

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

The Bologna ringway dataset: Improving road network conversion in SUMO and validating urban mobility via navigation services

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

The Bologna ringway dataset: Improving road network conversion in SUMO and validating urban mobility via navigation services / Bedogni, Luca; Gramaglia, Marco; Vesco, Andrea; Fiore, Marco; Härri, Jérôme; Ferrero, Francesco. - In: IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY. - ISSN 0018-9545. - STAMPA. - 64:12(2015), pp. 7234936.5464-7234936.5476. [10.1109/TVT.2015.2475608]

Availability:

This version is available at: <https://hdl.handle.net/11585/561777> since: 2016-11-23

Published:

DOI: <http://doi.org/10.1109/TVT.2015.2475608>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Bedogni, Luca; Gramaglia, Marco; Vesco, Andrea; Fiore, Marco; Härri, Jérôme; Ferrero, Francesco. (2015). The Bologna Ringway Dataset: Improving Road Network Conversion in SUMO and Validating Urban. IEEE Transactions on Vehicular Technology, 64(12), pp. 5464-5476.

The final published version is available online at: <http://dx.doi.org/10.1109/TVT.2015.2475608>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

The Bologna Ringway Dataset: Improving Road Network Conversion in SUMO and Validating Urban Mobility via Navigation Services

Luca Bedogni, Marco Gramaglia, Andrea Vesco, *Member, IEEE*, Marco Fiore, *Member, IEEE*, Jérôme Härrı, *Member, IEEE*, Francesco Ferrero

Abstract—The current lack of reference datasets of road traffic mobility for network simulation jeopardizes the reliability and reproducibility of vehicular networking research. We contribute to the ongoing effort to develop dependable and publicly available mobility traces, by (i) implementing an original version of the SUMO road network conversion tool that allows importing OSM data in a neat, automated fashion, (ii) generating an original dataset of road traffic in Bologna, Italy, (iii) providing a novel validation methodology that builds on open data provided by navigation service, which we leverage to assess the quality of the proposed Bologna dataset. These three contributions are expected to benefit the whole research community, since they do not only provide a new, ready-to-use, realistic dataset that can be input to network simulators, but also ease the generation and validation of further vehicular mobility traces for networking research.

I. INTRODUCTION

Connected vehicles are at the center of a telecommunication revolution. On the one hand, long-foreseen vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication are expected to hit the market soon: dedicated frequency bands have been allocated in Europe, US and Japan [1], standards have been finalized [2]–[5], and political actors and regulators are already crafting proposals enforcing all new vehicles to embed V2V radio interfaces [6]. On the other hand, the connected vehicle vision is a reality: cellular connectivity is an increasingly popular feature in cars, even budget models [7].

Many networking solutions and mobile services building on both V2V/V2I communication and connected vehicles are expected to operate at very large scales. This is the case of, e.g., opportunistic offloading [8], floating car data management [9], delay-tolerant message dissemination [10], or multi-hop routing [11],

The research leading to these results was supported by the People Programme (Marie Curie Actions) of the European Unions Seventh Framework Programme (FP7/2007-2013) under REA grant agreement n.630211 ReFlex. This work was also supported by the COLOMBO Project, which is funded, in part, by the European Commission under the Seventh Framework Program, FP7-ICT-2011-8, Grant 318622, as well as the support of its industrial members, namely, BMW Group Research and Technology, IABG, Monaco Telecom, Orange, SAP, SFR, ST Microelectronics, and Symantec. L. Bedogni (luca.bedogni4@unibo.it) is with the Department of Computer Science and Engineering, University of Bologna, Bologna, Italy. M. Gramaglia (marco.gramaglia@ieiit.cnr.it) is with CNR-IEIIT, Torino, Italy. M. Vesco and F. Ferrero ({vesco,ferrero}@ismb.it) are with Smart City Strategic Program, Istituto Superiore Mario Boella, Torino, Italy. M. Fiore (marco.fiore@ieiit.cnr.it) is with CNR-IEIIT, Torino, Italy and Inria, Villeurbanne, France. J. Härrı (jerome.haerri@eurecom.fr) is with EURECOM, Sophia Antipolis, France.

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Despite the deployment of a few first large-scale field tests [12], [13], the cost and complexity of experimental studies forces the performance evaluation of the aforementioned solutions to rely on simulation studies. Within this context, a reliable modeling of road traffic has been repeatedly proven to be paramount to the dependability of the results [14]–[16]. However, the availability of vehicular mobility datasets that are designed for networking purposes and that are publicly available is today limited to a few scenarios [17]–[19]. In turn, this inconvenient situation risks to severely limit the reliability and reproducibility of networking results [20].

The main reasons for the lack of reference mobility datasets are that (i) the tools to generate realistic road traffic are complex to configure and operate, and (ii) real-world input data to be fed to such tools is hard to retrieve. Concerning the first point, specialized software, such as the widely adopted SUMO [22], typically have a steep learning curve, especially when it comes to large-scale scenarios that involve high levels of detail and realism. This induces vehicular networking researchers to favor simple mobility scenarios in their simulations, at the cost of reliability. As for the second item above, a dependable road traffic simulation requires precise information on a number of aspects, including, e.g., the road network (down to per-lane settings or traffic light parametrization) and the trajectories of vehicles over it. These data are not easy to collect and correctly implement in the mobility simulation.

In this paper, we propose a three-fold contribution to ease the development of a reference set of realistic, ready-to-use datasets of vehicular mobility for networking studies.

- We improve the usability of software for the generation of road traffic datasets. Specifically, our focus is on a key phase of the generation process, i.e., the conversion of road networks from map databases to mobility simulators. We consider state-of-the-art open-source tools, i.e., SUMO [22] and OpenStreetMap [21], and propose some modifications to the NETCONVERT tool that allows linking the two.
- We present an original methodology to validate synthetic road traffic datasets using easily accessible real-world information provided by navigation services. Our approach provides a rapid way to verify if fundamental features of the synthetic mobility, i.e., routing paths and travel times, are realistic enough. We demonstrate a specific implementation of the methodology, based on Google Maps APIs [27].
- We apply our original road network converter and vali-

dation methodology to a use case scenario, i.e., a vast portion of the city of Bologna, Italy. The generation process employs realistic information on the macroscopic traffic flows in the region, collected by the iTetris European collaborative project [26]. The resulting dataset represents the mobility of more than 22,000 vehicles in a 25-km² area during one traffic peak hour in a typical morning, and its features are shown to match well those inferred from navigation services. Coherently with the objective of building a reference set of mobility scenarios for vehicular networking purposes, we make this novel dataset open to the research community¹.

These contributions foster the development and verification of new vehicular mobility datasets for networking research, by providing novel tools, guidelines and methodologies that can be largely reused. We provide a proof-of-concept of their practical usage, by generating and analyzing an original dataset, which is made available to the research community.

The document is organized as follows. After an introductory discussion of the generation process of synthetic road traffic, in Sec. II, we present our novel road network converter and compare it to legacy solutions in Sec. III. We apply our converter to a specific use case, i.e., the Bologna Ringway scenario, in Sec. IV. Using a realistic traffic-peak travel demand allows generating an original dataset. We validate such a scenario by leveraging a reusable methodology based on navigation service data, in Sec. V. Finally, Sec. VI presents related works, and Sec. VII concludes the paper.

II. SYNTHETIC VEHICULAR MOBILITY: A PRIMER

The generation of dependable datasets of road traffic for networking studies requires bringing together several different tools and data sources.

Road network. The first critical component is a comprehensive representation of the road network, which is not limited to the street layout, but includes detailed information on the nature of road segments (e.g., number of lanes per direction, or speed limits) and road junctions (e.g., presence and synchronization of traffic lights, or priorities among incoming roads). **Microscopic mobility models.** The second element is represented by validated models of the driver's behavior, which describe, at a microscopic level, his decisions in terms of acceleration, deceleration, lane changing, and, generally, his reactions to the surrounding environment.

Macroscopic traffic flows. The third major ingredient is a faithful model of the road traffic flows in the target region, so that the resulting mobility is realistic not only at a microscopic level but also from a macroscopic viewpoint. That implies gathering information on the traffic demand (i.e., the start time, the origin and the destination of each vehicle's trip, stored in an Origin-Destination, or O-D, matrix), as well as running appropriate traffic assignment techniques that determine the routes followed by each driver to reach his destination.

Validation. It is important that the synthetic mobility is validated by comparison against real-world data, including, e.g., traffic counts collected by induction loops, traffic volumes

computed through cameras, or traffic flows from surveys or floating car data samples.

All such components need to be integrated via a federated software, which generates a trace of the movement of each vehicle in the considered scenario with a high spatial (order of meters) and temporal (order of hundreds of ms) accuracy. We refer the interested reader to the detailed descriptions in [17], [18] for a thorough presentation of the generation process of road traffic datasets for networking research.

In this work we focus on the current state-of-the-art federating tool, i.e., the Simulator of Urban Mobility (SUMO) [22]. The specific problems we tackle the first component above, i.e., the representation of the road network, and the last, i.e., the validation of the simulated road traffic.

III. IMPROVED ROAD NETWORK CONVERSION IN SUMO

SUMO can import road network information from a number of sources, the most popular being, by far, OpenStreetMap (OSM) [21]. The OSM project provides crowdsourced maps of cities worldwide contributed by a vast user community, which are commonly regarded as the highest quality road data publicly available to date.

However, OSM is not designed for vehicular mobility simulation, and the conversion process is not trivial, as discussed in Sec. III-A. Since an incorrect representation of the road layout and features risks to impair the whole generation process, a-posteriori manual corrections of conversion errors are often required. The latter are, however, extremely time-consuming, and, as a consequence, they are just overlooked most of the times. In order to address this issue, we propose modifications to *NETCONVERT*, which make OSM-SUMO linkage significantly easier, in Sec. III-B, and we provide guidelines on common errors within OSM data that can impair the conversion process, in Sec. III-C.

A. Limitations in the legacy conversion

As many other microscopic traffic simulators, SUMO models the road network topology as a graph composed by sets of edges (i.e., the road segments) and nodes (i.e., the intersections). Both edges and nodes feature complex internal structures: edges are structured into lanes, whereas nodes define all possible connections among the lanes they interconnect, as well as the parametrization of traffic lights that possibly regulate them. Manual configuration of such composite representations is unfeasible at scales larger than a few intersections, which calls for tools that can perform an automated configuration of the road network. In the case of SUMO, the reference tool is *NETCONVERT*, which builds on data from several cartographic sources, among which OSM.

Specifically, OSM data contains information that can be used by *NETCONVERT* to determine the internal structure of edges, i.e., the number, direction, and speed limit of lanes. However, data from OSM – or any other common map database – does not include the information required to build the internal structure of nodes, i.e., intersections, in the SUMO representation. There, *NETCONVERT* implements a dedicated

¹Available at <http://www.cs.unibo.it/projects/bolognaringway/>.

algorithm to figure out the correct connections among the lanes entering and exiting a road intersection.

The algorithm works on a per-intersection basis; its baseline operation is described in [23], while refinements are presented in [24]. In brief, it first orders all incoming and outgoing edges according to their priority, inferred using some of the OSM lane attributes, such as the road type or the speed limit. Then, it determines edge-to-edge, lane-to-edge and lane-to-lane associations, in this order, by serving the highest priority edges first. Finally, it rechecks the internal structure for flaws such as disconnected lanes. The same ordered list of edges is used also to assign cycles to traffic lights governing the intersection. We refer the interested reader to [23], [24] for a comprehensive description of the algorithm.

The solution sketched above aims at accommodating the conversion of many different types of intersection using a same heuristic, and, conceptually speaking, does so in a sensible way. However, after some hands-on experiments on different OSM datasets, we found that the `NETCONVERT` output is affected by inaccuracies that can drastically reduce the capacity of the road network. Namely, we spotted these common conversion errors in the scenarios we considered:

- The legacy `NETCONVERT` algorithm tends to disable many lane-to-lane connections in presence of multi-lane edges. For example, in a four-way intersection joining roads that feature three-lanes each, it may happen that only the middle lane of each road is allowed to go straight, whereas the outermost ones are used for left and right turns only. An example can be observed in Fig. 3a, where only one of the two lanes in the road heading northwest is allowed to go straight at the top part of the intersection. While this can be acceptable in some cases, it is a configuration that generally reduces the straight traffic throughput at the intersection.
- Turning is only allowed on a lane-to-lane basis, i.e., incoming traffic from one left-turn lane is not allowed to flow into multiple lanes of the outgoing road on the left. Again, an example is shown in Fig. 3a, where the rightmost lane on the east-heading road is connected to just one lane of the two-lane road heading northwest. Enabling lane-to-multilane turning would not generate conflicts in traffic flows, and would significantly increase the turning traffic throughput at intersections.
- Left turns are allowed almost everywhere, even in multiple-lane boulevards. When not supported by a coherent traffic light synchronization, this tends to reduce the intersection out-flow, up to the point where deadlocks (i.e., situations where vehicles are stuck inside the intersection) appear. In fact, in many real-world situations, left turns are forbidden on main arteries, where drivers are forced to use service roads or perform 270° turns to go left. Left turn rules are in general difficult to figure out, and depend on the underneath traffic demand, as recommended, e.g., by the U.S. Highway Capacity Manual [25].
- Traffic light synchronization is generally more complex than needed. The dedicated `NETCONVERT` algorithm builds on the composite connection structure above,

which leads, at times, to long traffic light cycles that create significant queuing at intersections. An example is shown in Fig. 3b, where the green traffic light phase allows concurrent crossing by conflicting flows.

The overlapping of all these problems makes the SUMO road network prone to become congested. In some cases, it induces severe deadlocks that the simulator tries to solve by teleporting stuck vehicles, often when the traffic flow stability is already compromised. We remark that similar problems were encountered under different configurations of the current release of `NETCONVERT`, within SUMO version 0.22, as also proven by the results of our comparative analysis in Sec. IV-B.

B. `NETCONVERT` modifications

Our work focuses on improving the quality of the conversion of intersections and roundabouts, as well as redefines the traffic light synchronization logic, as detailed next.

1) *Connections*: Determining the rules according to which vehicles can traverse road junctions, such as intersections and roundabouts, plays a critical role in a smooth microscopic-level simulation of road traffic.

Fig. 1a outlines the notation that is at the base of the algorithm used by our original `NETCONVERT`, henceforth named *trigonometric*, to perform intersection conversion. For each incoming edge², we calculate the angle α between the edge itself and each of the outgoing edges. Angle α is used to infer if an outgoing edge is going straight ($-45^\circ \leq \alpha \leq 45^\circ$), turning right ($-135^\circ < \alpha < -45^\circ$) or left ($45^\circ \leq \alpha < 135^\circ$) with respect to the selected incoming edge. We then apply the following rules to create the internal connections:

- the rightmost lane of an incoming edge is connected to all lanes of an outgoing edge going right;
- each lane i of an incoming edge is connected to lane i of an outgoing straight edge. Let n_{in} and n_{out} be the number of lanes of the incoming and outgoing edges, respectively. If $n_{in} > n_{out}$ then the exceeding lanes of the incoming edge are not connected to any lane of the outgoing edge. This rule forces lane changing before entering the intersection, and avoids collisions during the simulation: the microscopic driver models in SUMO cannot handle situations where two vehicles enter together a same lane. Instead, when $n_{in} \leq n_{out}$, all lanes of the incoming edge are connected to those of the outgoing edge;
- the leftmost lane of an incoming edge is connected to all lanes of an outgoing edge going left.

Fig. 1b shows an example of left- and right-turn connections at an intersection with 5 edges. The bright red arrows between the lanes of the edges represent the enabled connections, i.e., the paths a vehicle can take to cross the intersection. When considering the incoming edge A, edge B has a relative angle ($-135^\circ < \alpha_{AB} \leq -45^\circ$) with respect to A, and edge D has a relative angle ($45^\circ \leq \alpha_{AD} < 135^\circ$) with respect to A. Consequently, the right turn is enabled from A's right lane to B and the left turn from A's left lane to

²An edge maps to an unidirectional road segment. A single bidirectional road segment is thus mapped into two distinct unidirectional edges.




Fig. 1: Trigonometric NETCONVERT. (a) Straight, left and right directions with respect to an incoming road. (b) Example of conversion for left and right connections at an intersection. (c) Example of conversion for straight connections at an intersection.

D. Fig. 1c shows the configuration of connections between incoming and outgoing edges along the same direction. When considering the incoming edge C, edge D has a relative angle ($-45^\circ \leq \alpha_{CD} \leq 45^\circ$), thus the two edges are interconnected lane by lane, as discussed above.

The trigonometric NETCONVERT also modifies the conversion logic in presence of roundabouts. Specifically, our proposed algorithm leverages the OSM data so as to flag all edges identified as roundabout. The flag simplifies the rules when creating connections between adjacent roundabout edges, and between roundabout and outgoing edges. The following rules apply:

- each lane i of an incoming edge, i.e., a non-roundabout edge, is connected to the lane i of the roundabout edge;
- each lane i of a roundabout edge is connected to the lane i of the next roundabout edge;
- the n_{out} rightmost lanes of the roundabout edge are connected to the rightmost lanes of the outgoing edge, where n_{out} is the number of lanes of the outgoing ones.

The highest priority, as defined by SUMO, is assigned to all roundabout edges. This assignment is important to preserve the left precedence rule governing ingress in the roundabout. In our experience, when the legacy NETCONVERT is used, the precedence rule is not always satisfied at simulation run-time.

2) *Traffic lights*: Automated traffic light configuration is a much desirable feature in the road conversion process. A reliable algorithm can dramatically reduce the manual configuration time of the simulation, and, at the same time, ensure that the in- and out-flow capacity of intersections is fully exploited.

The traffic light synchronization logic implemented in the trigonometric NETCONVERT builds on the identification of all incoming edges at an intersection, including their angle with respect to a common reference system, as depicted in Fig. 2. The rationale is that edges with an angle between 0° and 90° and between 180° and 270° should share the same traffic light phases. The edges with an angle between 90° and 180° and between 270° and 360° also share the same phases among them, but in opposition to the previous ones. In other words, this procedure aggregates the incoming edges into two clusters and configures coherent but opposite traffic




Fig. 2: Trigonometric NETCONVERT. Example of traffic light.

light phases between clusters. With reference to the example in Fig. 2, edges A and D belong to a first cluster whereas C and F belong to another.

We remark that some borderline scenarios can lead to synchronization errors. For example, an edge with angle equal to 91° and another one with angle equal to 179° are aggregated into the same cluster, even if they are almost perpendicular. These cases could be solved by increasing the number of clusters at the cost of adding traffic light phases. However, in our experience two clusters are enough to handle the vast majority of cases, without adding traffic light phases that would reduce the fraction of time dedicated to each phase, and thus slow down the vehicle flow.

3) *Practical examples*: Several representative examples of OSM-to-SUMO road network conversions are portrayed in Fig. 3. We compare the result obtained with the legacy NETCONVERT (top) and the trigonometric version (bottom).

Fig. 3a shows an example of an intersection merging a dual carriageway road with a single carriageway road, as converted by the legacy NETCONVERT. The conversion is flawed because vehicles incoming from the north direction cannot turn left on the avenue going southeast. This type of conversion error leads to increased route lengths, as vehicles must undergo




Fig. 3: Examples of road conversion from OSM to SUMO using the legacy (top) and trigonometric (bottom) NETCONVERT. (a,d) Dual carriageway intersection. (b,e) Multiple parallel carriageway intersection. (c,f) Five-entry, five-exit roundabout.

long detours to enter the avenue. Fig. 3d shows the conversion of the same intersection by the trigonometric NETCONVERT. It is worth noting that all turns are correctly represented in this case. Moreover, vehicles can now employ both lanes of the outgoing edge towards north, while one lane was enabled by the legacy NETCONVERT. These improvements lead to significantly increased in- and out-flow throughputs.

Fig. 3b shows an example of a more complex intersection converted by the legacy NETCONVERT. In this case, one dual carriageway road merges with a multiple parallel carriageway road. Multiple evident errors affect the converted road network in Fig. 3b. The road along the northwest-to-southeast heading has two lanes in each direction, yet in both directions vehicles are not allowed to move straight through the junction; instead, they have to undergo multiple lane changes, and, in the end, the capacity of the intersection to manage straight traffic is halved. Moreover, the legacy NETCONVERT leaves the possibility for drivers to turn on their left once they are in the middle of the intersection; at simulation run-time, this leads to deadlocks among vehicles occupying the center of the intersection. Finally, the traffic light synchronization determined by the legacy NETCONVERT allows vehicles traveling along perpendicular incoming roads to cross the junction at the same time (note the green lights in the picture), which results in congestion and accidents. Fig. 3e depicts the result of the conversion of the same intersection, performed with the trigonometric NETCONVERT. Here, all lanes are coherently connected, which enables the utilization of the full capacity of the intersection, and makes the traffic simulation more fluid

and realistic. Also, left turns from the center of the intersection are now prohibited³, which avoids deadlocks. Last but not least, Fig. 3e shows that the traffic lights are now properly synchronized in all directions.

Fig. 3c shows an example of a roundabout with five inbound and five outbound roads, converted from OSM to SUMO with the legacy NETCONVERT. A major error affects in this case the conversion: there are no connections between consecutive edges along the roundabout, which basically means that no route exists from any entry point to any exit points of the roundabout. In other words, the roundabout is excluded from all routes computed by SUMO at run-time, which may induce huge detours in vehicle trips, and a critical loss of capacity in the road network. Fig. 3f shows the conversion results with the trigonometric NETCONVERT. In this case connections within the roundabout are correctly modeled, and the roundabout becomes part of simulated road network.

As a final remark, we mention that we ran the trigonometric NETCONVERT on the SUMO scenario test suite, and, in most cases, we obtained identical results to those produced by the legacy conversion algorithm, as in Fig. 4a and Fig. 4b. The only differences emerged in a few cases where the trigonometric NETCONVERT enabled lane-to-multilane turns, as shown in Fig. 4c and Fig. 4d. As discussed in Sec. III-A, the additional connections created by our modified tool do not create conflicts, and increase the intersection capacity.

³The trigonometric NETCONVERT allows the user to enable or disable left turns on the avenues, through a specific command-line option.




Fig. 4: Examples of test suite road conversions using the legacy (a,c) and trigonometric (b,d) NETCONVERT.

C. OpenStreetMap modifications

Despite the improvements it brings over the legacy version, our proposed NETCONVERT cannot make up for issues that pertain to the original OSM data. The latter need to be fixed before the OSM-to-SUMO road network conversion.

In order to correctly model large-scale urban scenarios, the input OSM data should be as close as possible to the actual road infrastructure, so as to let the simulated road traffic flow smoothly. This is especially important when the macroscopic travel demand mimics that observed in the real world, and inaccuracies in the representation of the road network make it impossible to accommodate a realistic demand.

Among all possible inconsistencies between the modeled road infrastructure and the real one, two are the main aspects that need to be taken into account, based on our experience, as follows.

- **Correct number of lanes.** Several roads do not report, or wrongly report, the number of lanes the corresponding edges are composed of. The OSM format assigns a `highway` tag to all roads, along with a value indicating the road category, e.g., motorway, trunk, primary, etc. Depending on the category it belongs to, each road is assigned a default number of lanes. However, the default configuration is, most of the times, different from the real-world one: this leads to an unrealistic, and typically underestimated, road capacity that is not compatible with the actual travel demand. The `lanes` tag in OSM data overrides the default category-based number of lanes, and its correct setting is thus paramount to obtain a converted road layout in SUMO that is coherent with the real one.
- **One-way only roads.** Roads in OSM data are assigned, at times, a wrong value of the `oneway` tag that determines whether the road is one way or not. Such a misconfiguration prevents real-world routes from being used in the simulation, forcing vehicles to travel much longer paths to reach their destination. Depending on the `highway` value used for the road, the `oneway` tag might be not necessary. However, it has to be used every time the road presents differences with the standard behavior.

We thus strongly suggest that OSM data is carefully polished before the conversion is performed. Particular attention should be paid to fixing all instances of the two issues above.

IV. A USE CASE: THE BOLOGNA RINGWAY DATASET

We assess the quality of the novel OSM-to-SUMO road network conversion, by targeting one specific use case scenario, i.e., the Bologna Ringway dataset. Bologna is a middle-sized city of around 380,000 inhabitants, located in central-northern Italy. As shown in Fig. 5a, the urban area is composed by a limited-access downtown, surrounded by a ringway that features three lanes per direction; the peripheral regions of the city lay outside the ringway. The typical morning traffic consists in commuting drives from residential areas in the city outskirts, towards downtown, where offices and commercial activities are located. Commuters tend to drive around the city using the fast-transit ringway, and only enter the inner part of the city once close to their destination. Thus, most of the vehicular traffic flows along the ringway.

A. Simulation configuration

Our use case scenario encompasses the downtown region, the ringway around it, as well as part of the periphery, and it is aptly named Bologna Ringway. The layout is in Fig. 5a.

1) *Road network:* Information on the road infrastructure is extracted from up-to-date OSM data, downloaded in February 2015. We had to perform a number of fixes in the OSM data, based on the recommendations in Sec. III-C. In particular, we corrected critical errors on the number of lanes and bi-directionality of road segments along the main ringway, which completely impaired the simulation of vehicular mobility. Overall, we fixed 163 roads which did not report the correct number of lanes, 7 roads which were defined two times in the same area, and 6 roads that did not exist in the real network.

We employ the trigonometric NETCONVERT presented in Sec. III to convert OSM data in the Bologna Ringway scenario into a SUMO road network. For the sake of completeness, we test other conversion approaches as well, which allows appreciating the impact of the different conversions on the microscopic mobility simulation. We consider the following four configurations.

- **iTetris.** The iTetris dataset is the original road traffic trace developed in the collaborative research project of the same name [26]. The dataset is obtained by feeding the iTetris travel demand presented above, as well as a polished road network map derived from OSM data, to SUMO version 0.12, i.e., the current version during the project execution, in 2010. Interestingly, the dataset also features optimized, hand-configured traffic light synchronization that mimic that in the real world.
- **Legacy.** The Legacy dataset is obtained by combining the iTetris travel demand and a polished OSM data with the latest release of SUMO, i.e., version 0.22. Therefore, it employs the legacy NETCONVERT, and the automated traffic light synchronization algorithm present in that version. We consider two flavors: **A** uses basic conversion options, whereas **B** adds the recommended NETCONVERT options `--roundabouts.guess` and `--junctions.join`.
- **Trigonometric.** The Trigonometric dataset is the same as the Legacy one, but for the trigonometric NETCONVERT.

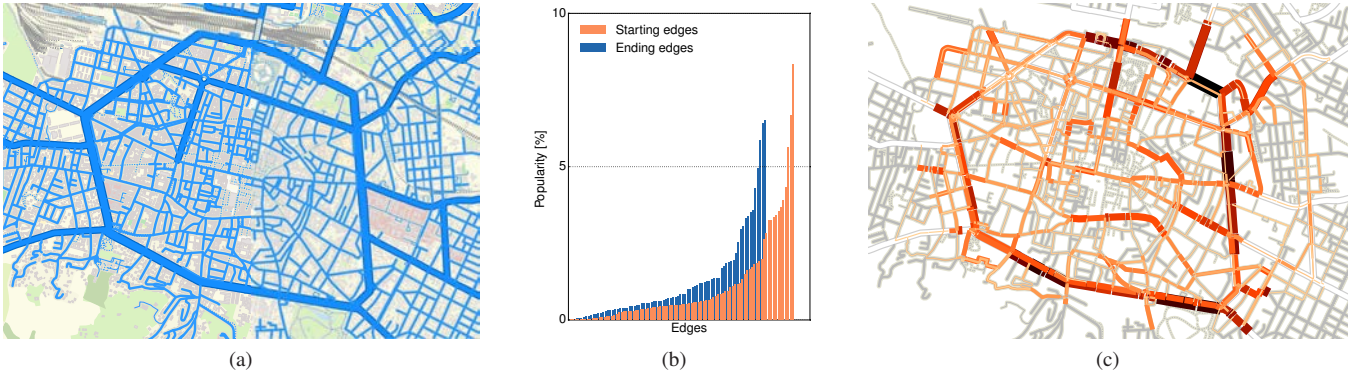


Fig. 5: Bologna Ringway. (a) Road network map. (b) Popularity of start and end edges. (c) Route distribution over the network.

It thus uses SUMO version 0.22, but it features the enhanced road network conversion and traffic light parametrization presented in Sec. III.

2) *Microscopic mobility models*: The microscopic-level models adopted are the default ones implemented by SUMO, i.e., Krauss' car-following model [28] and Krajzewicz's lane-changing model [29]. These control drivers' acceleration and overtaking decisions, respectively, by taking into account a number of factors (e.g., the distance to the leading vehicle, the traveling speed, the acceleration and deceleration profiles).

3) *Macroscopic traffic flows*: The travel demand we consider describes one hour of traffic in the Bologna conurbation, during the morning rush hour. The dataset was provided by the iTetris project [26], and represents the most challenging conditions for our performance evaluation: indeed, if the road network can support the 8 am traffic peak, it will also take in the travel demand observed during any other time of the day. The travel demand is structured as an Origin-Destination (O-D) matrix, converted into individual trips of vehicles: it thus describes the departure and arrival locations of each vehicle in the region during the simulated hour. A total of 22,213 individual trips is present in the dataset, starting at 93 different edges, and ending at 81 edges around the city.

Fig. 5b shows the distribution of trips over the starting and ending edges: we observe a limited number of edges that gather most of the vehicle departures and arrivals, with roughly the same four starting edges and five ending edges present in more than 20% of the total trips. The vast majority of edges are featured in a small percentage of trips, 2% or less each. Fig. 5c portrays instead the distribution of traveled routes over the street layout, as recoded in the synthetic trace. Thicker, darker road segments are part of the routes of a larger number of vehicles. We observe the importance of the ringway for traffic flows in the area, as the relative segments support most of the travel demand. The main entry points to downtown can also be spotted, as they feature significant traffic.

In order to perform a more general analysis, we generate mobility datasets under four different road traffic loads⁴: (i) the complete travel demand, i.e., 22,213 trips, which maps to the traffic peak hour between 8 am and 9 am; (ii) 90% of the travel demand, i.e., 19,940 trips, which can be mapped to

traffic between 7 am and 8 am; (iii) 80% of the travel demand, i.e., 17,806 trips, which can be mapped to traffic between 9 am and 10 am; (iv) 70% of the travel demand, i.e., 15,483 trips, which can be mapped to the hourly traffic from 10 am through 12 am. Considering different traffic loads allows simulating conditions other than the complete saturation observed during the morning rush hour.

In all settings, the final mobility is obtained by iterating Gawron's traffic assignment algorithm [30], until a user equilibrium is reached. Gawron's algorithm determines the least expensive (i.e., fastest) route for each vehicle, and then computes a cost for each road segment based on its occupancy level. At each iteration, it then moves part of the traffic to alternate, less congested paths, and recomputes road segment costs, until new iterations do not bring any advantage, i.e., the road capacity is exploited to a maximum.

B. Dataset analysis

A first representative metric of the synthetic mobility is the time series of the number of vehicles circulating in the simulated region, in Fig. 6. Each plot refers to a different road traffic load, expressed as a percentage of the iTetris travel demand. The curves in each plot refer instead to the simulation tool configurations, i.e., iTetris, Legacy A and B, and Trigonometric.

The number of vehicles always starts at zero, since the road network is empty at simulation startup. The simulation reaches a steady state, where vehicles are spread over the whole road network, after 15 minutes from the start of the first trip. Also, a peak is reached after one hour: then, vehicles stop being injected into the simulation, which runs until all cars already present in the road network reach their destination.

By comparing plots in Fig. 6a through Fig. 6d, we note that the number of running vehicles decreases as the travel demand is reduced: heavier traffic loads induce more severe congestion, which, in turn, increases travel times and forces a larger number of vehicles to co-exist in the road network.

When comparing instead the results obtained with different tool configurations, we observe that, in all travel demand settings, the original iTetris configuration yields a larger number of running vehicles than the Legacy and Trigonometric configurations. The enhancements introduced by SUMO version 0.22 over the version 0.12 originally used by iTetris allow

⁴The mapping between the load percentages and the day hours is inferred from non-public statistics provided by the transportation agency in Bologna.

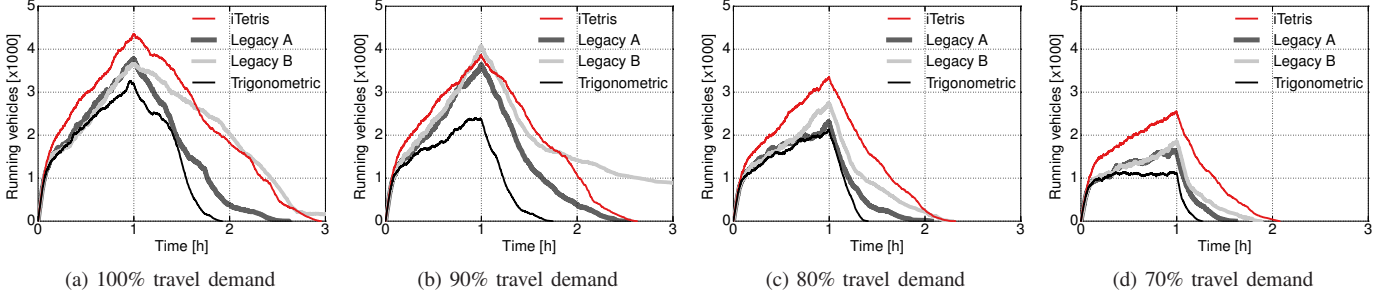


Fig. 6: Bologna Ringway. Concurrently traveling vehicles under different road traffic loads.

better accommodating the road traffic, reducing congestion and thus the number of concurrently traveling vehicles. A direct comparison between the Legacy and Trigonometric configurations highlights how the latter further reduces the number of vehicles present in the road network at the same time. The improvement is typically the most evident at the traffic peak occurring after around one hour of elapsed time, and it ranges between 10% and 40%, depending on the travel demand settings. It is also interesting that, under the Trigonometric configuration, simulations tend to complete much faster, again with gains ranging between 15% and 40%. The very long tails in the Legacy B case are due to a few intersections, unable to manage the traffic load and generating important queues that take a long time to be served.

In terms of road network conversion, these results demonstrate that the trigonometric `NETCONVERT` can bring significant advantages over the legacy approaches in the current SUMO release. Indeed, the SUMO representation of the road infrastructure provided by the trigonometric `NETCONVERT` makes traffic significantly smoother, which reduces congestion and allows discharging the road network much more fluidly once no new vehicles are injected.

Confirmations come from other road traffic performance metrics. Fig. 7 details the mean travel time of vehicles in the Bologna Ringway scenario, as recorded over simulation time. We observe that, in all cases, travel times tend to grow over time: this is coherent with the results in Fig. 6, showing that the road network becomes increasingly congested as simulation time elapses, and, thus, vehicles travel at lower speed and take more time to reach their destination. This effect becomes less pronounced as the travel demand is reduced, since traffic becomes then more fluid. The very low travel time recorded when close to the x-axis origin is a simulation artifact, due to the fact that the road network is initially empty, and thus only vehicles traveling over short distances are included in the computation. Again, we can remark that the simulation reaches a steady state after around 15 minutes, which corresponds to a change in the slope of all curves.

The comparison between the iTetris, Legacy and Trigonometric configurations underscores the advantage yielded by the latter, with average travel times that are reduced by up to two minutes. The gain is significant, considering the size of the scenario, and the fact that trips in the area typically last between five and ten minutes.

Fig. 8 portrays the mean time spent by vehicles in standstill

situations, which may be generated by queuing at intersections and ramps, or by generic congestion within the urban area. Different plots refer to varying road traffic loads. Clearly, it is desirable that such a waiting time is the lowest possible, since this implies a smoother road traffic; at the same time, the waiting time cannot be zero in a realistic scenario, since some queuing is unavoidable due to, e.g., traffic lights.

A cross-comparison of the four plots indicates that the waiting time is reduced along with the travel demand, which is intuitive, since we already observed that a lower demand implies less congestion in the road network. Also, both the Legacy and Trigonometric configurations outperform the original iTetris settings, attaining a significantly lower mean waiting time. Under heavy load, the iTetris configuration forces vehicles to spend, on average, 40% more time in standstill situations than what happens with the Legacy and Trigonometric configurations. When comparing the Legacy and Trigonometric cases, performance are identical in presence of heavy traffic conditions, in Fig. 8a. However, as the traffic load is reduced, some diversity emerges, with the Trigonometric setting accommodating the demand more smoothly than the Legacy configuration. The gain can reach peaks of 60%, in terms of waiting time reduction, when 70% of the original travel demand is considered.

Overall, we observe that, throughout all the results presented in this section, the trigonometric `NETCONVERT` presented in Sec. III produces a SUMO road network that is more efficient than that generated by the legacy `NETCONVERT`. The more consistent road infrastructure representation has an important impact on the simulated microscopic-level mobility, resulting in more fluid traffic and reduced queuing and congestion.

In particular, we remark that the Trigonometric configuration is the only capable to attain stable traffic conditions when considering 70% of the iTetris travel demand, in Fig. 6d, Fig. 7d, and Fig. 8d: there, the number of running vehicles is steady throughout the whole simulation, with nearly constant mean travel time and very low (10-20 seconds) average waiting times. This result is especially meaningful, since, as discussed in Sec. IV-A3, a 70% traffic load corresponds to the moderate traffic observed during the morning hours in the Bologna area we considered. Unlike what happens during the rush hours represented by the 80%-100% demands, where some congestion is unavoidable, the road network shall be able to accommodate the ordinary load observed during most of the day. In fact, the Trigonometric configuration is the only




Fig. 7: Bologna Ringway. Mean travel time under different road traffic loads.

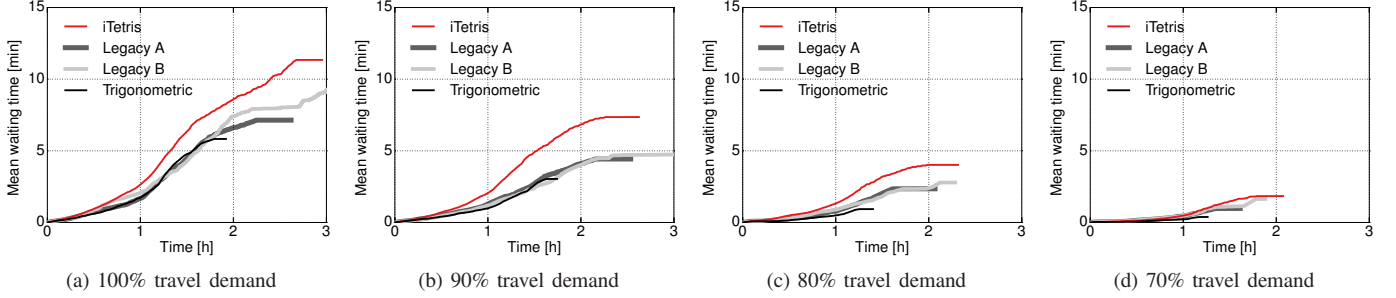


Fig. 8: Bologna Ringway. Mean waiting time under different road traffic loads.

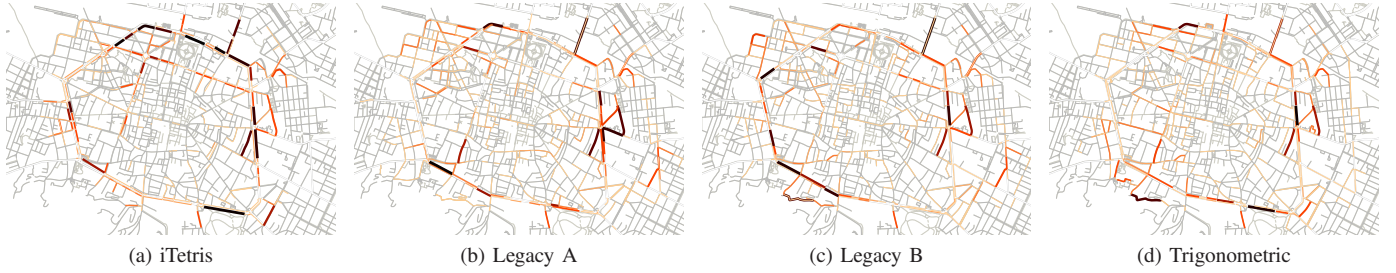


Fig. 9: Bologna Ringway. Road traffic density over the network, at the simulation peak traffic time.

capable of reproducing that realistic behavior.

Finally, Fig. 9 sketches the road traffic density recorded in the target area, one hour after the simulation start. Each plot refers to a different tool configuration, which allows appreciating how differences are not only related to the aggregate traffic statistics in the previous figures, but also to the geographical distribution of congestion. For instance, the iTetris configuration leads to a large traffic jam along the whole northeastern section of the ringway around downtown Bologna. The Legacy A and B configurations generate instead congestion in the eastern and southwestern regions of the city, respectively. Coherently with the previous results, the Trigonometric case is that yielding the mildest problems in terms of road traffic.

V. DATASET VALIDATION VIA NAVIGATION SERVICES

A major issue with synthetic mobility traces, such as the Bologna Ringway dataset presented above, is that they are the result of a simulation, and thus need to undergo some validation, as discussed in Sec. II. This aspect is often overlooked, and only a few previous works proposed validation

techniques that are, however, based on visual inspection [18] or city-specific hard-to-retrieve data [31].

A. Navigation-service-based validation

We propose an original approach to the validation of synthetic road traffic traces. Our methodology builds on publicly accessible data that is available for the vast majority of urban areas worldwide, and can thus be replicated for the validation of any mobility dataset.

Specifically, we leverage navigation services that also offer routing and travel time estimation functions, such as Google Maps. The results returned by such services are based on a number of real-world sources, including statistics produced by transportation authorities, floating car data generated by in-vehicle navigation systems, and crowdsourced data collected by GPS-enabled smartphones that run dedicated applications. The resulting information is extremely accurate, and can be retrieved via dedicated Application Program Interfaces (APIs). In the following we will consider Google Maps as our reference navigation service, due to its popularity.

Routing and travel time information from navigation services can be used as a ground truth for the equivalent data

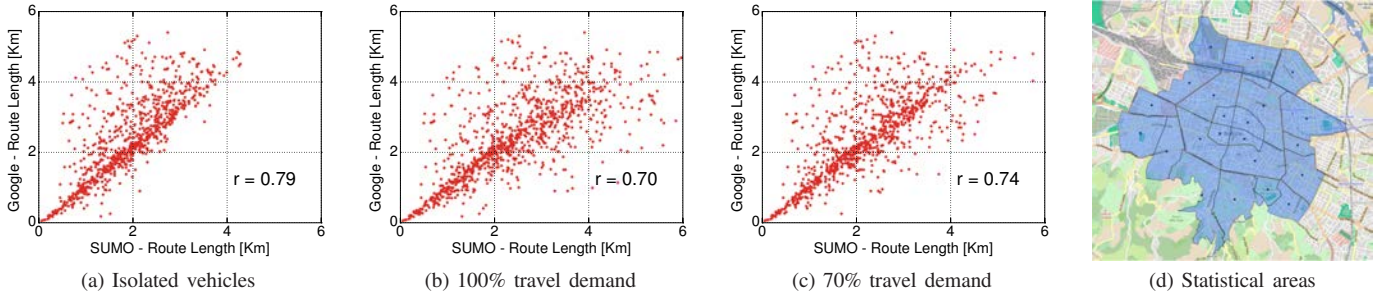


Fig. 10: Navigation-service-based validation. (a-c) Routing path lengths provided by Google Maps versus those recorded in simulation: one simulated probe vehicle (a), complete iTetris travel demand, corresponding to the morning road traffic peak (b), and 70% of the iTetris travel demand, corresponding to typical moderate morning traffic (c). (d) The 22 statistical areas of Bologna, as an OSM overlay; dots map to area centroids.

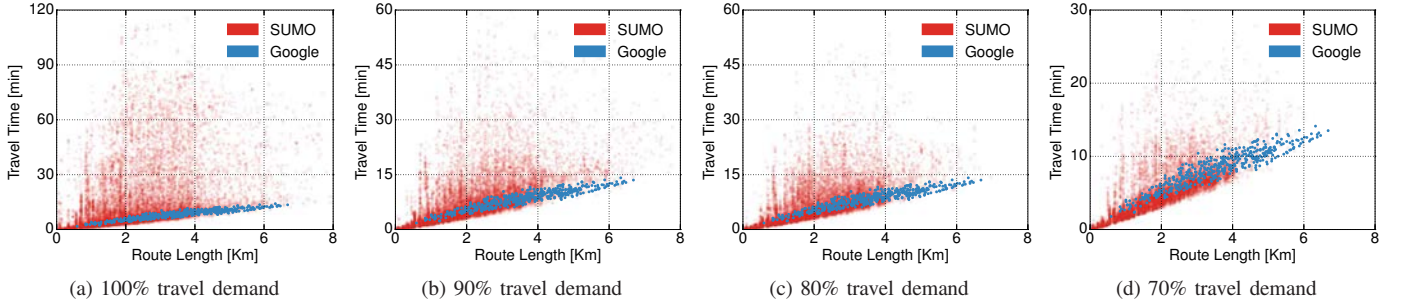


Fig. 11: Navigation-service-based validation. Travel time versus road lengths, as provided by Google Maps and recorded in simulation. Different plots refer to diverse road traffic loads.

extracted from a synthetic mobility dataset. More precisely, navigation service data can be leveraged for the validation of two key performance metrics of road traffic.

- **Routing paths.** The first step of our proposed navigation-service-based validation focuses on the analysis of the routes traveled by vehicles in simulation. By using dedicated Google Maps APIs, one can input the origin and destination pair of each trip in the synthetic dataset to the Google Maps routing function, and retrieve the corresponding path indicated by the navigation service. The latter can be compared to that followed by vehicles in simulation, so as to verify if traveled routes are realistic.
- **Travel times.** The second phase of the validation process concerns travel times. The procedure is similar to that adopted for the routing paths, but involves querying Google Maps APIs for trip duration estimates. This information can be then used as a benchmark for the equivalent data recorded in the synthetic dataset.

B. Validation of the Bologna Ringway dataset

The two-phase validation process above was run on the Bologna Ringway dataset, under the Trigonometric configuration. Once more, we remark that the methodology is based on open-access data, and it is thus not specific to this scenario, but it is replicable in most other urban areas.

The results of the routing path validation are portrayed in Fig. 10a. In all scatterplots, each point represents one trip: its x-axis value is the corresponding route length measured in simulation, whereas its y-axis value is the same route length retrieved from the Google Maps service. In the ideal case,

the routes assigned in simulation are the same suggested by the navigation service, and all points shall lie on the bisector line. This would map to a Pearson product-moment correlation coefficient r equal to one.

Fig. 10a refers to a simplified scenario, where only one vehicle is traveling at a time, between a selected pair of origin and destination locations. Thus, each point in the plot corresponds to one independent simulation. The rationale for this setting is that having a lone vehicle travel in the road network eliminates all bias due to queuing and traffic jams; as a matter of fact, the latter phenomena can induce vehicles to modify their routes, according to Gawron's algorithm sketched in Sec. IV-A3, so as to avoid congested areas – an aspect that Google Maps APIs do not take into account⁵. We observe a good match between navigation service data and simulation results: the vast majority of points lies along the bisector, and the correlation is rather strong at $r = 0.79$. Minor discrepancies are mainly related to the fact that, in some cases, Google Maps suggests routes longer than those computed in simulation. Such an effect is induced by a slightly stronger tendency of Google Maps routing to prefer paths over high-speed roads, when compared to SUMO routing. This comes at some cost in terms of traveled distance, but only occurs on a small set of cases, which explains the limited dispersion in the scatterplot.

Fig. 10c and Fig. 10b assess the reliability of the simulated routing paths once they are determined by Gawron's traffic assignment algorithm. To that end, we consider the complete travel demand provided by the iTetris project, in Fig. 10b, as

⁵Routing and travel time estimation in presence of traffic information are only available to business customers, through Google Maps APIs.

well as the case where only 70% of the demand is injected in the road network, in Fig. 10c. In both scenarios, the presence of heavy and moderate background road traffic triggers path balancing through Gawron’s algorithm iterations. We note that an increasing traffic load on the road network reduces the correlation, as it forces simulated vehicles to take longer routes to avoid congested road segments and intersections. Still, we remark that Gawron’s algorithm does a good job of accommodating the demand without massively re-routing vehicles: indeed, the correlation remains strong enough, at $r=0.7$ and $r=0.74$, respectively, which implies that comparable paths are traveled in simulation and suggested by the navigation service based on real-world data.

For the second validation phase, concerning travel times, we take a slightly different approach, which allows further demonstrating the generality of our methodology. We consider the 22 statistical areas of the city of Bologna, depicted in Fig. 10d: they are the basis for common statistical analyses by public administrations, and are typically available for most urban areas. We employ the centroids of the statistical areas as origin and destination points for queries on travel times to the Google Maps service; specifically, one query is issued for each combination of origin-destination centroids, generating a total of 462 statistical trips. We then compare the relationship between the travel time and the route length, as obtained by statistical trips from Google Maps, and by trips simulated in SUMO using the iTetris travel demand.

Fig. 11 shows the travel time as a function of the route length, as obtained from Google Maps and in simulation. Plots refer to the usual different traffic demands described in Sec. IV-A3. We remark that the simulated travel times overlap for the most part with the values provided by the navigation service, under all road traffic loads. More precisely, the overlap increases as the demand is reduced. Indeed, the dispersion of simulation points towards the top of the plot is due to congestion, which is neglected in Google Maps data: thus, reducing the injected traffic leads to trips whose duration is closer to that anticipated by the navigation service. It is also worth noting that the minimum travel times recorded in simulation are coherent with the minimum values provided by the Google Maps service, for any route length. Therefore, the speed profile of the vehicles in simulation well approximates the real-world ones.

In conclusion, all of the results presented in this section confirm the quality of the Bologna Ringway dataset. In particular, both the road network conversion through the trigonometric NETCONVERT presented in Sec. III and Gawron’s traffic assignment algorithm implemented by SUMO mimic real-world routing behaviors, inferred from navigation services.

VI. RELATED WORK

Our work primarily concerns the microscopic-level simulation of road traffic for vehicular networking purposes. Several tools have been proposed in the past that are dedicated to this task. Many of them are commercial software that ensure a very high level of detail at the cost of configuration complexity: examples include PARAMICS [32], CORSIM [33], or

VISSIM [34]. The economic cost and the steep learning curve of these tools have however favored the diffusion of simpler, open-source simulators of vehicular mobility, which better fit the needs of a vehicular network simulation.

Popular tools adopted by the networking research community were originally based on simplistic stochastic models, such as those implemented by the IMPORTANT framework [14]. The latter were then replaced by a random travel demand over realistic road topologies [35]. Later on, microscopic models of driver’s behavior, developed within transportation research, were included in such tools [36]. Ever since, open-source road network simulators have become increasingly complex, and nowadays allow simulating accurately the movement of individual vehicles over real-world road networks: widely adopted software include SUMO [22], VanetMobiSim [37], or STRAW/SWANS [38]. Among them, SUMO is recognized as the current state of the art, and is part of a number of federation efforts with network simulators, such as those carried out by iTetris [26], Veins [39], [40], or VSimRTI [41]. Yet, improvements to parts of SUMO are still possible, as in the case of the trigonometric NETCONVERT we present in Sec. III.

The availability of dependable simulation tools is critical to the generation of reliable synthetic datasets of vehicular mobility. However, it is not the only component required to that end. A major challenge lies today in the realistic modeling of the movement of large traffic flows across the road network. Without that piece of information, one has to resort to random trips [14], [35] or intuition-based traffic assignment [42], which are not representative of the actual movement of vehicles within urban regions. Clearly, a biased distribution of road traffic dramatically affects the topology of a vehicular network built on top of vehicle mobility, and thus puts the significance of results at stake [43].

A very limited number of datasets is available today, which rely on origin-destination (O-D) matrices that describe real-world traffic flows. Those describe vehicular mobility in the Canton of Zurich [44], in Luxembourg [17], in the city of Cologne [18], and in suburban areas of Bologna [45]. Thus, the the Bologna Ringway dataset we present in Sec. IV is a valuable contribution to the current set of dependable mobility traces available to the research community.

Finally, the validation of synthetic road traffic datasets is a last topic that is relevant to the work presented in this paper. This is a quite overlooked aspect in vehicular mobility simulation, and the only methodologies adopted to date within the networking community are limited to visual inspection against maps provided by live traffic services [18], comparison against traffic count data that are however specific to a city and generally hard to retrieve [31]. The approach we present in Sec. V advances those attempts, since it is based on data that is publicly available for most urban areas worldwide, and allows assessing the quality of the routing and travel times in the synthetic data.

VII. CONCLUSIONS

We presented several contributions to the ongoing effort of enabling dependable vehicular network simulation through

adoption of more realistic road traffic representations. Specifically, we introduced an original Bologna Ringway dataset, describing the movement, during the morning rush hour, of more than 22,000 vehicles in a 25-km² area that covers the center and outskirts of Bologna, Italy. It is our intention to make the dataset publicly available to the research community.

In order to generate the Bologna Ringway dataset, we employ an original version of the OSM-to-SUMO road network conversion tool, NETCONVERT, and show how it allows for a more reliable representation of road intersection and a better synchronization of traffic lights. These improvements lead to significant gains in terms of road network capacity, with reduced travel and waiting times in simulation.

Finally, we validated the routing paths and travel times in the Bologna Ringway dataset, by means of a novel methodology, based on data retrieved from navigation services. Our approach demonstrates the quality of the proposed mobility dataset, and is general enough to be re-used for the validation of other road traffic traces.

VIII. ACKNOWLEDGMENTS

We would like to thank Gian Marco Toso for his help in implementing and maintaining the trigonometric NETCONVERT. The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Unions Seventh Framework Programme (FP7/2007-2013) under REA grant agreement n.630211. EURECOM acknowledges the support of the COLOMBO project, which is partially funded by the European Commission under the 7th Framework Program, FP7-ICT-2011-8, grant agreement no. 318622, as well as the support of its industrial members: BMW Group Research and Technology, IABG, Monaco Telecom, Orange, SAP, SFR, ST Microelectronics, Symantec.

REFERENCES

- [1] J.B. Kenney, "Standards and Regulations," in H. Hartenstein, K.P. Laberteaux (Editors), *VANET – Vehicular Applications and Inter-Networking Technologies*, Wiley, 2010.
- [2] IEEE, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," *IEEE 802.11-2012*, 2012.
- [3] IEEE, "IEEE 1609 – Family of Standards for Wireless Access in Vehicular Environments (WAVE)," *U.S. Department of Transportation*, 2006.
- [4] OSI, "Communications access for land mobiles (ITS-CALM-M5)," OSI 21215 Standard.
- [5] ETSI, "Intelligent Transportation Systems (ITS)," *EN 302 665*, ETSI Standard.
- [6] Reuters, "Obama Backs Highway Fund Fix, Touts 'Talking' Cars," *The New York Times*, 2014.
- [7] G. Araniti, C. Campolo, M. Condoluci, A. Iera, A. Molinaro, "LTE for Vehicular Networking: A Survey," *IEEE Communications Magazine*, 51(5), 2013.
- [8] J. Whitbeck, Y. Lopez, J. Leguay, V. Conan, M. Dias de Amorim, "Relieving the wireless infrastructure: when opportunistic networks meet guaranteed delays," *IEEE WoWMoM*, Lucca, Italy, June 2011.
- [9] R. Stanica, M. Fiore, F. Malandrino, "Offloading Floating Car Data", *IEEE WoWMoM*, Madrid, Spain Jun. 2013.
- [10] C. Lochert, B. Scheuermann, M. Mauve, "Information Dissemination in VANETs," in H. Hartenstein and K. Laberteaux (Editors), *VANET – Vehicular Applications and Inter-Networking Technologies*, Wiley, 2010.
- [11] K.C. Lee, U. Lee, M. Gerla, "Survey of Routing Protocols in Vehicular Ad Hoc Networks," in M. Watfa (Editor), *Advances in Vehicular Ad-Hoc Networks: Developments and Challenges*, Information Science Reference, 2010.
- [12] Ann Arbor Safety Pilot, <http://safetypilot.umtri.umich.edu/>.
- [13] sim^{TD}, <http://www.simtd.de>.
- [14] F. Bai, N. Sadagopan, A. Helmy, "The IMPORTANT framework for analyzing the Impact of Mobility on Performance Of Routing protocols for Adhoc NeTworks", *Elsevier Ad Hoc Networks*, 1, 2003.
- [15] M. Fiore, J. Härrä, "The Networking Shape of Vehicular Mobility," *ACM MobiHoc*, Hong Kong, PRC, May 2008.
- [16] W. Viriyasitavat, F. Bai, O.K. Tonguz, "Dynamics of Network Connectivity in Urban Vehicular Networks," *IEEE Journal on Selected Areas in Communications*, 29(3), 2011.
- [17] Y. Pigné, G. Danoy, P. Bouvry, "A vehicular mobility model based on real traffic counting data," *Nets4Cars/Nets4Trains*, 2011.
- [18] S. Uppoor, O. Trullols-Cruces, M. Fiore, J.Barcelo-Ordinas, "Generation and analysis of a large-scale urban vehicular mobility dataset," *IEEE Transactions on Mobile Computing*, 13(5), May 2014.
- [19] M. Gramaglia, O. Trullols-Cruces, D. Naboulsi, M. Fiore, M. Calderon, "Vehicular Networks on Two Madrid Highways," *IEEE SECON*, Singapore, Jul. 2014.
- [20] S. Joerer, C. Sommer, F. Dressler, "Toward reproducibility and comparability of IVC simulation studies: a literature survey," *IEEE Communications Magazine*, 50(10), 2012.
- [21] OpenStreetMap, <http://www.openstreetmap.org>.
- [22] D. Krajzewicz, J. Erdmann, M. Behrisch, L. Bieker, "Recent Development and Applications of SUMO - Simulation of Urban MOBility," *International Journal On Advances in Systems and Measurements*, 5(3), 2012.
- [23] D. Krajzewicz, G. Hertkorn, J. Ringel, P. Wagner, "Preparation of digital maps for traffic simulation; part 1: Approach and algorithms," *Industrial Simulation Conference*, 2005.
- [24] D. Krajzewicz, J. Erdmann, "Road intersection model in SUMO," *SUMO User Conference*, 2013.
- [25] Highway Capacity Manual, <http://hcm.trb.org/>.
- [26] M. Rondinone *et al.*, "ITETRIS: a modular simulation platform for the large scale evaluation of cooperative ITS applications," *Simulation Modelling Practice and Theory*, 34:99–125, 2013.
- [27] Google Maps JavaScript API v3, "Directions Service", <https://developers.google.com/maps/documentation/javascript/directions>.
- [28] S. Krauß, P. Wagner, C. Gawron, "Metastable States in a Microscopic Model of Traffic Flow", *Physical Review E*, 55(304):55–97, May 1997.
- [29] D. Krajzewicz, "Kombination von taktischen und strategischen Einflüssen in einer mikroskopischen Verkehrsflusssimulation", in T. Jürgensohn, H. Kolrep (editors), *Fahrermodellierung in Wissenschaft und Wirtschaft*, VDI-Verlag, 104–115, Berlin-Adlershof, Germany, 2009.
- [30] C. Gawron, "An Iterative Algorithm to Determine the Dynamic User Equilibrium in a Traffic Simulation Model", *International Journal of Modern Physics C*, 9(3):393–407, 1998.
- [31] S. Uppoor, M. Fiore, "Characterizing pervasive vehicular access to the cellular RAN infrastructure: an urban case study," *IEEE Transactions on Vehicular Technology*, to appear.
- [32] Paramics: Microscopic Traffic Simulation, <http://www.paramics-online.com>.
- [33] CORSIM: Microscopic Traffic Simulation Model, <http://mctrans.ce.ufl.edu/featured/tsis>.
- [34] Ptv simulation – VISSIM, http://www.english.ptv.de/cgi-bin/traffic/traf_vissim.pl.
- [35] A. K. Saha, D. B. Johnson, "Modeling Mobility for Vehicular Ad Hoc Networks", *ACM VANET*, Philadelphia, PA, USA, Oct. 2004.
- [36] S. Jaap, M. Bechler, L. Wolf, "Evaluation of Routing Protocols for Vehicular Ad Hoc Networks in City Traffic Scenarios", *IEEE ITSC*, Vienna, Austria, Sep. 2005.
- [37] J. Härrä, M. Fiore, F. Filali, C. Bonnet, "Vehicular Mobility Simulation with VanetMobiSim," *Transactions of The Society for Modeling and Simulation*, 87(4), 2011.
- [38] D. Choffnes, F. Bustamante, "An Integrated Mobility and Traffic Model for Vehicular Wireless Networks," *ACM VANET*, 2005.
- [39] C. Sommer, R. German, F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing*, 10(1), 2011.
- [40] L. Bedogni, L. Bononi, M. Di Felice, A. D'Elia, R. Mock, F. Montori, F. Morandi, L. Roffia, S. Rondelli, T.S. Cinotti, F. Vergari, "An interoperable architecture for mobile smart services over the internet of energy," *IEEE WoWMoM*, Madrid, Spain Jun. 2013.
- [41] B. Schuenemann, "V2X Simulation Runtime Infrastructure VSimRTI: An Assessment Tool to Design Smart Traffic Management Systems," *Elsevier Computer Networks*, 55(14):3189–3198, 2011.

- [42] C. Barberis, G. Malnati, "Epidemic information diffusion in realistic vehicular network mobility scenarios," *IEEE ICUMT*, 2009.
- [43] D. Naboulsi, M. Fiore, "On the instantaneous topology of a large-scale urban vehicular network: the Cologne case," *ACM MobiHoc*, 2013.
- [44] B. Raney, N. Cetin, A. Völlmy, M. Vrtic, K. Axhausen, K. Nagel, "An agent-based microsimulation model of Swiss travel: First results," *Networks and Spatial Economics*, 3(1), 2003.
- [45] L. Bieker, D. Krajzewicz, A. Morra, C. Michelacci, F. Cartolano, "Traffic simulation for all: a real world traffic scenario from the city of Bologna," *SUMO*, 2014.



Andrea Vesco received the M.Sc. degree in Telecommunication Engineering and the Ph.D. in Computer and System Engineering from the Politecnico di Torino in 2003 and 2009 respectively. After one year of post-doc with the Control and Computer Engineering Department as a member of the Computer Networks Group at the Politecnico di Torino, he joined the Networking Lab at the Istituto Superiore Mario Boella (ISMB) in 2010. His research interests were on Quality-of-Service (QoS) over packet switched networks and wireless access networks. Moreover he carried out research activities on QoS over Network-on-Chip (NoC) in collaboration with the System-Level Design Group at the Columbia University of the city of New York. After ten years of experience and several publications in the networking field, he had the opportunity to focus the researches on the Smart City subject in 2013. His research effort is currently spent on the role of ICTs for the Sustainable Development of Smart Cities.



Luca Bedogni received the degree in computer science from the University of Bologna, Italy, where he also got his PhD in 2015. Currently he is a research assistant at the Department of Computer Science and Engineering at the University of Bologna, Italy. His research interests include the study, modeling and performance evaluation of future wireless networks operated in TV bands, cognitive wireless networks, and vehicular networking.



Jérôme Härrri is Tenured Assistant Professor at the Mobile Communication Group at EURECOM, France, and conducting research in wireless vehicular networks. Previously, he led the Traffic Telematics Junior Research Group at the Institute of Telematics of the Karlsruhe Institute of Technology (KIT), Germany. His research interests are related to the optimization of the vehicular wireless channel usage, to the investigation of cooperative ITS strategies and to the characterization of the mutual relationship between vehicular mobility and vehicular communication. He has authored and co-authored over 50 international journal and conference papers, and is involved in various National and European research projects related to Connected Vehicles. He holds a M.Sc. degree and a Dr. s sc. degree in telecommunication from the Swiss Institute of Technology (EPFL), Lausanne, Switzerland.



Marco Gramaglia is a postdoctoral research fellow at CNR-IEIIT, Turin, Italy. Previously, he was a researcher at Istituto Superiore Mario Boella, Turin, Italy. Between 2008 and 2012 he was a research assistant at Institute IMDEA Networks and a Ph.D. candidate at the telematics department of University Carlos III, Madrid, Spain. He holds a M.Sc. and a Ph.D. in telematics engineering from University Carlos III of Madrid ('09,'12) and a M.Sc. in Computer Science from Politecnico di Torino ('08). His research interests are mobility aspects of wireless networking and vehicular ad-hoc networks.



Francesco Ferrero is the Head of the Smart City Strategic Program at Istituto Superiore Mario Boella, Italy, where he leads a multi-disciplinary team of researchers conducting applied research in the field of economically, environmentally and socially sustainable cities. He has recently edited an Handbook of Research on the argument (<http://bit.ly/1DGZAPt>) and is involved in several National and European research projects related to smart and sustainable cities and mobility. He holds a M.Sc. degree in physics from the Università di Torino, Italy.



Marco Fiore (S05, M09) is a researcher at CNR-IEIIT, Italy, and an associate researcher at Inria, France. He received a PhD degree from Politecnico di Torino, in 2008, and a habilitation à diriger des recherches (HDR) from University of Lyon, in 2014. He was an associate professor at INSA Lyon, France, and a visiting researcher at Rice University, USA, and Universitat Politècnica de Catalunya, Spain. His interests are on mobile traffic analysis and vehicular networking.