On the Evaluation of Caching in Vehicular Information Systems

Nicholas Loulloudes and George Pallis and Marios D. Dikaiakos
Department of Computer Science, University of Cyprus, Nicosia, CY1678, Cyprus
{loulloudes.n, gpallis, mdd}@cs.ucy.ac.cy

Abstract—**VANETs have been envisioned as an infrastructure for deploying Vehicular Information Systems (VIS) that among others provide drivers with an up-to-date view on the prevailing traffic conditions. In this work we evaluate the benefits of caching vehicular information obtained from such VIS through VITP, a location-aware, application-layer communication protocol that we extend to support caching. We present an evaluation study of our approach conducting extensive simulation on large-scale vehicular networks under different realistic urban traffic conditions. Our results identify the critical parameters that affect information quality in VANETs as well as demonstrate the viability and effectiveness of the cache-enabled VITP.**

I. Introduction

Inter-vehicle communication has emerged as a promising field of research and development, where advances in wireless and mobile ad-hoc networks, GPS and sensor technologies can be applied to real-life problems and result to great market potential. The uptake of such technologies by the transportation industry is expected to result in VANETs. Due to vehicular mobility, VANETs are characterized by highly dynamic topologies, short-lived links, and frequent network disconnections [1]. Vehicular Information Systems (VIS) generate a substantially overhead on the VANET in order to obtain and maintain a global view of the prevailing conditions in the vehicular environment. This can lead to the saturation of the already limited network capacity, which in turn is bound to degrade the quality of VANET services (lower response times and loss of information quality) [2]. Therefore, the provision of efficient, robust, wide-area information services over VANETs remains an open challenge.

In this work, we investigate the implications of caching on the efficiency of VANETs and on the accuracy of VANET-based VIS. We examine the implications of caching in a number of different vehicular traffic scenarios that are more realistic and challenging. Building upon our prior work on VITP [3], we provide answers to the key question: *Does the co-existence of a proactive, location-aware, communication protocol and caching maintain acceptable levels of vehicular information quality while sustaining network performance?* Our main contributions are summarized as follows: a) We extend the VITP architecture [3] to support cache-based location aware services. b) Through extensive simulation, we identify the critical parameters that affect vehicular information quality as well as demonstrate the viability and effectiveness of the cache-enabled VITP.

II. Enabling Caching support in VITP

To support caching operations, we extend the VITP message syntax with a set of *cache-control* headers. These act as directives to VITP-peer caching decisions (Table I). The reader can find more details concerning the main design concepts of VITP architecture in [3]. We choose to extend VITP since it allows the dissemination of query messages proactively, thus providing better control in the number of messages injected in the VANET.

III. Simulation Testbed Setup

For simulating the effects of vehicle movement, mobility traces were generated using TrafficModeller [4] and SUMO[1]. For simulating the behaviour of VITP in a VANET, ns-2[2], was employed.

A. Vehicular Mobility Generation

Two sets of mobility traces, each one representing traffic on roads of a real city with different topographical layout. *Region 1* follows a Manhattan-like city and *Region 2* follows a common urban layout. Both regions were extracted from real-world maps through OpenStreetMap[3]. A number of "hot-spots" within these regions were defined, so as to simulate the traffic conditions that arise when people drive from/to their workplaces or houses. Vehicle top speed was bounded by the road speed limit. Mobility traces for 970 and 875 vehicles that move in Region 1 and 2 respectively were generated. These were later translated to ns-2 traces using TraceExporter[4]. For the network simulation, each vehicle is

[1]SUMO - Simulation of Urban Mobility, http://sumo.sourceforge.net/
[2]ns-2 - http://www.isi.edu/nsnam/ns/
[3]OpenStreetMap - http://www.openstreetmap.org
[4]TraceExporter - http://www.auto-nomos.de/

Directive	Value	Description
cacheable	Boolean	The reply can be cached.
expires	time-stamp	The time after the reply is considered expired (if cached).
private	Boolean	The cached reply can be reused only from the original peer that requested it.
public	Boolean	The cached reply can be reused from any peer.
validate-after	seconds	A peer must validate the cached reply with the target area after "n" seconds from the reply generation time.
p-validate-after	Boolean	A peer must validate the cached reply with the neighbour peers after "n" seconds since the reply generation time.
p-validate	Boolean	A peer must validate the cached reply with the neighbour peers before serving a request.
retransmit	Boolean	Serve request first from cache and retransmit it to target area also.

TABLE I
CACHE CONTROL DIRECTIVES USED IN A VITP REPLY MESSAGE

assumed to be equipped with IEEE 802.11 WiFi interface with a radio coverage of 200m and unlimited cache size. Simulations run for 1000s with a 200s warm up period.

B. Evaluation Scenarios and Query Generation

For the evaluation of the cache-based VITP we have set up the following scenarios:

Scenario 1. Each vehicle is aware of the whole road-network topology. Since travel times are heavily influenced by the prevailing traffic conditions, vehicles want to identify the road-paths towards their destination that will result in reduced the travel-times. To achieve this we introduce a query scheme implementing a "forward-scan-radar" traffic information system. Through its knowledge of the road network each vehicle can calculate all possible road-paths from its current position to the destination. *LookAhead (L)* queries are issued along all these possible road paths and are propagated to a certain depth in the path to obtain the conditions of the roads up to the specified depth. **Scenario 2** follows the paradigm of Scenario 1. The exception being here that unscheduled events take place (e.g., vehicle break-downs) in the aforementioned road-paths. Such events block road-paths and influence the normal traffic. **Scenario 3** assumes that vehicles would like to discover the availability of road-side facilities such as parking places. A stationary Road Side Unit (RSU) responsible for a facility broadcasts information concerning the facility and at certain time intervals. We randomly place 5 RSUs on the road-network each broadcasting information every 60 seconds.

In the above scenarios, VITP queries are issued with a $ReturnCondition = 5$. A Return Condition as specified in [3] determines the sampling size of the requested information a query must obtain before a VITP reply can be generated. For Scenarios 1 and 2, vehicles issues queries with $LookAhead = 2$. The values for the above parameters were selected by sampling the parameter space having in mind that a high *LookAhead* value can cause the saturation of wireless network bandwidth thereby causing a significant amount of query drops and on the other hand a high $ReturnCondition$ value increases vehicular information accuracy. A total of 29263 queries for *Region 1* and 16557 queries for *Region 2* were generated. For Scenario 3, 12231 and 7562 queries were generated for *Region 1 and 2* respectively. Replies are generated with *cache-control* header = *[cache-control: cacheable, expires=t, public]*. Consequently, information contained in each VITP reply message is cached to all intermediary nodes on its way towards the source node. Furthermore, we consider that all nodes use a *TTL-based cache replacement policy*. According to this policy, messages are removed from the cache as soon as their *TTL* (Time to Live) expires. A *TTL* specifies the maximum time for which a cached copy is considered valid.

To describe the performance of cache-enabled VITP in inter-vehicular networks, we employ the following metrics: 1) **Query Recall:** *the number of replies received while issuing queries towards a specific location of interest, over the number of replies that should have been received from that location*;
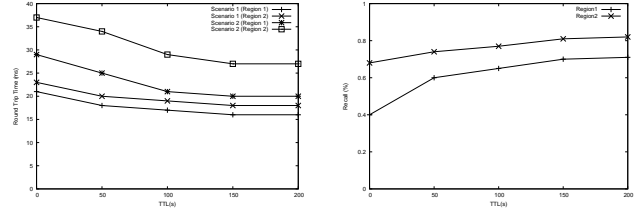
(a) Query Response Time   (b) Query Recall

Fig. 1.   Response Time and Query Recall vs. TTL

2) **Response Time:** *is the average Round Trip Time (RTT) of a successful VITP transaction*; 3) **Information Accuracy:** *measures how close the received value describing some vehicular information is to the actual value at the location of interest*; 4) **Number of Exchanged Messages:** *is defined as the total number of exchanged messages, including geographic routing messages and VITP query resolution messages throughout the whole simulation period*.

C. Caching Evaluation - Querying Road Traffic Conditions

For the three aforementioned scenarios we evaluate the performance of the cache-based VITP on both regions, under different *TTL* values assigned to query replies. $VITP_{TTL=0}$ emulates the original VITP where caching is not supported. The maximum value of $VITP_{TTL=200}$ denotes that traffic information is cached on VITP peers indefinitely. Response time (RTT) of the *LookAhead* traffic queries is measured by varying the *TTL* values of cached reply messages. According to the *TTL*-based replacement policy, when the *TTL* of a cached message expires, it is discarded by all vehicles' caches. Results are reported in Fig. 1(a).

For Scenario 1, longer *TTLs* result to a better diffusion of information throughout the vehicles' caches and, consequently, to an increased probability that the requested information is found nearby the requesting vehicle. For Scenario 2, due to the use of geographic routing, queries traverse a greater number of hops to overcome the break-down and reach the target location. This increase in the number of hops leads to an increase in RTT. Fig. 1(b) denotes the traffic information query recall for Scenario 1 under both regions. As *TTL* increases, the number of replicated information in the network increases and a query can be served also from information stored in other vehicles cache. This allows for vehicles to obtain the desired information faster and in addition reduces the amount of queries that should be re-generated in order to obtain information for the target locations of previous unresolved queries. The above observation is also reflected by examining the number and geographic distribution of replicated objects in the network as depicted by Figure 2. The x and y-axis denote each zone boundary in the road network while the z-axis (on the right) denotes the percentage of information completeness at each zone.
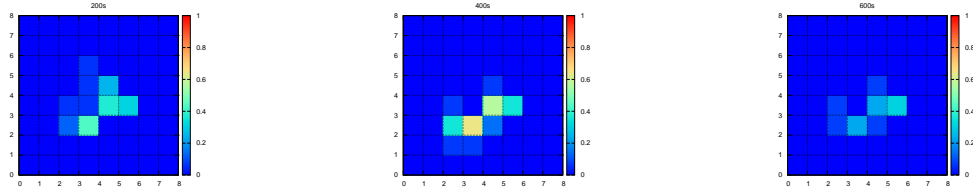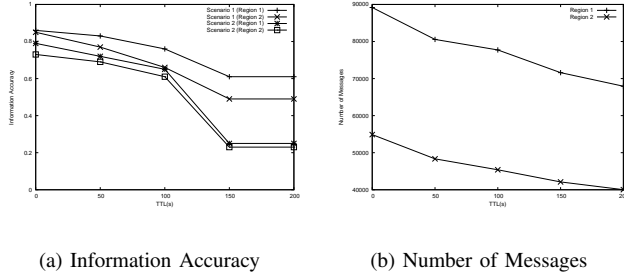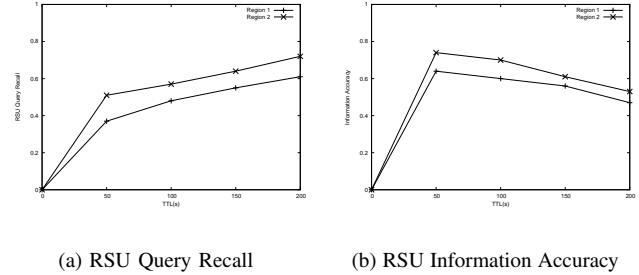
Fig. 2. Geographic distribution of Replicas in respect to simulation time

(a) Information Accuracy

(b) Number of Messages

Fig. 3. Information Accuracy and Number of Messages vs. TTL

(a) RSU Query Recall

(b) RSU Information Accuracy

Fig. 4. Caching implications on RSU discovery.

Information query accuracy for both scenarios is given on Fig. 3(a). Accuracy drops as the value of *TTL* increases denoting that there is high probability that VITP queries will be answered from other vehicle's cache instead of source location. The drop in information accuracy is expected since cached information does not accurately reflect the traffic conditions on the queried roads. Also, the road-network topology influences the rate of change of information accuracy between the two regions. We observe that the TTL of the cached information is directly influenced by the rate of change of traffic information in queried roads. The rate of change is influenced by the road length, vehicle density and traffic light existence. Therefore, caching policies should be adaptive and take into consideration these factors in order to adjust TTL of traffic information. Furthermore, the results depicted on Fig. 3(a) denote that information accuracy is heavily influenced by the presence of vehicle break-downs in the road-network. Fig. 3(b) depicts the total number of exchanged VITP messages for the resolution of traffic information queries with respect to *TTL* and the reduction in the network overhead achieved through the cache-based VITP is illustrated.

D. Caching Evaluation - Querying Road-Side Facilities Availability

Fig. 4(a) depicts the query recall of VITP messages for the discovery of road-side facilities. It is evident that in the lack of any information diffusion mechanism, as is in the original VITP, there is a very low probability that vehicles will locate and retrieve RSU broadcasted information. On the other hand, as Figure 4(b) denotes, increasing *TTL* values has a negative effect on the accuracy of the information received by querying vehicles.

To sum up our evaluation, the following remarks provide answers to the key question we defined in section I: 1) the cache-based VITP achieves information accuracy≥65% while at the same time reduces network overhead up to 12% under both normal traffic conditions and unscheduled traffic events; 2) The dissemination of RSU information without any caching mechanism is infeasible.

IV. CONCLUSIONS

VANETs have been envisioned to be useful in VIS. In this work, we explore the utilization of caching in VANETs by extending the architecture of VITP to support cache-based location aware services. Simulation results have shown that the use of a TTL-based cache replacement policy in urban environments, can achieve significant improvements under both normal traffic conditions and unscheduled traffic events.

REFERENCES

[1] G. Pallis, D. Katsaros, M. D. Dikaiakos, N. Loulloudes, and L. Tassiulas, "On the structure and evolution of vehicular networks," in *Proceedings of 17th Annual Meeting of the IEEE/ACM International Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems*, September 2009, pp. 502–511.

[2] B. Scheuermann, C. Lochert, J. Rybicki, and M. Mauve, "A fundamental scalability criterion for data aggregation in vanets," in *Proceedings of the 15th annual international conference on Mobile computing and networking*, 2009, pp. 285–296.

[3] M. D. Dikaiakos, A. Florides, T. Nadeem, and L. Iftode, "Location-aware services over vehicular ad-hoc networks using car-to-car communication," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 8, October 2007.

[4] L. Papaleontiou and M. D. Dikaiakos, "Trafficmodeler: A graphical tool for programming microscopic traffic simulators through high-level abstractions," in *Proceedings of the IEEE 69th Vehicular Technical Conference (VTC2009-Spring)*, April 2009.