A Simulation-based Decision Support System for Urban Traffic Management

Alessandro Calvio*, Anish Jindal†, Armir Bujari*, Gagangeet Singh Aujla† and Luca Foschini*
* Department of Computer Science and Engineering, University of Bologna, Italy
† Department of Computer Science, Durham University, United Kingdom
*{name.surname}@unibo.it,
†{name.surname}@durham.ac.uk

***Abstract*—One common challenge faced by smart cities is traffic congestion, caused also unintentionally by city planners, which significantly impacts urban life and the environment. To address these issues, various strategies and approaches are explored, including advanced traffic management methods, communication systems, and predictive modeling. In this study, we introduce a Digital Twin decision support system that focuses on accurately modeling a city's road network and replicating traffic patterns for specific time frames. Central to this project is SUMO, a microscopic vehicular traffic simulator used to create the reference road model. Our research presents a flexible and viable methodology for generating SUMO-compliant simulations that leverage real vehicular data from the city to analyze urban traffic and its conditions. To illustrate our approach, we apply it to a specific case study centered on examining environmental emissions in the city of Bologna, Italy. This approach shows promise in enhancing traffic management and overall urban efficiency within the broader context of smart cities.**

***Index Terms*—Smart City, Digital Twin, Simulation, SUMO**

I. INTRODUCTION

The Smart City paradigm advocates for the seamless integration of physical and digital realms, facilitated by the utilization of Internet of Things (IoT) and Information and Communication Technology (ICT). This integration is pivotal for achieving real-time data access and leveraging this data to enable precise control and the delivery of intelligent services for its residents [1]. In this context, Digital Twin technology is regarded as an innovative foundation for the concept of a smart city. It is envisioned as a digital representation of an object, process, or service, comprehensive enough to serve as the foundation for informed decision-making [2]. Within a smart city, traffic is one of the most prominent problems. According to the 2022 Global Traffic Scorecard by the mobility analytics firm Inrix [3], 58% of urban areas analyzed saw increased traffic delays over 2021, with the average American person losing 51 hours a year due to traffic congestion costing nearly \$869 per person in lost time. In the UK, these figures are 80 hours and \$926, respectively. The congestion is even worse in developed metropolitan cities such as London, which has an average of 156 hours and New York, which has an average of 117 hours lost in congestion. The transportation sector accounts for about 29% of total US greenhouse gas emissions [4], and traffic congestion would further deteriorate the climate and human health.

Moreover, with the increase in traffic congestion, significant problems arise in underdeveloped and developing countries. For example, emergency vehicles such as ambulances, police cars, and fire engines struggle to reach their destination on time. They must wait at intersections, which increases their delays and leads to the loss of lives and property [5].

Many methods and technologies are currently being used to make the transportation system more accessible and flexible. For instance, Meszaros *et al.* [6] proposed a Markov-based method for the arrival time and service time of vehicles near an intersection. The authors used the Markov arrival process and Markov Decision Process to calculate the time of vehicles near the intersection with traditional traffic control systems. In another work [7], researchers used Radio Frequency (RF) module on the vehicles instead of InfraRed (IR) module as the IR signal can only travel a short distance as compared to RF. The authors in [8] proposed an arrival time-based traffic signal optimisation solution that outperformed an arrival rate-based system in intelligent transportation environments in moderate and low traffic scenarios. Other approaches, such as information sharing [9] and machine learning [10], can be used to predict the city traffic flow. However, each approach has its own drawbacks, such as requiring pre-trained models on data, real-time information sharing, and robust communication infrastructure. Moreover, for deriving real-time traffic flow models, there is no uniform platform to combine the scattered information available and fill out the missing data.

On this front, Digital Twin technology plays a pivotal role, seamlessly integrating data from diverse sources to provide a comprehensive and real-time understanding of traffic conditions, enabling intelligent, data-driven optimizations. In this context, SUMO [11] is a well-known tool for simulation of urban mobility, equipped with a wealth of features such as demand generation, traffic lights, multimodal traffic, vehicle communication, traffic management, to name a few. However, mapping traffic from a real-city to SUMO to make it learn the realistic traffic patterns of the city for future modelling is a difficult task. This is mainly because of the errors introduced while converting data from one model or tool to another. Therefore, it is important to accurately map such data within the SUMO to precisely implement various traffic models for improving traffic awareness in the actual physical environment.

To this end, we present a Digital Twin decision support system which relies on (real) vehicular data, allowing city stakeholders to acquire timely information on traffic dynamics. The data are sourced from black boxes embedded on vehicles and can be fed to SUMO for further analysis on traffic flows. In this work, we present a generic methodology to effectively create a SUMO-based simulation scenario for a city and show its effectiveness for modelling the traffic demand of Bologna, Italy. The major contributions of the paper are:

- To develop a SUMO-based road network traffic generator for modelling the traffic demand for the city of Bologna.
- To accurately remove the tagging errors, invalid trips, and conversion errors from OpenStreetMap geographic information to SUMO.
- To make the Origin-Destination (OD) metric for different traffic assignment zones more precise.
- To show the implementation of the proposed scheme to model the traffic demand of Bologna for its environmental performance assessment.

The rest of the paper is organised as follows. Section II provides some background on Digital Twins. Section III discusses the methodology for the traffic demand modelling, while Section IV highlights some preliminary results. The paper is concluded in Section V.

II. BACKGROUND AND SYSTEM DESIGN

Digital innovation has revolutionized urban management through the emergence of Digital Twin platforms. These digital representations of physical systems and processes offer a powerful tool for real-time analysis, monitoring, and simulation of complex urban settings. However, defining the role of Digital Twins as mere mirrors of real-world situations would be reductive, given that, one of the core features of these data-driven platforms is also the ability to retroactively interact with assets present in the real world so that policies and strategies can be implemented entirely autonomously [12].

Although there are no widely accepted standard approaches to building these systems, there is a need for specific functional components that are essential for their operation. The first ingredient of data-driven Digital Twins is the continuous collection of data from the real world at different levels of precision and granularity. Our pursued system employs vehicular data, made available by the IPPODAMO system, consisting of vehicular positional information sensed and announced by embedded black box devices [13]. The data continuously enter the system and are fed to a tailored processing pipeline for some refinement, after which they are fed to dedicated Kafka topics and persisted elsewhere. The data ca be retrieved by third-party authorized providers and processed to reconstruct city traffic flows in (quasi) real-time.

The second component concerns the set of computational models that the Digital Twins exploit to mimic the monitored phenomenon. These models allow the characteristics and properties underlying a specific application domain to be encapsulated, enabling analysis and studies of specific phenomena that can support decision-making. In this context, the use of simulation tools is a popular practice to obtain a deeper level of understanding [14]. Using physical asset models makes it possible to perform analyses and test new scenarios before they are realized, and to plan the best strategies based on the derived outputs [13]. In our solution, we rely on SUMO as the main simulation tool, an open-source, portable, and microscopic multimodal traffic simulation package representing the state-of-the-art in this field and widely used in research. Among the various tools available, SUMO provides the ability to model and modify large road networks and populate them with a traffic demand generated randomly or from external data sources. Building on the experience of other works in the scientific community [15], in this paper, we approached the study of traffic management within the city of Bologna and the model generated with SUMO is the heart of the system. The data from the previous ingestion pipeline are used to build the proper traffic models for the simulation. Meanwhile, open-source data are utilized for network creation. Thanks to this design choice, it is possible to have a unified view of the traffic conditions present in a specific area to conduct analysis and identify critical issues in the transportation system about specific metrics, such as environmental emissions and their relationship to vehicle waiting times on the roads.

Finally, every Digital Twin must be an open system that can also be queried and used by external users or services. In this sense, a set of APIs and interfaces is an important requirement, effectively enabling interaction between heterogeneous systems. SUMO offers considerable flexibility in how it interacts with the simulation. On the one hand, it provides an intuitive graphical interface that allows users to easily visualize and interact with the simulation in a visual environment. This graphical interface is handy for real-time visualization of traffic dynamics and for making immediate observations about the behavior of vehicles and road infrastructure.

On the other hand, SUMO offers complete control of the simulation through the use of system commands and XML configuration files. This approach allows different simulations to be instantiated based on input parameters that can come from users without any manual intervention allowing a high degree of customization. In addition, SUMO offers the ability to interact with the simulation at runtime through the use of the TraCI library. This library allows users to write scripts or programs that can communicate directly with the simulation, dynamically changing parameters, obtaining real-time data, and influencing the behavior of vehicles or agents in the process.

III. SIMULATION CREATION METHODOLOGY

This section discusses the devised methodology used to build the SUMO simulation scenario at the core of the system. Starting with an area selection, the process includes the generation of topological information about the selected road network and the creation of traffic models using realistic floating car data, culminating in the actual simulation run and retrieval of the results. The process flow is depicted in Fig. 1, emphasizing the individual process stages that will be discussed in the following.

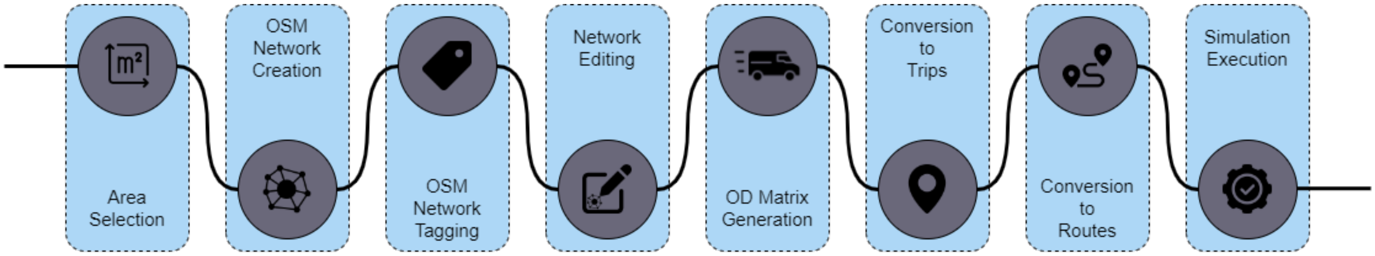

Fig. 1: SUMO Pipeline Structure

A. Network Generation

The starting point for a simulation is creating a network, a graph representation of the road network of interest. This information is encoded through XML files containing the definition of the main elements of the network (roads and intersections) as well as the way they are connected, thus defining the overall topology. In addition, through a process of network tagging, each of the elements must be enriched with a range of useful information (e.g. speed limits, presence of traffic lights) in order to correctly simulate the behavior of vehicles. SUMO provides several ways to generate a road network for simulations. To this end, we decided to create the road networks starting from the OpenStreetMap (OSM) geographic information service, specifically employing a tool provided by SUMO called OSMWebWizard (OWW).

OWW provides a web interface to interact with OSM maps so to create vehicular simulations with randomized traffic demand for specifically selected areas. Behind the scenes, the tool takes care of downloading network information from OSM and converting it into a format understandable by SUMO through the *NETCONVERT* tool. Then the traffic is randomly generated based on the traffic specifications, about volumes and types, entered by the user, using the *randomTrips.py* script.

Once the user selects the area to be simulated through the interface, OWW will transparently take care of generating a correct simulation for SUMO. However, far from being ideal, the conversion process introduces several structural errors that must be carefully considered and resolved for a truthful simulation of the interested scenario. Errors at the level of network structure can plague the way vehicles travel and the roads they may take, leading to executions that are offset from what occurs in reality.

Two error types were identified during the network generation process which must be carefully addressed. The first can be attributed to data import from the chosen geographic system (i.e., OSM) and tagging errors within it. In this case, SUMO imports all metadata and road information that is present within OSM, so if some information is erroneous or misrepresented in the source system, it is imported as is. These types of errors typically require manual intervention to correct them. In the case of Bologna, which is the focus of our evaluation, tagging errors concerned speed limits and lane counts on city streets. Some speed limit information was incorrect or missing. To address this, we used SUMO typemaps, files specifying default values for elements like roads and nodes with missing information. Hence, we set maximum speed limits of 50km/h for primary roads and 30km/h for residential and interior areas. For lane counts, default values were impractical due to Bologna's irregular topology, requiring manual corrections on specific roads. Additionally, OSM provided traffic light placement but not timing information. No external datasets on traffic light timing in Bologna were available for integration into the simulation. Thus, in this initial version of our work, we used SUMO's default settings, which define traffic light cycles with a period of 90 seconds.

On the other hand, another category of errors concerns those introduced into the network due to the OSM to SUMO conversion process that exploits the NETCONVERT tool. In all those cases where information is not explicitly entered, the tool exploits rely on assumptions to construct the missing parts. This happens with connections that describe how the various incoming and outgoing wools of a node are connected. This is particularly important for indicating the directions that can be used by vehicles within intersections. At the end of the conversion process, the network has a persistent error: whenever a road with two opposing lanes is present, a connection is created between one run and the other. This is unacceptable for a traditional road network because it would mean that, at any time, a car can make a so-called U-Turn on the same road it is traveling. To solve this problem, a U-TurnRemoval script was created that, by doing a scan of the network, identifies these situations and creates a set of rules to force the deletion of these erroneous links. At the end of the process, the rule set, encoded in XML format, is passed back to the netconvert command, which removes all the erroneous connections based on the directions provided. This script has been integrated within the OWW tool so that it can be handled uniformly with the other scripts that are already being called. At the end of the process, a network is obtained that is not particularly critical to the proper flow of vehicular traffic.

B. Traffic Demand Modelling

Once the network defining the spatial area for vehicular movement and the governing traffic regulations is established, it is necessary to populate it with appropriate traffic flows. This process is called Traffic Demand Modelling. As evident from the previous section, OWW already takes care of generating traffic within the selected area, albeit through a random process. However, this randomness presents constraints when conducting simulations aimed at analyzing and collecting metrics on what is happening within a city area. Therefore, it is necessary to create traffic flows that model the situation within the city, eventually in specific time slots of interest.

SUMO provides several methods for generating realistic traffic within a network ranging from aggregated traces, such as the use of inductions loops data, to more precise ones. In our case, we focused on the latter methodology, basing the traffic generation process on the use of OD matrices. An OD matrix is a data structure that highlights the number of vehicles flowing from one area of the network to another in both directions. Based on this information, which should model the traffic density in those areas, SUMO can generate a sufficient number of vehicular routes such that the volumes involved are reproduced.

Both temporal and spatial dimensions must be considered when defining an OD matrix. The former defines the reference period of the matrix and in SUMO is defined in terms of seconds. On the other hand, the spatial dimension refers to the definition of the areas that are considered by the matrix. In SUMO terminology, these areas are called Traffic Assignment Zones (TAZs) and are identified arbitrarily as appropriate. For the definition of TAZs in the Bologna metropolitan area, we decided to use census cells defined in 2011 and still in effect. In more detail, there are more than 2,000 geographic areas covering the entire Bologna municipality with a very fine spatial granularity. This feature also allows coarser levels of cell aggregation to be considered as needed.

The dataset is available in the open data section of the Bologna municipality in different formats (Shapefile, GeoJSON, etc.), however, there is a need for it to be converted to the TAZ format understandable to SUMO. This conversion is accomplished in two steps: first, the geographic information expressed in terms of latitude and longitude is transformed into the Cartesian reference system of the chosen network, in (x, y) coordinates, using the *polyconvert* tool and, each of them is associated with a unique ID. The areas thus obtained are saved as TAZs within a proper XML file, and for each of them, the *edgesInDistrict.py* script locates all the constituent edges that fall within it.

Once the TAZs for traffic demand modeling have been defined, the subsequent step involves creating the Origin-Destination (OD) matrix that will serve as the basis for generating the ultimate traffic flows. The matrix is derived from the floating car data stored in the system database as a result of the ingestion pipeline. One characteristic of these data is that the sampling period is not regular but only significant locations related to specific driving style events (acceleration, braking, exceeding speed limit) are taken into account, and in any case, at least one location per km is considered. When delving into the particulars of individual detections, the pipeline revolves around the following information fields: ID - anonymized vehicle identifier, DateTime - instant of the detection in yyyy-MM-dd HH:mm:ss format, Latitude - latitude coordinate of the detection, Longitude - longitude coordinate of the detection.

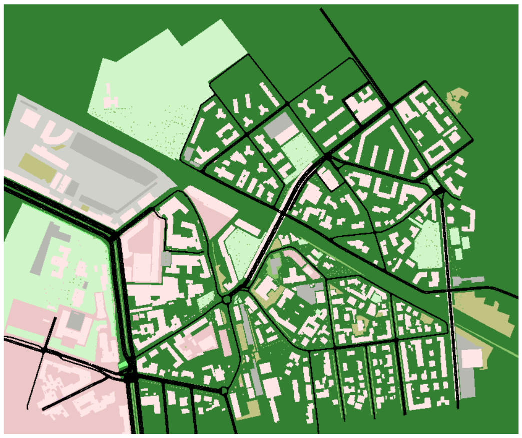

Fig. 2: SUMO network representation area of the simulation

The previously created TAZs, together with the available vehicular data, are the entry point for the creation of the O/D matrices. The spatial data's granularity is a crucial factor in this process, as it should enable accurate reconstruction of vehicle routes within the selected TAZs. In an ideal scenario, having at least one reference point within each TAZ is necessary. However, insufficiently fine sampling might hinder meeting this requirement, potentially resulting in less precise OD matrices. For the generation of the matrices, an algorithm was developed that associates each vehicular detection with the TAZ in which it is located and reconstructs the direction of travel. Going into more detail, through a spatial join operation each vehicular point is associated with its respective TAZ. Subsequently, the obtained data is organized based on time intervals relevant to the OD matrix. Within each band, a grouping by ID allows individual vehicle detections to be obtained so that the path taken can be analyzed. If two consecutive points in time are in different TAZs then the number of vehicles that have passed from the departure TAZ to the arrival TAZ considered is incremented by 1. Finally, the matrix thus obtained with the final counts is saved to an XML file in the format used by SUMO so that it can be used in the simulation. The OD matrices are the starting point for generating the actual traffic flows that will be input into the simulation. the *od2trips* tool enables the division of these matrices into individual vehicular trips. These trips are characterized in terms of departure and end trip edge and departure time instant, randomly chosen within the time range considered by the matrix. The trips thus generated are not validated, meaning that at runtime some of them may be incorrect and there may be no path connecting the start and end edges. At the simulation level, this results in the vehicle being stuck within the starting edge until it is removed from the simulation. Despite this, as long as the vehicle remains within the simulation it is counted as a valid vehicle and actually contributes to any statistics collected at the end of it as well.

To correct this problem, the *duarouter* offers a solution. Given a trip, it can reconstruct the path of the vehicle in question, in terms of roads to be traversed. The tool is also able to notice if a trip is invalid and possibly discard it. In this way, it is possible to filter the trip file generated by *od2trips* to remove all invalid trips and have a cleaner simulation. The cons of this approach is to end up with fewer trips than would be needed to meet the traffic volumes defined within the OD matrix. One solution to this was to use *od2trips'* scale parameter, which allows the traffic volumes to be scaled by an arbitrarily defined factor. By generating a surplus of trips beyond strict necessity, it becomes feasible to still adhere to the desired traffic volume while compensating for the potential reduction due to trip validation and filtering.

IV. EVALUATION

This section will present preliminary results and validation of the proposed methodology against scenarios related to the study of environmental emissions.

A. Experimental Settings

To validate and show the advantages of the approach concerning case studies related to environmental performance assessment, we instantiated a concrete scenario related to the city of Bologna. Specifically, the use case involved the study of the level of emissions during the peak hour, 12-14, within an area of about 1 km^2 around the Porta San Donato area and its surroundings, as shown in Figure 2. The reasons for choosing this area were to have a heterogeneous road network that included both main roads, such as the primary boulevards of Bologna and the San Donato road, and secondary roads, such as all the residential ones in the surroundings.

Given the limited area, the simulation was configured to directly use census cells as illustrated in the previous section, without considering any level of aggregation thus resulting in 68 distinct TAZs that were used to spatially filter the vehicular dataset. Regarding the latter, specifically, we considered data from 10 days of April 2019. Using more recent data would have resulted in bias due to the COVID-19 period in which traffic volumes altered. The data were then also temporally filtered to consider the period of interest for the study, namely the 12-14 time slot. The proposed scheme was run iteratively on each of these days, and the volumes obtained for each TAZ pair in each matrix were then averaged. Finally, to overcome the problem of invalid trips being generated during the process, a scaling factor of 1.1 was used.

The execution of the simulation was carried out on a node with SUMO, with 32GB RAM and 14 CPU cores. Table I shows statistics regarding both the size of the network and vehicles generated and the stability of the simulation itself.

In general, the chosen configuration resulted in the generation of 386 trips of which 329 are valid due to the lack of path validation of the od2trips tool. The execution does not show any particular criticality and finishes without creating queuing phenomena, which means that the network created by the procedure described in the previous section can dispose of all the traffic generated and that there are no significant problems to the network that jeopardize its entire simulation. In fact, at the end of it, no vehicles were teleported due to excessively long queues or colliding with other vehicles within intersections.

TABLE I: Statistics for the simulation execution

Simulation Statistics	
Nodes	79
Edges	133
Traffic Assignment Zones	68
Generated Trips	386
Valid Trips	329
Final Waiting vehicles	0
Teleports	0
Collisions	0

B. Preliminary Results

To validate the traffic volume modeling process, we made a comparison between the vehicle volumes generated by the simulation and those in the source dataset in terms of quantity and spatial distribution. For this comparison, tile aggregations of 150m sides were considered. Figure 3 shows the comparison between the spatial distribution present on average in the actual dataset (Figure 3a) and that generated in the two hours of simulation considered (Figure 3b). The comparison shows a difference in the distribution of traffic volumes due to simulation edge effects. While in the real dataset, the traffic is mainly concentrated along the avenues of Bologna, with volumes of about 80 cars, in the simulation case the traffic is shifted to the main street via San Donato with peaks of over 100 cars, as can be seen from the denser tiles crossing the network. This is mainly attributable to the fact that the boulevards are located along the edges of the considered network and, therefore, the trip generation algorithm takes them into account marginally, compared to the main roads crossing the network. In addition, the adopted methodology, which makes use of the od2trips script and an applied scaling factor greater than 1, leads to the generation of a higher number of vehicles than those detected from the actual dataset.

To assess the potential of this type of simulation with respect to the study of environmental emissions, Figure 4 shows the average levels of CO2 emitted, per second, by cars with the same level of aggregation in 150m tiles. To derive these kinds of results, SUMO provides several emission models that can be used as simulation output. In our case, the HBEFA3/PC_G_EU4 model for Euro norm 4 gasoline cars was used. The output files, thus generated, report for each second and for each car indications of the levels of CO2 (and other pollutants) emitted. As a validation of the process, it can be seen that the largest emission values follow the same distribution as in graph 3b, with the main areas distributed along the avenues and via San Donato. In addition, three emission peaks can be highlighted in the graph, values above 1000 mg/s, corresponding to the main intersections within the area. These results are reliable because, in these areas, waiting times due to traffic jams and traffic light waits are higher than in any other area. In this sense, the simulation helps to understand the points and reasons where it is useful to intervene to improve the overall conditions.

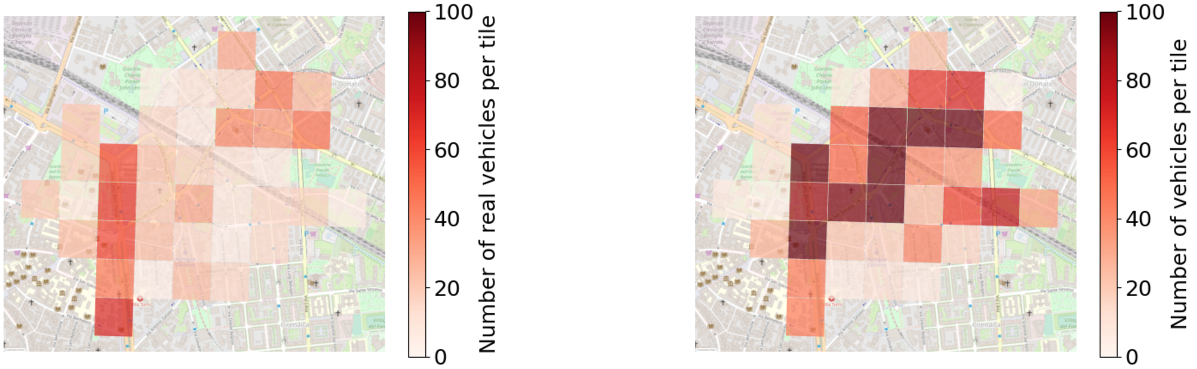

(a) Traffic density per Tile (Real Dataset)

(b) Traffic density per Tile (Simulated)

Fig. 3: Traffic density related metrics

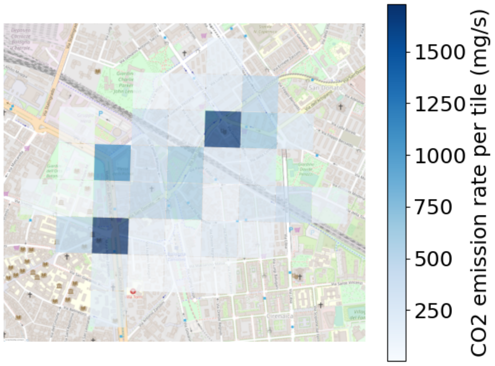

Fig. 4: CO2 emitted per Tile

V. CONCLUSIONS

In this paper, we present a SUMO-based methodology for creating simulations of real road networks, highlighting their value as integrated models in modern Digital Twin platforms. To demonstrate the effectiveness of this methodology, we applied it to a specific case study concerning the city of Bologna, focusing on the analysis of CO2 emissions produced by vehicles. Our traffic model is based on real data from reliable sources in the automotive industry. Currently, our system supports quasi-real-time data, but the primary focus is to evolve it towards a live data integration model. This transition will empower the platform with dynamic data feeds, enabling real-time visualization and affording operators the opportunity for informed and intelligent decision-making. Also, considering that modern Digital Twin platforms are mainly distributed, we intend to explore the implementation of this methodology in a distributed cloud continuum-based context and expand the range of possible analyses based on this work.

ACKNOWLEDGEMENT

This project has received funding from the NRRP project and The Royal Society's International Exchanges project under grant agreement number IES/R2/222187.

REFERENCES

[1] G. Mylonas, A. Kalogeras, G. Kalogeras, C. Anagnostopoulos, C. Alexakos, and L. Munoz, "Digital Twins From Smart Manufacturing to Smart Cities: A Survey," *IEEE Access*, vol. 9, pp. 143 222–143 249, 2021.

[2] Y. Wang, Z. Su, S. Guo, M. Dai, T. H. Luan, and Y. Liu, "A Survey on Digital Twins: Architecture, Enabling Technologies, Security and Privacy, and Future Prospects," *IEEE Internet of Things Journal*, vol. 10, pp. 14 965–14 987, 9 2023.

[3] B. Pishue, "2022 Global Traffic Scorecard: Congestion is Up Despite High Oil Prices," *INRIX Inc.*, 2023. [Online]. Available: https://inrix.com/blog/2022-traffic-scorecard/

[4] "Carbon Pollution from Transportation," *US Environment Protection Agency*, 2023. [Online]. Available: https://www.epa.gov/transportation-air-pollution-and-climate-change/carbon-pollution-transportation

[5] K. Nellore and G. P. Hancke, "Traffic Management for Emergency Vehicle Priority based on Visual Sensing," *Sensors*, vol. 16, no. 11, p. 1892, 2016.

[6] A. Mészáros and M. Telek, "Markov decision process and linear programming based control of map/map/n queues," in *Proc. of Computer Performance Engineering*. Springer, 2014, pp. 179–193.

[7] G. Beri, P. Ganjare, A. Gate, A. Channawar, and V. Gaikwad, "Intelligent Ambulance with Traffic Control," *International Journal of Electrical, Electronics and Computer Systems*, 2016.

[8] V. Paruchuri, S. Chellappan, and R. B. Lenin, "Arrival time based traffic signal optimization for intelligent transportation systems," in *Proc. of IEEE AINA*. IEEE, 2013, pp. 703–709.

[9] W. Liu, S.-W. Kim, K. Marczuk, and M. H. Ang, "Vehicle motion intention reasoning using cooperative perception on urban road," in *Proc. of IEEE ITSC*, 2014, pp. 424–430.

[10] M. Ozbayoglu, G. Kucukayan, and E. Dogdu, "A real-time autonomous highway accident detection model based on big data processing and computational intelligence," in *Proc. of IEEE Big Data*. IEEE, 2016, pp. 1807–1813.

[11] P. A. Lopez *et al.*, "Microscopic Traffic Simulation using SUMO," in *Proc. of IEEE ITSC*, 2018.

[12] A. Fuller, Z. Fan, C. Day, and C. Barlow, "Digital twin: Enabling technologies, challenges and open research," *IEEE Access*, vol. 8, 2020.

[13] A. Bujari, A. Calvio, L. Foschini, A. Sabbioni, and A. Corradi, "A Digital Twin Decision Support System for the Urban Facility Management Process," *Sensors*, vol. 21, no. 24, 2021.

[14] S. Ivanov, K. Nikolskaya, G. Radchenko, L. Sokolinsky, and M. Zymbler, "Digital twin of city: Concept overview," in *Proc. of IEEE Global Smart Industry Conference (GloSIC)*, 11 2020, pp. 178–186.

[15] M. Rapelli, C. Casetti, and G. Gagliardi, "TuST: from Raw Data to Vehicular Traffic Simulation in Turin," in *Proc. of IEEE/ACM DS-RT*, 10 2019, pp. 1–8.

Calvio, A., Jindal, A., Bujari, A., Aujla, G. S., & Foschini, L. (2023). A Simulation-based Decision Support System for Urban Traffic Management. In 2023 IEEE 28th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD) (128-133). https://doi.org/10.1109/camad59638.2023.10478428