Evaluating the Robustness of the DGT Approach for Smartphone-based Vehicular Networks

M. Picone, M. Amoretti, F. Zanichelli
Distributed Systems Group
Università degli Studi di Parma
Parma, Italy
Email: picone@ce.unipr.it

***Abstract*—To cope with millions users becoming increasingly connected to Internet on the move, location-based services may be better supported by decentralized infrastructures enabling improved scalability, access rate and resiliency. In this context, our previous work introduced the Distributed Geographical Table (DGT), an overlay scheme that builds and maintains virtual neighborhood relationships between peers with heterogeneous connections.**

In this paper we illustrate a smartphone-based vehicular network that uses the DGT, and we show its robustness against disconnections caused by the unavailability of connectivity/coverage (mostly occurring in rural areas), as well as overlay reconnections due to vertical handovers (mostly occurring in highly serviced urban areas). The simulative analysis of sample scenarios based on experimental measurements of coverage and connection throughput, carried out across/around Parma urban area, gives us valuable insights for defining an integrated model that will combine the DGT, user/vehicle mobility and connectivity/coverage types.

***Index Terms*—peer-to-peer; location-based services; mobile computing; vertical handover**

I. Introduction

Whereas car navigation has reached a certain level of maturity, traffic information services are still limited in coverage and grain. An open platform that collects, processes, analyzes, displays and stores real-time traffic status information *from a variety of independent data sources* has not been implemented and deployed, yet. The reason, from a technical viewpoint, is that a system able to manage a very high number of users would require abundant resources to provide an adequate QoS. If the whole information is managed by centralized servers, system scalability can be hardly guaranteed, unless practically unlimited computational and network resources can be deployed (as feasible for Google, TomTom and few other global service providers). Besides representing single points of failure, centralized servers would hardly scale for highly variable contexts, where up to millions mobile nodes change their location very quickly, in a wide and highly populated area. Indeed, search results would be incomplete or outdated with high probability.

To overcome these issues, a promising approach is to build the traffic information system on a smartphone-based vehicular network, where nodes interact and share information in decentralized, peer-to-peer fashion. In this context, our previous work introduced the Distributed Geographical Table (DGT), a P2P overlay scheme that allows nodes to maintain virtual neighborhood relationships with peers located around any physical location [17]. The DGT is maintained collectively by peer nodes sharing their location information with neighbors, thus potentially achieving improved performance. As all P2P approaches, the DGT allows to update and publish information directly, at low cost and with high scalability. It may also simplify the process of joining the virtual community and publishing new location-based services.

In this paper we illustrate a smartphone-based vehicular network that takes advantage of the DGT, and we show its fault-tolerance and robustness against disconnections caused by the unavailability of connectivity/coverage (mostly occurring in rural areas), as well as overlay reconnections due to vertical handovers (mostly occurring in highly serviced urban areas). The simulative analysis of sample scenarios based on experimental measurements of coverage and connection throughput, carried out across/around Parma urban area, gives us valuable insights for defining an integrated model that will combine the DGT, user/vehicle mobility and connectivity/coverage types.

The rest of the paper is organized as follows. Section II presents a survey and discussion of state-of-art research on location-based search in P2P networks and underlying space representation techniques. Section III recalls the main concepts and data structures of the Distributed Geographic Table, including the discovery and consistency preservation processes. Section IV describes the adopted model for vertical handover which considers the presence of overlapped areas with different network coverage. Section V presents significant simulation results showing the robustness of DGT-based localization considering different connectivity scenarios. Finally, we conclude with a discussion of our contributions and with some ideas for future work.

II. Related Work

The need for collaborative work with geospatial data has escalated in recent times also due to events such as terrorist activities and natural disasters. Distributed localization, a service that support *geo-collaboration* applications, is usually implemented by recursively dividing the 2D space into smaller areas in order to assign responsibilities for region of space to peers. The idea of hierarchical partitioning comes from the indexing of data structures for multidimensional data-sets such as R-tree [10] that is widely used in centralized databases. Globase.KOM (Geographical LOcation BAsed SEarch) [12] adopts a tree-based overlay where supernodes, *i.e.* the most powerful nodes, which tend to stay online for a long time, are responsible for indexing all nodes/services in one clearly defined geographical area. The normal (non-super) nodes in the network simply offer and consume services without having additional responsibilities. The idea is that the world projection is divided into disjoint, non-overlapping zones. Each zone is assigned to a supernode located inside the zone which keeps overlay/underlay contact addresses for all nodes in that zone. Supernodes form a tree where node A is called the parent of node B when Bs zone is inside As zone.

The main drawback of the hierarchical approach is that peers representing higher level zones may become bottlenecks for query routing, and possible points of failure for the whole system. For this reason, the most advanced distributed localization systems, instead of employing a number of centralized servers (either dedicated or selected among participating nodes) to carry the query load of the entire network, implement the peer-to-peer paradigm, for which every node in the network shares the load of indexing and searching data that refers to its area. GeoP2P [2] still performs a hierarchical partitioning of the 2D geographic space, but adopts a fully decentralized peer-to-peer overlay scheme, with overlay maintenance and query routing performed without super or special peers. Other peer-to-peer hierarchical schemes are HZSearch [24], DPTree [13], DiST [16]. These and other works [23], [14], [11] propose strategies for supporting complex queries over multi-dimensional data, such as select five available buildings closest to the airport.

None of these state-of-art solutions takes into account mobile peers, that would introduce several new challenges. Here we focus on the effects of *vertical handover*, that refers to the possibility that mobile terminals allow users to freely move and to switch connections among different access networks [22]. For example, a smartphone might be able to use both a high speed wireless LAN and a 2G/3G cellular technology for Internet access. The actual trend is to integrate complementary wireless technologies with overlapping coverage - beyond 3G (B3G) - to provide the expected ubiquitous coverage and to achieve the Always Best Connected (ABC) concept [9]. The ABC concept allows the user to use the best available access network and device at any point in time. Vertical handover between wireless technologies is being widely studied, with the objective of reducing delay [18], [7], [4], [5].

III. Proactive Localization with Distributed Geographic Table

By Distributed Geographic Table (DGT) we refer to an overlay scheme where each participant can efficiently retrieve node or resource information (data or services) located near any chosen geographic position. The responsibility for maintaining information about the position of active peers is distributed among nodes, in a way that a change in the set of participants causes a minimal amount of disruption. Compared to centralized solutions, the distributed approach is more scalable and robust against Denial Of Service (DOS) attacks and massive node disconnection. Moreover, it gives easier access to the market, since the platform does not need expensive infrastructure to efficiently work with a huge number of clients. Last but not least, the decentralized approach avoids the risks related to centralized control of information.

In our previous work [17], we defined a DGT-based strategy, inspired by Kademlia [15], for the localization of all nodes that are geographically (rather than virtually) close to the search initiator. Such a DGT-based localization scheme could be used to build a distributed vehicular application able to discover and inform drivers that are potentially interested in specific traffic messages or to data acquired by vehicle sensors.

In the following we briefly recall the basic concepts, data structures and protocols behind DGT-based localization strategy.

A. General concepts and data structures for DGT

We assume that every node knows its global position (GP), retrieved by means of a GPS receiver or other localization technologies. The same hypothesis has been made in the Mobile Millennium project [21] to implement a pilot traffic-monitoring system that uses GPS to gather traffic information, process it, and distribute it back to the phones in real time.

Each DGT node knows a set of real neighbors organized in a specific structure based on the distance with respect to the peer. The PeerDescriptor of a node contains contact information (IP address, port, proxy etc.), its geographic location and the list of message types for which it is interested. This structure is periodically updated by each single active peer with a specific discovery function described below.

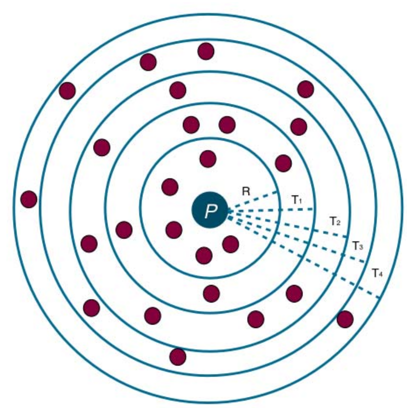

Fig. 1: In the DGT-based proactive localization protocol, each node's concentric data structures (GeoBuckets) are centered on the node itself.

Geographic position of the peer is the middle of K different concentric circles (each of them), each having a different (application-specific) radius R_i and thickness r_i, with i integer $\in [1, K]$, for which R_i is the sum of previous r_i from 1 to i. In other words, each peer stores a set of lists (called *GeoBuckets*) of neighbors, each list being sorted according to distance. These lists are regularly updated in order to have available the latest information about node positions.

B. DGT-based node discovery

When a new peer joins into peer-to-peer network, it receives an initial list of up to I references of other active nodes near its actual position. This list represents the initial knowledge of the new peer that can be used to start the first discovery procedure, in order to learn about other active users in the system. During each discovery process, a generic node P picks up the closest α known neighbors (if available), and sends simultaneous `FIND_NODE` requests *centered in its geographic location* (while, in general, the DGT allows for discovering nodes located around any point of the map). A node that receives this type of message extracts the closest known peers near the specified location from its GeoBuckets, and returns their references to the requesting node. After P receives answers from queried neighbors, it adds new received references to a temporary list and sorts it according to distance. Then P picks up again the closest α nodes from this temporary list that were not previously contacted, and repeats the same procedure with these other peers. If at the end of each iteration no new nodes are retrieved, P selects (if available) the closest not previously contacted K nodes and sends them the last turn of `FIND_NODE` requests. When finally P receives the answers from these neighbors, it adds new peers' references to the appropriate GeoBuckets. At the end of each run of the discovery procedure, P counts the number of new neighbors and upon its value it schedules the next activation of the discovery procedure. The general idea is that soon after the bootstrap or when neighbor peers show high dinamicity, the period of the discovery process may be very small (near T_{min}), increasing with time (towards T_{max}), when the knowledge becomes sufficiently distributed among active peers.

C. Consistency Preservation

Each active peer in the network may change its geographic position very frequently, periodically sending its updated GP to neighbors, in order to improve the accuracy of their knowledge base. As a measure to limit the overhead of this procedure, *i.e.* the message rate, each peer communicates its position updates to a neighbor only if the displacement is less or equal to ϵ (an application-specific distance). If during this message exchange a peer receives a node's update confirming that the new position is out of its area of interest, the neighbor's reference is removed from the appropriate GeoBucket and a `REMOVE` message is sent to the peer.

Peers are allowed to build and maintain an open-ended DGT, where each node has high knowledge of geographically close neighbors and a reduced view of the outer world. Moreover, peers are able to find new connected nodes that are entering the target area, incurring in limited computational and transmission costs.

D. Content Dissemination

The distributed knowledge about geographic neighborhood, periodically maintained by peers according to the DGT protocol, enables the dissemination of traffic information messages. In the vehicular context, each node can generate (manually or automatically) different types of messages to inform other participants about specific events such as traffic queues, dangerous situations or free spaces in a parking area. We recently started using an approach based on the publish/subscribe paradigm, that allows message distribution to all interested receivers in an area, by keeping messages alive in that zone for a specified period of time. Each message includes the *type* of notification (for example traffic events, sensor data. etc.), the *location* associated with the information, the *range* that represents the area surrounding the message source location that the notification should reach, the *time to live* of the event, and the *payload* containing - when necessary - additional and detailed information about the event. When a new message is generated, the publisher picks up from its GeoBuckets the closest known nodes within the notification's range that are interested in the particular information type (by reading their PeerDescriptor), and sends them the new message, trying to avoid duplications. When a notification is received, the node checks if it still matches the user interests or if it does not (in case of dynamic subscription) or if it is already known. When a peer receives the reference of a new node in its area of interest, it checks if in its knowledge base there are notifications not yet expired that may be useful for that peer. If the target peer has not yet been contacted for the same reason, the node sends the message. During this dissemination process, it is necessary to check if some messages have expired, and consequently to remove them and their references from the vehicle knowledge base, thus avoiding the distribution of an obsolete notification.

IV. Mobility Model with Vertical Handover

We want to evaluate the robustness of our DGT-based localization algorithm in a dynamic urban scenario with mobile devices, considering several (possibly overlapping) regions characterized by different types of network coverage (see figure 2). To this purpose, we use one model to describe the mobility of vehicles, and another model for vertical handover.

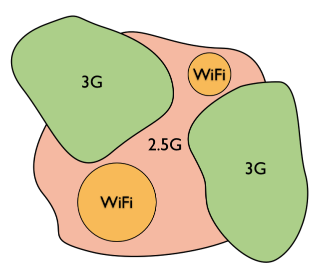

Fig. 2: Example of area with different, partially overlapping network coverages.

The realism of the mobility model is one of the fundamental elements in the performance evaluation of simulated V2V and V2I networking applications. Here we assume vehicles follow the Fluid Traffic Model (FTM) [19], that fits very well with scenarios characterized by different speed limits for each virtual path. The FTM describes speed as a monotonically decreasing function of vehicular density, forcing lower speed when the traffic congestion reaches a critical point. In our case the desired speed of a car moving along the points of a path p is evaluated according to the following equation:

$$v^{des} = \max\left[v_{min}, v^{p}_{max}\left(1 - \frac{k}{k_{jam}}\right)\right] \tag{1}$$

where v^{des} is the evaluated desired speed, v_{min} is the minimum car speed according to vehicle characteristics, v^{p}_{max} is the speed limit related to the path, k is the current density of the road given by $k = n/l$ (n represents the number of cars on the road and l its length), and k_{jam} is the vehicular density for which a traffic jam is detected.

Regarding vertical handover, different algorithms for reducing the delay and the packet loss rate have been proposed. The Always Best Connected (ABC) concept has been introduced in [9] and [28] to achieve seamless connectivity between WLAN and UMTS. The idea of using the vehicle speed as assessment criterion for vertical handover has been presented in [26] and [27]. The vertical handover algorithm we adopt in our analysis is based on the approach presented by Esposito *et al.* [25] that bases the handover decision both on vehicle speed and handover latency. Our version of the model considers a vehicle V moving with speed v^{des} in an environment characterized by several heterogeneous and overlapping access network at the same time. Defining SN (with bitrate B_{SN}) as the *serving network* to which the user is connected, CN (with bitrate B_{CN}) the *candidate network* and L as the handover latency (the time interval during which the peer does not receive any data due to the socket switching) the network switch is performed only if the time that the vehicle will spend in the area covered by the cell with higher bitrate (ΔT) is long enough to compensate for the data loss due to the switch overhead ($L < \Delta T$). Handover condition is defined as:

$$B_{CN} > \frac{B_{SN}}{1 - \frac{L}{\Delta T}} + \delta \tag{2}$$

where $\delta \in \mathbb{R}^+$ is an hysteresis factor used to avoid handover if the two competing networks have negligible bitrate difference. Since as previously described we are using a different mobility model and real vehicular city traces, instead of the Manhattan mobility model road composed of straight lanes used by the original authors, we need to redefine ΔT as follows:

$$\Delta T = \frac{\Delta x}{|v^{des}|} = \frac{R - d(V, CN)}{|v^{des}|} \tag{3}$$

where R is the radius candidate network station and $d(V, CN)$ is the geographic distance between the vehicle and the candidate network. The more the user is close to a cell site, the more he/she will stay within the coverage region of that cell site.

After each handover execution, the algorithm enters in idle mode for an inter-switch waiting period, T_w in order to avoid a high handover frequency that may happen when the vehicles travel on a border line between two different cells (ping-pong effects [29]) .

The types of wireless connection we take into account are 2.5G and 3G mobile telephone technologies, as well as WiFi and WiMAX. In the first case we consider that connectivity is provided through horizontal handover where available cells located in the area allow the communication, and do not involve changing the technology used to access the network at the data link layer.

V. Simulation Results

The analysis of the robustness of DGT-based localization, considering the vertical handover model defined in previous section has been performed by means of DEUS, a general-purpose tool for creating simulations of complex systems [1]. DEUS provides a Java API which allows to implement

- nodes (*i.e.* the parts which interact in a complex system, leading to emergent behaviors.. humans, pets, cells, robots, intelligent agents, etc.);
- events (e.g. node births/deaths, interactions among nodes, interactions with the environment, logs, etc.);
- processes (stochastic or deterministic, they regulate the timeliness of events).

Once simulation classes have been implemented, the dynamics of their instances in a specific simulation can be defined by means of a XML document (using the DEUS XML Schema), or a visual editor that generates the XML schema.

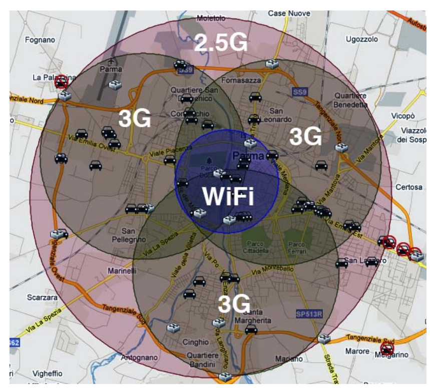

Fig. 3: Simulated network regions in Parma urban area.

We have simulated the DGT overlay in a 10km^2-squared area centered on the city of Parma, taking into consideration hundreds of vehicles which move over a set of road paths generated by means of the Google Maps API. Each simulated vehicle selects a random path in the set and starts moving over it according to the FTM mobility model until it reaches its first destination. From there the vehicle departs again towards a new destination, and so on and so forth until the simulation completes. Using the features provided by the Google API, we have created a simple HTML&Javascript control page that allows us to monitor the temporal progression of the simulated system with the possibility to select any node and visualize its neighborhood (some videos are available at [6]). We have considered a DGT overlay where available nodes have $K = 3$ different GeoBuckets, with a thickness of 1Km and a dynamic discovery period ranging from 1.5 min to 6 min, depending on the number of discovered nodes (if the latter decreases, then the period increases). Simulations cover ten hours of system life (10000 virtual time units) and have been averaged over several execution runs with different seeds. Additional parameter values are $T_w = 20[s]$ to reduce pingpong effects and $\delta = 1[Mbps]$ as hysteresis value to avoid handover if the two competing networks have negligible data rate difference. Road accident events are scheduled during the simulation according to a Poisson process with mean inter-arrival value of 1000 VTs. These and other events are sensed by vehicles and disseminated over the DGT through different message types.

In the evaluation, we have considered the following performance metrics:

- PMN (adimensional): Percentage of Missing Nodes in the GeoBuckets of a peer, with respect to those actually present in the area.
- PMN_{GB0} (adimensional): global distribution of PMN in the inner GeoBucket.
- r_m (dimension: [Kbyte/min]): Average DGT message rate - per peer, per minute.
- $RT(x)$ (dimension: [s]): Reconnection Time, *i.e.* the average time required by a temporary disconnected peer to recover the knowledge of its neighborhood and minimize the PMN value under $x\%$.
- $PacketLoss/min$: average number of packets that fail to reach the destination per minute per peer. It takes in account both DGT and content dissemination packets.
- $\%Coverage$: Estimated coverage percentage of traffic information messages at a certain time of the simulation. It is evaluated as the number of peers that actually received a specific message over those that should have it.

With the aim of improving the accuracy and the realism of simulated models, we performed multiple field measurements using Android smart-phones (HTC Desire and Samsung Galaxy) on a vehicle moving along several Parma streets. In this way, we obtained experimental data about:

- *Uplink and Downlink Rates*: Real-world communication rates for uplink and downlink channels.
- *Cell Tower Information*: Information about cell towers located in the area of interest, such as geographic location, provider, connection type, measured distance to cell tower and RSSI value.

All measurements have been carried out with different smartphones and SIM modules of three italian providers (TIM, 3 ITA, Vodafone). Tower locations were used to build a map of available cell towers (the overall coverage is schematically shown in figure 3), each one characterized by a specific connectivity type and a coverage area with a $1.5Km$ radius. A WiFi region with a radius of $2.0Km$ is also located in the city center (which likewise approximates the coverage that is actually available in Parma center) to provide higher data rates

#	Type	Uplink min-max [Kbit/s]	Downlink min;max [Kbit/s]
1	2.5G	30-90	60-170
2	3G	35-1150	91-2650
3	WiFI	100-2000	2000-10000

TABLE I: Types and performance of the available network regions in the simulated urban area.

	Disc.	2.5G	3G	WiFi
Disc.	0 s	1 s	1 s	5 s
2.5G	1 s	0 s	0 s	5 s
3G	1 s	1 s	0 s	5 s
WiFi	5 s	5 s	5 s	0 s

TABLE II: Vertical handover latency timetable.

where the density of peers is very high and consequently larger messages are exchanged among nodes to share the information about neighborhood. For each type of connectivity table I reports the ranges of data rates experimentally obtained on the field and thus the limits for the values considered by the simulation.

A first simulative analysis has been carried out to evaluate the robustness of the DGT overlay with respect to vertical handover in five scenarios with different network coverage, considering the vertical handover latencies reported in Table II. Simulated coverage distributions are summarized in Table III, starting from a fully operational Scenario 1 and proceeding with a progressive decrease of available connection types for WiFi and 3G networks.

Graph (c) in figure 4 reveals how significantly the presence of a WiFi region in the first scenario and the related expensive vertical handover effect influences the reconnection time needed by a peer to recover the PNM under the 10% of missing nodes. In scenarios where there is no WiFi area the required period is quite smaller. This behavior affects only marginally the number of packet lost per minute (b) and the global PMN distribution (a) - the latter results slightly higher in the first scenario compared to the others. Furthermore, the percentage of missing node in the inner GeoBucket (b) remains really low in all simulated scenarios, allowing for a high coverage of traffic information disseminated among nodes. Those results are a consequence of the efficiency and robustness of the DGT approach that allows to quickly identify new available nodes close to peer's geographic location, by means of periodic discovery and maintenance procedures, and at the same time to detect disconnected nodes.

A second simulative analysis aimed at measuring the robustness of the DGT overlay with respect to increasing values

Scenario	2.5G Regions	3G Regions	WiFi Regions
1	100%	100%	100%
2	100%	75%	0%
3	100%	50%	0%
4	100%	25%	0%
5	100%	0%	0%

TABLE III: Simulated scenarios with different connectivity coverage of the urban area.

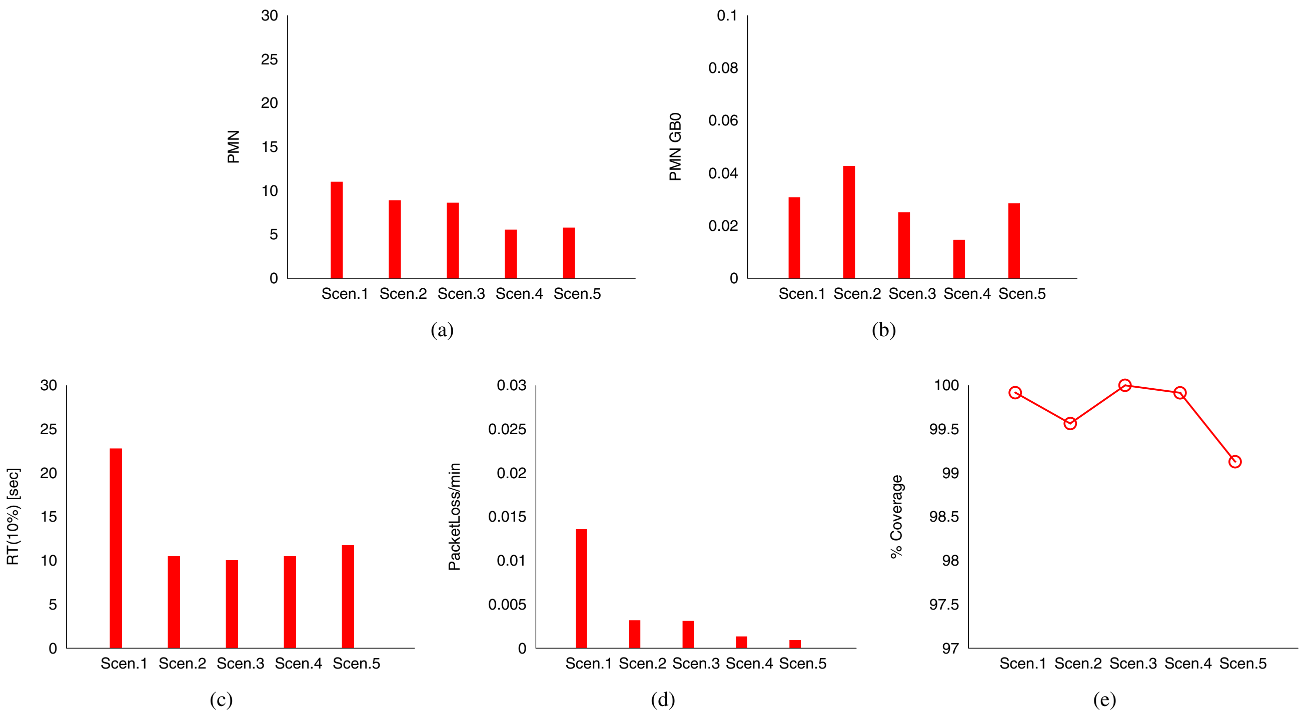

Fig. 4: Results related to different simulated scenarios. (a) PMN, (b) PMN in GB0, (c) RT(10%), (d) PacketLoss / min, (e) % Coverage of traffic information messages

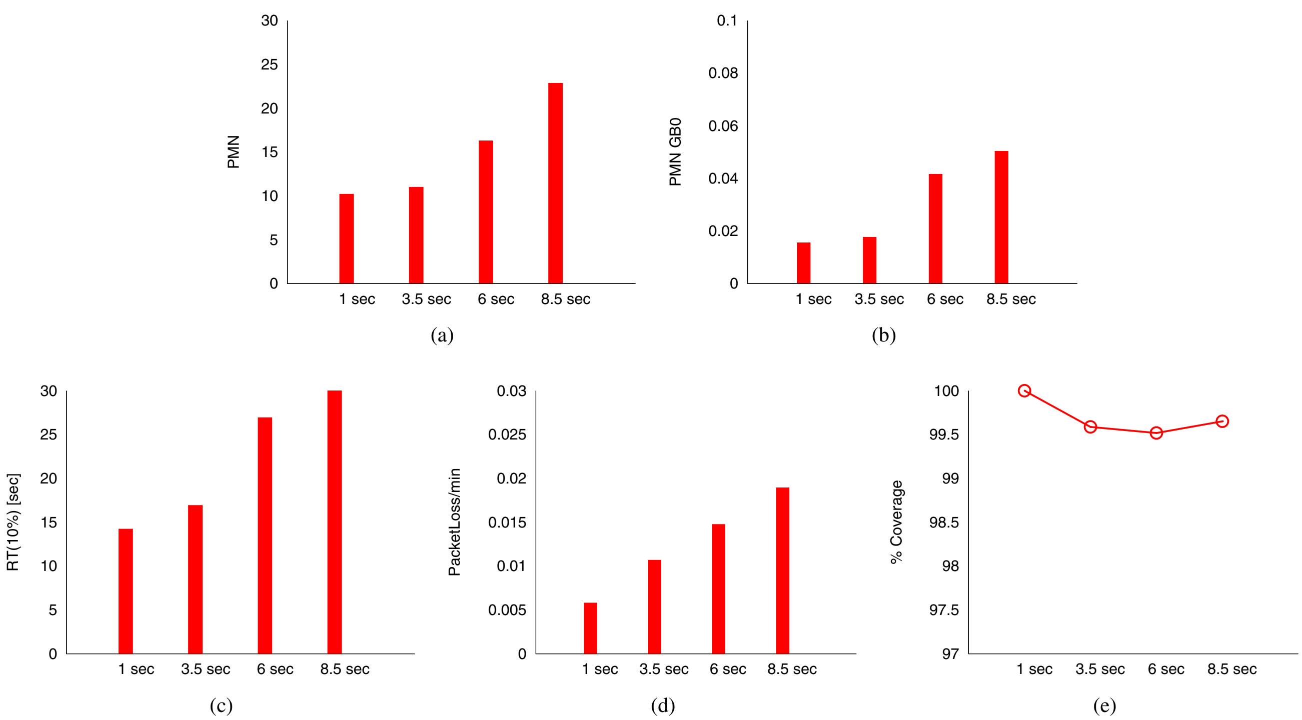

Fig. 5: Results related to different latency values. (a) PMN, (b) PMN in GB0, (c) RT(10%), (d) PacketLoss / min, (e) % Coverage of traffic information messages

of vertical handover WiFi latency. We considered a rather pessimistic range ($L \in [1s; 8.5s]$) of latency values (if compared to the ones used in other papers, *e.g.* [25]), in order to heavily test our approach, also taking into account that the latency is constant for all the duration of the simulation and for all peers. In terms of network coverage the simulated scenario is again Scenario 1, where all types of connectivity are available at the same time in the area of interest. As expected, an increased value for the WiFi vertical handover latency heavily affects RT values (as shown in figure 5(c)), and consequently the global percentage of missing nodes (figure 5(a)), that proportionally grows with the latency. The same behavior can be observed for the number of lost packets per minute (figure 5(d)) that however remains globally small. As presented in the analysis related to the variation of the connectivity coverage, one of the main important metrics for evaluating the robustness of the DGT approach is the PMN evaluated in the first GeoBucket(s). In fact, a high knowledge in the inner container and a gradually reduced value in the others mean that in any case the peer can perform successfully the discovery procedure, keeping the neighborhood updated. Most importantly, the peer can properly disseminate traffic information messages to its neighbors. Results presented in figure 5(b) and (c) confirm how the design of this peer-to-peer inter-vehicular DGT network is robust also in presence of a high values of latency, being able to correctly deliver messages with a percentage always higher than 99%.

VI. CONCLUSION

In this paper we have presented a preliminary investigation on the robustness of our DGT-based localization protocol to vertical handover in highly serviced urban areas. We have illustrated some significant simulative scenarios whose results evidence the independence of the percentage of missing nodes in the inner GeoBucket from peer disconnections due to vertical handovers as well as the short time required subsequently to recover knowledge about most neighbors consequently allowing to correctly deliver

Presented results are relevant also to show how a traffic information system and/or an inter-vehicular network based on mobile devices and a peer-to-peer approach is feasible, and how in the near future those kind of systems could be really and massively be used by end users.

We plan to extend the current model defining a new one that will combine the DGT overlay characteristics, user/vehicle mobility and connectivity/coverage type and we are also working on a real implementation of our DGT overlay, to be initially tested in a controlled environment, then in a real urban scenario. To this purpose, we use Sip2Peer [20], a SIP-based API for robust connection and communication among peers.

REFERENCES

[1] M. Amoretti, M. Agosti, F. Zanichelli, *DEUS: a Discrete Event Universal Simulator*, Proc. of the 2nd ICST/ACM International Conference on Simulation Tools and Techniques (SIMUTools 2009), Roma, Italy, March 2009.

[2] S. Asaduzzaman, G.v. Bochmann, *GeoP2P: an Adaptive and Fault-tolerant Peer-to-peer Overlay for Location Based Search*, CoRR 2009.

[3] A. Balasubramanian, B. N. Levine, A. Venkataramani, *DTN Routing as a Resource Allocation Problem*, in Proc. of SIGCOMM07, Kyoto, Japan, August 2007.

[4] S. Busanelli, M. Martaló, G. Ferrari, G. Spigoni, N. Iotti, *Experimental investigation of vertical handover algorithms between WiFi and UMTS networks*, Proc. of the International Conference on Future Generation Communication Networks (FGCN 2010), Jeju Island, Korea, December 2010.

[5] J. Choi, H. Lee, *Supporting Handover in an IEEE 802.11p-Based Wireless Access System*, Proc. of VANET10, September 24, 2010, Chicago, Illinois, USA.

[6] Distributed Systems Group, *DGT project home page*, http://dsg.ce.unipr.it/dgt

[7] A. Durantini, M. Petracca, *Performance Comparison of Vertical Handover Strategies for PSDR Heterogeneous Network*, IEEE Wireless Communications, Vol.15, No.3, June 2008.

[8] Google, *Google Latitude*, http://www.google.com/latitude

[9] E. Gustafsson, A. Jonsson, *Always Best Connected*, IEEE Wireless Communications, vol. 10, pp. 49-55, 2003.

[10] A. Guttman, *R-Trees: A Dynamic Index Structure for Spatial Searching*, in SIGMOD '84, pages 47-57, June 1984.

[11] A. Harwood, E.Tanin, *Hashing Spatial Content over Peer-to-Peer Networks*, in ATNAC, p.5, 2003.

[12] A. Kovacevic, N. Liebau, R. Steinmetz, *Globase.KOM - A P2P Overlay for Fully Retrievable Location-Based Search* , Proc. of the Seventh IEEE International Conference on Peer-to-Peer Computing, September 2007.

[13] M. Li, W. Lee, A. Sivasubramaniam, *DPTree: A Balanced Tree Based Indexing Framework for Peer-to-Peer Systems*, in 14th IEEE ICNP, pp.12-21, November 2006.

[14] B. Liu, W. C. Lee, D. L. Lee, *Supporting Complex Multi-Dimensional Queries in P2P Systems*, June 2005.

[15] P. Maymounkov, D. Mazieres, *Kademlia: A Peer-to-peer Information System Based on the XOR Metric*, First International Workshop on Peer-to-peer Systems, Cambridge, MA, USA, March 2002.

[16] B.Namand, A.Sussman, *DiST: fully decentralized indexing for querying distributed multidimensional datasets*, in 20th IPDPS, p.10, April 2006.

[17] M. Picone, M. Amoretti, F. Zanichelli, *Proactive Neighbor Localization Based on Distributed Geographic Table*, 8th International Conference on Advances in Mobile Computing & Multimedia (MoMM2010), Paris, France, November 2010.

[18] S. Salsano, L. Veltri, G. Martiniello, A. Polidoro, *Seamless vertical handover of VoIP calls based on SIP Session Border Controllers*, in IEEE International Conference on Communications (ICC '06), Istanbul, Turkey, June 2006.

[19] I. Seskar, S. Marie, J. Holtzman, J. Wasserman, *Rate of location area updates in cellular systems*, IEEE Vehicular Technology Conference (VTC'92), Denver, Colorado, USA, May 1992.

[20] Distributed Systems Group, *Sip2Peer*, http://code.google.com/p/sip2peer

[21] Mobile Millenium Project, http://traffic.berkeley.edu/

[22] E. Stevens-Navarro, U. Pineda-Rico, and J. Acosta-Elas, *Vertical Handover in beyond Third Generation (B3G) Wireless Networks*, International Journal of Future Generation Communication and Networking, vol. 1, no. 1, pp. 51-58, December 2008.

[23] E. Tanin, A. Harwood, H. Samet, *Using a Distributed Quadtree Index in Peer-to-peer Networks*, The VLDB Journal, 16(2):165-178, 2007.

[24] D. A. Tran, T. Nguyen, *Hierarchical Multidimensional Search in Peer-to-peer Networks*, Computer Communications, Vol.31, pp.346-357, 2008.

[25] F. Esposito, A. Vegni, I. Matta, A. Neri *On modeling speed-based vertical handovers in vehicular networks*, IEEE Global Telecommunications Conference (GLOBECOM10), Miami, Fla, USA, 2010.

[26] Z. Yan, H. Zhou, H. Zhang, and S. Zhang, *Speed-Based Probability-Driven Seamless Handover Scheme between WLAN and UMTS*, in MSN, Wuhan, China, Dec. 10-12, 2008.

[27] D. Kwak, J. Mo, and M. Kang, *Investigation of handoffs for IEEE 802.11 networks in vehicular environment*, in ICUFN, Hong Kong, June 7-9, 2009, pp. 8994.

[28] J.A. Olivera et al., *VANBA: a simple handover mechanism for transparent, always-on V2V communications*, in Proc. on IEEE 69th Vehicular Technology Conference (VTC2009-Spring), April 26-29, 2009.

[29] S. Balasubramaniam and J. Indulska, *Vertical Handover Supporting Pervasive Computing in Future Wireless Networks*, Computer Communication Journal, vol. 27, no. 8, pp. 708719, 2003.