Computer Communications 78 (2016) 28–44

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Mobility and connectivity in highway vehicular networks: A case study in Madrid

Marco Gramaglia [a,b,*], Oscar Trullols-Cruces [c], Diala Naboulsi [d], Marco Fiore [e,f], Maria Calderon [b]

[a] IMDEA Networks, 28918 Madrid, Spain
[b] UC3M, 28911 Madrid, Spain
[c] UPC, 08034 Barcelona, Spain
[d] INSA Lyon and Inria, 69621 Villeurbanne, France
[e] CNR-IEIIT, 10129 Torino, Italy
[f] Inria, 69621 Villeurbanne, France

A R T I C L E I N F O

Article history:
Received 11 August 2015
Revised 22 October 2015
Accepted 31 October 2015
Available online 2 December 2015

Keywords:
Vehicular networks
Highway traffic
Synthetic traces
vehicle-to-vehicle communication
Complex networks

A B S T R A C T

The performance of protocols and architectures for upcoming vehicular networks is commonly investigated by means of computer simulations, due to the excessive cost and complexity of large-scale experiments. Dependable and reproducible simulations are thus paramount to a proper evaluation of vehicular networking solutions. Yet, we lack today a reference dataset of vehicular mobility scenarios that are realistic, publicly available, heterogeneous, and that can be used for networking simulations straightaway. In this paper, we contribute to the endeavor of developing such a reference dataset, and present original synthetic traces that are generated from high-resolution real-world traffic counts. They describe road traffic in quasi-stationary state on three highways near Madrid, Spain, for different time-spans of several working days. To assess the potential impact of the traces on networking studies, we carry out a comprehensive analysis of the vehicular network topology they yield. Our results highlight the significant variability of the vehicular connectivity over time and space, and its invariant correlation with the vehicular density. We also underpin the dramatic influence of the communication range on the network fragmentation, availability, and stability, in all of the scenarios we consider.

1. Introduction

A key enabling technology of future Intelligent Transportation Systems (ITS), vehicle-to-vehicle (V2V) communication is envisioned to interconnect vehicles into distributed, self-organized networks. The latter are expected to complement today's mobile access architecture, and support services such as cooperative awareness, collision avoidance, or data dissemination.

The emergence of large-scale vehicular networks requires that a large fraction of vehicles is equipped with dedicated radio interfaces. Such a pervasive deployment of V2V communication is closer than one would imagine: standards for V2V communication, such as IEEE 802.11–2012 [1], IEEE 1609 [2],OSI CALM-M5 [3] and ETSI ITS-G5 [4] are now finalized, and regulators in the USA plan to enforce V2V radio interfaces on all new vehicles by 2017 [5]. Early large-scale field tests are also in progress, e.g., within the sim^{TD} project in Germany, or the Ann Arbor Safety Pilot in Michigan, USA.

These notwithstanding, experimental trials of vehicular networking solutions remain an exception, due to their costs and complexity. The vast majority of applications, protocols and architectures for upcoming vehicular networks is evaluated via computer simulation. The dependability of results is then conditional on the level of realism of the models assumed, and the representation of the mobility of individual vehicles is often the single feature that introduces the largest bias [6].

For that reason, during the past decade, significant efforts have been made to gather real-world road traffic data [7,8], develop effective tools for the simulation of vehicular movement [9–12], and generate realistic synthetic mobility traces [13–15]. Still, a reference set of realistic, publicly shared, heterogeneous road traffic scenarios for networking simulation is not yet available. This situation, originated by a manifest scarcity of mobility traces featuring the required level of realism and spatiotemporal granularity, is raising questions on the dependability and reproducibility of research results [16]. Within such a context, this paper puts forward several major contributions, as follows.

* Corresponding author at: Marco Gramaglia, IMDEA Networks Institute Avenida Mar Mediterraneo 22, 28918, Leganes, Madrid, Spain.

E-mail address: marco.gramaglia@imdea.org, mgramagl@it.uc3m.es (M. Gramaglia).

http://dx.doi.org/10.1016/j.comcom.2015.10.014
0140-3664/

(a) Highway locations (b) M30 - Aerial view (c) M40 - Aerial view (d) A6 - Aerial view

Fig. 1. (a) Geographical location of the measurement points on the three highways considered in our study, near Madrid, Spain: M30 (A), M40 (B) and A6 (C). (b),(c),(d) Close-by views of measurement points on M30, M40 and A6.

First, we take a step forward in the direction of dependable and reproducible vehicular networking research, by providing to the community multiple novel realistic highway traffic traces for network simulation. The traces are based on real-world traffic count measurements that feature an unprecedented level of detail, and are representative of heterogeneous motorway segments and road traffic conditions, as discussed in Section 2.

Second, we outline a detailed methodology to generate synthetic mobility traces of unidirectional highway traffic starting from road traffic counts. The traces model road traffic in quasi-stationary conditions, where macroscopic features such as the average vehicular density, speed, and out-flow observed on each highway lane are invariant over the full span of the simulated road segment. To that end, we leverage inherent properties of the real-world data for the per-vehicle calibration of well-known car-following and lane-changing microscopic models. Details are provided in Section 3.

Third, we characterize the vehicular network connectivity resulting from the proposed synthetic traces. To that end, we perform a network protocol-independent study, by adopting an instantaneous topology model, as discussed in Section 4. We investigate the impact of a wide range of parameters, including time (i.e., hour of the day, day of the week), highway settings (i.e., number of lanes, speed limits), road traffic conditions (i.e., free flow or synchronized traffic), and V2V communication range. Our results, presented in Section 5, underscore, in all of the scenarios we considered, the following properties: (i) the dramatic impact that relatively small communication range variations have on the network structure; (ii) the prevalent role of the vehicular density in driving network connectivity via three-phase dynamics; (iii) the limited availability and stability of long-range multi-hop vehicular networks, (iv) the fact that the highway vehicular network is difficult to navigate.

Finally, a comparative review of the related literature is provided in Section 6, before we draw conclusions in Section 7.

2. Source measurement data

The synthetic traces we present in this paper are based on empirical data that comes from real-world measurements carried out in the region of Madrid, Spain. The data, kindly provided to us by the Spanish office for the traffic management (Dirección General de Tráfico, DGT) and the Madrid City Council, details the vehicular traffic conditions on the following three arterial highways.

M30. With an average distance of 5.17 km from the city center, M30 is the inner part of the Madrid city beltway system, which also comprises the outermost M40 and M50. The data employed in this study comes from measurements along the northbound direction, close to the junction with the A-2 Motorway and marked as A in Fig. 1a. There, M30 features 4 lanes in the main carriageway, as it can be observed in the aerial view of Fig. 1b. The speed limit along M30 is 90 km/h.

M40. Motorway M40 is a part of the intermediate layer of the Madrid city beltway system. It has an average distance of 10.7 km from the city center, and traverses both the most peripheral areas of the municipality and several surrounding minor cities. The measurement point, marked as B in Fig. 1a, is at the 12.7-km milepost, where M40 traverses the suburb of San Blas and the town of Coslada. The measures cover the southbound carriageway, in Fig. 1c, which includes 3 lanes with a speed limit of 100 km/h.

A6. Autovía A6 is a motorway that connects the city of A Coruña to the city of Madrid. A6 enters the urban area from the northwest, collecting the traffic demand of the conurbation built along it. The data collection point is placed around the 11-km milepost in the Madrid direction, depicted with a C in Fig. 1a, where A6 features 3 lanes, as per Fig. 1d. The speed limit is 120 km/h.

2.1. Collecting fine-grained traffic count data

The sensors deployed on the three highways are induction loops, i.e., loops of wires buried under the concrete layer and creating a magnetic field. When a vehicle transits on the vertical axis of the loop, it induces a variation in the magnetic field. If two loops are placed close to each other, other metrics, e.g., the vehicle speed and length, can be also determined.

Usually, these devices are programmed to supply coarse-grained data, since public transportation authorities are generally interested in aggregate measures on, e.g., the number of vehicles transiting on a road, their average speed, or the percentage of heavy vehicles [1], so as to detect major alterations of traffic conditions [17,18]. The loops used in this paper are normally configured to supply data averaged over 60 s, but their setup was changed specifically for our study, so as to provide fine-grained information on each transiting vehicle.

Not only the level of detail, but also the timing and duration of the measurements are critical aspects of the data collection. Indeed, vehicular traffic presents significant daily variability, and rush hours yield diverse traffic conditions than off-peak hours, especially on main arterial roads like those we consider. In order to capture such temporal heterogeneity, and compatibly with the limitations imposed by the dedicated setup needed at the induction loops, we collected the following datasets.

One day-long dataset, collected on M30 during 24 h of a typical weekday in May 2010. This dataset features variable conditions, from very sparse traffic at night to heavy congestion during the morning rush hours. It thus provides a rather complete view of the possible traffic scenarios met on a real-world highway.

Sixteen 30-min datasets, collected on M40 and A6. These datasets were recorded on multiple weekdays of May 2010, during the morning traffic peak (from 8:30 a.m. to 9 a.m.), and during off-peak hours (from 11.30 a.m. to 12 p.m.). The rationale for these shorter datasets is that they allow us to generalize our study, by investigating the effects induced by different roads (e.g., number of lanes, speed limits and proximity to the city center) and different weekdays.

[1] As an example, Dirección General de Tráfico provides elaborations of the traffic count data via the Infocar web service at http://infocar.dgt.es, with visualizations of the historical aggregate data at the observation points.