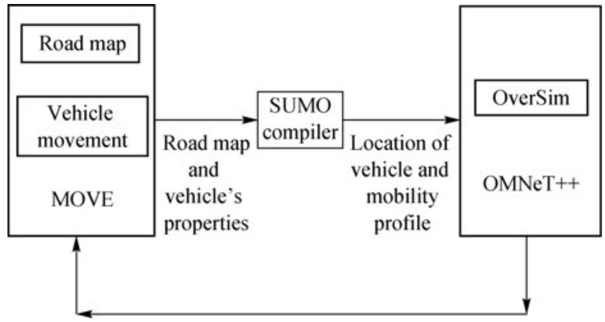

Fig. 6 Traffic and network simulators

mobility (SUMO) [12] with OMNET++ to simulate vehicle to vehicle communications. SUMO uses mobility model generator for vehicular networks (MOVE) [13] as the mobility simulator. The two components of MOVE are a road map editor and a vehicle movement editor. The road map editor is used to generate a road map from topologically integrated geographic encoding and referencing (TIGER) database or Google earth files. Movement of vehicles can be generated automatically or manually using the vehicle movement editor. It specifies the properties of each vehicle such as vehicle speed, duration of trip, origin and destination of vehicle, vehicle departure time, etc. The road map and the vehicle's properties are sent to the SUMO compiler which generates a trace file for the network simulator. This trace file includes the location of each vehicle at every time instant for the entire simulation time and their mobility profiles. This trace file is used in the simulations on OMNeT++. We have performed many sets of experiments to evaluate the performance of our approach in different settings. In the following subsection, we describe, in detail, the different simulation settings and process of traffic simulation in our approach.

5.1 Traffic simulation

We have simulated traffic scenarios in two cities of India as our reference. To simulate the traffic of an urban area we have chosen South Delhi whereas for semi-urban area we have chosen Allahabad. The following steps were taken to develop the map and simulate traffic using SUMO:

Fixed nodes and edges To simulate the traffic we use the Google earth maps of South Delhi and Allahabad. We mark the major roads and use all the major intersections as fixed nodes (info-stations). We also mark the edges between the fixed nodes.

XML files for coordinates of nodes and edges We formulate the coordinates of all fixed nodes and also identify the edges connecting these nodes. In general, we convert the given information into three files, which are used by SUMO as input data. Two of the files contain the network information, to be converted into node and link information in SUMO. Figure 7 depicts the map of South Delhi which is taken as reference. Further, Fig. 8 depicts the snapshot of traffic simulation for the map of South Delhi.

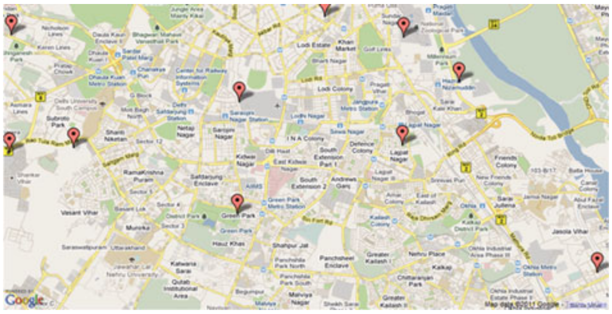

Fig. 7 Map of South Delhi showing fixed nodes and edges

Fig. 8 Snapshot of traffic simulation for South Delhi

We divided a simulated day into three traffic types: 07:00–10:00 and 17:00–19:00 (peak traffic), 10:00–17:00 and 19:00–00:00 (moderate traffic), and 00:00–7:00 (low traffic). The simulation settings for vehicles are summarized as follows:

South Delhi (urban area) This is a high vehicle density scenario with an organized four-lane road network. Maximum 1 000 vehicles may be present concurrently. Maximum speed of vehicles is 60 km/h and near intersections (info-stations) the average speed is 30 km/h. The vehicle density is dependent upon time of day. Peak traffic density, 800 to 1 000 vehicles, moderate traffic density, 300 to 500 vehicles, and low traffic density, 100 to 300 vehicles.

Allahabad (semi-urban area) This is a somewhat low vehicle density scenario with a less organized two-lane road network. A maximum of 600 vehicles can be present concurrently. Maximum speed of vehicles is 50 km/h and near the intersections (info-stations) average speed is 20 km/h. The vehicle density is again dependent upon time of day. Peak traffic, 400 to 600 vehicles, moderate traffic, an average of 300 vehicles, and low traffic, an average of 100 vehicles.

In both of the above settings, publishers and subscribers are distributed randomly among all the vehicles. The maximum percentage of subscriber vehicles is 10% whereas maximum 20% act as publishers. Vehicles utilize a predefined set of matching publications and subscriptions from a file to generate at fixed rate. Simulation results are collected for both the deployment scenarios with two further variations: 1) info-stations are connected to internet; 2) info-stations exchange information through vehicles moving between them. In the next subsection we present our simulation results and their analysis. Further, the distribution of the number of vehicles is also of two types: in one set of experiments, vehicles are uniformly distributed across the map whereas in other set vehicle distribution replicates real life situations where some roads near hot spots like metro and rail stations, shopping malls, airports, and some popular junctions experience heavier traffic than other roads.

5.2 Simulation results

In this subsection we present our simulation results. Two sets of results are presented, one with uniform distribution of vehicles and other with traffic skewed around hot spots. These results are gathered for simulation environments of both South Delhi (urban area) and Allahabad (semi-urban area), and for both deployment scenarios, info-station with internet connectivity (connected) and without connectivity (isolated). In all the results, we measure the delivery ratio expressed as a fraction of subscribers successfully receiving the notification with respect to the number of info-stations.

5.2.1 Uniform distribution of vehicles

Figures 9(a)–9(d) depict the results where a uniform distribution of vehicles is considered. Figures 9(a) and 9(b) depict the results obtained for high and low densities of vehicles, respectively, in the South Delhi setting. In Fig. 9(a), it can be observed that when vehicle density is high then the performance in the isolated info-station case is somewhat comparable to the connected info-station case. It may also be noted that as the number of info-stations increases, the delivery ratio increases and reaches a maximum of 95% and 89% for connected and isolated info-stations respectively. This suggests that if vehicle density is high then with more info-stations, we can achieve a similar delivery ratio with isolated info-station to the connected info-station case.

The situation is completely different when the density of vehicles is low. Here, the connected info-station scenario

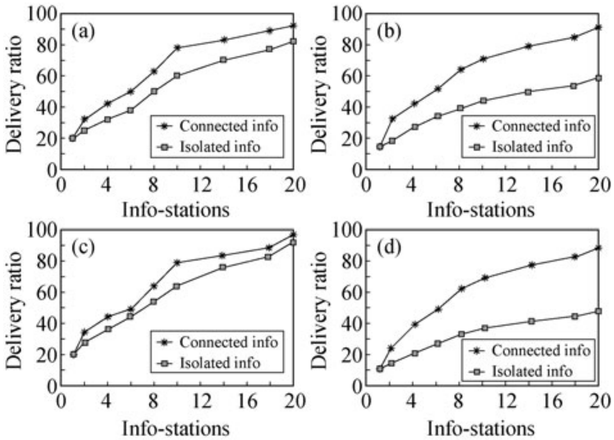

Fig. 9 Delivery ratio: uniform distribution of vehicles. (a) urban area, high vehicle density; (b) urban area, low vehicle density; (c) semi-urban area, high vehicle density; (d) semi-urban area, low vehicle density

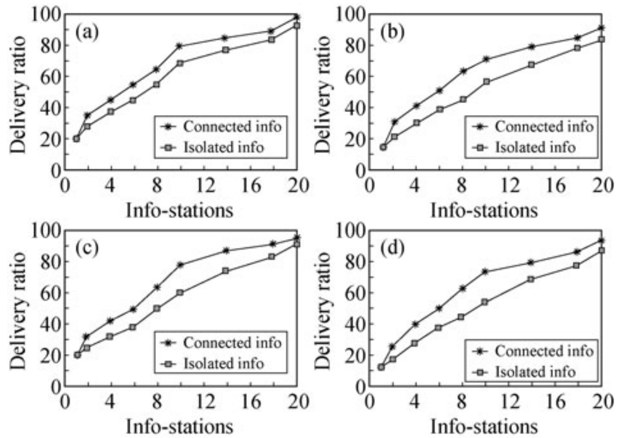

Fig. 10 Delivery ratio: skewed distribution of vehicles. (a) urban area, high vehicle density; (b) urban area, low vehicle density; (c) semi-urban area, high vehicle density; (d) semi-urban area, low vehicle density

performs much better than the isolated info-station case. It may be observed in Fig. 9(b) that the maximum delivery ratio is not degraded much for the connected info-station case in low density but for isolated info-stations it is reduced by almost 50%. This is due to partitioning of the DHT overlay as there are very few or no vehicles moving between them.

Figures 9(c)–9(d) depict the results obtained for high and low density of vehicles respectively, in the Allahabad setting. Here, a similar pattern of results is obtained. In the case of connected info-stations very similar performance is recorded as in the South Delhi setting for both high and low densities whereas for isolated info-stations, performance is much reduced.

5.2.2 Skewed distribution of vehicles

Figures 10(a)–10(d) present the results for the skewed distribution of vehicles. Figures 10(a) and 10(b) show the results for urban area whereas Figs. 10(c) and 10(d) provide the results for semi-urban settings. We have marked some areas in the map as hot spots with a vehicle distribution such that roads and info-stations around them experience higher traffic density.

Here, we recorded substantial improvement in the performance of the isolated info-station scenario, both for high and low densities. The improvement is more remarkable in the case of low density, both in the case of urban and semi-urban areas where we observed that delivery ratios almost double with respect to the uniform distribution setting.

The reason behind this remarkable improvement is justifiable. We have taken several snapshots of running simulation to closely observe the DHT formation process among the info-stations. In the chord DHT protocol, every node periodically runs a stabilization procedure to discover node joining or leaving. In this stabilization procedure, nodes ask their successor to tell the ID of its predecessor. In isolated info-station case, the underlying multi-hop communication among vehicles is the only way to send these stabilization messages in DHT layer. If there are no vehicles between two neighbors of the DHT ring, then this results in timeout of the DHT heartbeat messages. Nodes which are not able to reply the stabilization messages due to underlying intermittent connectivity are considered to be out of the DHT ring. As the density of vehicles is skewed around a few clusters of the info-stations, we observe formation of small DHT rings of info-stations around the hotspots. Further, publications and subscriptions happen around those hot spots only due to the presence of most of the vehicles around them. Consequently, we observe increase in delivery ratio as notifications are routed easily towards subscribers due to good connectivity through dense vehicle population and notifications have to travel lesser distance to reach to subscriber.

5.2.3 Delay in notification delivery

Figures 11(a) and 11(b) depict the simulation results for delay in delivering notifications to interested subscribers. These results are collected for the urban scenario where delay in notification delivery is plotted against vehicle density. Figure 11(a) provides the results for uniform distribution of vehicles whereas Fig. 11(b) shows the results for a skewed distribution. It may be noted that when info-stations are potentially connected to the Internet then vehicle density causes very little effect on delay in notification delivery. In the case of isolated info-stations, when uniform distribution of vehicles is considered then delay increases when vehicle density increases. On the other hand, when vehicle distribution is skewed then the performance of system is remarkably satisfactory. This is due to both publishers and subscribers may be found in the vicinity of hot spots where smaller DHTs get formed due to vehicle concentration around them.

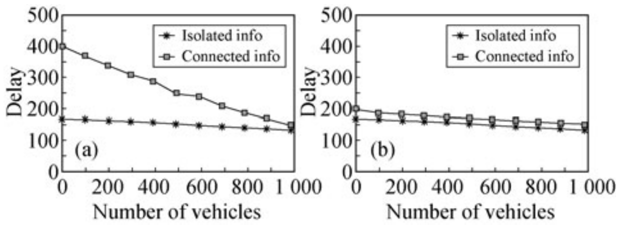

Fig. 11 Delay in notification delivery. (a) urban area, uniform distribution; (b) urban area, skewed distribution

6 Related work

Recently, some approaches [14–17] have been proposed which use the publish/subscribe paradigm for information dissemination in VANET like settings. These approaches have contributed significantly toward understanding the applicability of publish/subscribe over VANET.

In these approaches [14–17], a hybrid setup is assumed where there are stationary info-stations and moving vehicles communicating in cooperative manner. These info-stations are assumed to be connected to the Internet for timely information spreading. Vehicles are assumed to be installed with navigation system and GPS and they behave like mobile sensors that collect information about traffic condition, parking slots, etc. Then with the help of GPS and navigation system a publish/subscribe middleware is used to disseminate information geographically.

The main goal in these approaches [14–17] is to design a publish/subscribe middleware for vehicular networks that considers location and time in its design objectives. This middleware enables the application developers to publish notifications towards specific locations taking advantage of the information that can be extracted from the vehicle's navigation systems (location, map, destination of the driver, etc.) to generate subscriptions. The navigation system decides if a vehicle is interested in receiving a specific notification or not. Mobile vehicles transfer this information to the info-station on its way. All the info-stations are directly accessible to each other with the help of the Internet. A centralized system combines the gathered information and generates traffic warnings. Traffic warnings are sent to the nearest info-station and from there they are routed towards the affected road segment by vehicle-to-vehicle communication.

Authors of [18,19] introduced an approach for segmented DHTs of moving cars. They assume that a city is divided into segments and every segment has its own DHT where each vehicle can communicate to other directly. Vehicles may cross segments while moving and they leave the previous segment (or DHT) and join the next. Though their approach looks feasible but the effect of frequent joining and leaving of vehicles to different DHTs is yet to be inverstigated. The authors have not provided any supporting simulation results. Moreover, if vehicles can communicate directly with each other, all the vehicles can overhear the messages in a segment and thus formation of DHT in such small segment appears to be a redundant feature.

6.1 Discussion

The approach closest to ours is presented in [16]. Although the proposed approach is effective, it depends heavily upon GPS, navigation system and digital maps of cities. They also require highly sophisticated vehicles equipped with these devices. Also, in these approaches info-stations are assumed to be connected to the Internet and a centralized server is required to gather information to build notifications. This requires the pre-deployment of infrastructure to make the design workable. Further, it is desirable to store the relevant information zone-wise and maintain given zone for vehicles to be notified. The subscription has to be routed towards a zone to get matching notifications. This is achieved through the network of info-stations assumed to be connected to the Internet.

We argue that for different application scenarios, the proposed solution in [16] has limitations even if this extensive infrastructure is there. For example, in any urban locality, there are parking spots for cinemas, bars etc. Some parking spots are near to points of interest and some are several kilometers away. If vehicles are only notified once they arrive in the vicinity of a point of interest that parking is not available then it would be a significant obstacle for drivers to find alternative parking. If drivers can get information in a timely manner, they can choose another parking location or another cinema altogether. Further, if info-station infrastructure is not there, then vehicles maintain the information by working in collaboration (ad-hoc persistence). In other words, vehicles crossing the affected area or moving nearby keep on transferring the desired publication to each other to maintain it there. This is hugely dependent on vehicle density. For instance, if a vehicle having desired information crosses the area and does not find any other vehicle for some time going towards that area, then some important notifications will be missed.

Finally, the approach in [16] assumes that there is a centralized server attached to the Internet which collects information and then generates notifications. These notifications must be generated by a centralized entity because they also include the location information of persistence area where they must be stored, routed and maintained for a given time. They require infrastructure of info-stations because the publications have to be routed to the centralized server in a timely fashion, notifications have to be generated and then must be forwarded towards the affected area in time. Their approach cannot be solely dependent on an ad-hoc mode of dissemination.

Our approach does not assume that vehicles are fitted with sophisticated devices like GPS-navigation systems. Instead of generating notifications in centralized manner, the underlying DHT is used to route and store related publications and subscriptions close to rendezvous info-stations. In our approach, we relax the requirement of info-stations to be connected to the Internet. Info-stations are just a wireless devices with an omni-directional antenna and some memory. They are not required to have any infrastructure to talk to other info-stations. Instead of the continuous transfer of notifications from one vehicle to another, these low cost info-stations can store the information for a given time in a given area. For [16], one cannot find info-stations everywhere (because that approach assumes that info-stations are connected to the Internet). For our approach, it is just as simple as putting a stationary node into ad-hoc mode, which is less expensive to install and maintain. Our contributions can be listed as follows:

- Results are gathered for two scenarios; info-stations connected to the Internet and isolated info-stations where links are formed in an ad-hoc manner through vehicles moving among them.
- For the connected info-stations case, our approach is somewhat similar to [16] with respect to performance issues like delivery ratio and delay in delivery. This is justifiable as when the info-stations are not leaving and joining dynamically and the link between info-stations are stable then whether they communicate and build notifications in the DHT manner (as in our approach) or by any other manner (as assumed in [16]), there would not be much difference performance wise. This is also evident in our results as compared to performance achieved by [16] in similar simulation scenarios.
- The distinguishing feature of our approach is the solution provided for the scenario where info-stations are not connected. In this case there is an enforced dynamism in the DHT structure of info-stations due to dynamism of ad-hoc links formed by moving vehicles between them. This dynamism is dependent upon the density of vehicles which changes according to the time of day. Accordingly, the DHT of info-stations can be split or different DHTs can be merged together. Interestingly, this merge and split does not have a major impact on the performance when the skewed distribution of vehicles is used in our simulations which models the density of vehicles near hotspots as it happens in real life. This is due to the fact that the vehicles that need information and vehicles that provide information are found to be close to these hotspots.

7 Conclusion

We have presented our approach for information dissemination in a hybrid VANET of moving vehicles and stationary info-stations. Our approach utilizes the publish/subscribe communication paradigm and a DHT based overlay network. We do not use GPS to locate vehicles on the road. The location of vehicles is determined by info-stations in a cooperative manner. We have simulated two deployment scenarios, both connected to the Internet and an isolated scenario that utilizes vehicles to disseminate information. We have used traffic scenarios for two Indian cities (South Delhi and Allahabad) as a reference for urban and semi-urban areas. The simulation results suggest that even when the info-stations are not connected, the performance of our approach is acceptable. This is true even for the situation when vehicle density is low if the skewed distribution of vehicles is considered.

References

1. Kenney J. Dedicated short-range communications (dsrc) standards in the united states. Proceedings of the IEEE, 2011, 99(7): 1162–1182
2. Hadaller D, Keshav S, Brecht T, Agarwal S. Vehicular opportunistic communication under the microscope. In: Proceedings of the 5th International Conference on Mobile Systems, Applications and Services, MobiSys'07. 2007, 206–219
3. Eugster P, Felber P, Guerraoui R, Kermarrec A. The many faces of publish/subscribe. ACM Computing Surveys, 2003, 35: 114–131
4. Mishra T, Garg D, Gore M. A publish/subscribe communication infrastructure for vanet applications. In: Proceedings of the 2011 IEEE Workshops of International Conference on Advanced Information Networking and Applications, WAINA'11. 2011, 442–446
5. Ratnasamy S, Francis P, Handley M, Karp R, Shenker S. A scalable content-addressable network. SIGCOMM Computer Communication

Review, 2001, 31: 161–172

6. Rowstron A, Druschel P. Pastry: scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In: Proceedings of the IFIP/ACM International Conference on Distributed Systems Platforms Heidelberg, Middleware'01. 2001, 329–350
7. Stoica I, Morris R, Liben-Nowell D, Karger D, Kaashoek M, Dabek F, Balakrishnan H. Chord: a scalable peer-to-peer lookup protocol for internet applications. IEEE/ACM Transaction on Networking, 2003, 11: 17–32
8. Oversim website. http://www.oversim.org/
9. OMNET++ website. http://www.omnetpp.org/
10. Wessel K, Swigulski M, Köpke A, Willkomm D. Mixim: the physical layer an architecture overview. In: Proceedings of the 2nd International Conference on Simulation Tools and Techniques, Simutools'09. 2009, 78:1–78:8
11. Wegener A, Piórkowski M, Raya M, Hellbrück H, Fischer S, Hubaux J P. Traci: an interface for coupling road traffic and network simulators. In: Proceedings of the 11th Communications and Networking Simulation Symposium, CNS'08. 2008, 155–163
12. Krajzewicz D, Hertkorn G, Rössel C, Wagner P. Sumo (simulation of urban mobility); an open-source traffic simulation. In: Proceedings of the 4th Middle East Symposium on Simulation and Modelling, (MESM2002). 2002, 183–187
13. Karnadi F K, Mo Z H, Lan K C. Rapid generation of realistic mobility models for vanet. In: Proceedings of the WCNC. 2007, 2506–2511
14. Leontiadis I, Mascolo C. Opportunistic spatio-temporal dissemination system for vehicular networks. In: Proceedings of the 1st International MobiSys Workshop on Mobile Opportunistic Networking, MobiOpp'07. 2007, 39–46
15. Leontiadis I. Publish/subscribe notification middleware for vehicular networks. In: Proceedings of the 4th International Doctoral Symposium on Middleware, MDS'07. 2007, 12:1–12:6
16. Leontiadis I, Costa P, Mascolo C. A hybrid approach for content-based publish/subscribe in vehicular networks. Pervasive and Mobile Computing, 2009, 5: 697–713
17. Leontiadis I, Costa P, Mascolo C. Extending access point connectivity through opportunistic routing in vehicular networks. In: Proceedings of the 29th Conference on Information Communications, INFOCOM'10. 2010, 486–490
18. Prinz V, Brocco M, Woerndl W. Peer-to-peer algorithms for c2c communication systems. In: Proceedings of the 6th International Conference on Advances in Mobile Computing and Multimedia, MoMM '08. 2008, 376–381
19. Prinz V, Eigner R, Woerndl W. Cars communicating over publish/subscribe in a peer-to-peer vehicular network. In: Proceedings of the 2009 International Conference on Wireless Communications and Mobile Computing: Connecting the World Wirelessly, IWCMC'09. 2009, 431–436

Mrs. Tulika is an assistant professor in the Department of Computer Science and Information Technology, Sam Higginbottom Institute of Agriculture, Technology & Sciences, Allahabad, India. She received her MSc in Information Technology from Birla Institute of Technology, Mesra, Ranchi, India. She is currently pursuing her PhD in the Department of Computer Science and Engineering, Thapar University, Patiala, India. She has more than 10 publications to her credit.

Dr. Deepak Garg's research interests lie in the area of efficient algorithm design. He has 15 years of experience in teaching undergraduate and postgraduate students. He is currently an associate professor in the Department of Computer Science and Engineering, Thapar University, Patiala, India. He has more than 80 publications to his credit. He is Chair of the ACM North India SIGACT and Secretary of the IEEE CS Chapter. He has organized several workshops and conferences on advanced computing.

Dr. Manoj Madhava Gore received his PhD in Computer Science and Engineering from the Indian Institute of Technology, Kanpur, India. Dr Gore has been on the faculty rolls of Motilal Nehru National Institute of Technology, Allahabad, since January 1991 where he is currently working as a professor in the Department of Computer Science and Engineering. He also served in the Indian Institute of Information Technology of Allahabad, as an associate professor for one academic year. Dr Gore has research interests and publications in the areas of database modeling, transaction recovery, and networks. He is a fellow of IETE (India), a senior member of the Computer Society of India, and a member of IE (India), the Indian Society of Technical Education, IEEE, and ACM. He has more than 60 publications to his credit.