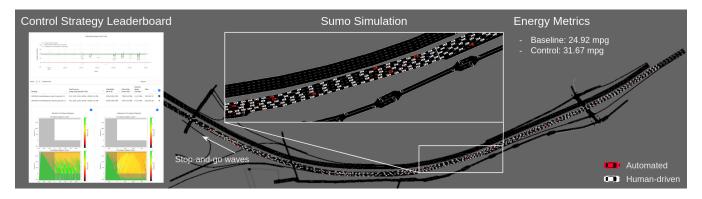
# Integrated Framework of Vehicle Dynamics, Instabilities, Energy Models, and Sparse Flow Smoothing Controllers

Jonathan W. Lee<sup>†</sup>, George Gunter<sup>‡</sup>, Rabie Ramadan\*, Sulaiman Almatrudi<sup>†</sup>, Paige Arnold<sup>§</sup>, John Aquino<sup>†</sup>, William Barbour<sup>‡</sup>, Rahul Bhadani<sup>||</sup>, Joy Carpio<sup>†</sup>, Fang-Chieh Chou<sup>†</sup>, Marsalis Gibson<sup>†</sup>, Xiaoqian Gong<sup>§</sup>, Amaury Hayat<sup>§</sup>, Nour Khoudari\*, Abdul Rahman Kreidieh<sup>†</sup>, Maya Kumar<sup>‡</sup>, Nathan Lichtlé<sup>¶</sup>, Sean McQuade<sup>§</sup>, Brian Nguyen<sup>†</sup>, Megan Ross\*, Sydney Truong<sup>§</sup>, Eugene Vinitsky<sup>†</sup>, Yibo Zhao<sup>†</sup>, Jonathan Sprinkle<sup>||</sup>, Benedetto Piccoli<sup>§</sup>, Alexandre M. Bayen<sup>†</sup>, Daniel B. Work<sup>‡</sup>, Benjamin Seibold\* circles.consortium@gmail.com



**Figure 1.** Traffic waves generated by human driving increase the energy consumption of traffic flow. A small fraction of well-controlled automated vehicles can smooth the flow and the reduce energy consumption.

#### **Abstract**

This work presents an integrated framework of: vehicle dynamics models, with a particular attention to instabilities and traffic waves; vehicle energy models, with particular attention to accurate energy values for strongly unsteady driving profiles; and sparse Lagrangian controls via automated vehicles, with a focus on controls that can be executed via existing technology such as adaptive cruise control systems.

- \*-Temple University
- †–University of California, Berkeley
- <sup>‡</sup>–Vanderbilt University
- §-Rutgers University-Camden
- "-University of Arizona
- ¶-École normale supérieure Paris-Saclay, Paris-Saclay University



This work is licensed under a Creative Commons Attribution International 4.0 License.

DICPS '21, May 18, 2021, Nashville, TN, USA © 2021 Copyright held by the owner/author(s). ACM ISBN 978-1-4503-8445-2/21/05.

https://doi.org/10.1145/3459609.3460530

This framework serves as a key building block in developing control strategies for human-in-the-loop traffic flow smoothing on real highways. In this contribution, we outline the fundamental merits of integrating vehicle dynamics and energy modeling into a single framework, and we demonstrate the energy impact of sparse flow smoothing controllers via simulation results.

CCS Concepts: • Computer systems organization  $\rightarrow$  Embedded and cyber-physical systems; • Computing methodologies  $\rightarrow$  Modeling and simulation; • Software and its engineering  $\rightarrow$  Software creation and management.

**Keywords:** automated vehicles, traffic control systems, microsimulation, energy models, fuel economy

#### 1 Introduction

An emerging research direction is the development of concepts and technologies that enable energy optimization of traffic flow using automated vehicle technologies to control the overall traffic, by exploiting the influence of these vehicles on the bulk traffic flow. In this endeavour it is crucial to correctly model system-level impacts when sparse controllers (*i.e.*, automated vehicles) close the loop in the

real-world application. Moreover, because most of the agents remain human-controlled vehicles, the human-in-the-loop aspects (particularly the reactions of humans to the controllers) are critical for successful control design.

Because such automated vehicles may have different objectives compared to human drivers or classical *Adaptive Cruise Controlled* (ACC) vehicles, a critical enabler of real-world deployment is an integrative simulation/software tool that combines an accurate capture of vehicle dynamics, vehicle energy modeling, as well as controllers, in a way that both micro-scale and macro-scale flow features are reproduced with suitable fidelity. Such tools are needed but not readily available to support research and development.

To illustrate the importance of human driving behavior on traffic flow, the seminal work [20] showed experimentally that human-controlled traffic can exhibit "phantom jams" in which a traffic jam emerges, not from an outside influence such as an accident or lane-reduction, but from the collective behavior of drivers. Phantom jams are an important area of traffic control research because they represent a true inefficiency in the flow in which both system throughput is degraded and average fuel efficiency declines [19, 26], and frequently arise in non-automated (human) traffic flows [20]. It was subsequently shown [19] that these jams could be effectively "smoothed out" with a single automated vehicle on a ring of 20+ other human drivers. The result was that throughput, and notably, energy efficiency of the system were both improved.

Since then, a number of theoretical works have expanded on techniques for using sparsely adopted *connected and automated vehicles* (CAVs) to smooth traffic [15, 25]. Despite the potential for CAVs to significantly improve operation conditions for congested traffic flow, currently available vehicles in the form of ACC do not seem to display such a tendency. Across [11, 12, 14], several commercially available ACCs were tested experimentally, and all such systems actually contributed to phantom jams. As a result, there is a gap between where traffic flow theory and experimentation suggest CAVs could be and where they are in practice.

In order to enable addressing this disconnect, this work develops a high fidelity simulation environment that allows users to quickly implement and test CAV control schemes in complicated traffic environments to assess their capacity to smooth waves and improve energy consumption. A number of microsimulation packages exist, both proprietary and open-source. We build on the widely used open source traffic simulation environment *Simulation of Urban MObility* (SUMO) [13], but recognize it does not directly produce stop-and-go waves (a core feature of traffic that leads to wasted energy). Additionally, we also recognize that suitable models to estimate energy from these trajectories are needed. On one hand, aggregated models such as the US EPA MOVES [1] model produce energy estimates, but based on

average speeds on road segments. On the other extreme, models such as the US DOE Autonomie [18] are highly detailed but computationally costly and not open-source compatible. Recognizing the gap between the tools, here we outline our preliminary work to producing a set of software tools that enables researchers concerned with the development of vehicle controllers to improve the energy efficiency of traffic to develop, test, and benchmark their ideas. The resulting tools take care of the intricacies of non-equilibrium traffic flow theory (e.g., producing phantom jams) while also appropriately approximating vehicle energy dynamics compatible with the traffic dynamics. We demonstrate this energy modeling on generic vehicle types and a Toyota RAV4 vehicle, which is the vehicle platform we intend to use later for large scale traffic control. We also incorporate clear metrics for benchmarking through a leaderboard. The leaderboard provides common traffic scenarios and reports metrics (e.g., systemlevel energy and flow metrics) to assess controllers' effect on the flow dynamics on a system-level scale. While the tools presented here are in a preliminary stage, they are already being used by our multi-institution consortium that includes more than 50 researchers building out these tools [17].

The remainder of the article is as follows. In Section 2 we discuss the challenges to integrate stop-and-go traffic dynamics into SUMO. In Section 3 we outline our work to build simplified energy models that approximate Autonomie models and thus are suitable for integration with SUMO. In Section 4 we discuss the integration efforts and a proof-of-concept case study. Conclusions and future perspectives are given in Section 5.

### 2 Traffic Dynamics Modeling

Traffic is primarily simulated and understood via two different paradigms: *macroscopic* and *microscopic*. In macroscopic modeling, aggregate states of traffic flow are modeled over space and time. Relevant quantities are the *flow rate*, *speed*, and *density* of traffic, and typically partial differential equations are used for simulation, resembling fluid dynamics.

In contrast, the microscopic view models interactions between individual vehicles. These techniques use the *speed*, *spacing-gap*, and *speed-difference* of each vehicle in the network to advance the simulation as a system of ordinary differential equations. Since the primary focus of this work is to understand how controllers implemented on a few CAVs can improve traffic flow, the microscopic paradigm is a natural choice. It allows for high fidelity modeling of individual vehicles, both from a control design perspective and from an energy consumption perspective.

Of particular interest for using automated vehicles to smooth traffic flow is the car-following behavior. Typically, *car-following models* (CFMs) are either first-order or secondorder ODEs, corresponding to either modeling the speed of a vehicle directly or prescribing its acceleration. Suitable for traffic waves are second-order models of the form

$$\dot{v}(t) = f_{\text{CFM}}(s(t), v(t), \Delta v(t)) 
\dot{s}(t) = \Delta v(t)$$
(1)

where  $\dot{v}(t)$  is the acceleration of an ego vehicle, governed by the equation  $f_{\text{CFM}}$ , s(t) is the space-gap to the lead vehicle, v(t) is the vehicle's speed, and  $\Delta v(t)$  is the inter-vehicle speed-difference. Generally, CFMs are used to describe the driving behavior of human drivers, with some popular choices being the Gipps Model [10], the Optimal-Velocity model [16], and the Intelligent Driver Model (IDM) [22]. In this work we use the IDM for human driving, which is of the form

$$f_{\text{IDM}}(s(t), v(t), \Delta v(t)) = a \left[ 1 - \left( \frac{v}{v_0} \right)^{\delta} - \left( \frac{s^*(v, \Delta v)}{s} \right)^2 \right]$$

$$s^*(v, \Delta v) = s_0 + vT + \frac{\max\{0, v\Delta v\}}{2\sqrt{ab}},$$
(2)

and a, b,  $v_0$ ,  $\delta$ , T, and  $s_0$  are model parameters that influence driver behavior. Moreover, the acceleration (2) is replaced by zero whenever  $f_{\text{IDM}} < 0$  and v = 0.

Of central interest in this work are phantom jams, in which small disturbances to an otherwise steady traffic flow can translate into large slow-downs that move move backwards in the flow. Mathematically, phantom jams can be interpreted as dynamic instabilities that grow into nonlinear waves [9]. In suitable scenarios, the (in)stability of a uniform flow is equivalent to the concept of *string-stability* [23]. To determine whether a CFM can exhibit phantom jams due to stringinstability, a linear stability analysis using partial derivatives can be applied to derive a transfer function:

$$\begin{split} \ddot{y}_i(t) &= \alpha_1(y_{i-1} - y_i) - \alpha_2 \dot{y}_i + \alpha_3 \dot{y}_{i-1} \;, \\ \text{where } \alpha_1 &= \frac{\partial f}{\partial s} \,, \; \alpha_2 = \frac{\partial f}{\partial (\Delta v)} - \frac{\partial f}{\partial v} \,, \; \alpha_3 = \frac{\partial f}{\partial (\Delta v)} \;, \\ \text{and } F(\omega) &= \frac{\alpha_1 + \alpha_3 \omega}{\alpha_1 + \alpha_2 \omega + \omega^2} \;. \end{split}$$

Here  $F(\omega)$  represents the amplification of a perturbation of frequency  $\omega$ , and y(t) is a linearization of the CFM. The system is string stable if  $|F(\omega)| \leq 1 \ \forall \omega \in i\mathbb{R}$ , which is equivalent to the simple algebraic criterion

$$\lambda \ge 0$$
 where  $\lambda = \alpha_2^2 - \alpha_3^2 - 2\alpha_1$ . (3)

Traffic modeled with a given CFM is then able to exhibit phantom jams when (3) is violated (see [5] for more details).

Leveraging this stability theory, a microsimulation environment is set up in SUMO based on the IDM (2) to model traffic streams of human drivers which exhibit phantom jams, with the goal of designing CAV controllers to then smooth those waves. To do this, IDM parameters are chosen such that (i) a realistic fundamental diagram is reproduced; and (ii) at/near the critical density, a transition from stability to instability occurs with growth rates that grow waves from small perturbations on a time scale of 10s–30s. We choose  $a=1.3 \text{ m/s}^2$ ,  $b=2 \text{ m/s}^2$ ,  $v_0=30 \text{ m/s}$ ,  $\delta=4$ , T=1 s, and  $s_0=1 \text{m}$ .

The integrator time step for the simulation is chosen small enough so that the stability criterion (3) and wave behavior from the IDM are well-captured, but that it still allows for efficient simulation. Moreover, strict adherence to speed/acceleration/gap bounds of the discrete time-stepping system is guaranteed via standard fail-safes. Random noise is added to the prescribed IDM acceleration values to trigger any dynamic instabilities (the noise effects are properly filtered out when feeding accelerations into the energy models below). Finally, lane changing, geometries, vehicle inflow spawning, and routing are all handled using SUMO's existing capabilities, while the aforementioned custom car-following logic is implemented using the software package Flow [24], and further implementation details are described in Section 4. This framework leverages existing modeling work of the microsimulation to handle "real" traffic, while at the same time allowing for high fidelity and principled choice/development of a CFM that satisfies the critical requirement that it produces systematic instabilities and waves that resemble real stop-and-go traffic.

# 3 Energy Models

This project requires vehicle energy models that are (a) accurate for highly non-constant velocities; (b) representative of a variety of vehicle types; yet also (c) structurally simple enough to allow for fast evaluation, use in optimization/control, and open source implementation. To meet requirements (a) and (b), the Autonomie software [3] is employed on a selection of representative vehicles (Table 1), and a systematic model-reduction procedure is devised to generate simple fitted models that meet requirement (c).

Vehicle portfolio: The selected vehicles are chosen to satisfy diversity and prevalence in the US. The top 5 vehicles in Table 1 represent generic vehicles that are representative for current market vehicles of that specific type [3]. Toyota RAV4 is chosen because this is the consortium's [17] intended primary controller vehicle used for flow smoothing. Each vehicle model represents a class of vehicles that have comparable weight (with load assumed half full) and fuel consumption characteristics. The road share in Table 1 is obtained as follows: (i) share of trucks (Class3 PND) vs. passenger vehicles (rest) from TN DoR [21] vehicle registration data (as future field tests are expected in TN); (ii) distribution within passenger vehicles from CNCDA [4] sales data.

Autonomie: The simulation software *Autonomie Rev 16SP7* includes detailed vehicle dynamics, energy models, and a library for several types of vehicles, which can be used for estimating fuel/energy consumption and other vehicle performances such as emissions, regenerative braking, etc. [3]. Each vehicle model is composed of detailed plant and controller models for its components, including engine, drivetrain, driver, and environment. The RAV4 model is based on a small SUV template, adapted to the specifications of the

RAV4 [7]. Autonomie works with MATLAB and Simulink, where its blocks and files can be modified and customized.

Virtual chassis dynamometer: In order to compute vehicle performance maps in a full parameter space of driving situations, the Autonomie model framework is modified as follows: (1) in the driver model, braking is deactivated, and the accelerator is forced to a prescribed test pattern; (2) in the environment model, a PI controller is inserted to artificially adjust road grade so that actual vehicle speed matches the target speed as the accelerator test pattern varies power delivered to the wheels; (3) in the gearbox demand and transient controllers, gear shifting is deactivated; (4) in the engine transient controller, the idle speed controller logic is bypassed to simply pass through the engine torque demand; (5) in the wheel plant, additional load is introduced to cancel the artificial road grade signal.

The thus modified model is then run, gear-by-gear, through a complete velocity—load phase space, where the load represents various acceleration values. Each set point is held for 10 seconds to allow the model to reach local equilibrium. The resulting maps (vehicle speed to engine speed, vehicle speed and wheel force to engine torque, and engine speed and torque to fuel-rate) are then used in the next step.

**Semi-principled model:** We formulate a simplified, semi-principled, model that has a physics-based part (using vehicle mass, road load coefficients, gear ratios, final drive ratio, tire diameter, maximum engine torque, maximum engine speed, and engine idle speed), but it also relies on the maps obtained from the above virtual chassis dyno. The model takes as inputs the instantaneous vehicle speed v, acceleration a, and road grade  $\theta$ , and outputs engine speed, engine torque, fuel consumption, gear, transmission output speed, wheel force, wheel power, and feasibility of the given  $(v, a, \theta)$  with respect to engine speed and engine torque. Gear scheduling is based on choosing the (feasible) gear that yields the minimal fuel consumption. In contrast to the original Autonomie model, this simplified model now yields the fuel consumption rate f (and other outputs) as a direct function  $f = F(v, a, \theta)$ .

The model version used herein is simplifying real vehicle dynamics by assuming that the torque converter is always locked (in reality, an open torque converter bypass facilitates vehicle launch and mitigates driveline vibration). This simplification results in slight underpredictions of the fuel rate in lower gears, and it will be improved in future versions.

**Tuning parameters:** Beyond the physics-based vehicle parameters, the simplified model also uses a few tuning parameters. These are extracted in an automated fashion from the original Autonomie vehicle model run on test cycles: minimum engine torque after gear shifting, fuel cut speed, upshifting engine speeds, and downshifting vehicle speeds.

**Fitted polynomial model:** To generate even simpler models, a further simplification step is applied. For each vehicle, the semi-principled model is evaluated on a grid in a region of the feasible (v, a)-space. To these data, a (capped)

| vehicle model<br>(engine type) | model<br>error SP | model<br>error FP | HV<br>share |
|--------------------------------|-------------------|-------------------|-------------|
| Compact Sedan (SI)             | -2.0%             | -2.7%             | 23.59%      |
| Midsize Sedan (SI)             | 2.7%              | 2.9%              | 32.92%      |
| Midsize SUV (SI)               | -2.2%             | -2.3%             | 17.56%      |
| Midsize Pickup (SI)            | -3.1%             | -3.0%             | 10.32%      |
| Class3 PND (CI)                | 0.2%              | -2.1%             | 15.61%      |
| 2019 RAV4 (SI)                 | -1.6%             | -1.0%             | _           |

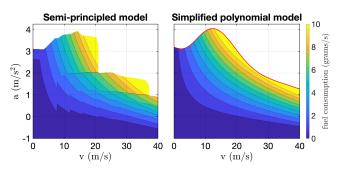
**Table 1.** Vehicle models. Error SP (FP) is the relative error of the semi-principled (fitted polynomial) model vs. Autonomie baseline. (HV=human vehicle; SI=spark-ignition; CI=compression-ignition; PND=pickup-and-delivery)

degree 3 bivariate polynomial (in v and a) is fitted (least squares sense with non-negativity constraints on the 10 parameters) of the form

$$f(v,a) = \max \left\{ C_0 + C_1 v + C_2 v^2 + C_3 v^3 + p_0 a + p_1 a v + p_2 a v^2 + q_0 a_+^2 + q_1 a_+^2 v, \beta \right\},$$

where  $a_+ = \max\{a, 0\}$ , and  $\beta$  is the minimum fuel rate, which is not necessarily zero because different vehicles have different criteria for enacting a fuel cut.

Moreover, a fit for the boundary of the feasibility region is produced, in the form of a function g(v), above which (v, a)-pairs are infeasible. Figure 2 shows a plot of the semi-principled and fitted polynomial models for a midsize SUV.



**Figure 2.** Semi-principled fuel consumption model (left) and simplified polynomial model (right) for midsize SUV. The red curve is g(v), the boundary of the feasibility region.

**Model validation:** For each vehicle, the accuracy of both model simplifications is validated against the Autonomie ground truth based on standard drive cycles (for that vehicle type). Used are: (a) the EPA cycles UDDS, FTP-75, US06, 505 for light-duty vehicles [8]; (b) the globally (UNECE) approved WLTC [2]. Table 1 summarizes, for each vehicle, the relative error of both model types relative to Autonomie, averaged over all respective drive cycles. One can see that on average, all simplified models are safely within 4% relative error.

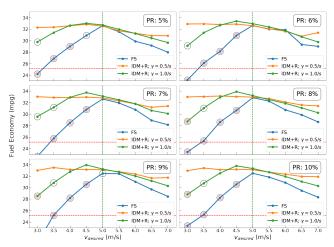
# 4 Integration and Simulation Results

The major components of microsimulation that SUMO handles are *car-following*, *lane-changing*, and *routing*. We have built an interface layer on top of SUMO using Flow [24], which serves as the backbone infrastructure of our integrated framework. Flow builds on SUMO's underlying logic to allow both for custom definition of car-following logic, and for the implementation of vehicle level control schemes, such as in the form of *reinforcement learning* (RL). Through Flow, microsimulations with custom network geometries are instantiated in SUMO and interfaced with different vehicle controllers and behavior models, road networks, traffic environments, vehicle energy models, etc. This framework is modular in its components and features, thereby enabling a systematic assessment of different choices of model design.

Simulation outputs are post-processed in a custom data pipeline implemented in Amazon AWS cloud. Output data are uploaded to AWS S3 storage as CSV files, which then triggers automatic post-processing using AWS Lambda and AWS Athena services. A load-balancing mechanism is implemented to handle parallel uploads at any scale. Centralized storage allows for standardized post-processing that is decoupled from simulating vehicle dynamics; i.e., simulations do not need to be rerun as post-processing procedures and models are updated. SQL queries are used to compute several key performance indicators (KPIs). Resultant table schemata are intentionally designed to facilitate quick file ingestion and plot rendering by an online benchmarking leaderboard for ranking controller strategies. The leaderboard, built with DataTables & Plotly. js, features standardized scenarios, evaluation metrics and plots, and vehicle distributions.

The primary KPI of this study is fuel economy (*i.e.*, *miles per gallon* (mpg)) of the whole traffic flow, leveraging the energy models from Section 3 to estimate fuel consumption<sup>1</sup>. Further, the evaluation pipeline has implemented several vehicle distributions by mapping vehicles to energy models. CAVs are mapped to the RAV4 model, and others are mapped onto the "human" vehicles according to proportions shown in Table 1. Several other KPIs (*e.g.*, network speed, network inflow rate) and plots (*e.g.*, time-space diagrams, mpg histograms, various telemetry vs. relative time or distance) further illustrate each controller's performance relative to a standardized baseline of uncontrolled traffic flow with waves.

**Experimental setup:** The benchmarking pipeline is designed for a broad customizable range of scenarios, parametric studies, and vehicle distributions. The scenario discussed herein features a network model of CA SR-134 (~1-mile stretch from the I-210 Pilot project [6]). The road is assumed flat, and ramps are disengaged to isolate the effects of waves due to congestion. All human-driven vehicles are dynamically equivalent, utilizing the IDM parameters



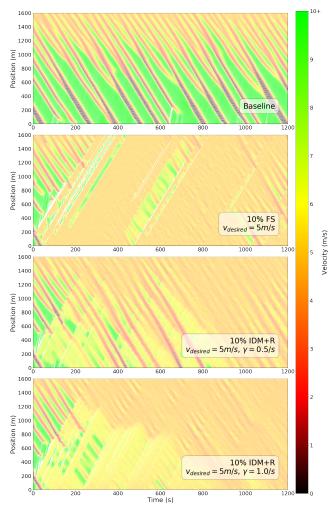
**Figure 3.** Fuel economy results of three vehicle controllers over parameter  $v_{\rm desired}$  (x-axis) and penetration rates (subplots). Red (black) circles denote decreased inflow rate (network speed) by more than 10%. Red dashed line denotes baseline fuel economy, and green dashed line highlights the true average traffic speed downstream.

from Section 2 and SUMO's lane-changing. User-defined, longitudinally-controlled CAVs are evenly injected at prescribed penetration rates. Edges are added upstream and downstream to remove boundary effects from metrics evaluation. Unless unsafe, vehicles are injected at 2050 vehicles per hour (a value that was found to lead to both reasonable congestion and desirable wave formation). Congestion is achieved by imposing a 5 m/s speed limit immediately downstream from the simulation domain. The simulator uses a ballistic integration scheme with a time step of 0.4 s. Zeromean Gaussian noise of standard deviation 0.1 m/s² is added to the acceleration to provoke waves, as previously discussed. The simulation is warmed up for 720 s to let waves establish throughout the domain and then run for 1200 s.

**Sparse Lagrangian traffic controllers:** As a demonstration of the framework, two parameterized vehicle controllers are submitted through the pipeline at various penetration rates (5%–10%): (1) FollowerStopper (FS) [19] with varied  $v_{\text{desired}}$  (3 m/s–7 m/s) and (2) IDM with Relaxation (IDM+R) [5], given by  $\dot{v} = f_{\text{IDM}}(s, v, \Delta v) + \gamma (v_{\text{desired}} - v)$ . For simplicity, the human CFM parameters in (2) are assumed precisely known here. The control gain  $\gamma$  is a tunable parameter. Here we study varying  $v_{\text{desired}}$  (3 m/s–7 m/s) and  $\gamma$  (0.5 /s & 1.0 /s).

**Results:** As shown in Figure 3, as long as  $v_{\rm desired}$  is a reasonable estimate of the true average speed of traffic, all tested controllers result in significant energy improvements, with the exception of the FS for  $v_{\rm desired} < 3.5$  m/s. As the flagging thresholds in Figure 3 show, the FS sacrifices on network speed and throughput when  $v_{\rm desired}$  is too low, while the IDM+R controller generally appears more robust in those regards, yet still performs well in fuel economy. Furthermore,

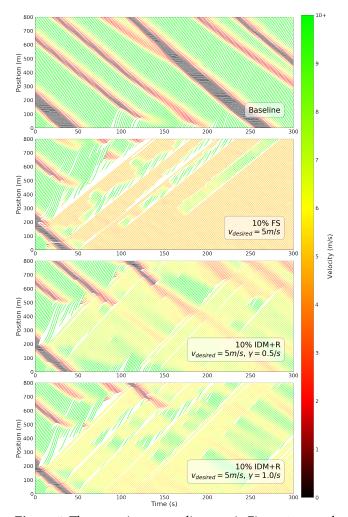
<sup>&</sup>lt;sup>1</sup>The energy models of Section 3 are implemented twice in: (1) Flow for RL optimization; (2) AWS for post-processing evaluation.



**Figure 4.** Time-space diagrams of controllers at 10% penetration rate and  $v_{\rm desired} = 5$  m/s, which yield similar fuel economy. The uncontrolled baseline shows unimpeded waves; controlled cases show some wave dissipation.

the IDM+R controller consistently matches or outperforms the FS. The IDM+R controller has a free parameter  $\gamma$  (how aggressively to steer v towards  $v_{\rm desired}$ ). The results suggest that more aggressive control (i.e., larger  $\gamma$ ) has equal or positive impact if  $v_{\rm desired}$  does not significantly underestimate the average speed. Otherwise, the more aggressive controller suffers from similar drawbacks as the FS. Conversely, overestimating the true average speed results in less efficient wave dampening.

Figure 4 & 5 shows how the controllers act to interrupt the progression of waves via their increased tendencies to drive at uniform speeds. It is apparent that no controller removes all waves, and that there are qualitative differences across the different controllers. Most prominently, the FS conducts noticeably more wave removal than the IDM+R controllers; however, it does so at the expense of allowing



**Figure 5.** The same time-space diagrams in Figure 4 zoomed in on both axes to highlight individual trajectories.

more lane changes. These aspects result in an overall similar performance in terms of energy metrics.

# 5 Conclusions and Outlook

The presented framework provides a standardized benchmarking tool for researchers to develop and assess Lagrangian traffic-smoothing controllers. Neither of the presented controllers were constructed to be optimal; rather they are simple recipes taken from existing work that prove the concept. The authors present this work as a vision of how future optimized controllers may be evaluated against one another.

## Acknowledgments

The authors would like to thank Kenneth Butts (Toyota) for helpful discussions and comments. This material is based upon work supported by the National Science Foundation under Grants CNS-1837244/CNS-1837652/CNS-1837481/OISE-1743772. This material 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

- Environmental Protection Agency. 2020. MOtor Vehicle Emission Simulator (MOVES). https://www.epa.gov/moves.
- [2] UK Vehicle Certification Agency. 2020. The Worldwide Harmonised Light Vehicle Test Procedure (WLTP). From https://www.vehicle-certification-agency.gov.uk/fuel-consumption-co2/the-worldwide-harmonised-light-vehicle-test-procedure/.
- [3] Argonne National Laboratory (ANL). 2020. Autonomie Compiled Vehicles. A technical manual for training in Autonomie.
- [4] CNCDA. 2020. Comprehensive information on the California vehicle market. Technical Report 1. CA New Car Dealers Association. https://www.cncda.org/wp-content/uploads/Cal-Covering-4Q-19.pdf.
- [5] S. Cui et al. 2017. Stabilizing traffic flow via a single autonomous vehicle: Possibilities and limitations. In 2017 IEEE Intelligent Vehicles Symposium (IV). 1336–1341. https://doi.org/10.1109/IVS.2017.7995897
- [6] F. Dion et al. 2015. Connected Corridors: I-210 Pilot Integrated Corridor Management System, Concept of Operations.
- [7] EPA. 2019. 2018 Toyota 2.5L A25A-KS Engine Tier 2 Fuel ALPHA Map Package. Version 2019-03.
- [8] EPA. 2020. Dynamometer Drive Schedules. From https://www.epa.gov/vehicle-and-fuel-emissions-testing/dynamometer-drive-schedules.
- [9] M. R. Flynn et al. 2009. Self-sustained nonlinear waves in traffic flow. Phys. Rev. E 79, 5 (2009), 056113.
- [10] P. G. Gipps. 1981. A behavioural car-following model for computer simulation. Transp. Research B: Methodological 15, 2 (1981), 105–111.
- [11] G. Gunter et al. 2019. Model-based string stability of adaptive cruise control systems using field data. *IEEE Transactions on Intelligent Vehi*cles 5, 1 (2019), 90–99.
- [12] G. Gunter et al. 2020. Are Commercially Implemented Adaptive Cruise Control Systems String Stable? *IEEE Transactions on Intelligent Trans*portation Systems (2020), 1–12. https://doi.org/10.1109/TITS.2020. 3000682
- [13] D. Krajzewicz et al. 2012. Recent Development and Applications of SUMO - Simulation of Urban Mobility. *International Journal On Advances in Systems and Measurements* 3&4 (December 2012). https://elib.dlr.de/79485/
- [14] V. Milanés et al. 2013. Cooperative adaptive cruise control in real traffic situations. *IEEE Trans. Intel. Transp. Systems* 15, 1 (2013), 296–305.
- [15] M. L. Delle Monache et al. 2019. Feedback control algorithms for the dissipation of traffic waves with autonomous vehicles. In Comput. Intell. & Optim. Methods for Control Eng. Springer, 275–299.
- [16] A. Nakayama et al. 2016. Quantitative explanation of circuit experiments and real traffic using the optimal velocity model. New Journal of Physics 18, 4 (2016), 043040.
- [17] B. Nguyen and J.W. Lee. 2020. CIRCLES. https://circles-consortium. github.io/
- [18] Department of Energy. 2020. Vehicle Energy Consumption and Performance Analysis. From https://www.autonomie.net/expertise/ VehicleEnergyConsumptionAndPerformanceAnalysis.html.
- [19] R. E. Stern et al. 2018. Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments. *Transp. Res. Part C Emerg.* 89 (2018), 205–221.
- [20] Y. Sugiyama et al. 2008. Traffic jams without bottlenecks experimental evidence for the physical mechanism of the formation of a jam. New Journal of Physics 10, 3 (2008), 033001.
- [21] TDOR. 2020. Active Motor Vehicle Registrations in TN. From https://www.tn.gov/revenue/title-and-registration/vehicle-titling---

#### registration/vehicle-registration.

- [22] M. Treiber et al. 2000. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E* 62, 2 (2000), 1805.
- [23] R. E. Wilson and J. A. Ward. 2011. Car-following models: fifty years of linear stability analysis—a mathematical perspective. *Transportation Planning and Technology* 34, 1 (2011), 3–18.
- [24] C. Wu et al. 2017. Flow: Architecture and benchmarking for reinforcement learning in traffic control. arXiv preprint arXiv:1710.05465 (2017), 10.
- [25] C. Wu et al. 2018. Stabilizing Traffic with Autonomous Vehicles. In 2018 IEEE International Conference on Robotics and Automation (ICRA). 6012–6018. https://doi.org/10.1109/ICRA.2018.8460567
- [26] F. Wu et al. 2019. Tracking vehicle trajectories and fuel rates in phantom traffic jams: Methodology and data. *Transp. Res. Part C Emerg.* 99 (2019), 82–109.