A holistic approach to the energy-efficient smoothing of traffic via autonomous vehicles

Amaury Hayat, Xiaoqian Gong, Jonathan Lee, Sydney Truong, Sean Mcquade, Nicolas Kardous, Alexander Keimer, Yiling You, Saleh Albeaik, Eugene Vinistky, et al.

A holistic approach to the energy-efficient smoothing of traffic via autonomous vehicles

Amaury Hayat, Xiaoqian Gong, Jonathan Lee, Sydney Truong, Sean McQuade, Nicolas Kardous, Alexander Keimer, Yiling You, Saleh Albeaik, Eugene Vinistky, Paige Arnold, Maria Laura Delle Monache, Alexandre Bayen, Benjamin Seibold, Jonathan Sprinkle, Dan Work, Benedetto Piccoli

Abstract The technological advancement in terms of vehicle on-board sensors and actuators, as well as for infrastructures, open an unprecedented scenario for the management of vehicular traffic. We focus on the problem of smoothing traffic by controlling a small number of autonomous vehicles immersed in the bulk traffic stream. Specifically, we aim at dissipating stop-and-go waves, which are ubiquitous and proven to increase fuel consumption tremendously and reduce. Our approach is holistic, as it is based on a large collaborative effort, which ranges from mathematical models for traffic and control all the way to building infrastructures capable of measuring energy efficiency and providing real-time data. Such an approach allows to clearly set and measure a metric for success in the form of a reduction of at least 10% of fuel consumption using 5% of autonomous vehicles immersed in bulk traffic. The chapter illustrates the overall approach and provides simulation results on a tuned microsimulator for the California I-210.

1 Introduction

1.1 Stop-and-go waves

It is commonly known that having many drivers at the same speed on the road and the same distance between vehicles can lead to jam [49, 46]. Stop-and-go waves can seem to appear without any particular reason (no road reduction, no roadwork ahead, etc.). There is however a simple explanation to this phenomena: steady states in traffic flows are sometimes unstable [7]. Stop-and-go waves have a large impact both on the economy and sustainability of traffic. Not only can they reduce the

Arizona State University (XG), Ecole des Ponts Paristech (AH), INRIA Grenoble - Rhône Alpes (MLDM), Rutgers University Camden (PA,SMQ,BP,ST), Temple University (BS), University of Arizona (JS), University of California at Berkeley (SA, AB, NK, JL, EV, YY), Vanderbilt University (DW)

outflow compared to a steady-state situation where all vehicles are equally spaced and have the same speed, but in addition the fuel consumption and CO2 emissions associated with constantly braking and accelerating are much higher than the fuel consumption and CO2 emission of the corresponding steady-state [39].

This chapter introduces an initiative aiming at reducing stop-and-go waves by incorporating a small number of autonomous vehicles (AV) in the traffic and using it as a system controller. After presenting an overview of the project in Section 1, we focus in Sections 2 to 3 on a control design approach derived from theory-based model, and in Section 4 we present some results when using the resulting controller in simulations on a highway model. In particular, we show that in a highly congested traffic this controller allows for a strong reduction of fuel consumption for the same throughput, even with a very low proportion of AVs in the traffic.

1.2 An initiative to increase energy efficiency

Finding ways to increase energy efficiency, i.e., the miles traveled per gallon of fuel (denoted "fuel economy" in the US) or distance traveled per Joule of electric power, through traffic control is not a new idea. Numerical simulation results and stability analysis to regulate traffic flow via autonomous vehicles are available, see [10, 20, 47, 51]. Other approaches used variable speed limits, see [2, 21, 52], or jam absorption [23, 34]. Our project is based on evidence shown by the experimental results on a ring road, see [44, 45, 53].

However, the energy efficiency of today's vehicular mobility relies on the combination of two elements: control via static assets (traffic lights, metering, variable speed limits, etc.); and onboard vehicle automation (adaptive cruise control (ACC), ecodriving, etc.). These two families of controls were not co-designed and are not engineered to work in coordination. Recent studies have shown limitations of controls, and even sometimes negative impacts of ACC [30]. Our general approach focuses on the technology development, implementation and prototyping, and validation of Mobile Traffic Control (MTC). In other words using AVs as mobile controllers in the traffic flow.

This can be viewed as an extension of classical traffic control (in which static infrastructure actuates traffic flow). In the MTC paradigm, automated vehicles impact the entire surrounding traffic via their behavior, offering enhanced possibilities to optimize the energy footprint of traffic, if designed correctly.

We want to demonstrate for the first time that considerably reduced fuel consumption of all vehicles in traffic can be achieved via distributed control of a small proportion of controlled autonomous vehicles (CAV)s. The level of autonomy expected for the CAV is level 2 (according to the SAE taxonomy J3016_201806). Compared to baseline vehicular technologies, our work offers a significant design departure: control algorithms for the CAV consider the impact one vehicle can have on overall traffic, improving resulting overall fuel consumption. We focus on using a few vehicles as traffic controllers (via CAV technology) to improve the energy ef-

ficiency of traffic flow to further optimize energy efficiency. The target is to achieve energy gains exceeding 10% in average for all vehicles on the road when the traffic is congested, through automation of less than 5% of the vehicles in the flow (called penetration rate in the following).

This ambitious goal is motivated by prior field experiments that demonstrated fuel consumption reductions of up to 40% with a single CAV on a single lane track under ideal conditions [45]. Of course, for real life multilane traffic many new difficulties have to be taken into account such as lane-changing, drivers' responses to actuation, interaction between CAV, topography of the road, etc.

To achieve our aim, several challenges need to be addressed on the technical level:

- Establishing the minimum sensing and connectivity required for eliminating traffic waves with mobile actuation;
- Investigating control requirements to dampen stop-and-go traffic;
- Designing simulation models representing reliably the multilane stop-and-go traffic and the relevant behaviors;
- Deducing a precise and realistic energy model for vehicles consumption, taking into account the different types of vehicles;
- Designing sensing systems on a highway as well as estimation algorithms to detect the traffic state using on-board vehicle sensing and/or infrastructure sensor networks;
- Finding efficient control algorithms either from accurate mathematical models and tools or from machine learning methods.

In the following, we present the main methods and achievements related to each of these issues, before focusing on the design of control algorithms from microscopic mathematical models.

Testbed Development

Research and development for a testbed began with evaluating modern Close-Circuit Television (CCTV) camera systems for possible deployment on roadside poles for traffic observation. A determination of high-resolution (4K) cameras, capable of individual pan-tilt-zoom control, was made. Roadside poles at a height of 110ft were selected to minimize optical occlusion between vehicles. Six cameras could be housed on a single roadside pole with their fields of view adjusted for complete coverage of 500 linear feet of roadway. These developments are done on a portion of the I-24 (Tennessee), in complement to the test site of the I-210 California presented in this chapter. Motivations for this site were roadside site layout, traffic patterns, and existing data sources. Testing of this proposed configuration was conducted on an existing roadside pole owned by the Tennessee Department of Transportation (TDOT). The test also provided valuable video data for computer vision algorithm development. The poles carry 18 4K resolution cameras that can provide continuous coverage of traffic on I-24, capturing trajectories from $\approx 150,000$ vehicles daily [17].

Fig. 1: Installation of a roadside pole on the I-24

Fast car tracking computer vision methods

To meet the needs of our project, we required a new method for object tracking. Existing methods, while fairly accurate, fall far short of the testbed needs in terms of speed; a fast method is necessary to process video data as it is streamed, and to provide high-fidelity estimates of vehicle positions. The core innovation of this new tracking method is to rely on a convolutional neural network for object localization within a small crop of a video frame, rather than utilizing object detectors that specialize in finding all object locations within an entire video frame. The lightweight localization network is able to run much faster than the object detector at only a small cost in accuracy. Computer vision algorithms have been continually developed on video data from cameras on roadside traffic monitoring poles, mimicking the setup of the I-24 testbed. This has enabled algorithms to be tuned for performance with large numbers of objects in view, which exhibit realistic vehicle dynamics within the field of view.

Sensing and hardware development

Modern vehicles (since 1995 in the United States) operate through an system of control units and sensors that are interconnected through a Controller Area Network (CAN) architecture. Messages sent between electronic control units throughout a vehicle are sent through a CAN bus (or in some cases, on multiple buses), which provides an opportunity to detect the vehicle state directly by recording these messages. While there exist a bevy of off-the-shelf solutions to inspect CAN messages or (in some cases) record these data, such solutions do not reliably deliver the messages in a timely way that permits for closed-loop control of the system.

To address these limitations in data fidelity and timeliness, we developed a custom library called `libpanda` that interfaces with vehicle's CAN bus through optimized software and off-the-shelf hardware for data acquisition in real-time [5, 50]. `Libpanda` is a multi-threaded C++ based library for custom Panda interface applications (such as those in [41]) using the observer based software design pattern (i.e. callbacks). Also featured in the software area are pre-made data recording utilities, simple data visualizers, startup services, and network handling services.

`Libpanda` utilities save CSV-formatted data that are of interest to vehicle state estimation, control, and in some cases estimation of surrounding traffic. Recorded data values include instantaneous velocity, acceleration, wheel-speeds, steering angle (and rate), accelerator angle, brake angle, cruise control settings and state, as well as other vehicle states. Localization and information on the surroundings include radar trace information from vehicles equipped with adaptive cruise control systems, permitting estimation of following distance and lead-vehicle dynamics. Additional contextual information is provided by synchronized GPS sensors that come standard with the off-the-shelf CAN-bus hardware. `Libpanda` records, for each drive, a pair of files that are synchronized to GPS time, and which include the CAN data in one file, and GPS data in the other.

Downstream analysis of these data files are performed by our data-analytic tool `strym` [4], written in Python, for further downstream analysis of acquired data. `Strym` uses existing standards for decoding CAN messages files that are CSV-formatted, to produce time-series data representing desired signals for a particular drive.

Validation of CAN data through offline analysis using `strym` permitted us to produce runtime analysis and computation on the data through a software bridge that joins the Robot Operating System (ROS) [38] with streaming data produced by `libpanda`. Through this bridge, additional software components can make runtime decisions on data coming from the CAN bus—and could potentially inject control commands to vehicles in the future.

Using these software and hardware interfaces it is possible to imagine how to put a human-in-the-loop with sensed data from the vehicle that a driver cannot perceive in real time. Inspired by this possibility we created the CAN Coach, which is a system that continuously feeds time gap sensor information from the CAN bus back to the driver in near real time. Three sets of preliminary experiments are conducted in which the study vehicle follows a lead vehicle driving a specified driving profile to assess the potential of the CAN Coach to modify driver behavior. The experiments consider three modes: *Normal driving* (the driver is given no prompt and no feedback), *Instructed driving* (driver is given a prompt to drive at a two second time gap, but is not given any feedback from the CAN Coach), and *Coached driving* (two second prompt and CAN Coach feedback). The mean time gap errors from the 2 second target are 0.39s (Normal Driving), 0.09s (Instructed Driving), and 0.01s (Coached Driving). The standard deviation of the time gap error with the CAN Coach reduced by 72% and 68% from Normal Driving and Instructed Driving, respectively. Given this reduction of mean and standard deviation of the time gap error, we conclude that it is possible to "coach" drivers using only data from the CAN. Below shows how CAN Coach fits in the loop:

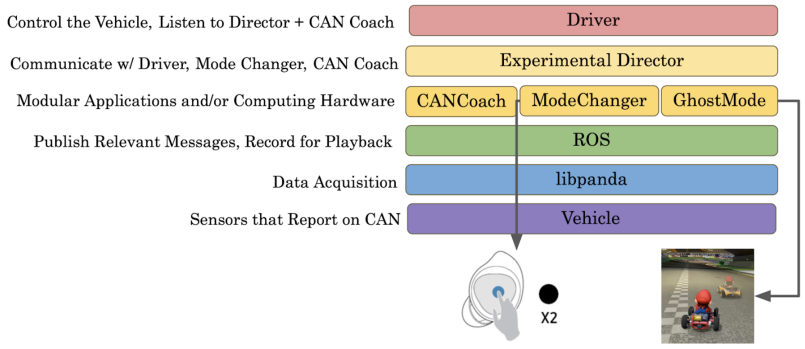

Fig. 2: A diagram showing where `libpanda`, ROS, and the CAN Coach fit in the hardware/software stack.

Traffic Flow Modelling

To train and test wave smoothing control strategies, we have been working on developing high fidelity traffic simulation models that explicitly produce non-equilibrium phenomena. Target traffic patterns, such as stop-and-go waves and traffic congestion, are evaluated and carefully tested for consistency with known phenomena on real roads. We developed a model to capture bulk traffic on a stretch of I-210 in California. We investigated various calibration routines and associated metrics [43]. This was done to prepare for the development of a model on I-24, and so while not complete, it served as a starting point to better design micro simulations with the macro statistics of bulk traffic flow.

Energy Modelling

Based on two fitted energy models, calibrated to measurements conducted by Toyota, we have produced two simplified models that provide an easy structural adaptability to the control and simulation framework employed below. Specifically, the models are of the form $P(v,a)$, i.e., the instantaneous energy consumption rate at time t is given by $P(\dot{x}_i(t), \ddot{x}_i(t))$, where $\dot{x}_i(t)$ is the speed, and $\ddot{x}_i(t)$ the acceleration, of a vehicle at that same time t. The quantity P can be fuel usage per time (g/s or ℓ/s), or the power equivalents (kW) of the battery depletion per time. One model (PriusEV 1.0) is for a hypothetical vehicle that possesses all the characteristics of the 2017 Prius, except for being a fully electric vehicle rather than a hybrid vehicle. While the (proprietary) measurements/models by Toyota do cover real battery properties, particularly the fact that the power consumption does depend on the battery's state of charge, this simplified fit in Fig. 3 is restricted to a fixed battery state of charge of 60%. The other model (Tacoma 1.0) represents an internal combustion engine, the

2017 Tacoma. These simplified models, shown in Fig. 3, are represented by explicit equations (thus admit rapid evaluation), and are piecewise smooth, convex functions. This last property is an approximation of reality, as a real vehicle's gear switching dynamics tends to introduce non-convexity effects that may render a vehicle-specific unsteady driving profile more energy efficient than a uniform velocity profile. The motivation for this simplification is that we wish to have energy metrics representing vehicle types in an averaged sense, rather than results that are fine-tuned to a specific single vehicle. In particular, the simple structure facilitates optimization and machine learning.

Figure 3. Visualization of both fitted simplified models for power consumption as a function of instantaneous vehicle speed and acceleration. Left: PriusEV 1.0 model for an electric vehicle. Right: Tacoma 1.0 for an internal combustion engine vehicle.

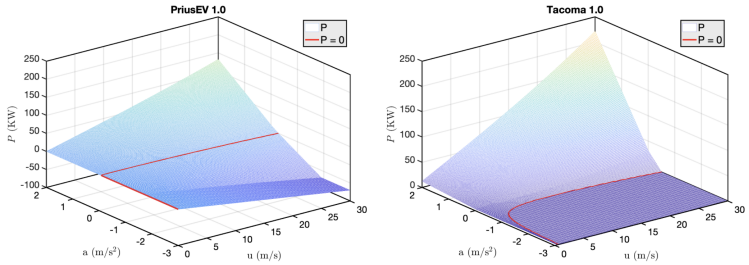

Fig. 3: Visualization of simplified models for energy consumption rate P (in power equivalent) as a function of instantaneous vehicle speed u and acceleration a. **Left:** PriusEV 1.0 model for an electric vehicle. The negative values of P reflect that the electric vehicle can recuperate some of the energy when braking. **Right:** Tacoma 1.0 for an internal combustion engine vehicle.

Safety

With real-life implementation in mind, ensuring safety of the control algorithms is paramount. We seek to establish a method of validating safety via the use of reachability analysis, i.e. looking at the set of reachable states for the AV given admissible controls. We aim at producing safe sets that represent a set of states from which, when evolved forward for all time, our controller's action will not result in a collision. To solve the reachability problem, we formulate the problem as a two player game with a leading vehicle and subject (following) vehicle. The two player game is an optimization problem optimizing over control inputs of both vehicles. The cost function of the optimization problem is the minimum distance between two vehicles at all times. The subject vehicle is governed by a controller, while the leading vehicle's input is some input that minimizes the cost function. Both vehicles have a maximal acceleration and deceleration. The safety set is determined as a positive value set of the optimal value function. An example is shown in Fig. 4.

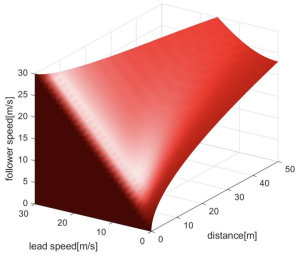

Fig. 4: An example of safe set. The safe set is defined in the domain where distance is positive, lead speed is between 0 and 30, and follower speed is between 0 and 30. The red surface represents the boundary of the safety set. The safe set is the set of states bounded by the surface (i.e. all states below the surface).

Control design using Reinforcement Learning

Due to the complexities of the driving and lane changing model, one approach pursued is to model-free reinforcement learning (RL) methods which can return effective controllers via optimization. At a high level, the RL formulation consists of defining a reward function (in this case, maximizing average miles per gallon) and then searching for a controller that optimizes the sum of the reward over a given horizon. By casting the problem as an optimization problem, we can find controllers for systems whose system dynamics are challenging to analyze. RL works by running the controller repeatedly in an environment, i.e. a simulator or a real world deployment. Figure 5 demonstrate the process of running RL to acquire sequences of state, action, reward pairs. Over these repeated runs, we acquire an estimate of the expected cumulative reward and then use optimization schemes to update the controller parameters to increase the expected cumulative reward, often using first order / gradient-based methods. By running this process repeatedly, we eventually return a controller that is close to a local maximum of the cumulative reward function.

Formally, since the autonomous vehicles are decentralized and their observations are local, we cast the problem as a Decentralized Partially Observed Markov Decision Process (DEC-POMDP), a formalism in which we have decentralized actors trying to each optimize their reward functions while having only partial access to the true global system state. In particular, each of our cars observe only their speed, the speed of their leading car, and the distance to the lead car; these are all states that can easily be acquired by radar but are insufficient. RL studies the problem of how an agent can learn to take actions in its environment to maximize its cumulative discounted reward. Specifically each agent, indexed by i, tries to optimize $J^{\pi_i} = \mathbb{E}_{\rho_0,\ p(s_{t+1}|s_t,a_t)}\left[\sum_{t=0}^{T} r_t(s_t, a_t^i) \mid \pi(a_t^i|o_t)\right]$ where r_t is the reward at time t and the expectation is over a distribution of initial states ρ, the probabilistic dynamics $p(s_{t+1}|s_t, a_t)$, and the probabilistic controller π_i which depends on a local observa-

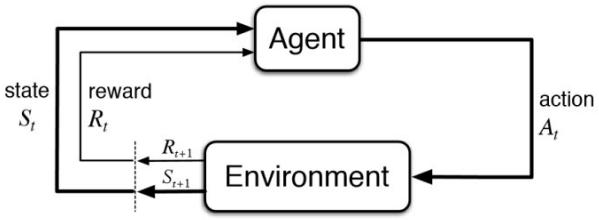

Fig. 5: The Reinforcement Learning loop. At each time-step, an agent (the controller) receives a state and reward from the environment/simulator, uses the state to compute a desired action, then executes that action to receive its next state and reward. This loop continues until the system horizon is reached.

tion o_t rather than the global state o_t. Here we have each agent optimize a local reward function rather than a global reward function.

As a preliminary example of this approach, we use *Proximal Policy Optimization* [42] to approximately optimize the following intuitive reward function: $r(s_t, a_t) = -P(v_t, a_t)$ i.e. the negative of the instantaneous consumed energy. Additionally, a bonus of 5 is added to the reward every time the vehicle completes 50 meters to ensure that it makes forward progress. After several hundred iterations, the controller generates the space-time diagram shown in Fig. 6 (note, only one lane showed) at a penetration rate of 10%.

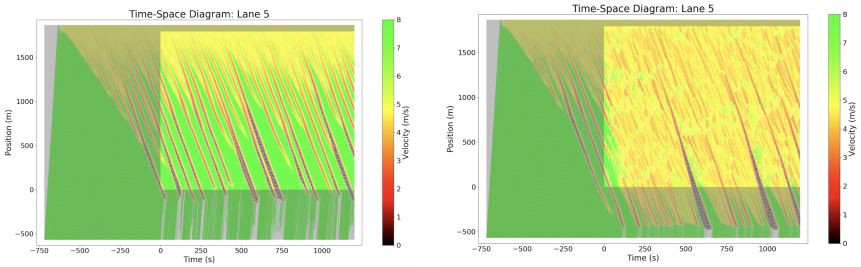

Fig. 6: Left-most lane of the I-210 without control (left) and under control of AVs at a 10% penetration rate (right).

Control design using mathematical models

Control strategies were designed in multi-line multi-population microscopic models by using a small number of autonomous vehicles (AV) (less than 5% penetration rate) to represent Lagrangian control actuators that can smooth stop-and-go waves in

multilane traffic flow. A microscopic multilane model is typically composed of two components: longitudinal dynamics for each lane and a lane-change mechanism. The parameters used on the lane changing mechanism were shown to have a potentially large impact on the macroscopic behaviour of the systems [28]. This implies that the control algorithm has to be robust with respect to the lane-changing mechanism. Two other approaches were considered by looking at the traffic at two higher scales: a macro-model consisting of a PDE coupled to several ODEs by a flux relation, which represents the interaction between the global traffic flow and the AV, identified as moving bottlenecks. And a new mean-field model that allows us to take the behavior of the well-known and well-studied microscopic models (such as the IDM) to a mesoscopic scale. This is obtained by considering the microscopic equations as equations on concentrated distributions that are the individual vehicles and showing a convergence when the number of vehicles goes to infinity. In the following, we focus on a way to design an efficient controller from a microscopic model. This is detailed in Sections 3 to 4.

2 Control system studied

Depending on the scale at which they represent vehicular traffic, mathematical traffic models usually can be classified into different categories: microscopic, mesoscopic, macroscopic, and cellular. We refer to the survey papers [6, 3], and reference therein, for general discussions about the models at various scales in the literature. In this section, we only focus on microscopic and macroscopic models.

Microscopic traffic models represent traffic by looking at each single vehicle. The aim is to simulate each single vehicle via the variables of position and velocity. Each vehicle is seen as a particle or agent, and its trajectory usually evolves according to the behaviour of other vehicles in front. The dynamics of all vehicle can then be caused by a system of ordinary differential equation (ODE). In this section, we first focus on two microscopic traffic models:

- The Bando-Follow-the-Leader (Bando-FTL) model.
- The Intelligent Driver Model (IDM).

Then, we introduce novel multilane and multiclass traffic by dividing the vehicle population into human-driven vehicles and autonomous vehicles. The autonomous vehicles are influenced by external policy makers with controlled dynamics. The multilane traffic is distinguished by its hybrid nature: continuous dynamics on each lane and discrete events due to lane-changing. For the lane-changing mechanism, we consider three components: safety (cars do not become too close), incentive (cars benefit from changing lanes with larger prescribed acceleration) and cool down time (cars cannot change lanes too rapidly). To this end, we study an optimal control problem related to the controlled microscopic hybrid system to model, for instance, the minimization of energy cost.

2.1 Microscopic traffic models: Bando-FTL model and IDM model

Consider a population of $P \in \mathbb{N}$ vehicles on an open stretch road of a single lane. For each vehicle $i \in \{1,\dots,P\}$, let $i_{\mathrm{L}} \in \{1,\dots P\}$ and $i_{\mathrm{F}} = \{1,\dots,P\}$ be the indices of the leading vehicle and following vehicle i, respectively. In other words, we associate with vehicle i an indices vector $\iota(i) = (i, i_{\mathrm{L}}, i_{\mathrm{F}})$ where i_{L} is index of the leader and i_{F} the index of the follower of vehicle i. To fix notation, we assume that $i_{\mathrm{L}} = 0$ if there is no other vehicle in front of vehicle i and $i_{\mathrm{F}} = 0$ if there is no other vehicle following vehicle i. Let the considered time horizon $T \in \mathbb{R}_{>0}$ be fixed. For each vehicle $i \in \{1,\dots,P\}$, let $(x_i(t), v_i(t)) \in \mathbb{R} \times \mathbb{R}_{\geq 0}$ be the position-velocity vector of vehicle i at time $t \in [0,T]$.

The Bando-FTL model

The Bando-Follow-the-Leader (FTL) model is a combination of the Bando model and the Follow-the-Leader model. In this section, we will introduce the Bando model and Follow-the-Leader model separately before giving the formulation of the Bando-FTL model.

The Bando model, which is also called optimal velocity (OV) model, is a traffic model proposed in [1]. One key feature is that each vehicle adjusts its acceleration or deceleration according to the difference between the optimal and their own "optimal" of preferred velocity. Specifically, the dynamics are given by the following system of second order ODEs

$$\begin{aligned} \dot{x}_i(t) &= v_i(t), && t \in [0,T], \quad i \in \{1,\dots,P\}, \\ \dot{v}_i(t) &= a(V(h_i(t)) - v_i(t)), && t \in [0,T], \quad i \in \{1,\dots,P\}. \end{aligned} \tag{1}$$

In the above formula, the dot represents the time derivative and $a \in \mathbb{R}_{>0}$ is a sensitivity constant. For every $t \in [0,T]$, $h_i(t) \coloneqq x_{i_{\mathrm{L}}}(t) - x_i(t)$ is the headway of vehicle i. The function $V : \mathbb{R} \to \mathbb{R}$ is called "OV-function" and describes the **o**ptimal **v**elocity determined by the headway. In general, V is monotonically increasing with respect to the headway, i.e. the optimal velocity is "smaller" if the headway is small and increases as the headway increases. In addition, if the headway is large enough for a sufficiently large time, the vehicle's velocity should arrive at the maximum velocity $v_{\max} \in \mathbb{R}_{>0}$, which may depend on the driver. One example of OV-function is $V\colon \mathbb{R}_{\geq 0} \mapsto \mathbb{R}_{\geq 0}, h \mapsto V(h)$, with

$$V(h) = v_{\max} \frac{\tanh(h-l-d)+\tanh(l+d)}{1+\tanh(l+d)}, \tag{2}$$

where $l \in \mathbb{R}_{>0}$ is the length of the vehicles and $d \in \mathbb{R}_{>0}$ is the minimum safe distance between vehicles. Note that the equilibrium point of the Bando model is obtained when all vehicles travel at constant velocity and have the same headway. The Bando model has been extended in [33] by considering interactions with the following vehicle (beside the leading) to stabilize the traffic flow.

The follow the leader (FTL) model was introduced in [16]. It is based on the idea that the acceleration of a vehicle is determined only by the vehicle in front. In particular, the acceleration of a vehicle is directly proportional to the difference between its leader's velocity and its own velocity, and is inversely proportional to the vehicle's headway. The main dynamics are given by

$$\begin{aligned} \dot{x}_i(t) &= v_i(t), && t \in [0,T], \quad i \in \{1,\dots,P\}, \\ \dot{v}_i(t) &= b\tfrac{v_{i_\mathrm{L}}(t)-v_i(t)}{(h_i(t))^2}, && t \in [0,T], \quad i \in \{1,\dots,P\}, \end{aligned} \tag{3}$$

where $b \in \mathbb{R}_{>0}$ reflects the sensitivity of the driver. In the case that there is no leading vehicle for vehicle i, i.e., $i_\mathrm{L} = 0$, the dynamics of vehicle i is given by

$$\dot{x}_i(t) = v_{\max} \qquad \forall t \in [0,T]. \tag{4}$$

We would like to mention the following drawback of the FTL model: the acceleration for vehicle i is zero as long as vehicle i and its leader have the same velocity. This implies that an extremely small headway is allowed even with high speeds.

One way to handle the above mentioned drawback of the FTL model is to combine the Bando model and the FTL model in the following way:

$$\begin{aligned} \dot{x}_i(t) &= v_i(t), && t \in [0,T], \quad i \in \{1,\dots,P\}, \\ \dot{v}_i(t) &= a(V(h_i(t)) - v_i(t)) + b\tfrac{v_{i_L}(t)-v_i(t)}{(h_i(t))^2}, && t \in [0,T], \quad i \in \{1,\dots,P\}. \end{aligned} \tag{5}$$

We point out that the Bando-FTL model justifies the fact that drivers adjust their acceleration or deceleration based on their own velocities, the optimal velocities and the velocities of their leading vehicles. This model allows formation and persistence of stop-and-go waves, and can be well calibrated following the experimental results of [46], see [11].

The intelligent driver model

The intelligent driver model (IDM) introduced in [48] is a time-continuous microscopic car-following model, which is widely used in the traffic engineering community. It assumes that each vehicle-driver decides to accelerate or to brake depending only on their own velocity, and on the position and velocity of the leading vehicle immediately ahead. To simplify notations, we set the velocity difference or approaching rate of vehicle i at time $t \in [0,T]$ as $\Delta v_i(t) \coloneqq v_i(t) - v_{i_L}(t)$ and the net distance of vehicle i at time $t \in [0,T]$ as $s_i(t) \coloneqq x_{i_L}(t) - x_i(t) - l$, where $l \in \mathbb{R}_{>0}$ is again the vehicle length. The dynamics of vehicles are then described by the following system of ODE:

$$\begin{aligned} \dot{x}_i(t) &= v_i(t), && t \in [0,T], \quad i \in \{1,\dots,P\}, \\ \dot{v}_i(t) &= a\left(1 - \left(\tfrac{v_i(t)}{v_0}\right)^\delta - \left(\tfrac{s^*(v_i(t),\Delta v_i(t))}{s_i(t)}\right)^2\right), && t \in [0,T], \quad i \in \{1,\dots,P\}, \end{aligned} \tag{6}$$

where the desired minimum gap of the vehicle i represented by the function $s^* : \mathbb{R}\times\mathbb{R}$, is given by

$$s^*(v_i(t), \Delta v_i(t)) \coloneqq s_0 + v_i(t)\tau + \frac{v_i(t)\Delta v_i(t)}{2\sqrt{ab}} \qquad t \in [0,T]. \tag{7}$$

The parameters' meaning and choice are as follows:

- $v_0 \in \mathbb{R}_{>0}$ is the desired velocity the vehicle would drive in free traffic (m/s)
- s_0 is the minimum desired net distance
- τ is the safety time gap
- a is the maximum vehicle acceleration
- b is the comfortable braking deceleration
- the exponent $\delta \in \mathbb{R}_{>0}$ a tuning parameter, usually set to 4.

The desired minimum gap depends on the safety time gap, the vehicle acceleration, deceleration and the velocity difference. Specifically, Eq. (7) contains three terms:

1. The minimum distance s_0 in congested traffic;
2. The safety gap that the follower must have with its leader, $v_i(t)\tau$;
3. The term $\frac{v_i(t)\Delta v_i(t)}{2\sqrt{ab}}$ which is designed to stabilize the platoon vehicle in terms of velocity.

Furthermore, the acceleration of vehicle i can be separated into a free road term

$$a_{\text{free}}(t) = a\left(1 - \left(\frac{v_i(t)}{v_0}\right)^{\delta}\right), \qquad t \in [0,T] \tag{8}$$

and an interaction term

$$a_{\text{int}}(t) = -a\left(\frac{s^*(v_i(t),\Delta v_i(t))}{s_i(t)}\right)^2 = -a\left(\frac{s_0+v_i(t)\tau}{s_i(t)} + \frac{v_i(t)\Delta v_i(t)}{2\sqrt{ab}s_i(t)}\right)^2, \qquad t \in [0,T]. \tag{9}$$

Note that in the case of a free road when the net distance of vehicle i, s_i, is large, the vehicle's acceleration is governed by the free road term Eq. (8) which vanishes as v_i approaches v_0. Thus a vehicle on a free road will gradually approach its desired velocity v_0. For large approaching rate Δv_i, the interaction term Eq. (9) is dominated by the term $-a\left(\frac{v_i(t)\Delta v_i(t)}{2\sqrt{ab}s_i(t)}\right)^2 = -\frac{(v_i(t)\Delta v_i(t))^2}{4bs_i(t)^2}$ which leads to a driving behavior that compensates velocity differences while trying not to brake much harder than the comfortable braking deceleration b. For negligible velocity differences and small net distance, the interaction term Eq. (9) is approximately equal to $-a\frac{(s_0+v_i(t)\tau)^2}{s_i(t)^2}$. This resembles a simple repulsive force such that small net distances are quickly enlarged towards an equilibrium net distance.

We point out that the IDM model has many drawbacks when it comes to the drivers' safety and the vehicles' real capability especially in the case of a collision, see [12]. Therein, a modified version of IDM model was proposed and tested in terms of string stabilization.

2.2 Lane changing conditions for multilane traffic

In the case of multilane traffic, frequent lane changing maneuvers could disrupt traffic flow and even worse, lead to accidents. In addition, lane changing behaviors have significant impact on formulation and propagation of the stop-and-go traffic waves. Recently, efforts to model lane change have rapidly increased, see [29, 26, 25]. In this subsection, we propose lane changing rules based on acceleration due to two main advantages, see [54]:

1. The lane changing decision-making process is dramatically simplified
2. One can readily calculate the accelerations based on an underlying microscopic traffic model

We consider an open stretch of road with $m \in \mathbb{N}_{\geq 1}$ lanes and assume that the number of vehicles on lane $j \in J := \{1, \dots, m\}$ is P_j. We associate each vehicle to $i \in \{1, \dots, P_j\}$, on lane $j \in J$, a vector of indices $\iota(i) = (i, j, i_{\mathrm{L}}, i_{\mathrm{F}})$, where i_{L} and i_{F} are defined as in Section 2.1.

We design lane changing rules based on three components: *cool-down time*, *safety* and *incentive for lane-changing* and assume that lane changing is performed instantaneously. Then, we propose that a vehicle performs a lane change if the following conditions are met:

1. The vehicle did not make a recent lane change. To prevent a vehicle changing lane too frequently, we assign each vehicle a cool-down time $\tau > 0$. A vehicle is allowed to perform lane changing at most once every τ time units. This constraint is motivated by two reasons: the lane change is not instantaneous (so there should be a fixed minimal duration between two lane-changes), a driver that has just changed lanes is less likely to change lanes again right away. From [24], only the 15% of the vehicles cross a lane while traveling the road section. In the simulation presented in [19], we choose the cool-down time $T_1 = 5s$ over the time interval $[0, 1000s]$. The simulation shows that our model with this cool-down time can handle the case when around 40% of vehicles performing lane-changing. In particular, one can model higher lane-changing frequency by taking smaller cool-down times.
2. The lane change can be performed safely into a target lane. This corresponds to having enough space in between the two vehicles in the target lane where the vehicle is moving to.
3. There is sufficient incentive to perform lane changing. In other words, if the expected acceleration of the vehicle in the new lane is sufficiently bigger than its acceleration on the current lane, the lane change will be performed (assuming that the other mentioned requirements are met).

We now instantiate the above rules in a mathematically sound and precise way. We associate to each vehicle i in lane $j \in J$ an internal time for lane change called $\tau_i^j : [0, T] \to \mathbb{R}_{\geq 0}$ with $T > 0$ being fixed, such that the following holds:

$$\begin{aligned} \dot{\tau}_i^j(t) &= 1, \qquad t \in [0,T] \\ \tau_i^j(0) &= \tau_{i,0}^j \in [0,\tau), \end{aligned} \tag{10}$$

where $\tau \in \mathbb{R}_{>0}$ is the cool down time. To avoid synchronous lane changes, we assume that all initial conditions of the internal times are distinct one from another. That is, we assume for $j, j' \in J$, $i \in \{1,\dots,P_j\}$, $k \in \{1,\dots,P_{j'}\}$, with $j \neq j'$ and $i \neq k$,

$$\tau_{i,0}^j \neq \tau_{k,0}^{j'}.$$

We assume that a vehicle does not lane change unless its internal time reaches the cool down time τ. In addition, we reset the internal time for each vehicle to be zero just after it reaches the cool-down time τ. Thus, we set

$$\begin{aligned} \dot{\tau}_i^j(t) &= 1, \quad t \in]k\tau,(k+1)\tau] \cap [0,T], \quad k \in \{1,2,\dots\}, \\ \tau_i^j(k\tau+) &= 0, \quad j \in J, \qquad\qquad\qquad\qquad i \in \{1,\dots,P_j\}, \end{aligned} \tag{11}$$

where $k\tau+$ represents the right hand side limit which is needed as τ_i^j is only piecewise continuous. Note that we do not allow two vehicles to change lane at the same time. This assumption is reasonable since it has been experimentally shown that lane changing is not frequent in a traffic flow, see [27].

Let us now describe the safety and incentive conditions. Let $\Delta \in \mathbb{R}_{>0}$ be fixed. Vehicle i on lane $j \in J$ will change to lane $j' = j+1$or $j-1 \in J$ at time $t \in [0,T]$, if the following conditions are met

$$\text{Safety: } \bar{a}_i^{j'}(t) \geq -\Delta \text{ and } \bar{a}_l^{j'}(t) \geq -\Delta; \qquad \text{Incentive: } \bar{a}_i^{j'}(t) \geq a_i^j(t) + \Delta,$$

where l is the index of the potential follower of vehicle i on the new lane j', $a_i^j(t)$ denote the acceleration of vehicle i on lane j at time t, and $\bar{a}_l^{j'}(t)$ and $\bar{a}_i^{j'}(t)$ are the expected acceleration of vehicles l and i on lane j' at time t, respectively.

2.3 A controlled hybrid system

In this subsection, we present a multi-population multilane model. First, we divide the whole traffic population into two classes: autonomous vehicles and human-driven vehicles. Let $M_j \in \mathbb{N}_{\geq 1}$ and $N_j \in \mathbb{N}_{\geq 1}$ represent the number of autonomous vehicles and human-driven vehicles on lane $j \in J$, respectively. Due to the fact that external policy makers can influence the dynamics of the autonomous vehicles, we only add controls to the autonomous vehicles instead of controlling the whole traffic population.

The dynamics of vehicles on each lane are continuous and are governed by a system of ODEs consisting of, for instance Bando-FTL model or IDM model described in Section 2.1. Roughly speaking, for any microscopic model, the acceleration of an individual vehicle is determined by its desired velocity, and the position-velocity

vectors of its leader and its own. Let $\left(x_i^j(t), v_i^j(t)\right) \in \mathbb{R}\times\mathbb{R}_{\geq 0}$ be the position-velocity vector of human-driven vehicle i and $\left(y_k^j(t), w_k^j(t)\right) \in \mathbb{R}\times\mathbb{R}_{\geq 0}$ be the position-velocity vector of autonomous vehicle k on lane $j \in J$ at time $t \in [0,T]$. Then, the dynamics of vehicles on the j-th lane can be written for $t \in [0,T]$ as

$$\begin{aligned}
\dot{y}_k^j(t) &= w_k^j(t), && (j,k) \in J \times \{1,\dots,M_j\},\\
\dot{w}_k^j(t) &= \mathrm{ACC}(y_k^j(t), w_k^j(t), z_{k_{\mathrm{L}}}^j(t), v_{k_{\mathrm{L}}}^j(t), v_k^{j,*}(t)) + u_k^j(t), && (j,k) \in J \times \{1,\dots,M_j\},\\
\dot{x}_i^j(t) &= v_i^j(t), && (j,k) \in J \times \{1,\dots,N_j\},\\
\dot{v}_i^j(t) &= \mathrm{ACC}(x_i^j(t), v_i^j(t), \tilde{z}_{i_{\mathrm{L}}}^j(t), \tilde{v}_{i_{\mathrm{L}}}^j(t), v_i^{j,*}(t)), && (j,k) \in J \times \{1,\dots,N_j\},
\end{aligned} \tag{12}$$

where $\mathrm{ACC}\colon \mathbb{R}\times\mathbb{R}_{\geq 0}\times\mathbb{R}\times\mathbb{R}_{\geq 0}\times\mathbb{R}_{\geq 0} \mapsto \mathbb{R}$ is the general formulation of a vehicle's acceleration based on a microscopic traffic model, $(z_{k_{\mathrm{L}}}^j(t), v_{k_{\mathrm{L}}}^j(t))$ and $(\tilde{z}_{i_{\mathrm{L}}}^j(t), \tilde{v}_{i_{\mathrm{L}}}^j(t))$ are the position-velocity vectors of the leaders of autonomous vehicle k and human-driven vehicle i on lane j, respectively, at time $t \in [0,T]$ and $v_k^{j,*}(t)$ and $v_i^{j,*}(t)$ are the desired velocities of autonomous vehicle k and human-driven vehicle k on lane j, respectively. Notice that the dynamics of autonomous vehicle k on lane j are distinguished by the control term $u_k^j\colon [0,T] \mapsto \mathbb{R}$.

In addition, the lane changing mechanism of the vehicles introduced in Section 2.2 generates discrete events for the multilane, multi-class traffic system. The presence of both time-dependent continuous dynamics in Eq. (12) and discrete events lead to a system of hybrid nature, see [35, 36]. Such systems are characterized by the presence of continuous dynamics, which, at discrete times – so called switching times – are effected by logic variables. The latter, in turn, change their values depending on the values of the continuous variables at switching times. We call such system a hybrid system or a switched system. In our example, the switching times are precisely the lane changing times. The discrete variables are the lane indices of each vehicle, which may change depending on position, speed and accelerations of vehicles, which represent the continuous variables. Because of the cool down time assumption, we avoid the famous Zeno phenomenon, which impacts strongly the behavior of hybrid systems [9].

Furthermore, we formulate an optimal control problem related to the mentioned controlled hybrid system to minimize, for instance, the energy cost. Let $\mathcal{U} := \{\mathbf{u}\colon [0,T] \mapsto \mathbb{R}^M\}$ be the set of admissible controls, where $M = \sum_{j\in J} M_j$ is the total number of the autonomous vehicles. We define a cost functional $F\colon \mathcal{U} \mapsto \mathbb{R}$, by

$$F(\mathbf{u}) = \sum_{j\in J} \int_0^T \left\{ L_j\left(\mathbf{x}^j(t), \mathbf{v}^j(t), \mathbf{y}^j(t), \mathbf{w}^j(t), \mathbf{u}^j(t)\right) + \sum_{k=1}^{M_j} \frac{|u_k^j(t)|}{M_j} \right\} \mathrm{d}t \tag{13}$$

where for each $j \in J$, the Lagrangian function $L_j\colon \mathbb{R}^{N_j}\times\mathbb{R}_{\geq 0}^{N_j}\times\mathbb{R}^{M_j}\times\mathbb{R}_{\geq 0}^{M_j}\times\mathbb{R}^{M_j} \mapsto \mathbb{R}$ is sufficiently smooth. The hybrid optimal control problem is to minimize the cost functional Eq. (13) over the set of admissible controls $\mathcal{U}$, where $\mathbf{x}, \mathbf{v}, \mathbf{y}, \mathbf{u}$ satisfies the hybrid system Eq. (12). For the existence of optimal control, we refer the interested reader to [13].

2.4 A mean-field approach

Macroscopic traffic models describe vehicular traffic as fluid flow by assuming a sufficiently large number of vehicles on a road. By capturing and predicting the main phenomena of microscopic dynamics, macroscopic models provide overall and statistical views of traffic. There are three important variables in a macroscopic traffic model:

- the traffic density $\rho : \mathbb{R} \times [0,T] \to \mathbb{R}_{\geq 0}$
- the speed $v : \mathbb{R} \times [0,T] \to \mathbb{R}_{\geq 0}$ of the traffic
- the flow rate $q : \mathbb{R} \times [0,T] \to \mathbb{R}_{\geq 0}$ passing through a fixed point

The three variables satisfy the following relationship: for $(x,t) \in \mathbb{R} \times [0,T]$,

$$q(x,t) = \rho(x,t)v(x,t). \tag{14}$$

The evolution of the traffic density is governed by partial differential equations (PDEs), for example the Lighthill-Whitham-Richards (LWR) model in [31, 40] consists of a single conservation law for the density, and closing the system by assuming that the speed depends only on the density. Models with more equations and different "closure relationships" were developed, see [14, 15] for an extensive discussion. For the multi-population traffic system introduced in Section 2.3, one can use a coupled ODE-PDE system to model the dynamics of a small number of autonomous-vehicles with ODEs and the large number of human-driven vehicles with PDEs. This is a combination of the microscopic and macroscopic models using multiple scales.

One may use a mean-field approach to relate the two different scales, i.e. microscopic and macroscopic, both formally and rigorously by letting the number P of vehicles go to infinity. For a population of $P \in \mathbb{N}$ vehicles on an open stretch of a single lane, we again let for $t \in [0,T]$ the vector $(x_i(t), v_i(t)) \in \mathbb{R} \times \mathbb{R}_{\geq 0}$ be the position-velocity of vehicle $i \in \{1,\dots,P\}$. We assume that drivers adjust their acceleration based on the position-velocity vectors of several vehicles ahead (instead of only one vehicle as it is usually assumed in classical car following models). To realize that, we introduce a locally Lipschitz convolutional kernel function $H : \mathbb{R} \times \mathbb{R}_{\geq 0} \mapsto \mathbb{R}$ which is of sub-linear growth and describes the dynamics of the P vehicles as

$$\begin{aligned} \dot{x}_i(t) &= v_i(t), && t \in [0,T], \quad i \in \{1,\dots,P\}, \\ \dot{v}_i(t) &= H * \mu_P(x_i, v_i)(t), && t \in [0,T], \quad i \in \{1,\dots,P\}, \end{aligned} \tag{15}$$

where $\mu_P(x_i, v_i) = \frac{1}{P} \sum_{i=1}^{P} \delta_{x_i, v_i}$ is the probability measure obtained using Dirac distributions placed at a cars' position-speed location, also called the empirical measure. Then, the rigorous mean-field limit of the finite-dimensional ODE system in Eq. (15) consists of a Vlasov-Poisson type PDE

$$\partial_t \mu + v \nabla_x \mu = \nabla_v \cdot [(H * \mu)\mu], \qquad (t,x,v) \in [0,T] \times \mathbb{R} \times \mathbb{R}_{\geq 0} \tag{16}$$

which describes the evolution of the density distribution $\mu : (0,T)\times\mathbb{R}\times\mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ of infinitely many vehicles in (x,v) space, see [13].

For multi-population and multilane traffic with a small number of autonomous vehicles and a large number of human-driven vehicles, we can proceed as follows. The mean-field limit is computed only for the human-driven vehicles by taking their number to infinity, while the number of autonomous vehicles is kept fixed. The lane-changing maneuvers of the human-driven vehicles generate in the limit a source term for the Vlasov-Poisson type PDEs Eq. (16) (one equation per lane).

The autonomous vehicles' lane-changing behaviour (discrete dynamics) can be considered as controls, thus leading to a controlled hybrid system as in Section 2.3. Similarly, in the limit one obtains a coupled hybrid system of controlled ODEs (for the autonomous vehicles) and Vlasov-Poisson type PDEs with source terms. It turns out that the optimal controls for the finite-dimensional hybrid ODE system converge to the optimal controls of the hybrid ODE-PDE system when the number of human-driven vehicles approaches infinity. This is proved rigorously by taking advantage of Γ-convergence [8] and using a generalized Wasserstein metric [37]. For details on the limit process see [18].

2.5 Numerical implementation and difficulties about modelling

A traffic smoothing algorithm, or "controller" will be selected among candidate controllers that we first test in a simulation in the traffic simulator "sumo" [32]. We use this platform to evaluate the controller's potential real world performance in reducing the energy consumption of the bulk traffic. This step allows us to select the best performing controllers for the eventual road test. To meaningfully evaluate the road performance, the bulk traffic in the simulation must be similar to real life bulk traffic. There are features of the bulk traffic that are especially important in evaluating the controllers, such as stop-and-go traffic waves.

We designed a microscopic traffic model in sumo to simulate the bulk traffic. The simulation will include stop-and-go waves in expected traffic density regimes using the Intelligent Driver Model (IDM), augmented with a small amount of additive noise in the velocity to trigger instabilities that can grow into full nonlinear stop-and-go waves if traffic flow is in the dynamically unstable regime. This microscopic model must be parameterized to have macroscopic features such as decreasing velocity with increasing density and congestion formation on a segment of road. The simulation will have a rate of cars spawning at the upstream boundary, called the inflow, and a rate of cars leaving the segment, called the outflow. When the inflow is higher than the outflow congestion will form. The vehicle accelerations from the microscopic model are used to calculate emissions and energy usage for the bulk traffic flow. In order to test that our bulk traffic model was representing the type of traffic we require, we started with a model of a small section of highway and no on ramps or off ramps. We needed a way to force different levels of congestion and observe the presence or lack of stop-and-go waves. An initial challenge was the dramatic effect of simulation

step size on the presence and quality of stop-and-go waves. Unexpected traffic waves would be present for large time steps.

Designing this bulk traffic model on a small section of highway involves intricate car-spawning conditions at the beginning of the highway segment. In order to control congestion levels, we designed a downstream "ghost cell," which is a section of road used only to control conditions on the road segment of interest, where vehicles obey a lower speed limit which propagates upstream to traffic in the highway segment. The vehicles in the downstream ghost cell approach this new speed with a relaxation term so as to smooth out unexpected acceleration artefacts and prevent them from propagating into the segment of interest. Before this segment, we also had to implement an upstream ghost cell in order to circumvent issues due to vehicles spawning logic in a high congestion setting. These two ghost cells provided a way to control the traffic density and make sure we were achieving stop-and-go waves in expected density and throughput regimes. This design of a highway segment represents the minimum testbed to meaningfully evaluate the potential road performance of controllers. This was done before including other necessary features, such as inclusion of trucks, more sophisticated lane change models, and highway topology, such as on/off ramps.

Fig. 7: Visualization of a simulation on the stretch of the I-210 (top), zoom on a start of congestion (bottom). Regular vehicles are represented in white, AVs are represented in red but disabled in this simulation.

2.6 Energy model used

The results presented in Section 4 assume that all vehicles on the roadway (uncontrolled as well as controlled vehicles) are all combustion engine vehicles with identical fuel consumption characteristics. We use a simplified, regularized, energy function that is calibrated to measurements conducted by Toyota on a 2017 Tacoma vehicle. The energy demand is expressed via a power equivalent function $P(v, a)$ that expresses the instantaneous power as a function of the vehicle speed v and acceleration a. The simplified function P, shown in Fig. 3, is designed to structurally resemble physical power functions, augmented by structurally simple correction terms. It is of the form

$$P(v,a) = \max\left(m\,av + C_0 + C_1 v + C_2 v^2 + C_3 v^3,\ 0\right) + \max\left(p_1 a + p_3 av,\ 0\right), \quad (17)$$

where $m = 2041\,\mathrm{kg}$, $C_0 = 3405.54\,\mathrm{W}$, $C_1 = 83.1239\,\mathrm{kg\,m/s^2}$, $C_2 = 6.76507\,\mathrm{kg/s}$, $C_3 = 0.70413\,\mathrm{kg/m}$, $p_1 = 4598.71\,\mathrm{kg\,m/s}$, and $p_3 = 975.127\,\mathrm{kg}$. To obtain fuel consumption (in volume per time), the fitted conversion factor $15.09\,\mathrm{KW} = 1\,\mathrm{gallon/h}$ is used, which incorporates the engine's efficiency. Of the two terms in Eq. (17), the first represents a physics-equivalent power function (on flat roads) of the form $P = vF$, where the force F is composed of the acceleration force ma, and a generalized friction and drag force function (the $C_j v^j$ terms). The second term is a correction term that accounts for engine and powertrain inefficiencies. In both terms, the maximum with zero expresses the fact that no energy is recuperated during braking, which is the typical situation for combustion engine vehicles. In contrast, electric vehicles do generally recuperate a certain amount of energy during braking (see Fig. 3), albeit not the full amount possible during the braking that tends to occur in strong stop-and-go traffic waves. Hence, flow smoothing will still improve the energy efficiency of electric vehicles, but not as significantly as the results shown in Section 4.

In the simulator, the vehicle energy is evaluated as follows. Given that the integration time step Δt is chosen suitably small, energy is computed via a simple first order quadrature rule: for each vehicle, evaluate the energy rate $f(v(t), a(t))$ at time t, with a version of the acceleration devoid of noise, and add $f(v(t), a(t))\Delta t$ to the cumulative energy.

3 Control design

As stated in the introduction, our ideal goal is to suppress stop-and-go waves in the system, namely bringing it to a situation where all vehicles have the same velocity. Given the models' dynamics in Eq. (5) or in Eq. (6), this corresponds to bringing all the vehicles to a steady state. This steady-state is uniquely defined by the velocity, and in particular a given velocity imposes all the vehicles headway. In this section we first present a simple and ideal controller. We explain why this controller works in ideal

situations, and can prevent the appearance of new stop-and-go waves, but struggles to cope with dissipating existing stop-and-go waves in practical situations. Then we present an approach to make this controller more robust and adapt it to realistic situation with strong stop-and-go waves. We illustrate these points by providing numerical simulations on an ideal one lane straight road. Numerical simulations on a realistic mutlilane setting as well as quantitative results are the subject of the next section.

3.1 Ideal controller

Let us assume that, given the inflow and outflow conditions, there exists a steady-state with velocity $v_{\text{opt}} \in \mathbb{R}_{>0}$ to the system. A natural idea to force the traffic flow into this steady-state would be to set the velocity of each AV to be v_{opt}, using the following proportional control law:

$$\dot{v}_{\text{AV}}(t) = -k \left(v_{\text{AV}}(t) - v_{\text{opt}} \right), \tag{18}$$

where $k \in \mathbb{R}_{>0}$ is a design parameter to be chosen. To understand why such a simple idea might work, one has to see stop-and-go waves as propagating backwards because all vehicles in the jam alternate between decelerating – when they are at the top of the wave – and accelerating – when they are at the bottom of the wave. Usually such acceleration brings them to a higher speed than the steady-state speed, which therefore inevitably leads to a further deceleration. Few AVs in the bulk traffic dampening stop-and-go waves are like rocks on a shore braking sea waves. Hence, this control law has already been used in several works of research, in both theoretical and experimental settings. In [11] for instance the authors study this control law implemented on a single AV on a ring-road, i.e., with the i-th vehicle being also identified as the $N + i$-th vehicle and with the Bando-FTL (see Eq. (5)) traffic model. They show theoretically that this controller guarantees the local exponential stability of the system for up to 9 vehicles (which correspond to a penetration rate of 11%). They also show numerically that this stability holds reasonably well up to 20 vehicles (penetration rate 5%). In [22], the authors show that this controller guarantees in fact the local exponential stability for any number of vehicles on the road (penetration rate as low as desired), moreover they show the exponential decay rate and the optimal k to not depend on the number of vehicles. However, and as it could be expected, the bound we are able to obtain on the basin of attraction decreases strongly with the number of vehicles. This means that in order to stabilize an entire ring-road with a single vehicle, the vehicle's velocities must be very close to the steady-state value if the number of vehicles becomes large. In other words, this means that this controller might not be able to ensure the stability of motion in heavy stop-and-go sequences. And this, as we will see in the next subsection, will be the main problem. Nevertheless, this means that this simple controller is good at preventing the emergence of stop-and-go waves in a steady flow.

The controller was also tested in a field experiment, on a single lane ring-road with 20 vehicles and a single AV. Several variants of the controller where tested. Among them a variant where[1]

$$\dot{v}_{\mathrm{AV}}(t) = -\beta \left(v_{\mathrm{AV}(t)} - v_1(t) \right), \tag{19}$$

where β is a positive coefficient, $\alpha = \min\left(\min\left(\frac{\Delta x - \Delta x^s}{\gamma}, 0\right), 1 \right)$ with Δx^s being a safety distance and

$$v_1(t) = \alpha(t) v_{\mathrm{target}}(t) + (1 - \alpha(t)) v_{\mathrm{lead}}(t), \tag{20}$$

$$v_{\mathrm{target}}(t) = \frac{1}{\tau} \int_0^\tau v_{\mathrm{AV}}(\tau) d\tau + \min\left(\max\left(\frac{\Delta x - l_1}{l_2}, 0 \right), 1 \right), \tag{21}$$

with τ a characteristic time, and l_1 and l_2 two distances, typically $l_1 = 7m$ and $l_2 = 23m$. This amounts to changing the speed to be reached v_{opt} by v_1, an objective speed that depends on time. This time dependency is important, as we will see later on in Section 3.2. Despite the fact that a single lane ring-road is a very favorable situation, the results were astonishing with a nearly complete disappearance of stop-and-go waves and up to 40% fuel consumption reduction.

3.1.1 Limits in practical situations

While the controller in Eq. (18) works perfectly when the system is close to the target steady-state, it suffers several limitations when the speed variance is high, all related to the following fact: if an AV follows the controller in Eq. (18) while the system undergoes large stop-and-go waves, it might crash in its leading vehicle at the top of the wave. A natural quick fix would consist of adding a safety mechanism that overrides the controller and brakes when the AV is too close from its leading vehicle. However, if the AV is itself decelerating when at the top of the stop-and-go wave, it propagates the stop-and-go wave and this destroys the stabilizing effect of the controller. This is where the difficulty lies, and the reason why we consider these two simple options.

- The first option is to replace v_{opt} by $\min(v_{\mathrm{opt}}, v_{\mathrm{lead}})$ where v_{lead} is the velocity of the vehicle in front of the AV. However, this does not prevent it from being caught in a stop-and-go wave in numerical experiments. In some cases, the stop-and-go wave reduces the average velocity in a manner that v_{lead} is below v_{opt} most of the time. This, in turn, leads the AVs to propagate the stop-and-go wave instead of dissipating it. This can be observed in Fig. 8 representing the speed variance of vehicles with time of a straight stretch of road with 3 lanes. The road is 2000m and 5% of vehicles are AVs, while the 95% remaining are regular vehicles modeled with Bando-FTL model. For the first 500s the AVs obey the same dynamics as the regular vehicles and then start following the controller dynamics at $t = 500$s.

[1] In [11], this variant is given as a discretized control law. We state it its continuous version.

The simulation is done in `MATLAB` using a fourth order explicit scheme. We see on Fig. 8 that a little before $t = 500s$, the speed variance is high, indicating heavy stop-and-go waves, and this does not seem to change after the AVs are turned on. The switch at $t = 500s$ does not lead to any apparent reduction of the overall speed variance of each lane, indicating that the AVs do not seem to have a real smoothing effect in this case.

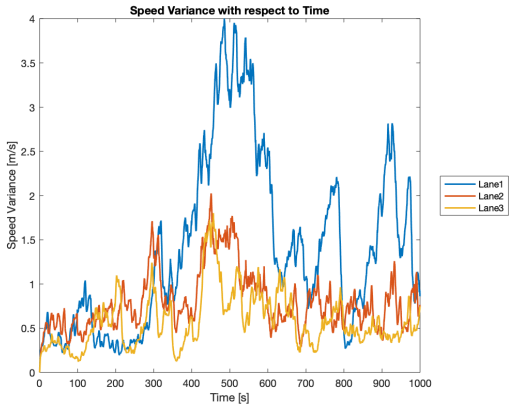

Fig. 8: Speed variance with respect to time in 3 lanes of a straight road, using the controller with $\min(v_{\text{opt}}, v_{\text{lead}})$ instead of v_{opt}. AVs are turned on at $t = 500s$. The different lanes are represented in different colors, the proportion of AVs in inflow is 5% and the same for each lane

- Another approach consists of adding to the control law a term that mimics a regular vehicle behavior. For instance for a traffic flow modelled with Bando-FTL (see Eq. (5)), this would give the new control law

$$\begin{aligned}\dot{v}_{\text{AV}}(t) = a(V(x_{\text{lead}}(t) - x_{\text{AV}}(t)) - v_{\text{AV}(t)}) \\ + b\frac{v_{\text{lead}}(t) - v_{\text{AV}}(t)}{(x_{\text{lead}}(t) - x_{\text{AV}(t)})^2} - k(v_{\text{AV}}(t) - v_{\text{opt}}(t)),\end{aligned} \tag{22}$$

  and the parameter $k \in \mathbb{R}$ properly chosen as well as the parameter $a, b \in \mathbb{R}$ introduced in Eq. (5). When too close from the leading vehicle, the FTL term is dominant and is expected to avoid crashes. However this approach was shown in [7] to have some limitations: not only with the noise inherent to practical situations there is a inferior bound to the penetration rate that can achieve the stabilization of the system, the k needed for the controller to stabilize efficiently the system increases exponentially when the penetration rate decreases linearly and could become quickly unpractical.

Finally, another limit comes from the estimation of the steady-state speed v_{opt} and the possible subsequent errors. While underestimating v_{opt} would mean diminishing

the throughput, overestimating v_{opt} could render the system unstable. Indeed, in this case the AV tries to reach a higher speed than what the overall system can reach. As a consequence it will always end up being too close to its leading vehicle, triggering a safety brake and starting a wave again. This effect was already seen in field experiments on a ring road [45].

3.2 Robust approach

Our approach to tackle the limitations mentioned in Section 3.1.1 and have a robust controller is trying to stabilize another speed $v_{\text{cmd}}(t)$ instead of v_{opt}. Motivated by the field experiment in [11, 45], this speed should be time-dependent and the controller then becomes

$$\dot{v}_{\text{AV}}(t) = -k\big(v_{\text{AV}}(t) - v_{\text{cmd}}(t)\big) + \dot{v}_{\text{cmd}}(t), \tag{23}$$

The term $\dot{v}_{\text{cmd}}$ is for compensating that v_{cmd} might depend on t. Indeed, we recover then $(v_{\text{AV}}(t)-v_{\text{cmd}}(t))' = -k(v_{\text{AV}}(t)-v_{\text{cmd}}(t))$ so that the velocity of the AV converges to $v_{\text{cmd}}(t)$. From this point, the goal is to find a v_{cmd} that would conciliate two opposite goals: being as large as possible to avoid decreasing the outflow, while making sure that the AVs are decelerating as little as possible.

Our proposed algorithm to determine v_{cmd} is the following, inspired by the TCP algorithm used to reduce congestion in communication networks. The general philosophy is to select a target speed, reduce it if the AV has to brake at some point (meaning that the target speed was chosen too high compared to the stop-and-go wave), and increase it slowly again. Besides, if the AV starts to have a headway significantly larger than the usual headway in the stop-and-go wave, the speed is increased to the target speed. Finally, if the AV is really too far from its leading vehicle, v_{cmd} is set to a free flow speed to catch up. This leads v_{cmd} to have an hysteresis behavior that allows to adapt to the current stop-and-go wave.

More precisely, we define three distances, d_0, d_1 and $d_2 \in \mathbb{R}_{>0}$ (typically $d_0 = 4.5$, $d_1 = 25m$ and $d_2 = 100m$) and two speeds:

- $v_{\text{des}} = 0.95v_{\text{opt}}$, slightly slower than the steady-state speed v_{opt} to compensate the potential errors and overestimation of v_{opt} which would make the system unstable, as seen in the previous subsection.
- v_{fast} a speed of free-flow, typically significantly higher than v_{opt}.

Then, denoting the headway $h \equiv x_{\text{AV}} - x_{\text{lead}}$, we act for $t \in [0,T]$ as follows:

- If $h(t) < d_0$, then
$$v_{\text{cmd}}(t) = v_{\text{lead}}(t) \quad \text{(safety measure)} \tag{24}$$
- If $h(t) \in (d_0, d_1]$, then we denote by t_1 the last time at which $h(t)$ crossed the threshold d_0 (if it does not exist we set $t_1 = 0$ and have
$$v_{\text{cmd}} = v_{\text{des}} \min(c_1 + c_2(t - t_1), 1) \quad \text{(TCP-type hysteresis)}, \tag{25}$$

where c_1 and c_2 are design parameters with $b = 0.8$ typically.

- If $h(t) \in [d_1, d_2)$, then $v_{\text{cmd}} = v_{\text{des}}$
- If $h(t) \geq d_2$ then

$$v_{\text{cmd}} = v_{\text{fast}} \text{ (catching-up)} \tag{26}$$

Note that, by setting initially $t_1 = 0$, we ensure that the AV will start by stabilizing the system to a speed that is significantly smaller than v_{des}. This will ensure that the AV has some room to absorb a stop-and-go wave that would be already present in the system.

4 Results

We present the implementation of our controller on a model representing a stretch of the I-210 California described in Section 2.5. We consider a situation where too many vehicles arrive from inflow compared to the outflow speed and the system gets rapidly congested. After a time of 600s the road is entirely congested with strong stop-and go waves (see Fig. 9 left-hand side). The AVs then are activated and follow the control algorithm described in Section 3.2. They represent 5% of the total number of vehicles and are spanned randomly among lanes in inflow with a uniform distribution. In these simulations, the outflow speed is imposed at 5.5 $m.s^{-1}$ while vehicles are arriving with an inflow speed of 25.5 $m.s^{-1}$. The road has five lanes and is 1430 m long and represents a part of the State Route 134 as a segment of the I-210. We observe an exceptionally large diminution of fuel consumption and a large increase of the energy efficiency (+40% reduction of fuel consumption) while having a comparable throughput (+1% vehicles per minute).

These numbers are quite impressive and, of course, have to be taken with precaution for several reasons. First and foremost, the car-following model used for the simulations (the IDM with the chosen parameters) produces, in the given traffic flow regime, strong stop-and-go traffic waves throughout the domain, so the uncontrolled baseline state is governed by persistent strong braking and acceleration phases. Real traffic flow may only rarely be in such an extreme strong waves state all over the road; and some drivers may drive less aggressively. Second, it is assumed that the uncontrolled vehicles always drive according to the chosen model dynamics and accept the control vehicles' dynamics without any adverse reactions. That being said, the results are not completely unrealistic. In the field, experiments on a ring-road using an AV as a sparse controller [45], a 40% fuel reduction was obtained, which is in the same ballpark.

Besides the absolute value of these numbers, it is interesting to compare to the ideal baseline, namely the steady-state behavior where all vehicles keep the same velocity and headway. This ideal baseline represents the lower possible bound on fuel consumption (equivalently the higher bound on energy efficiency). In this case the throughput is still roughly the same, despite an increase, (+5.8% vehicles per hour) while the energy efficiency shows a reduction of fuel consumption by 55%

with respect to the uncontrolled case. This means that the energy efficiency achieved with our controller applied to 5% AVs is already quite close to the ideal situation in terms of reduction of fuel consumption. Increasing the proportion of AVs logically improves the energy efficiency, but with a decreasing marginal effect as it is bounded anyway by the ideal case. This is summarized in Table 1.

Quantity	Uncontrolled	2% AV	5% AV	15%AV	30% AV	ideal
Throughput (vehicles per hour)	7319	7365	7353	7286	7262	7750
Throughput vs. baseline	0%	+0.63%	+0.46%	−0.45%	−0.77%	+5.8%
Energy efficiency (miles per gallon)	18.1	27.4	32.0	35.5	36.5	40.4
Fuel consumption reduction	0%	34%	43%	49%	50%	55%

Table 1: Effect of AVs on the traffic energy efficiency

To visualize the effect on the stop-and-go waves, we present Fig. 9, which contains the time space diagrams for each lane with and without the controller. The x axis represent the time, the y axis the positions of the different vehicles at the considered time, and the color the actual velocity where red stands for low velocities and green for high velocities. On the left hand side of Fig. 9 there are only regular vehicles, while on the right-hand side there are 5% of AVs equipped with the controller of the previous section. The AVs are activated in the portion of time that is highlighted (upper-right side). This means that, until reaching this time, they imitate human normal human drivers (modelled by the Bando-FTL model) and as soon as they reach this time, they start following the controller. We see that without any controllers strong stop-and-go waves appear and propagate backward on the road. This is illustrated by the many parallel lines of darker colors.

When adding the controllers these stop-and-go waves are partially dissipated even though some still manage to remain. One could do the same with a much larger proportion of AVs and see that, logically, even less waves appear as seen in Fig. 10 when one has 15% of AVs.

5 Conclusion

In this chapter we presented an initiative aiming at dissipating stop-and-go waves by using a small number of autonomous vehicles (AVs) as sparse controllers for traffic flow in a complex setting: a multilane highway. Because of the changes of lane, the resulting system is of hybrid nature with boundary conditions which make it difficult to study theoretically. We proposed several approaches to obtain control algorithms, from Reinforcement Learning controllers to model-based controllers. We have focused on finding robust model-based controllers when the traffic is highly congested. In a highly congested situation the model-based controller exhibits a high energy reduction compared with the uncontrolled case, even with a small proportion of AVs also called penetration rate. Nevertheless, this controller has

limitations and the research is still ongoing to ensure suitability for application to an experimental setting. When increasing the proportion of AVs, the marginal gains in energy efficiency are decreasing. Therefore, an interesting question is the existence of an an optimal trade-off between penetration rate and energy efficiency.

Acknowledgements The authors would like to thank Maria Teresa Chiri for her remarks. This research is based upon work supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Vehicle Technologies Office award number CID DE-EE0008872. The views expressed herein do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References

1. Masako Bando, Katsuya Hasebe, Akihiro Nakayama, Akihiro Shibata, and Yuki Sugiyama. Dynamical model of traffic congestion and numerical simulation. *Physical review E*, 51(2):1035, 1995.
2. L.D. Baskar, De Schutter, and H. Hellendoorn. Model-based predictive traffic control for intelligent vehicles: Dynamic speed limits and dynamic lane allocation. *Proceedings of the 2008 IEEE Intelligent Vehicles Symposium (IV'08)*, pages 174–179, 2008.
3. Nicola Bellomo and Christian Dogbe. On the modeling of traffic and crowds: A survey of models, speculations, and perspectives. *SIAM review*, 53(3):409–463, 2011.
4. Bhadani, Rahul and Sprinkle, Jonathan. *Strym: A data-analytic tool for CAN-bus messages*. Department of Electrical & Computer Engineering, The University of Arizona, 2020. 0.3.1.
5. Matt Bunting, Rahul Bhadani, and Jonathan Sprinkle. Libpanda - a high performance library for vehicle data collection. In *The Workshop on Data-Driven and Intelligent Cyber-Physical Systems*, page Submitted, 2021.
6. José A Carrillo, Massimo Fornasier, Giuseppe Toscani, and Francesco Vecil. Particle, kinetic, and hydrodynamic models of swarming. In *Mathematical modeling of collective behavior in socio-economic and life sciences*, pages 297–336. Springer, 2010.
7. Shumo Cui, Benjamin Seibold, Raphael Stern, and Daniel B Work. Stabilizing traffic flow via a single autonomous vehicle: Possibilities and limitations. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 1336–1341. IEEE, 2017.
8. G. Dal Maso. *An introduction to* Γ *convergence*. Springer, New York, 1993.
9. Sergey Dashkovskiy and Petro Feketa. Zeno phenomenon in hybrid dynamical systems. *PAMM*, 17(1):789–790, 2017.
10. L. C. Davis. Effect of adaptive cruise control systems on traffic flow. *Phys. Rev. E*, 69:066110, Jun 2004.
11. Maria Laura Delle Monache, Thibault Liard, Anaïs Rat, Raphael Stern, Rahul Bhadani, Benjamin Seibold, Jonathan Sprinkle, Daniel B Work, and Benedetto Piccoli. Feedback control algorithms for the dissipation of traffic waves with autonomous vehicles. In *Computational Intelligence and Optimization Methods for Control Engineering*, pages 275–299. Springer, 2019.
12. Oussama Derbel, Tamas Peter, Hossni Zebiri, Benjamin Mourllion, and Michel Basset. Modified intelligent driver model for driver safety and traffic stability improvement. *IFAC Proceedings Volumes*, 46(21):744–749, 2013.
13. Massimo Fornasier, Benedetto Piccoli, and Francesco Rossi. Mean-field sparse optimal control. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2028):20130400, 2014.

14. M. Garavello and B. Piccoli. *Traffic flow on networks*, volume 1 of *AIMS Series on Applied Mathematics*. American Institute of Mathematical Sciences (AIMS), Springfield, MO, 2006. Conservation laws models.
15. Mauro Garavello, Ke Han, and Benedetto Piccoli. *Models for vehicular traffic on networks*, volume 9 of *AIMS Series on Applied Mathematics*. American Institute of Mathematical Sciences (AIMS), Springfield, MO, 2016.
16. Denos C Gazis, Robert Herman, and Richard W Rothery. Nonlinear follow-the-leader models of traffic flow. *Operations research*, 9(4):545–567, 1961.
17. Derek Gloudemans, William Barbour, Nikki Gloudemans, Matthew Neuendorf, Brad Freeze, Said ElSaid, and Daniel B Work. Interstate-24 motion: Closing the loop on smart mobility. In *2020 IEEE Workshop on Design Automation for CPS and IoT (DESTION)*, pages 49–55. IEEE, 2020.
18. X. Gong, B. Piccoli, and G. Visconti. Mean-field limit of a hybrid system for multi-lane multi-class traffic. Preprint arXiv:2007.14655, 2020.
19. Xiaoqian Gong, Benedetto Piccoli, and Giuseppe Visconti. Mean-field of optimal control problems for hybrid model of multilane traffic. *IEEE Control Systems Letters*, 2020.
20. Maxime Guériau, Romain Billot, Nour-Eddin [El Faouzi], Julien Monteil, Frédéric Armetta, and Salima Hassas. How to assess the benefits of connected vehicles? a simulation framework for the design of cooperative traffic management strategies. *Transportation Research Part C: Emerging Technologies*, 67:266 – 279, 2016.
21. Youngjun Han, Danjue Chen, and Soyoung Ahn. Variable speed limit control at fixed freeway bottlenecks using connected vehicles. *Transportation Research Part B: Methodological*, 98:113 – 134, 2017.
22. Amaury Hayat, Benedetto Piccoli, and Sydney Truong. Dissipation of traffic jams using a single autonomous vehicle on a ring road. preprint, 2020.
23. Z. He, L. Zheng, L. Song, and N. Zhu. A jam-absorption driving strategy for mitigating traffic oscillations. *IEEE Transactions on Intelligent Transportation Systems*, 18(4):802–813, 2017.
24. M. Herty and G. Visconti. Analysis of risk levels for traffic on a multi-lane highway. *IFAC-PapersOnLine*, 51(9):43–48, 2018.
25. Wen-Long Jin. A kinematic wave theory of lane-changing traffic flow. *Transportation research part B: methodological*, 44(8-9):1001–1021, 2010.
26. Wen-Long Jin. A multi-commodity lighthill-whitham-richards model of lane-changing traffic flow. *Procedia-Social and Behavioral Sciences*, 80:658–677, 2013.
27. E Kallo, A Fazekas, S Lamberty, and M Oeser. Microscopic traffic data obtained from videos recorded on a german motorway. *Mendeley Data, v1*, 7, 2019.
28. Nicolas Kardous, Amaury Hayat, Sean McQuade, Xiaoqian Gong, Sydney Truong, Paige Arnold, Alexandre Bayen, and Benedetto Piccoli. A rigorous multi-population multi-lane hybrid traffic model and its mean-field limit for dissipation of waves via autonomous vehicles. Preprint, 2020.
29. Jorge A Laval and Carlos F Daganzo. Lane-changing in traffic streams. *Transportation Research Part B: Methodological*, 40(3):251–264, 2006.
30. Chi-Ying Liang and Huei Peng. Optimal adaptive cruise control with guaranteed string stability. *Vehicle system dynamics*, 32(4-5):313–330, 1999.
31. Michael James Lighthill and Gerald Beresford Whitham. On kinematic waves ii. a theory of traffic flow on long crowded roads. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 229(1178):317–345, 1955.
32. P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wiessner. Microscopic traffic simulation using sumo. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 2575–2582, 2018.
33. Akihiro Nakayama, Yūki Sugiyama, and Katsuya Hasebe. Effect of looking at the car that follows in an optimal velocity model of traffic flow. *Physical Review E*, 65(1):016112, 2001.
34. Ryosuke Nishi, Akiyasu Tomoeda, Kenichiro Shimura, and Katsuhiro Nishinari. Theory of jam-absorption driving. *Transportation Research Part B: Methodological*, 50:116 – 129, 2013.

35. Benedetto Piccoli. Hybrid systems and optimal control. In *Proceedings of the 37th IEEE Conference on Decision and Control (Cat. No. 98CH36171)*, volume 1, pages 13–18. IEEE, 1998.
36. Benedetto Piccoli. Necessary conditions for hybrid optimization. In *Proceedings of the 38th IEEE Conference on Decision and Control (Cat. No. 99CH36304)*, volume 1, pages 410–415. IEEE, 1999.
37. Benedetto Piccoli and Francesco Rossi. Generalized wasserstein distance and its application to transport equations with source. *Archive for Rational Mechanics and Analysis*, 211(1):335–358, 2014.
38. Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, page 5. Kobe, Japan, 2009.
39. Rabie A Ramadan and Benjamin Seibold. Traffic flow control and fuel consumption reduction via moving bottlenecks. In *Transportation Research Board Conference*, 2017.
40. Paul I Richards. Shock waves on the highway. *Operations research*, 4(1):42–51, 1956.
41. Harald Schafer, Eder Santana, Andrew Haden, and Riccardo Biasini. A commute in data: The comma2k19 dataset, 2018.
42. John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
43. Sadman-Ahmed Shanto, George Gunter, Rabie Ramadan, Benjamin Seibold, and Dan Work. Challenges of Microsimulation Calibration with Traffic Waves using Aggregate Measurements. Preprint, 2020.
44. Raphael E. Stern, Yuche Chen, Miles Churchill, Fangyu Wu, Maria Laura [Delle Monache], Benedetto Piccoli, Benjamin Seibold, Jonathan Sprinkle, and Daniel B. Work. Quantifying air quality benefits resulting from few autonomous vehicles stabilizing traffic. *Transportation Research Part D: Transport and Environment*, 67:351 – 365, 2019.
45. Raphael E Stern, Shumo Cui, Maria Laura Delle Monache, Rahul Bhadani, Matt Bunting, Miles Churchill, Nathaniel Hamilton, Hannah Pohlmann, Fangyu Wu, Benedetto Piccoli, et al. Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments. *Transportation Research Part C: Emerging Technologies*, 89:205–221, 2018.
46. Yuki Sugiyama, Minoru Fukui, Macoto Kikuchi, Katsuya Hasebe, Akihiro Nakayama, Katsuhiro Nishinari, Shin-ichi Tadaki, and Satoshi Yukawa. Traffic jams without bottlenecks—experimental evidence for the physical mechanism of the formation of a jam. *New journal of physics*, 10(3):033001, 2008.
47. Alireza Talebpour and Hani S. Mahmassani. Influence of connected and autonomous vehicles on traffic flow stability and throughput. *Transportation Research Part C: Emerging Technologies*, 71:143 – 163, 2016.
48. Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical review E*, 62(2):1805, 2000.
49. Joseph Treiterer and Jeffrey Myers. The hysteresis phenomenon in traffic flow. *Transportation and traffic theory*, 6:13–38, 1974.
50. The University of Arizona. *Libpanda: A software library and utilities for interfacing with vehicle hardware systems.*, 2020.
51. M. Wang, W. Daamen, S. P. Hoogendoorn, and B. van Arem. Cooperative car-following control: Distributed algorithm and impact on moving jam features. *IEEE Transactions on Intelligent Transportation Systems*, 17(5):1459–1471, 2016.
52. Meng Wang, Winnie Daamen, Serge P. Hoogendoorn, and Bart van Arem. Connected variable speed limits control and car-following control with vehicle-infrastructure communication to resolve stop-and-go waves. *Journal of Intelligent Transportation Systems*, 20(6):559–572, 2016.
53. Fangyu Wu, Raphael E. Stern, Shumo Cui, Maria Laura [Delle Monache], Rahul Bhadani, Matt Bunting, Miles Churchill, Nathaniel Hamilton, R'mani Haulcy, Benedetto Piccoli, Benjamin Seibold, Jonathan Sprinkle, and Daniel B. Work. Tracking vehicle trajectories and fuel rates in phantom traffic jams: Methodology and data. *Transportation Research Part C: Emerging Technologies*, 99:82 – 109, 2019.

54. Zuduo Zheng. Recent developments and research needs in modeling lane changing. *Transportation research part B: methodological*, 60:16–32, 2014.

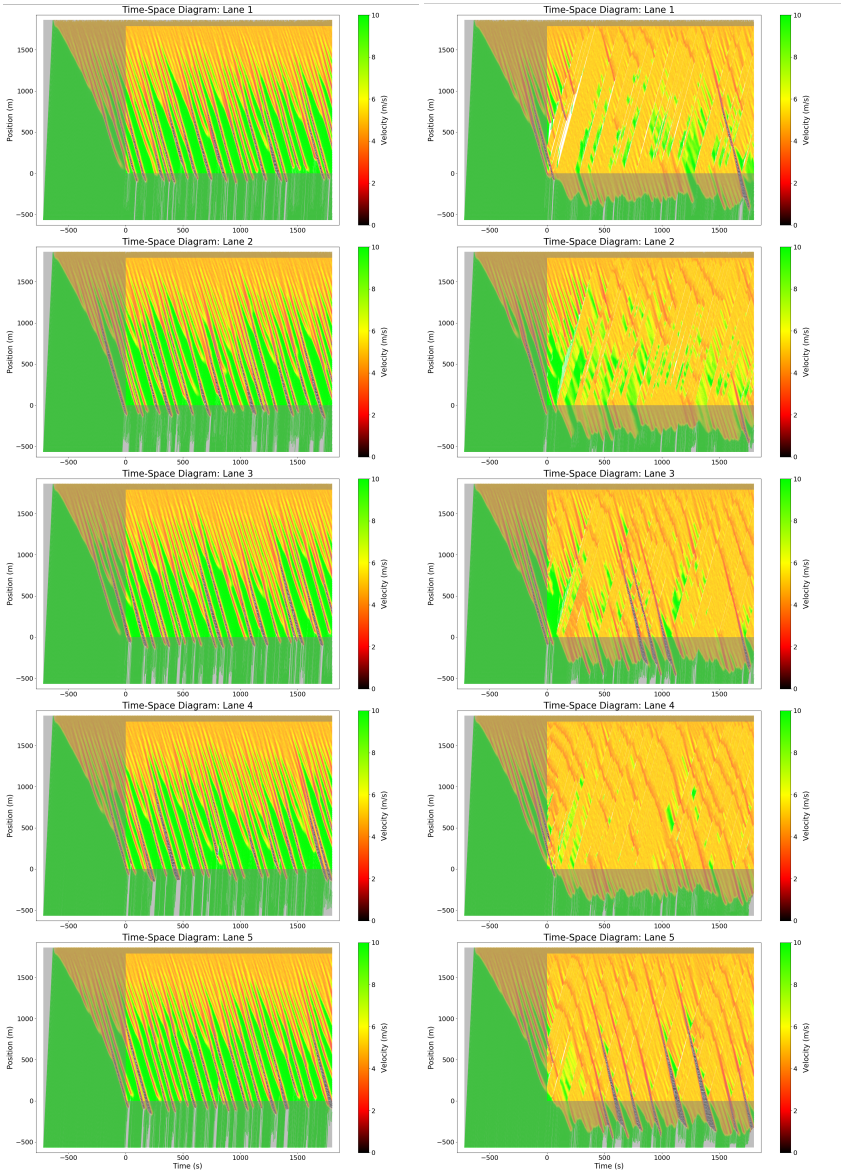

Fig. 9: Time-space diagrams with no control (left side) and 5% AVs (right-side). Colors represent the vehicle velocities.

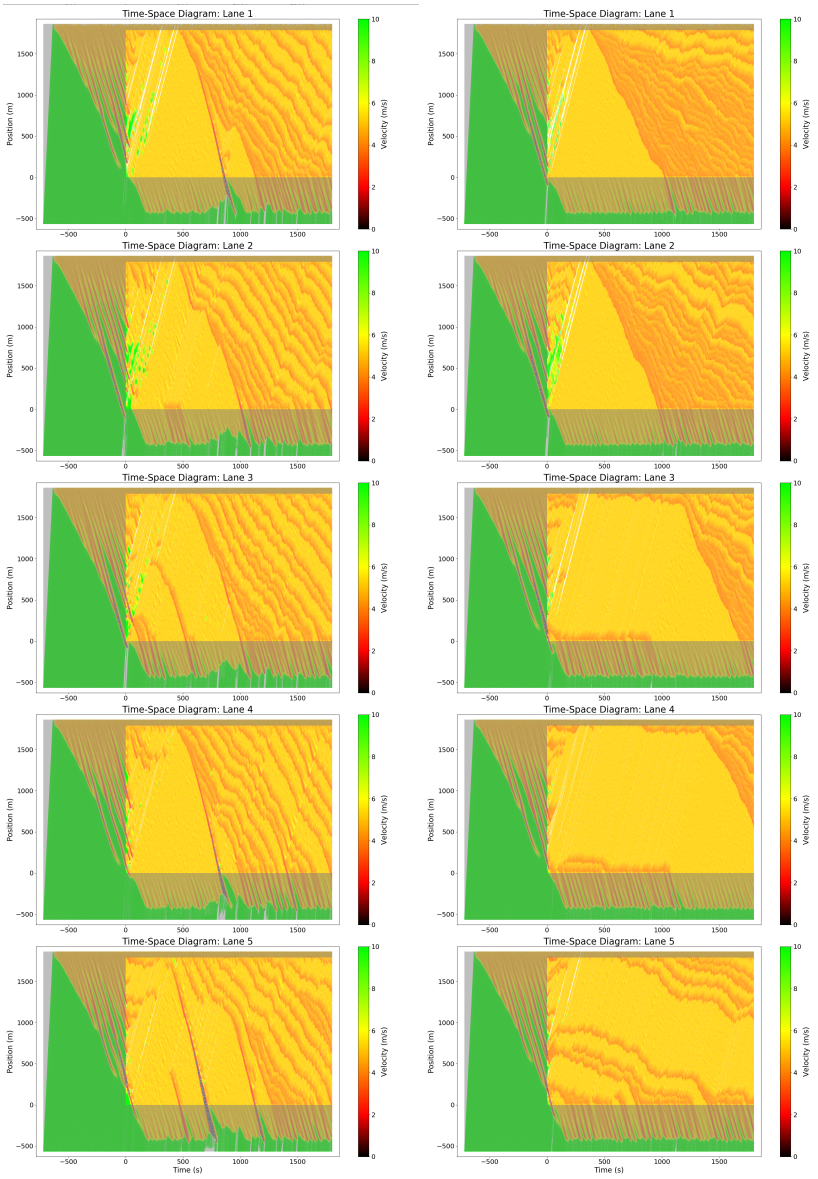

Fig. 10: Time-space diagrams with 15% AVs (left) and 30% AVs (right). Colors represent the vehicles velocities.