

# Lane-Free Microscopic Simulation for Connected and Automated Vehicles

Dimitrios Troullinos<sup>1</sup>, Georgios Chalkiadakis<sup>2</sup>, Diamantis Manolis<sup>1</sup>,  
Ioannis Papamichail<sup>1</sup>, and Markos Papageorgiou<sup>1</sup>, *Life Fellow, IEEE*

**Abstract**—This paper presents the ongoing development of the microscopic *TrafficFluid-Sim* simulator, aimed primarily for Connected and Automated Vehicles (CAVs) under a novel lane-free traffic paradigm. In particular, *TrafficFluid-Sim* builds on the SUMO simulation infrastructure to model traffic environments featuring two novel vehicle characteristics: (i) Vehicles can be located at any arbitrary lateral position within the road boundaries; and (ii) Vehicles may exert, based on their automated driving and connectivity capabilities, “vehicle nudging” to other surrounding vehicles. As such, *TrafficFluid-Sim* enables simulation of novel CAV movement strategies for various types of road infrastructure and is an indispensable tool for the design, testing and evaluation of the characteristics of a future CAV traffic flow as an efficient artificial fluid, as envisaged by the ongoing *TrafficFluid* ERC project.

## I. INTRODUCTION

The arrival of Connected and Automated Vehicles (CAVs) [1] promises the dawn of a new era of road transportation, that is linked to the emergence of novel, safer, and more efficient than existent traffic flow paradigms [2], [3]. Even though CAVS are still fledging, the need to enhance existing [4], [5], [6] or develop entirely new traffic simulators that will allow their design and study at the traffic level is obvious and well recognized in the traffic research community. This is due to the fact that certain key aspects, and characteristics of the prevalent pre-CAV traffic are linked to assumptions or constraints that are likely to become obsolete; while novel features, enabled by the superb CAV capabilities, are emerging and need to be captured in traffic simulation.

In particular, parallel road lanes emerged as a safety-fostering measure at the mid-20th century, when the density of automobiles increased strongly in road transportation. Traffic lanes simplify the driving operations and undoubtedly improve on safety and average driving speed, yet they introduce some disadvantages. A considerable proportion of traffic accidents (10%) occur due to lane changes [7], an operation that all drivers perform on a regular basis, either for overtaking or to follow a specific route. Another important limitation is that lanes have a fixed width according to a presumed maximum vehicle width, which results in lower

lateral road occupancy, and hence lower road capacity. Such compromises are unavoidable with human drivers, given their high reaction times and limited visual capabilities. However, in traffic environments with high CAV penetration rates, such limitations will soon be strongly mitigated, as CAVs have superb sensing capabilities, allowing continuous and reliable monitoring of their environment on a 360° basis; as well as efficient computer-based decision-making. Thus, in recent years, there is a drive to conceive and investigate novel traffic paradigms—e.g., notably, the *TrafficFluid* [8] concept—which entails the need for developing corresponding new traffic simulators to accommodate the novel characteristics of these paradigms.

## II. THE TRAFFICFLUID CONCEPT

*TrafficFluid* [8] is a novel paradigm for vehicular traffic, which is being investigated within its namesake ERC Advance Grant, and targets fully automated vehicles (at SAE Levels 4 or 5) [9], equipped with vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication systems. *TrafficFluid* is based on two key principles:

(1) *Lane-Free traffic*: Vehicles can be in any lateral position within the road boundaries, thus lifting the lane changing task.

(2) *Nudging*: Vehicles can exert a “nudging” effect, i.e., apply a pushing force, on front vehicles. This concept implies that slower downstream vehicles may be nudged to move aside so as to facilitate faster upstream vehicles to pass. More importantly, at a macroscopic level, nudging has a capacity increasing and stabilizing effect on traffic flow [8], [10].

Results obtained so far in *TrafficFluid* (e.g., in [8], [11], [12], [13]) showcase aspects of the concept’s potential. In order to further examine emerging approaches and, in particular, implications at the traffic level, there is an imminent need for a traffic simulator that incorporates these core principles, while retaining a flexible approach. The simulator’s structure has to be modular, so as to integrate and evaluate approaches originating from many different research fields, ranging from Control Theory [14] to Artificial Intelligence [15], Machine Learning [16], [17], [18], and Multi-Agent Systems [19], [20].

In this work, we present the development of *TrafficFluid-Sim*, a lane-free microscopic simulator aimed for V2V and/or V2I-enabled CAVs. *TrafficFluid-Sim* possesses the aforementioned desirable characteristics; while special care is taken during its development so that no inter-tool “communication” delays, which could affect efficiency and user experience,

The research leading to these results has received funding from the European Research Council under the European Unions Horizon 2020 Programme / ERC Grant Agreement no. 833915, project TrafficFluid, see: <https://www.trafficfluid.tuc.gr>

<sup>1</sup>Dynamic Systems and Simulation Laboratory, Technical University of Crete, 73100 Chania, Greece (e-mail: dtroullinos@dssl.tuc.gr, dmanolis@dssl.tuc.gr, ipapa@dssl.tuc.gr, markos@dssl.tuc.gr)

<sup>2</sup>School of Electrical and Computer Engineering, Technical University of Crete, 73100 Chania, Greece (e-mail: gehalk@intelligence.tuc.gr)

are introduced. The simulator builds on and extends the SUMO [4] simulation infrastructure. Importantly, we provide the user with a dynamic library for lane-free vehicle movement control, the incorporation of which allows one to avoid all commutation cost between SUMO and an external application. In what follows, we present in detail our design and implementation choices. Among other things, we explain how lane-free vehicle movement is achieved in *TrafficFluid-Sim*; and also how the simulator handles external traffic and appropriately inserts vehicles at entry points, by offering the user two related functionalities: (a) the option to laterally position the vehicles in such a way that the entire capacity of the road is exploited; and (b) allow the user to sample each vehicle’s desired speed from a distribution of her choice (offered by the API), and then assign a lateral position for the vehicle that is proportional to its respective desired speed. We also present envisaged extensions of the simulator, most of which actually correspond to work currently in progress. We note that a (previous) version of our simulator has already been used for obtaining the results of [11].

### III. DESIGN OF THE SIMULATOR

This section presents the *TrafficFluid-Sim*’s different design and implementation aspects. Specifically, we discuss how we built upon the microscopic traffic simulation tool SUMO, so as to provide enhancements for the incorporation of lane-free vehicle movement and related necessary or desirable functionalities; and detail the implementation of the lane-free *vehicle controller* dynamic library.

#### A. Simulator Overview

The purpose of this work is to build a simulator that enables various control strategies for lane-free vehicle movement, for CAVs, including the nudging effect. We choose SUMO [4] (Simulation of Urban MObility) as the traffic platform to build upon, since, as an open-source project, it can easily be adjusted to our needs.

Due to the popularity of SUMO, multiple frameworks (e.g. VEINS, iCS, VsimRTI) exist that can connect with SUMO via its traffic control interface (TraCI), in order to provide additional functionalities, such as enhanced vehicle communication (see Section V). Our primary endeavour is to create a *modular approach* regarding the road infrastructure characteristics (including space-time varying internal boundary separating the two traffic directions, see [21]) and the vehicles and their movement design (physical dimensions, controllers, connectivity with surrounding CAVs and infrastructure, desired speeds, etc.), thus providing a contained and flexible environment to systematically evaluate novel approaches at various levels (vehicle movement, infrastructure, traffic control).

TraCI enables communication with SUMO, giving access to information regarding the simulation environment, and managing many aspects of it online. However, TraCI imposes certain time limitations, which escalate as the number of controllable vehicles rises. This motivated us to directly extend the code-base of SUMO, rather than relying on TraCI. While

SUMO does not support connected vehicles or vehicles with custom controllers, we implement such functionalities by “overriding” its default behavior. Therefore, our contribution can be viewed as an additional plug-in to the existing SUMO implementation.

We should mention here that, besides TraCI, one alternative option would be to consider Libsumo, a low-level Application Programming Interface (API) that does not exhibit the computational limitations imposed by socket communication. However, this choice was deemed to be rather restrictive, and thus inappropriate for our requirements. Note that Libsumo does not support the GUI application, meaning that it works only with the command line version of SUMO. In addition, the functionalities of the APIs (TraCI or Libsumo) would be limiting for some of the extensions we provide; or would require additional computational tasks to monitor or control the traffic environment properly. By working directly with the open-source code-base of SUMO, we can meticulously exploit the existing infrastructure and extend it appropriately.

Thus, we retain the current functionality, and develop complementary components that can both monitor online a lane-free traffic environment, and also control vehicles considering the lane-free paradigm. This is achieved via an additional *dynamic library* we implemented, which allows development in C/C++.<sup>1</sup> In more detail, the library provides the user with the ability to follow a core coding structure which contains an *initialization* (and *finalization*) function. These functions, as the naming suggests, are executed once at the initialization (respectively, the termination) of the traffic environment. Also a function that is executed in every simulation time-step is introduced. There, a user can develop code in order to test lane-free vehicle controllers, emulating vehicle connectivity aspects via retrieving of information from SUMO, since we provide a concrete and easy to use API, enabling the direct communication with SUMO without any time-related overhead (that TraCI imposes). Fig. 1 shows an overview of the components that form *TrafficFluid-Sim*.

By means of this library, a user can utilise the available API in order to obtain information regarding:<sup>2</sup> the vehicles’ properties (type, length, width, etc.) or current status (position, speed, etc.); properties of a specific road segment

<sup>1</sup>.dll file (or .so file in Linux)

<sup>2</sup>Here we simply highlight the features, as they will be presented in more detail in the following sections.

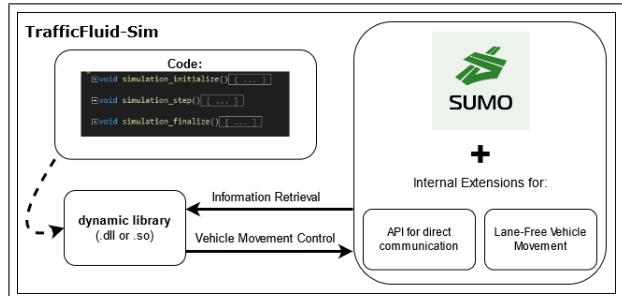


Fig. 1. An overview of TrafficFluid-Sim’s components.

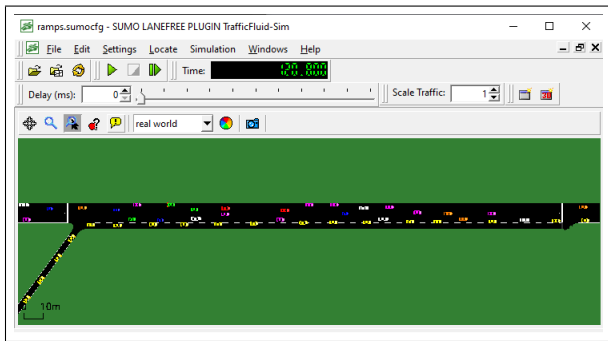


Fig. 2. A snapshot of the TrafficFluid-Sim application.

(ids of vehicles operating in this segment, length, width, additional information for given regions of the road); and online information from loop detectors. Additionally, one can control the vehicles through their unique ID with the API, where longitudinal and lateral acceleration values are supplied as control inputs to the relevant function.

Fig. 2 shows a snapshot of the application for a lane-free traffic scenario with an on-ramp. It is evident that the vehicles operate with no restrictions whatsoever regarding lateral placement. Lane-markings only exist to indicate the acceleration lane for the vehicles to merge on the lane-free highway.

In the following subsections, we outline key features and functionality offered; and present the different aspects of the implementation in more detail.

### B. Lane-Free Vehicle Movement

Networks in SUMO are composed by multiple connected road segments (edges), essentially depicted by a graph representation, and SUMO maintains local coordinates for each vehicle, with respect to the road segment the vehicle is residing. For the longitudinal coordinate, the value is essentially reflecting the distance from the respective segment’s starting point. In a lane-free environment, this value is suitable, so no substantial changes are needed. However, this is not the case for the lateral coordinate, where the value we can obtain from SUMO (besides information about the lane we are currently in), is the lateral distance from the respective lane’s center. Moreover, SUMO does not actually consider lateral dynamics, since it does not involve anywhere lateral speeds or acceleration. One can find information about lateral speeds only within the Sublane model<sup>3</sup> of SUMO, but this is only relevant to the lane-changing model, whereby the vehicles simply move by providing a new lateral distance deviation value.

These features are not adequate for our purposes, as, in order to model lane-free movement, we need to consider both *longitudinal* and *lateral* dynamics with constant respective accelerations for each simulation time-step. SUMO supports this only for the longitudinal motion of vehicles.<sup>4</sup> Thus, we added information about the lateral speed of the vehicles

<sup>3</sup>Sublane depicts more realistic lane-changing behavior.

<sup>4</sup>Through the ‘Ballistic-Update’ option of SUMO.

“internally” (i.e., by altering the existing SUMO simulator code), and modified the update process of the vehicles’ lateral position.

In more detail, the *local* coordinates  $(x,y)$  of the rectangular-shaped vehicles that we provide are as follows: The longitudinal position  $x$  of the vehicle is now the distance of its center point from the starting point of the road segment. Accordingly, the lateral position  $y$  measures the distance from the right road boundary to the vehicles’ center point. Hence, the vehicles observe a single lateral position, without knowledge about lateral placement w.r.t. to lane centers, as this information is not meaningful anymore for designing a vehicle movement control strategy under lane-free settings.

### C. Handling of Demand

In contrast to real traffic networks, any simulated network infrastructure has upstream boundaries, at which vehicles need to be inserted at a pre-specified demand rate. This is a delicate operation, because the simulated traffic conditions (e.g. dense traffic or congestion) may not allow for the whole demand to enter, as this might lead to unrealistically high vehicle densities just after insertion. Thus, vehicles should only be inserted if there is enough space to the vehicle in front; and with an insertion speed that would not lead to immediate strong deceleration, or even crashing with the front vehicle. Vehicles that cannot be inserted, due to such restrictions, are placed in a virtual queue to be considered for insertion later on. On the other hand, the vehicle insertion policy should not be too conservative, as this might lead to the creation of virtual queues even in cases where the demand is lower than the road capacity.

In a regular lane-based scenario, one normally specifies a demand of vehicles either for individual lanes, or for the whole edge (road segment). For the latter option, SUMO automatically distributes the demand to each lane. Contrary to a lane-based environment, where a simple time-gap or space-gap policy can be applied, such a policy is not appropriate for lane-free environments, since it would not exploit laterally the actual capacity of the road, which constitutes one of the intrinsic properties when operating under the lane-free paradigm. Recall that the space-gap ( $s_g$ ) is the distance from the front bumper of the vehicle attempting to enter, to the rear bumper of the front one, and the time-gap ( $t_g$ ) is defined as  $t_g = s_g/v_p$ , where  $v_p$  is the speed of the vehicle attempting to enter. Likewise, if we have any two quantities, we can calculate the remainder one using the equation above.

Therefore, we designed a new method that determines where (laterally) a vehicle shall be inserted in the road; as well as whether it should actually be inserted or stay in a virtual queue of vehicles in case of lack of sufficient space due to other vehicles in front. In the latter case, the vehicle has to wait for the next time-step before re-attempting insertion. To this end, we need to establish how the vehicles will spawn at an entry point, when defining a global cross-section, in a way that is organically tied to a lane-free environment. Essentially, we use a constant time-gap policy (provided by the user) customised for this setting, which also




Fig. 3. Vehicle is inserted randomly within the range of one of the available lateral regions (indicated with light blue).

accounts for the width of the vehicle when entering. As such, each vehicle attempting to enter (according to the demand) has a set of available lateral regions that it can spawn.

Initially, the vehicle considers a single lateral region that corresponds to the whole road width. Then, its insertion speed will be the minimum of the user-defined departure speed and the average speed of the front  $n = 5$  vehicles. This bounding serves to prevent collisions in case the region is congested, where the vehicles move with lower speed (or even in extreme cases they cannot move at all). Given the insertion speed and the pre-specified time-gap, the corresponding required space-gap in front of the vehicle is calculated, and the next step is to obtain a lateral position to spawn according to the downstream traffic.

To achieve that, we first scan (in ascending order, according to longitudinal position) the front vehicles until one is found that conforms to the calculated space-gap. Then, for all front vehicles that have a space-gap less than the required value, we restrict, according to their widths and lateral positions on the road, the available lateral regions. Finally, the vehicle randomly spawns in one of the remaining lateral regions, which guarantees that there is no vehicle in front with less than the required space-gap. Fig. 3 demonstrates visually the aforementioned procedure. Of course, if, at any point in the above procedure, the set of available lateral regions becomes empty (i.e., there is no lateral position to insert the vehicle with an acceptable space-gap in front of it), then the procedure is stopped, and the vehicle does not spawn, but is preserved in the virtual queue of SUMO. Vehicles within this virtual queue have the opportunity to enter in subsequent time-steps, when front vehicles will have advanced. It should be noted that in the aforementioned process, we impose a minimum lateral distance between the vehicles, that is selected by the user.

In Figs. 4 & 5, we illustrate instances of lateral vehicle placement for two different respective demands, according to the aforementioned insertion policy.

Simulated vehicles possess a desired speed that they strive to achieve when and where the traffic conditions allow. This desired speed is decided at insertion point, and accompanies the vehicles, until their exit from the network. The user may select desired speeds for the entering vehicles via the API, based on a range of desired speeds  $[v_{d,low}, v_{d,high}]$ , sampled according to a distribution (e.g. uniform). The desired speed distribution may be independent of the lateral vehicle insertion positions, but a specific consideration for

the selection of desired speeds may be to establish a mapping with the range of initial lateral placements. Specifically, some vehicle movement strategies favor a lateral vehicle position, while driving, that is proportional to the vehicle's desired speed, i.e., vehicles with higher desired speed tend to drive further left on the road, compared to vehicles with lower desired speed (e.g. trucks). In such cases, it is reasonable to relate accordingly the lateral insertion position of each vehicle with its assigned desired speed at the entry points, so as to avoid unrealistic weaving maneuvers after insertion. For instance, each desired speed  $v_d$  can be assigned based on a linear mapping from the lateral position range to the range  $[v_{d,low}, v_{d,high}]$  of desired speeds (accounting also for the width of the vehicle, e.g., when the vehicle has width  $w_v$ , and the road has width  $w_r$ , the vehicle considers the lateral region  $[w_v/2, w_r - w_v/2]$ , so if it spawns at  $w_v/2$ , then the selected desired speed will be  $v_{d,low}$ ).

The aforementioned insertion policy aims to fully exploit the available lateral space. However, notice that this method, along with mentioned linear mapping will not result in a uniform (or other desired) distribution of desired speeds. Therefore, for specific simulation scenarios requiring a specific probability distribution, we also developed an alternative insertion policy, in which we refrain from exploiting fully the lateral space.

In the alternative policy, we initiate the process of introducing a new vehicle to an entry point by drawing a sample  $v_d$  from the user-chosen distribution (provided by the API), within the range  $[v_{d,low}, v_{d,high}]$ . This sample value  $v_d$  will be the desired speed of the vehicle. Thereafter, the policy decides upon the lateral placement for the vehicle. This is done via a linear mapping from the range  $[v_{d,low}, v_{d,high}]$  to the range of available lateral positions  $[w_v/2, w_r - w_v/2]$ , according to the sampled  $v_d$  value. Next, the first vehicle in front of the vehicle to be inserted is identified, and a candidate insertion speed  $v_p$  is calculated according to the time-gap provided by the user and the space-gap between the two vehicles. Note that we consider that a vehicle is in front, if the lateral regions resulting from their lateral placements and width overlap. This is also noticeable in Fig. 6, where we show visually the above procedure. The vehicle will be inserted with insertion speed  $v_{insert} = \min\{v_p, v_d\}$ , unless  $v_{insert} < \min\{v_d, v_f\}$ , where  $v_f$  is the speed of the vehicle in front. In that case, the vehicle will remain in the virtual queue and attempt reentering with the same desired speed (and, consequently, the same lateral position) in the next time-step.




Fig. 4. Vehicle placement for a demand of 7500veh/h, time-gap of 0.5s, minimum lateral distance of 0.25m.




Fig. 5. Vehicle placement for a demand of  $12000\text{veh/h}$ , time-gap of  $0.5\text{s}$ , minimum lateral distance of  $0.25\text{m}$ .

In what follows, we showcase its use (assuming uniform distribution for the desired speeds), against that of the first method in a scenario with an extremely high value of demand, specifically, one of  $18,000\text{veh/h}$ . This is a demand clearly beyond the capacity of any lane-based approach for a road width as the one in Table I, corresponding to a road comprised of 3 lanes. We compare the two insertion policies in terms of the flow values achieved in each case.

We collect measurements of the flow with a sampling period of  $T = 2\text{min}$ . All relevant parameters are tuned according to Table I, where  $d_{lat,min}$  is the minimum lateral distance the vehicles are allowed to have upon entering;  $l_v, w_v$  are the length and width of the vehicles respectively (which are equal for all vehicles in this scenario, for simplicity); and  $w_r$  is the width of the road.

As observed in Fig. 7, where we display measured flows for the two different insertion policies, the first policy (indicated with method 1 in the figure) yields an average flow of  $12,998\text{veh/h}$ . This effectively sets an upper bound for the achievable maximum flow, given the parameter choices of Table I. The vehicle movement strategy used to obtain these results is an ad-hoc strategy according to the one presented in [8]. On the other hand, the alternative policy (indicated with method 2 in the figure), using a uniform distribution  $\mathcal{U}(25,35)$ , was observed to lead to under-utilisation of the road capacity. We now have vehicles spawning with desired speeds according to a uniform distribution within the range of  $[25,35]\text{m/s}$ , as desired, but this results into a lower average flow of  $8,312\text{veh/h}$ . Thus, maintaining such a distribution comes with a cost of a reducing by 36% for the average




Fig. 6. Vehicle has a predefined initial lateral position, and obtains an insertion speed according to the space-gap from the vehicle in the front.

TABLE I  
PARAMETER CHOICES FOR THE DEMAND EXAMINED.

time-gap	$d_{lat,min}$	$l_v$	$w_v$	$w_r$
$0.5\text{s}$	$0.25\text{m}$	$3.2\text{m}$	$1.6\text{m}$	$10.2\text{m}$




Fig. 7. Measured flows of vehicles from a detector placed 100 meters downstream of the entry point.

flow inserted (w.r.t. the one achieved by the first policy (method 1)). These results indicate the expected trade-off between demand insertion at the capacity of the road versus guaranteeing that the distribution of desired speeds will be exactly mapped with the lateral insertion position of entering vehicles.

#### D. Using a Dynamic Library for Lane-Free Vehicle Control

Normally, vehicle strategies in lane-based SUMO are structured according to two types of models, a ‘‘Car-Following’’ model, and a ‘‘Lane-Changing’’ one, that regulate the longitudinal speed and the lane-placement accordingly. This type of architecture is obviously not appropriate for lane-free vehicle controllers. Also, something crucial, that we need to address, is the fact that these models are pre-compiled, and constitute a part of the SUMO environment. This means that the users can select one of the available models, and ‘tune’ them via model-related parameters. There is no flexibility to easily incorporate different controller designs to this environment, unless we do so by accessing the open-source code-base of SUMO (or via TraCI).

Especially for a novel traffic paradigm, like *TrafficFluid*, a more flexible design, that allows for interdisciplinary research approaches and refined control of the vehicles, is fundamental. Therefore, we have made the necessary developments to make use of an external dynamic library (that is linked with the main application at execution time) in order to enable the development of any type of vehicle controller just by providing the longitudinal and lateral acceleration values for each simulation time-step. Importantly, this essentially means that *any* vehicle strategy can be incorporated without the need to re-compile SUMO; moreover, after its compilation, the library is essentially treated for all purposes as an integral part of the application, thus avoiding all communication costs between SUMO and external frameworks (that would have been otherwise unavoidable). In conclusion, we have the benefit of flexibility, but without the cost of time-delays.

#### E. Dynamic Library Structure and API features

Our library consists of 3 core functions that can be utilised:

- *simulation\_initialize*: Executed once before the first time-step; intended for initialization of variables, memory, files, etc.
- *simulation\_step*: Executed once in every time-step. Within this function, users can control the vehicles, and monitor online the traffic environment
- *simulation\_finalize*: Executed once before exiting the simulation; intended for deallocation of memory, saving log files, etc.

The user can build upon the development environment and communicate directly with SUMO via the API provided by our implementation (see Fig. 1). Regarding the memory structure that contains the vehicles for a given road segment, vehicles are always retained in ascending order, according to their respective (local) longitudinal position  $x$ . This information is updated dynamically, since vehicles have different speeds and, also, are following a route and may thus change road segments. Such ordering is quite convenient for many reasons. First, it allows us to perform collision checks much more efficiently; and to perform the dynamic allocation of vehicles to road segments much faster. Furthermore, this information can be provided to the user via the API.

For the remainder of this subsection, we showcase the current capabilities of our API:

1) *Vehicles*: A user can obtain information, such as: the ID's of all vehicles in the network; or only the ID's of vehicles in a specific road segment. In the latter case, they are obtained in ascending order according to their (local) longitudinal position. Given a specific vehicle's ID value, information relevant to this vehicle is easily available, such as: its longitudinal and lateral positions (either local coordinates, as specified in Section III-B, or global coordinates); longitudinal and lateral speeds, dimensions (length, width); the vehicle's type id (as defined when setting up the traffic environment); the vehicle's name (the one displayed by SUMO).

Besides information retrieval, one can *control* a vehicle via its ID, providing the longitudinal and lateral acceleration values (in  $m/s^2$ ) as inputs to the relevant function. Moreover, through the API, the user can manually *insert* a new vehicle either at the initialization phase, or at any arbitrary simulation time-step. In this case, the initial conditions and additional information should be provided by the user. That information involves: initial position and speed (longitudinal and lateral), a (valid) vehicle type, and a (valid) route for the vehicle to follow. We also retain internally a *desired speed* for each vehicle; and, through the API, we can obtain or change the existing value.

As mentioned, we use local coordinates, something quite convenient for designing a vehicle movement strategy, since the vehicle just follows a given route, and does not need to take into account the geometry of the road. However, this poses a limitation when a vehicle needs to observe downstream traffic in the next road segment, as the local coordinates of downstream vehicles correspond to a different road segment. For this reason, considering two vehicles that can potentially be located on different road segments, we pro-




Fig. 8. The appointed vehicle, located on an on-ramp and about to merge in the highway, has direct access to the vehicles located in the next segment.

vide information about the longitudinal and lateral distance from the perspective of one of the two vehicles. Complementary to the aforementioned functions is another feature that allows for each vehicle to easily observe *downstream* and *upstream traffic*, according to their routing scheme, i.e., they can observe vehicles that may be located in the upcoming or previous road segments. This essentially translates to obtaining an array of all the downstream (upstream) vehicles' ID values (up to a specified distance), in ascending order, based on the longitudinal distance from the ego-vehicle that request such information, even if the downstream (upstream) vehicles are beyond the current road segment. For example, a vehicle in an on-ramp, which is advancing to merge on the mainstream, can directly observe downstream traffic from the main highway and be properly informed about the distances, as illustrated in Fig. 8.

Moreover, we developed the necessary components to emulate a ring-road scenario using a single road segment. In this case, the vehicles are directly "transmitted" to the starting point of the segment instead of exiting. Longitudinal distance  $d_x(i, j)$  is properly adjusted for this feature, e.g., a vehicle towards the end of the road segment will properly observe vehicles at the starting point. In order to introduce vehicles in the ring-road, the recommended solution is to utilise the API, and insert the desired density of vehicles when initializing the simulation. When this functionality is extended to incorporate multiple road segments, the demand can be alternatively generated from an on-ramp, connected to the main highway that emulates the ring-road.

2) *Road Segments*: Similarly, users can obtain an array of the road segments' ID values in the network. For a given road segment, a user can obtain its properties (length, width), and additional online information for user-defined regions. Essentially, at any simulation step, we can directly gain information for a specific region of the road segment regarding the traffic density at that region, or the average speed of the vehicles. This is illustrated in Fig. 9, where such information will be provided only according to the vehicles within the appointed region.

In addition, there is the flexibility to use these functions only for a specific type of vehicles, i.e., we can independently monitor each type of vehicles in the network. For instance, in Fig. 9, we may be interested only in the type of vehicles indicated with white color, thus the information we will receive will be calculated from these 3 vehicles only.

The outlined functionalities related to road segments are valuable for several possible applications, including real-time

traffic management via emulated V2I communication.

3) *Detectors*: One of the features of SUMO, when designing a traffic scenario, is the support of multiple types of detectors, which can be placed anywhere throughout the network. The simple ‘E1’ Loop Detectors of SUMO are integrated with our internal extensions. Likewise, these loop detectors are represented by unique ID values. With this, we can monitor online the number of vehicles that have passed from the corresponding detector. Additionally, we retain the same flexibility here regarding monitoring specific types of vehicles.

Lastly, we have embedded a set of *event-based* functions, triggered when specific situations occur, such as: a new vehicle entering the network, a vehicle reaching a final destination; a collision between two vehicles; or a vehicle being located out of the road boundaries. Obviously, these functions are accommodated with information regarding the involved vehicles’ IDs.

#### F. Designing a Lane-Free Traffic Scenario

We should emphasise that we can design traffic scenarios exactly as we would normally do for the standard SUMO application. In order to design a traffic scenario, the core components that we need is a road network, specification of vehicle types along with their properties, and finally the demands.

Lane-free vehicles can now be incorporated through the definition of a vehicle type. We have included an additional ‘virtual’ model option (besides the ones provided by SUMO), that the assigned vehicle types can now employ. When a user specifies the model as “lane-free”, it will not conform to any standard lane-based model, but the user will be granted access to the control of all vehicles of this type through the dynamic library. Thus, these vehicles will now move according to the developed movement strategy. Any type of lane-free controller can be developed and adapted to the same structure, provided that it yields the longitudinal and lateral accelerations (in  $m/s^2$ ) in every simulation time-step.

#### IV. EXTENSIONS

Currently, vehicles are operating according to a predefined route, simply following a specified path from origin to destination. We rely on SUMO, which requires either just an origin and a destination segment (then a valid path is automatically generated upon initialization of the vehicles), or a (valid) full path consisting of the full sequence of road




Fig. 9. We can directly request information about density or average speed just by providing the upper and lower bounds of the region, e.g., if what is displayed is the total road segment, its length is 150 meters, and we are interested only in the last 50 meters, we define the region as [100, 150].

segments the vehicle should follow, if we wish more refined control. Also, SUMO provides some options regarding re-routing (either online or offline) of the vehicles, but none of them conforms with the flexibility we need in order to support the design of V2I applications for traffic management in a lane-free environment.

One impending extension is to provide the necessary functionalities for allowing the user to assign turn probabilities at bifurcation nodes of the road network, in a lane-free traffic environment, omitting the origin-destination routing of the vehicles for specific applications.

Another forthcoming extension is the incorporation of *jerk* [22] dynamics (both longitudinal and lateral), meaning that the vehicles will have the option to operate based on jerk commands from the vehicle control, where jerk is the time derivative of acceleration. To this end, we need to update the simulation of vehicle movement accordingly. Of course, the default option will still be to control through acceleration, and jerk control will constitute an alternative that the user will be able to switch to via either the API or through an additional parameter when specifying a type of vehicles.

#### V. RELATED WORK

In this section we provide a brief overview of some of the most popular external frameworks for SUMO, and further motivate our choice to directly extend this open-source project.

To begin with, Veins [23] is a network simulator that emphasizes on the realistic depiction of communication structures, simulating IEEE 802.11p and IEEE 1609.4 DSRC/WAVE network layers, and providing an extensive list of useful metrics. “iTETRIS Control System” (iCS) [24] delivers a platform for large-scale simulation and assessment of Cooperative Intelligent Transportation Systems (C-ITS). They additionally aspire to develop novel traffic control and communication strategies, along with new metrics that capture the overall performance of the network in terms of common traffic issues, such as: vehicle emissions, traffic congestion, travel time etc. VSimRTI [25] is aimed at providing a comprehensive and user-friendly environment. It embraces likewise realistic communication structures and energy consumption and is suitable for large-scale simulations. Flow [26] is a modular framework that facilitates the use of machine learning (in particular, reinforcement learning) methods to address the complexity of traffic dynamics. Likewise, Flow communicates with SUMO via TraCI, and is developed in Python. It also facilitates the creation of a generic simulator, appropriate for microscopic traffic modelling and the embedding of custom vehicle controllers.

All aforementioned frameworks are flexible and extensible, at least to a certain extent. Still, to accommodate our needs, considerable source code modifications and extensions would have been required, since no single framework provides our desired TrafficFluid capabilities, while any such framework would induce a substantial additional overhead (w.r.t. time and memory) for communicating with SUMO. This led us to the choice of developing a direct extension of

the SUMO framework, purpose-built for modeling the novel lane-free traffic paradigm. In addition, the new framework is flexible enough to allow its easy incorporation into existing code-bases, and also for future extensions.

## VI. CONCLUSIONS

We presented *TrafficFluid-Sim*, a new, modular, and extensible simulator, that supports a novel paradigm of lane-free traffic, such as the one coined by *TrafficFluid* [8]. We discussed the main components of the new simulator, as well as forthcoming extensions. The simulator is already employed for the evaluation and expansion of vehicle movement strategies [8], [11], [12] under the novel traffic paradigm. Videos illustrating the presented insertion policies, and a simulation example with one on-ramp when employing an ad-hoc movement strategy (see [8]), are available at <https://bit.ly/3twy4qc>. We intend to expand the existing *TrafficFluid-Sim*'s components and, following its evaluation and feedback by the community, to elevate it into a fully-fledged, feature-rich platform that will enable the rapid progress of novel lane-free traffic research approaches; and, more generally, serve the community's needs and demands.

## REFERENCES

- [1] D. Elliott, W. Keen, and L. Miao, "Recent advances in connected and automated vehicles," *Journal of Traffic and Transportation Engineering (English Edition)*, vol. 6, no. 2, pp. 109–131, 2019.
- [2] I. Papamichail, N. Bekiaris-Liberis, A. I. Delis, D. Manolis, K.-S. Mountakis, I. K. Nikolos, C. Roncoli, and M. Papageorgiou, "Motorway traffic flow modelling, estimation and control with vehicle automation and communication systems," *Annual Reviews in Control*, vol. 48, pp. 325–346, 2019.
- [3] P. Gora and I. Rüb, "Traffic models for self-driving connected cars," *Transportation Research Procedia*, vol. 14, pp. 2207–2216, 2016, transport Research Arena TRA2016.
- [4] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using sumo," in *21st IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2018.
- [5] J. Casas, J. L. Ferrer, D. Garcia, J. Perarnau, and A. Torday, *Traffic Simulation with Aimsun*. New York, NY: Springer New York, 2010, pp. 173–232.
- [6] K. Axhausen, K. Nagel, and A. Horni, Eds., *Multi-Agent Transport Simulation MATSim*. London: Ubiquity Press, Aug 2016.
- [7] M. Rahman, M. Chowdhury, Y. Xie, and Y. He, "Review of microscopic lane-changing models and future research opportunities," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1942–1956, Dec 2013.
- [8] M. Papageorgiou, K. S. Mountakis, I. Karafyllis, I. Papamichail, and Y. Wang, "Lane-free artificial-fluid concept for vehicular traffic," *Proceedings of the IEEE*, vol. 109, no. 2, pp. 114–121, 2021.
- [9] On-Road Automated Driving (ORAD) committee, *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*, 2018.
- [10] I. Karafyllis, D. Theodosis, and M. Papageorgiou, "Analysis and control of a non-local pde traffic flow model," *International Journal of Control*, 2020.
- [11] D. Troullinos, G. Chalkiadakis, I. Papamichail, and M. Papageorgiou, "Collaborative multiagent decision making for lane-free autonomous driving," in *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. International Foundation for Autonomous Agents and Multiagent Systems, 2021, pp. 1335–1343.
- [12] V. K. Yanumula, P. Typaldos, D. Troullinos, I. Papamichail, and M. Papageorgiou, "Optimal path planning for connected and automated vehicles in lane-free traffic," in *24th IEEE International Conference on Intelligent Transportation (ITSC)*, 2021.
- [13] I. Karafyllis, D. Theodosis, and M. Papageorgiou, "Two-dimensional cruise control of autonomous vehicles on lane-free roads," *arXiv preprint arXiv:2103.12205*, 2021.
- [14] R. Dorf and R. Bishop, *Modern Control Systems*, 13rd ed. Pearson, 2017.
- [15] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed. USA: Prentice Hall Press, 2009.
- [16] T. M. Mitchell, *Machine Learning*, 1st ed. USA: McGraw-Hill, Inc., 1997.
- [17] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. The MIT Press, 2018.
- [18] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [19] M. Wooldridge, *An Introduction to MultiAgent Systems*, 2nd ed. Wiley Publishing, 2009.
- [20] N. Vlassis, "A concise introduction to multiagent systems and distributed artificial intelligence," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 2007.
- [21] M. Malekzadeh, I. Papamichail, M. Papageorgiou, and K. Bogenberger, "Optimal internal boundary control of lane-free automated vehicle traffic," *Transportation Research Part C: Emerging Technologies*, vol. 126, no. 103060, 2021.
- [22] I. A. Ntousakis, I. K. Nikolos, and M. Papageorgiou, "Optimal vehicle trajectory planning in the context of cooperative merging on highways," *Transportation Research Part C: Emerging Technologies*, vol. 71, pp. 464–488, 2016.
- [23] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing (TMC)*, vol. 10, no. 1, pp. 3–15, January 2011.
- [24] M. Rondinone, J. Maneros, D. Krajzewicz, R. Bauza, P. Cataldi, F. Hrizi, J. Gozalvez, V. Kumar, M. Röckl, L. Lin, O. Lazaro, J. Leguay, J. Härrri, S. Vaz, Y. Lopez, M. Sepulcre, M. Wetterwald, R. Blokpoel, and F. Cartolano, "iTETRIS: A modular simulation platform for the large scale evaluation of cooperative its applications," *Simulation Modelling Practice and Theory*, vol. 34, pp. 99–125, 2013.
- [25] B. Schünemann, "V2X simulation runtime infrastructure VSimRTI: An assessment tool to design smart traffic management systems," *Computer Networks*, vol. 55, no. 14, pp. 3189–3198, 2011.
- [26] C. Wu, A. Kreidieh, K. Parvate, E. Vinitzky, and A. M. Bayen, "Flow: A modular learning framework for autonomy in traffic," *arXiv preprint arXiv:1710.05465*, 2017.