2014 IEEE 13th International Symposium on Network Computing and Applications

An Adaptive Peer-to-Peer Overlay Scheme for Location-based Services

Giacomo Brambilla[1], Marco Picone[1], Michele Amoretti[2], Francesco Zanichelli[1]
1: Dept. of Information Engineering, 2: SITEIA.PARMA — University of Parma, Italy
giacomo.brambilla@studenti.unipr.it, {marco.picone, michele.amoretti, francesco.zanichelli}@unipr.it

***Abstract*—One envisioned distinctive feature of smart cities is the interconnection among mobile users and vehicles, to support the fulfillment of location-based services. This can be obtained with centralized architectures, and with all the problems of scalability and robustness that such a solution involves. On the other hand, a more complex but more reliable, completely distributed approach can overcome this kind of problems. In this paper, we present the Adaptive Distributed Geographic Table (ADGT), a peer-to-peer overlay scheme suitable for the development of location-based services. In particular, the ADGT allows to efficiently retrieve peers or resources, to broadcast messages within any geographical region, and to be automatically notified about any type of information around any geographical location, following the publish/subscribe model. What mainly differentiates the ADGT from the other solutions in literature is the adaptivity of the overlay's topology to peers' mobility. Actually, the ADGT has the capability to adapt the neighborhood of each mobile peer depending on speed and direction. We have evaluated the ADGT by simulating different scenarios, and the results show that it acts well, ensuring high quality of messages dissemination and low cost in terms of data usage.**

***Keywords*-peer-to-peer computing; mobile computing**

I. Introduction

In recent years, the idea of a networked city has become one of the most promising development models. For example, users and vehicles are becoming more and more connected due to the rapidly increasing demand the mobile devices market has experienced. As a consequence, there has been a growing attention to location-based services, which, for example, allow to locate people on a map, discover nearby social events or receive alerts (such as warnings of traffic jams along the user route). To design reliable services that take into account geographical locations, while preserving the privacy of the users and supporting millions of connected devices, can be a very challenging and demanding task.

In this context, peer-to-peer (P2P) solutions have been proposed. The decentralized nature of peer-to-peer networks increases the robustness, because it removes the single point of failure that can be inherent in a client-server based system [1], but, on the other hand, increases the architectural complexity of the system itself, in particular from the perspective of routing and resource discovery.

In this paper, we propose a novel extension to the P2P overlay scheme known as Distributed Geographic Table (DGT) [2], which supports the development of location-based services. The resulting Adaptive Distributed Geographic Table (ADGT), in comparison with other peer-to-peer overlay schemes with similar purposes, fulfills all the requirements that are essential to location-based services, such as geographic broadcast or the retrieval of resources near any geographical location.

We have extensively evaluated the ADGT with OSMobility [3], which allows to simulate the motion of different entities in realistic geographical spaces and, adopting the software-in-the-loop simulation methodology, allows to test deployment software on simulated devices, immersed in simulated environments. Our performance analysis shows that the ADGT is cost-effective in terms of data rate, and therefore highly suited to mobile devices.

In Section II, we analyze and compare the main peer-to-peer overlay schemes for location-based services which have been proposed in literature, and we draw attention to their weak points. In Section III, we describe the architecture of the ADGT. This is followed, in Section IV, by the explanation of the performed simulations for the evaluation of the ADGT overlay and the related results, in Section V. Finally, in Section VI, we present our conclusions and future work.

II. Related Work

Traditional peer-to-peer overlay schemes, such as Kademlia [4] or Chord [5], are not particularly suitable for location-based services, since they completely ignore geographical distances between peers of the network. The geographical distance, however, assumes an aspect of remarkable importance for location-based services, since they consume and produce geolocated information. For these reasons, the design of peer-to-peer overlay schemes proper to the development of location-based services is an extensively discussed issue in literature. Table I lists the main location-aware peer-to-peer overlays, sorted chronologically, and compares them based on the main features that location-based services should have.

In GeoPeer by Araújo and Rodrigues [6], network peers arrange themselves to form a Delaunay triangulation, with the addition of long range contacts. GeoPeer is capable of providing some of the fundamental operations of location-based services, such as geographical multicast and queries.

978-1-4799-5393-6/14 $31.00 © 2014 IEEE
DOI 10.1109/NCA.2014.35

The main weaknesses of this work stem from the fact that only stationary peers are taken into account — which is a serious deficiency, since location-based services have to cope with devices that are mobile by definition.

A second notable solution is Globase.KOM by Kovačević, Liebau and Steinmetz [7]. This solution adopts a superpeer-based overlay, where each superpeer manages one of the rectangular zones that constitute the hierarchical layers in which the geographic space is partitioned. Clearly, making use of a not completely distributed architecture implies a great effort in terms of complexity and a significantly poor scalability. Moreover, similarly to GeoPeer, this solution considers only stationary peers. Thus, it completely ignores the drastic increase of the maintenance overhead that would result by applying this protocol to mobile nodes.

Picone *et al.* describe a structured overlay scheme where each participant can efficiently retrieve peers or resources located near any chosen geographical location [2]. In such a system, called Distributed Geographic Table (DGT), the main provided service is to route requests to find available peers in a specific area. The DGT has been designed with primary emphasis on peers' mobility: each peer maintains a set of logic concentric circles, around its own geographical location, and keeps them constantly updated, in order to have the latest neighbors' location. Although the DGT takes into account the mobility of peers, it does not consider speed and direction of the motion. In effect, the mechanism by which each peer of this overlay scheme maintains updated its neighborhood completely ignores the different ways a peer can move.

Geodemlia, by Gross *et al.*, is a peer-to-peer overlay that allows users to search for location-based information around specific geographic locations [8]. The Geodemlia overlay scheme is inspired by Kademlia [4] and provides geographical methods for search and store. Also, like the DGT, it uses concentric circles to divide the geographical space. However, this overlay completely overlooks peers that change their geographical location.

Heep *et al.* developed Overdrive, a peer-to-peer overlay very similar to the DGT, but considering also speed and direction of peers even if, actually, it uses information of speed and direction only to reduce the number of sent messages and not to enhance the overlay itself [9]. Differently from the DGT, Overdrive adopts a recursive approach for routing instead of the iterative one inspired by Kademlia that the DGT adopts.

Examining these works in literature, we have noticed that none of them really adapt the topology of the network based on peers' movements. We consider it an essential feature proper to a location-based service. For these reasons, we have extended the peer-to-peer DGT overlay in order to readjust itself considering the speed and direction of the peers. Moreover, we have adopted a recursive algorithm for message routing. Finally, we have provided a publish/subscribe mechanism by which each peer can be automatically notified about any type of information around any geographical location, by simply exhibiting interest. All these features characterize the ADGT.

Table I: Comparison of location-aware peer-to-peer overlay schemes.

	Mobility of peers	Speed and direction of peers	Geographic broadcast	Completely peer-to-peer	Publish/Subscribe mechanism	Adaptive overlay scheme
GeoPeer			✓	✓		
Globase.KOM			✓			
DGT	✓			✓		
Geodemlia			✓	✓		
Overdrive	✓	✓	✓	✓		
ADGT	✓	✓	✓	✓	✓	✓

III. ADGT Architecture

With the objective to fulfill all the requirements of a location-aware peer-to-peer overlay scheme, we have decided to redesign and improve the DGT implementation. In particular, to make the ADGT take fully into account peers' mobility, we have formalized a new data structure for the management of neighborhood, different from all those available in literature. Such a new data structure is based on the idea that a peer should directly connect to those peers from which it is most likely to obtain satisfactory contents. Doing so, an adaptive topology that reacts to peers' movements is obtained. Furthermore, we have switched from the traditional Kademlia-like iterative routing algorithm to a more efficient recursive one. To achieve this, we have changed the operations of discovery of other peers in the network. Finally, we have introduced the mechanism of subscription to any type of information around any geographical location, making the ADGT a complete peer-to-peer overlay scheme for location-based services.

A. General Concepts

The original conceptual model that characterized the first version of the DGT has not particularly changed. We define $\mathcal{P}$ as the set of peers, each one described by a unique $id \in$

$\mathcal{I}$, where $\mathcal{I}$ is the space of identifiers, and an ordered pair $w = (latitude, longitude)$. If $\mathcal{W}$ is the space of geographic coordinates and $w \in \mathcal{W}$, then we can define a generic peer $p \in \mathcal{P}$ as $\langle id_p, w_p \rangle$, where $id_p \in \mathcal{I}$ and $w_p \in \mathcal{W}$.

In the ADGT overlay scheme, the distance between two peers is evaluated as the great-circle distance, which is the shortest distance between two points on the surface of a sphere, measured along the surface of the sphere itself: $d : \mathcal{W} \times \mathcal{W} \rightarrow \mathcal{R}$.

The neighborhood of a geographical location is defined as the set of peers that are geographically close to that specific location. In other words, those peers which are located inside a given surrounding region. Defining $\mathcal{A}$ as the set of geographic regions delimited by a closed curve and $GB_w \in \mathcal{A}$ as a region centered in the geographic location w, it is possible to define the neighborhood as $\mathcal{N} : \mathcal{W} \times \mathcal{A} \rightarrow 2^{\mathcal{P}}$, where $2^{\mathcal{P}}$ represents the set of all the possible connections between peers of the network.

In particular, the neighborhood of a specified geographical location $l \in \mathcal{W}$, center of the region $GB_l \in \mathcal{A}$, can be defined as $\mathcal{N} = \{p \in \mathcal{P} | w_p \subseteq GB_l\}$, where w_p is the geographical location of peer $p \in \mathcal{P}$.

B. Data Structures

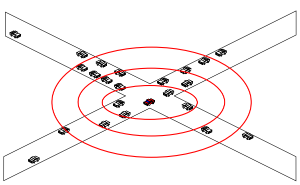

(a) Circular GeoBucket structure

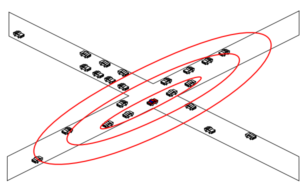

(b) Elliptical GeoBucket structure

Figure 1: Comparison between the traditional circular GeoBucket structure and the new adaptive elliptical GeoBucket structure.

In the first version of the DGT, and also in all the peer-to-peer overlay schemes discussed in Section II, every peer maintains its neighborhood within a circular region centered on its geographical location. Therefore, there is a lack of differentiation by the peers between the directions wherein the neighborhood is maintained. Imagine a vehicle that moves slowly along the streets of a small town: its neighborhood has to be maintained within an almost circular region because the vehicle is interested in receiving information from all the directions. But when the vehicle travels along a road at high speed, such as an highway, it is mainly concerned with the information that comes along its travel direction, rather than from other directions: its neighborhood has to be maintained mostly forward and backward, rather than laterally.

We believe that the real capability of a peer-to-peer overlay scheme that takes into account the peers' mobility is the adaptability of the neighborhood, depending on the peer's speed and direction.

Thus, in the ADGT, each peer stores a set of lists of neighbors, called GeoBucket, each list being sorted according to the distance from the center that the GeoBuckets have in common. Such lists are regularly updated in order to have the latest peers' positions. We have switched from the typical circular shape of DGT's GeoBuckets, to an elliptical one that is more general, where both the semi axes of the ellipse depend on the velocity of the peer, *i.e.*, depend both on the direction and speed of the peer. Formally, a GeoBucket structure is defined as a group of K different concentric ellipses, each having a different semi-major axis a_i and semi-minor axis b_i, with i integer $\in [1, K]$, for which

$$a_i = i \cdot t_i \cdot (1 + S \cdot \frac{v}{V_{max}})$$

$$b_i = i \cdot t_i \cdot \frac{V_{max}}{V_{max} + S \cdot v}$$

where v is the value of the speed of the peer, V_{max} is a constant value which represents the maximum speed, t_i is the thickness between each GeoBucket and S is a parameter that relates the speed of the peer to the shape of the GeoBucket. The idea behind such elliptical GeoBuckets is that the higher is the speed of the peer, the higher is the eccentricity of the ellipses: when the peer is stationary, its speed is 0 and the eccentricity of the ellipses is also 0, so the GeoBuckets are circular. On the other hand, when the peer reaches the maximum speed, the eccentricity of the ellipses is high, so the GeoBuckets have an elongated shape. Also, the direction of the semi-major axis coincides with the direction of movement of the peer.

More formally, we can say that the neighborhood of a peer $p \in \mathcal{P}$ is $\mathcal{N} = \{p \in \mathcal{P} | w_p \subseteq gb_p\}$, where $gb_p \in \mathcal{A}$ represents the region covered by the GeoBucket structure of the peer p. This region can be defined as $gb_p = \{w \in \mathcal{W} | d(w, f_1) + d(w, f_2) \leq 2 \cdot a\}$, where f_1 and f_2 are the two foci of the ellipse and a is the semi-major axis.

Such a kind of GeoBucket perfectly adapts to the different ways a peer can move, taking into account its direction and speed, in addition to its changes of position.

C. Routing Strategies

The original DGT implementation used an iterative algorithm, very similar to the one adopted by Kademlia, during the discovery process of new peers. We have decided to move to a recursive implementation, to improve the efficiency of the routing procedure [10].

In particular, we have defined a discovery message characterized by a geographical location around which to search for new neighbors. When a peer wants to discover new neighbors, it chooses the closest peer among those it knows and sends it a discovery message specifying the geographical location of interest. If a peer receives a discovery message, it replies to the requester with a set containing the β closest peers to the specified geographical location, and forwards the discovery request to the closest peer to the specified geographical location it knows. If the peer which receives the discovery message is the closest peer to the specified geographical location among its neighbors, the recursive discovery stops and the request is not forwarded.

D. Publish/Subscribe Mechanism

One of the most important features that a location-based service must provide is the possibility to notify the user about something of his/her interest near any geographical location. For example, the service should be able to notify the presence of friends nearby the user, or the presence of traffic jams along the route that the user is traveling. Nevertheless, all previous works have made the assumption that a peer has to know only its geographical neighborhood, *i.e.*, those peers whose geographical location is not too far from that of the peer itself. This means that the peer can be notified about any type of information only if it is located in its vicinity. Therefore, a user could not be able to monitor traffic conditions along a particular city thoroughfare, without necessarily being close to it.

In the ADGT, we have introduced a publish/subscribe mechanism through which peers can be notified about any information they are interested in around any geographical location [11]. Our significant improvement is apparently simple: while in the original DGT every peer had a unique GeoBucket structure, now it can have multiple GeoBucket structures, associated with different locations. In this way, it is possible, for example, to place a GeoBucket structure on a particular road junction to be automatically notified about any related warning.

Obviously, this brings out changes in the maintenance strategies of the neighborhood. Previously, it was possible to state that if peer B belongs to the neighborhood of peer A, then peer A must be in the neighborhood of B, that is, if a peer A adds B to its GeoBucket structure, then B must also do the same. However, with the introduction of the possibility of being able to keep more GeoBucket structures for each peer, such an assumption is no longer verified. It is thus required a mechanism to allow a peer to inform another peer about the interest in receiving location updates.

To do this, every peer maintains a limited set of references to the last peers which have expressed interest in receiving location updates. Periodically, each peer iterates over its set and transmits location update messages to interested peers. If a peer receives a location update message from another peer about which it is no longer interested in, all it needs to do is to inform the sender, that will remove it from its set.

Having multiple GeoBucket structures allows the peers to become aware of other peers near any geographical location, and so to query those peers about anything or to inform them about particular information they want to know, following the publish/subscribe model.

Actually, multiple GeoBucket structures make geographical broadcast possible. All that is needed is a GeoBucket structure that covers the area interested by the broadcast and disseminates messages to all the peers contained in the GeoBucket structure.

IV. Simulations

To evaluate the ADGT, we have used OSMobility [12], which allows to simulate the motion of different entities in realistic geographical spaces. The main advantage that derives from the use of OSMobility is the possibility to directly test deployment software on simulated devices, avoiding to write additional source code, different from that we actually use on mobile devices. In particular, we have provided OSMobility with a realistic mobility model, and a communication model based on ns-3's LENA LTE-EPC package [13]. A layered representation of the testbed is illustrated in Figure 2.

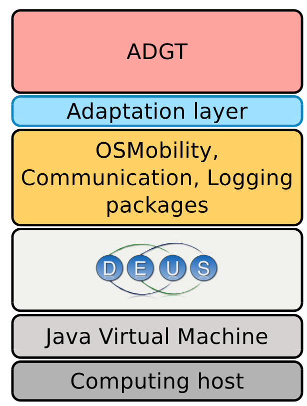

Figure 2: Layered representation of the testbed.

A. Mobility Model

OSMobility is integrated with OpenStreetMap, an open database which provides geographical data, and uses such data to compute trajectories. Thus, OSMobility allows to simulate vehicles running on highways or urban roads, pedestrian walking within limited traffic zones, bicycles moving on cycling lanes, and so on.

We have evaluated our overlay scheme in three completely different scenarios:

1) vehicles moving within our university campus;
2) many more vehicles moving within our city (Parma);
3) a lot of vehicles moving along the highways of our administrative region (Emilia-Romagna).

These scenarios differ not only for the number of vehicles, but also for the way in which they move. For example, within

(1) University campus (2) Parma (3) Highways in Emilia-Romagna

Figure 3: Scenarios adopted during simulations.

our university campus, the speed of vehicles is very low and there are very few roads. On the other hand, vehicles moving in Parma have a higher speed and can travel along many different paths. At last, vehicles that drive along the highways go at full speed and their direction does not change frequently.

In all these scenarios, we have tried to choose configurations as much realistic as possible. In particular, in the simulations within our university campus, each peer randomly chooses two points of interest among the ones we have defined (such as the entrances or the sport center), and then follows the shortest path calculated between them. Also the number of vehicles is not randomly chosen, but is based on realistic assumptions. With regard to the simulations within Parma, vehicles randomly select two geographical locations and move along the fastest path that joins them. In the third scenario, peers move along the highways between two toll booths, still randomly chosen. Vehicles are simulated so that the traffic density is comparable to the actual one, as reported by the web portal about mobility of Emilia-Romagna [14].

In addition, every vehicle adopts the Fluid Traffic Model as mobility model, where the speed is a monotonic decreasing function of the vehicle density [15].

B. Communication Model

Since the main application of the ADGT we have conceived is based on the use of mobile devices like smartphones, the modeling of communication delays plays a fundamental role. For this reason, to better characterize the communication among the ADGT peers in the urban environment, we have adopted the model illustrated by Amoretti *et al.*, using ns-3 with the Lena LTE-EPC package [16]. The latter provides the E-UTRA part of the Long Term Evolution (LTE) technology, dealing with PRY, MAC and Scheduler functionalities, and support for the LTE RLC and PDCP protocol, together with EPC data plane features, such as the S1-U interface and the SGW and PGW entities. Shortly, such an ns-3 package supports the detailed simulation of end-to-end IP connectivity over LTE-EPC.

C. Configuration Parameters

In order to evaluate the behavior and the performance of the ADGT overlay scheme throughout the simulations, we have identified the main configuration parameters that characterize it. In particular, they are

- β: maximum number of peers returned by a discovery response;
- ϵ: minimum distance (dimension: [km]) that must be travelled by a peer, before it notifies its neighbors about its change of location;
- K: the number of GeoBuckets that constitute each GeoBucket structure;
- t: the thickness (dimension: [km]) of each GeoBucket;
- L: the maximum number of peers contained in a GeoBucket;
- S: the value which relates the speed of the peer to the shape of the GeoBucket.

D. Evaluation Metrics

Both quantitative and qualitative performance evaluation metrics have been taken into account. Quantitative metrics show the cost of the ADGT in terms of transmitted data — this is an important aspect, since typically smartphones, and in general mobile devices, have traffic tariffs based on transmitted data amount. On the other hand, qualitative metrics show the behavior of the ADGT from a point of view of the quality of the service. So, they represent capabilities of the overlay scheme such as the warning responsiveness or the service accessibility.

In particular, we have defined

- Data Rate (DR), calculated as the average number of bits that are transmitted per unit of time by each peer (dimension: [kbit/s/peer]);
- Coverage Percentage (CP), calculated as the number of peers that have actually received a specific warning message, over those that should have received it;

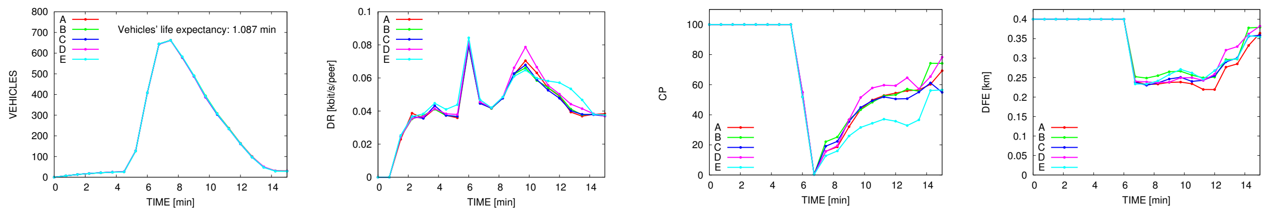

(a) Number of Vehicles (b) Data Rate (c) Coverage Percentage (d) Distance From Events

Figure 4: Simulation results of the scenario inside the university campus.

- Distance From Events (DFE), calculated as the average distance between the geographical location of peers which have not received yet a warning message and the geographical location of the associated event.

In order to compute CP and DFE values, a randomly chosen peer periodically produces a warning message related to its geographical location and disseminates it among its neighbors.

V. RESULTS

Taking into account our past experiments with the DGT, we have decided to set $\beta = 20$, $\epsilon = 1$, $K = 5$, $t = 0.4$, $L = 20$. Since the most characteristic parameter of the ADGT is S, we have defined five different configurations, denoted by the letters A, B, C, D and E, corresponding to increasing S values (as listed in Table II).

Table II: Configuration parameters of the simulations.

	A	B	C	D	E
S	0	1.5	2.5	5	10

The first simulation scenario has been defined to test the response of the ADGT in a situation of sudden high traffic density. We have modeled what typically happens in our university campus at the end of the day, *i.e.*, a significant number of vehicles leave the parking areas and move towards the points of interest. Simulation results presented in Figure 4 show that a too large value of S badly affects the qualitative performance (CP and DFE values) of the ADGT, because of the very small area of the university campus. The best S value, among those considered, is the one related to configuration D. Figure 4a shows the trend of the number of vehicles during the simulation. The warning message has been generated from a single randomly chosen peer, after 6 minutes from the beginning of the simulation. Figure 4b shows the DR of each peer. As the reader can see, the DR grows very slowly with the progress of the simulation, which is a very good result since it confirms the possibility to employ the ADGT on real mobile devices. Figure 4c shows the CP of the generated warning message, which grows from a very low value (when the warning message has been generated) to a high value, after a few minutes. Finally, one of the most interesting results is shown in Figure 4d, where the average of the DFE starts from a value of 0.2 km and rapidly arrives at the value of 0.4 km, which is the distance within the warning message has validity. In summary, the most noticeable result of this simulation is that a good quality of service can be guaranteed, despite the impossibility to set up a stable network topology, because of the very short routes the vehicles cover in the university campus.

On the other hand, the second scenario allows to evaluate the behavior of the ADGT in a situation of normal traffic condition. In particular, vehicles move on the roads of Parma, covering a period of five hours. During the first simulated hour, a vehicle is generated every 0.9 seconds. Every vehicle is removed, once it reaches its destination, and substituted by a new vehicle, placed at a randomly chosen location. In this way, after the first hour of simulation, there are constantly 4000 traveling vehicles. In addition, starting from the first hour and with a period of 30 minutes, a randomly chosen vehicle generates a warning message associated to the vehicle's geographical location, with a validity distance equals to 1.5 km. Figure 5 illustrates the simulation results of this second scenario. In particular, in Figure 5a the reader can see the number of vehicles. Figure 5b shows that also in this simulation the DR is very low. Compared with the evaluation of the DGT in a similar urban scenario, the data rate is reduced by one order of magnitude [2]. This remarkable result is mainly due to the recursive routing algorithm used by the ADGT. Figure 5c shows that the ADGT is highly responsive, considering that, if the generation of a new event implies a natural decrease of the CP value, this is quickly followed by a fast growth. In this case it is possible to observe another important result: the gap between the CP values of simulation A (where $S = 0$) and the others. This demonstrates that an adaptive overlay, *i.e.*, with $S > 0$, is much more efficient than a static overlay. Also the DFE trend in Figure 5d, confirms this feature, since, after each warning message generation, it remains almost constant. Figures 5e and 5f show more clearly that

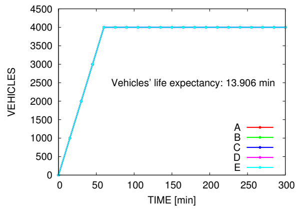

(a) Number of Vehicles

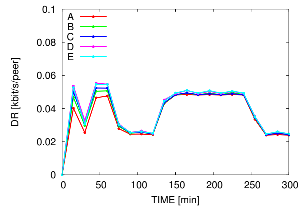

(b) Data Rate

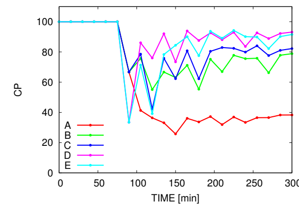

(c) Coverage Percentage

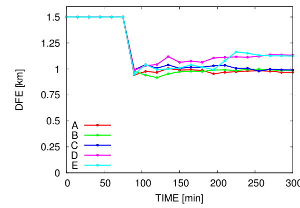

(d) Distance From Events

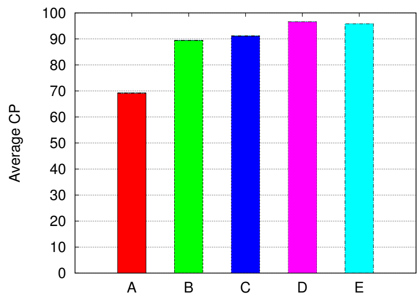

(e) Average Coverage Percentage

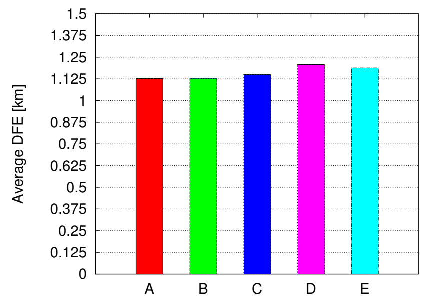

(f) Average Distance From Events

Figure 5: Simulation results of the scenario inside the city of Parma.

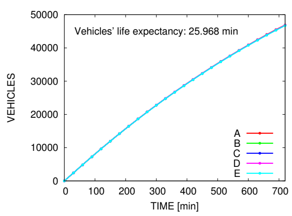

(a) Number of Vehicles

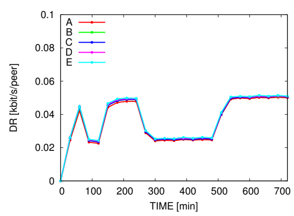

(b) Data Rate

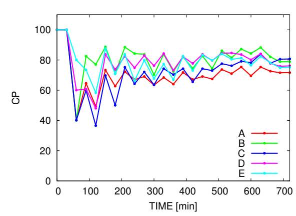

(c) Coverage Percentage

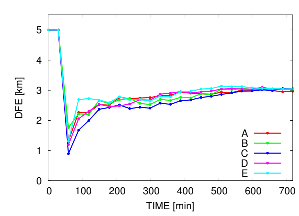

(d) Distance From Events

Figure 6: Simulation results of the scenario along the highways.

the performance quality increases from configuration A to configuration D. In particular, with respect to configuration A, configuration D provides a CP increment of about 27 percentage points. There is no clear advantage in using configuration E instead of D.

Finally, the last simulated scenario allows to evaluate the ADGT in a highly structured environment, where the direction of the peers does not change frequently. In particular, we have considered the highway segments that connect three cities in the Emilia-Romagna region, namely Piacenza, Parma and Reggio nell'Emilia. The simulation spans 12 hours and adopts an arrival period of the vehicles at the toll booths equals to 0.74 seconds. When a vehicle is generated, it travels towards another randomly chosen toll booth. Once the destination has been reached, the vehicle is removed from the simulation and not replaced. Furthermore, every hour a random vehicle generates a warning message with a range of validity equals to 5 km.

The results of this simulation are shown in Figure 6. In Figure 6a, the reader can see that the simulation ends before a constant number of vehicles is reached, due to the (realistic) mobility model parameters we have used. The Data Rate is presented in Figure 6b. Again, it is very low, compared with the current traffic data rates available to mobile devices. Finally, Figure 6c and Figure 6d show that, also in this scenario, the ADGT demonstrates a very good efficiency in the transmission of information.

VI. Conclusion

In this paper we have presented the ADGT, an adaptive peer-to-peer overlay scheme allowing the realization of location-based services for mobile peers, such as georeferenced information subscription and retrieval, as well as location-specific message broadcasting. Actually, the ADGT is the only peer-to-peer overlay scheme that effectively takes into account the speed and direction of the peers to adapt its topology. Moreover, with respect to the state of the art, the ADGT provides a higher number of location-based functionalities.

The ADGT has been evaluated using OSMobility, a discrete event simulator for mobile peers based on OpenStreetMap, which allows to test deployment software on simulated devices, immersed in simulated environments. Obtained results show that the ADGT is suitable for the realization of location-based services, due to high message dissemination quality and low costs in terms of data transmission.

Next works will investigate further methods for the adaptivity of the overlay scheme to peers' mobility, for example based on information of the road shape obtained from a map. Furthermore, aspects related to the preservation of the privacy of peers will be taken into account in the architecture of the overlay scheme.

REFERENCES

[1] E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, "A survey and comparison of peer-to-peer overlay network schemes," *Communications Surveys Tutorials, IEEE*, vol. 7, no. 2, pp. 72–93, Second 2005.

[2] M. Picone, M. Amoretti, and F. Zanichelli, "Proactive neighbor localization based on distributed geographic table," *Int.'l Journal of Pervasive Computing and Communications*, vol. 7, no. 3, pp. 240–263, 2005.

[3] DSG. (2013) Osmobility. [Online]. Available: https://code.google.com/p/osmobility/

[4] P. Maymounkov and D. Mazières, "Kademlia: A peer-to-peer information system based on the xor metric," in *Revised Papers from the First International Workshop on Peer-to-Peer Systems*, ser. IPTPS '01. London, UK, UK: Springer-Verlag, 2002, pp. 53–65.

[5] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," *SIGCOMM Comput. Commun. Rev.*, vol. 31, no. 4, pp. 149–160, August 2001.

[6] F. Araújo and L. Rodrigues, "Geopeer: a location-aware peer-to-peer system," in *3rd IEEE International Symposium on Network Computing and Applications (NCA 2004)*, August 2004, pp. 39–46.

[7] A. Kovacevic, N. Liebau, and R. Steinmetz, "Globase.kom - a p2p overlay for fully retrievable location-based search," in *7th IEEE International Conference on Peer-to-Peer Computing (P2P 2007)*, September 2007, pp. 87–96.

[8] C. Gross, D. Stingl, B. Richerzhagen, A. Hemel, R. Steinmetz, and D. Hausheer, "Geodemlia: A robust peer-to-peer overlay supporting location-based search," in *Peer-to-Peer Computing (P2P), 2012 IEEE 12th International Conference on*, September 2012, pp. 25–36.

[9] B. Heep, M. Florian, J. Volz, and I. Baumgart, "Overdrive an overlay-based geocast service for smart traffic applications," in *Wireless On-demand Network Systems and Services (WONS), 2013 10th Annual Conference on*, March 2013, pp. 1–8.

[10] B. Heep, "R/kademlia: Recursive and topology-aware overlay routing," in *Telecommunication Networks and Applications Conference (ATNAC), 2010 Australasian*, October 2010, pp. 102–107.

[11] P. T. Eugster, P. A. Felber, R. Guerraoui, and A.-M. Kermarrec, "The many faces of publish/subscribe," *ACM Comput. Surv.*, vol. 35, no. 2, pp. 114–131, June 2003.

[12] G. Brambilla, A. Grazioli, M. Picone, F. Zanichelli, and M. Amoretti, "A cost-effective approach to software-in-the-loop simulation of pervasive systems and applications," in *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2014 IEEE International Conference on*, March 2014, pp. 207–210.

[13] N. Baldo, M. Requena-Esteso, J. Nin-Guerreo, and M. Miozzo, "A new model for the simulation of the lte-epc data plane," in *5th ICST/ACM International Conference on Simulation Tools and Techniques (SIMU-Tools 2009)*, March 2012.

[14] Emilia-Romagna. (2011) Mobilità. [Online]. Available: http://mobilita.regione.emilia-romagna.it

[15] Transportation Research Board of the National Academies, "75 Years of the Fundamental Diagram for Traffic Flow Theory," in *Symposium on The Fundamental Diagram (Greenshields 75), Woods Hole, Massachusetts*, July 2008.

[16] M. Amoretti, M. Picone, F. Zanichelli, and G. Ferrari, "Simulating mobile and distributed systems with deus and ns-3," in *High Performance Computing and Simulation (HPCS), 2013 International Conference on*, July 2013, pp. 107–114.