



HAL
open science

Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments

Raphael E. Stern, Shumo Cui, Maria Laura Delle Monache, Rahul Bhadani, Matt Bunting, Miles Churchill, Nathaniel Hamilton, R'Mani Haulcy, Hannah Pohlmann, Fangyu Wu, et al.

► To cite this version:

Raphael E. Stern, Shumo Cui, Maria Laura Delle Monache, Rahul Bhadani, Matt Bunting, et al.. Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments. Transportation research. Part C, Emerging technologies, 2018, 89, pp.205-221. 10.1016/j.trc.2018.02.005 . hal-01614638

HAL Id: hal-01614638

<https://inria.hal.science/hal-01614638>

Submitted on 11 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments

Raphael E. Stern^{a,1}, Shumo Cui^{b,1}, Maria Laura Delle Monache^{c,1}, Rahul Bhadani^d, Matt Bunting^d, Miles Churchill^a, Nathaniel Hamilton^e, R'mani Haulcy^f, Hannah Pohlmann^g, Fangyu Wu^a, Benedetto Piccoli^{h,2}, Benjamin Seibold^{b,2}, Jonathan Sprinkle^{d,2}, Daniel B. Work^{a,2,*}

^aDepartment of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, 205 N. Mathews Ave, Urbana, IL 61801, USA.

^bDepartment of Mathematics, Temple University, 1805 North Broad Street, Philadelphia, PA 19122, USA.

^cInria, University Grenoble Alpes, CNRS, GIPSA-lab, F-38000 Grenoble, France.

^dElectrical and Computer Engineering, University of Arizona, Tucson, AZ 85721-0104, USA.

^eLipscomb University, 1 University Park Drive, Nashville, TN 37204, USA.

^fYale University, New Haven, CT 06520, USA.

^gPennsylvania State University, University Park, PA 16801, USA.

^hDepartment of Mathematical Sciences, Rutgers University – Camden, 311 N. 5th St, Camden, NJ 08102, USA.

Abstract

Traffic waves are phenomena that emerge when the vehicular density exceeds a critical threshold. Considering the presence of increasingly automated vehicles in the traffic stream, a number of research activities have focused on the influence of automated vehicles on the bulk traffic flow. In the present article, we demonstrate experimentally that intelligent control of an autonomous vehicle is able to dampen stop-and-go waves that can arise even in the absence of geometric or lane changing triggers. Precisely, our experiments on a circular track with more than 20 vehicles show that traffic waves emerge consistently, and that they can be dampened by controlling the velocity of a single vehicle in the flow. We compare metrics for velocity, braking events, and fuel economy across experiments. These experimental findings suggest a paradigm shift in traffic management: flow control will be possible via a few mobile actuators (less than 5%) long before a majority of vehicles have autonomous capabilities.

Keywords: Traffic waves, Autonomous vehicles, Traffic control

1. Introduction

1.1. Motivation

The dynamics of traffic flow include instabilities as density increases, where small perturbations amplify and grow into stop-and-go waves that travel backwards along the road [1, 2, 3, 4]. These so-called *phantom* traffic jams are an experimentally reproducible phenomenon, as demonstrated in different experiments [2, 5, 6]. Common wave triggers include lane changing [7, 8, 9], but they can even be generated in the absence of any lane changes, bottlenecks, merges, or changes in grade [2, 5]. Moreover, these waves can be captured in microscopic models of individual vehicle motion [10, 11, 12] (see also the reviews [13, 14, 15]) and macroscopic models described via solutions to continuum problems [3, 16, 17, 18, 19, 20]. Since these waves emerge from the collective dynamics of the drivers on the road, they are in principle avoidable if one could affect the way people drive. Recognizing the rapid technological innovations in traffic state estimation and control, this work provides experimental evidence that these waves can be reduced by controlling a small number of vehicles in the traffic stream.

*Corresponding author

Email address: dbwork@illinois.edu (Daniel B. Work)

¹R. Stern, S. Cui, and M.L. Delle Monache contributed equally to this work.

²B. Piccoli, B. Seibold, J. Sprinkle, and D. Work contributed equally to this work.

A necessary precursor to dissipating traffic waves is to detect them in real-time. Advancements in traffic state estimation [21, 22, 23] have facilitated high resolution traffic monitoring, through the advent of GPS smartphone sensors [24, 25, 26, 27] that are part of the flow—termed *Lagrangian* or *mobile* sensors. Now commercialized by several major navigation services, the use of a small number of GPS equipped vehicles in the traffic stream has dramatically changed how traffic is monitored for consumer-facing mobility services, which previously relied on predominantly fixed sensing infrastructure.

Currently, traffic control is dominated by control strategies that rely on actuators at fixed locations or are centralized. Such systems include *variable speed advisory* (VSA) or *variable speed limits* (VSL) [28, 29, 30, 31, 32], which are commonly implemented through signs on overhead gantries, and ramp metering [33, 34, 35], which relies on traffic signals on freeway entrance ramps. More recently, coordinated systems to integrate both ramp metering and variable speed limits have been proposed [36, 37, 38, 39]. A common challenge of VSL and ramp metering systems is the small flexibility of the systems due to the high cost of installation of the fixed infrastructure, which consequently limits the spatial resolution of the control input.

Recent advancements in vehicular automation and communication technologies provide new possibilities and opportunities for traffic control in which these smart vehicles act as Lagrangian actuators of the bulk traffic stream. When a series of adjacent vehicles on a roadway are connected and automated, it is possible to form dense platoons of vehicles which leave very small gaps. A key challenge for vehicle platoons is to design control laws in which the vehicle platoon remains stable, for which significant theoretical and practical progress has been made [40, 41, 42, 43, 44, 45, 46, 47, 48]. In contrast to the vehicle platoon setting, in which all vehicles are controlled, or the variable speed limit and ramp metering strategies which actuate the flow at fixed locations, this research aims to dissipate congestion-based stop-and-go traffic waves using only a sparse number of autonomous vehicles already in the flow, without changing how the other, human-driven, vehicles operate.

The notion to dissipate stop-and-go waves via controlling vehicles in the stream represents a shift from stationary to Lagrangian control, mirroring the transition to Lagrangian sensing that has already occurred. The key advantage in mobile sensing projects [24, 26, 27] is that a very small number of vehicles being measured (3-5%) suffices to estimate the traffic state on large road networks [25]. In the same spirit, our research experimentally demonstrates that a small number of Lagrangian controllers suffices to dampen traffic waves.

The ability of connected and automated vehicles to change the properties of the bulk traffic flow is already recognized in the transportation engineering community. For example, the works [49, 50, 51, 52] directly address the setting where a subset of the vehicles are equipped with automated and/or connected technologies, and then assess via a stability analysis or simulation the extent to which the total vehicular flow can be smoothed. Recently, several works have explored extensions to the variable speed limit control strategies in which connected or automated vehicles are used to actuate the traffic flow [53]. For example, the work [54] develops a VSL strategy that is implemented in simulation with connected vehicles where the traffic evolves according to the kinematic wave theory. It follows a similar strategy proposed in [55], where a coordinated VSL and ramp metering strategy is implemented via actuation of the entire vehicle fleet (i.e., 100% penetration rate). Although not explicitly designed as a variable speed limit controller, the article [56] advocates a “slow-in, fast-out” driving strategy to eliminate traffic jams, using a microscopic model also in line with kinematic wave theory. The work [57] proposes a similar jam absorbing strategy as [56] based on Newell’s car following theory, and its effectiveness is assessed in simulation.

Interestingly, an experimental test of the “slow-in, fast-out” strategy [56] is provided in [58], in which five vehicles are driven on a closed course. The lead vehicle in the platoon of five vehicles drives initially at a constant speed, then decelerates as if driving through a congestion wave, and then accelerates back to the cruising speed. The third vehicle in the platoon initially leaves a large gap, and due to the extra gap it is able to maintain the cruising speed and effectively absorb the jam. In contrast to the experiment [58], the present work fully replicates the setup of Sugiyama et al. [2, 5], in which the stop-and-go wave is generated naturally from the human drivers in the experiment, without an external cause. Moreover, the controllers proposed in the present work are distinct.

We also note some preliminary field experiments to harmonize speeds via connected vehicles are recently reported in [59, 60], in part to measure the impact of connected vehicles following an infrastructure-generated advisory speed on the traffic stream behind the connected vehicles. In the present article, we instead dampen waves on a closed ring, which simplifies the experimental setup, and facilitates detailed data collection on the performance of the controllers.

1.2. Problem statement and contributions

The present article is inspired by the work of Sugiyama et al. [2, 5] which are the first works to demonstrate via experiments that traffic waves can emerge without infrastructure bottlenecks or lane changing. A series of experiments is conducted where approximately 20 vehicles drive in a ring of fixed radius with each driver following the vehicle in front of them. The experiments [2, 5] are foundational because they demonstrate the emergence of traffic waves caused (unintentionally) by human driving behavior. However, they do not offer a solution for dampening these waves.

To address this gap, we design and execute a series of ring-road experiments which show that an intelligently controlled autonomous vehicle is able to dampen stop-and-go waves. The experimental setup (described in Section 2) follows the setting of Sugiyama et al. [2, 5], with the modification that one vehicle is an autonomous-capable vehicle which can run a variety of longitudinal control laws. Similar to the Sugiyama et al. [2] experiment, the position and velocity of each vehicle is tracked via a 360 degree camera. We additionally instrument each vehicle in the 22-car fleet (see Table 5 in Appendix 5.1 for a detailed description of the vehicle fleet) with an OBD-II scanner to log the real-time fuel consumption of each vehicle, such that the impact of the traffic waves and controllers on the bulk fuel consumption can be recorded.

We present three experiments (labeled as A, B, and C in Section 4) and two distinct control strategies (detailed in Section 3) that can be used to dampen stop-and-go waves created by human drivers. The first control strategy is to follow a fixed average velocity (selected based on observation) as closely as possible without collisions. It is implemented in Experiment A via an automatic control algorithm (called *FollowerStopper*) and in Experiment B via a carefully trained human driver. The second type of control strategy is a *proportional-integral (PI) controller with saturation*, which is a natural extension of the PI controller, a simple and widely used controller in industrial applications. The controller is only based on the knowledge of the autonomous vehicle speed over a time horizon. The control action is saturated at small gaps to avoid collisions, and long gaps to avoid slowing down of traffic. Compared to the average velocity controllers (Experiments A and B), the PI controller with saturation directly estimates the average velocity and thus needs no external input.

The results of each of the three experiments are presented and compared in Section 4. In each experiment, stop-and-go waves arise dynamically when all vehicles are under human control. Once one vehicle is activated to be autonomous (with the control algorithms described in Section 3), the traffic waves are dissipated. Compared to when waves are present, the Lagrangian control results in up to 40% less fuel consumption, and a throughput increase of up to 15%. Future perspectives for Lagrangian vehicular control are provided in Section 5.

2. Experimental methodology

We briefly describe the experimental setting in which stop-and-go waves are observed to develop and subsequently dampened via control of a single vehicle in the experiment (mimicking a low penetration rate on a long freeway stretch). The experiments follow the ring setting Sugiyama, et al. [2, 5]. A key advantage of the ring road experimental setup [2, 5] is that it removes other effects like boundary conditions, merging lanes, or intersections. To aid in interpretation of trajectory and fuel consumption datasets made available with this work, we concisely describe the experimental design and data collection methods in Section 2.1. The protocol for each experiment, including the specific instructions given to the drivers are presented in Section 2.2.

2.1. Experiment design

We consider a single-lane circular track of radius 41.4 meters to the center of the lane (260 meter circumference) with 21 to 22 vehicles depending on the experiment (see Figure 1). Small modifications to the Sugiyama et al. [2, 5] setup include a larger circumference of the circle and driving in counter-clockwise direction, to account for the larger average US vehicle size and the location of the steering wheel. An asphalt track is marked with small circular cones, and is otherwise nearly flat and uniform (no marking, light poles, parking barriers, or other potential obstacles). Short (3 cm) orange indicators are placed to mark the inside of the track ring.

A fleet of 22 passenger vehicles equipped with data acquisition hardware is used in the experiment. One of the 22 vehicles is the University of Arizona self-driving capable *Cognitive and Autonomous Test (CAT) Vehicle*, which can be transitioned between manual velocity control and autonomous velocity control. A trained human driver controls the steering wheel of the CAT Vehicle at all times during all experiments. We underscore that *only one* vehicle is ever



(a) Alignment of vehicles at start of Experiment A.



(b) Alignment of vehicles 93 seconds into Experiment A when wave is present in back right.



(c) Alignment of vehicles 327 seconds into Experiment A when the CAT Vehicle is actively dampening the wave.

Figure 1: Ring track with 21 vehicles during Experiment A.

controlled to dampen the traffic wave, either via automation of the vehicle velocity or through the trained driver. This setting in which a single vehicle is controlled on a ring road mimics a low penetration rate of automated vehicles on a long stretch of highway. All other vehicles are driven by University of Arizona employees that have completed safe driver training but received no other special driving training. Drivers are instructed to drive safely and are requested to attempt to close any widening gaps between their vehicle and the vehicle ahead (see Section 2.2 for the precise driver instructions).

Data from each experiment is collected via a video camera and OBD-II data loggers. The 360 degree camera is placed at the center of the track and used to record each experiment. The resulting video is processed via computer vision techniques to identify the center of each vehicle in each frame, which is then smoothed to generate the vehicle trajectories, following the approach described in [61]. Data from the in-vehicle devices are gathered through the OBD-II standard interfaces available on all US cars starting in 1996 [62]. The in-vehicle devices measure the instantaneous fuel consumption of each vehicle.

2.2. Experiment mechanics

Each experiment lasts between 7–10 minutes to limit driver fatigue and begins with all vehicles uniformly spaced around the track according to the position of their front-left tire. The CAT Vehicle begins each experiment in manual mode, and is switched into a control mode during the experiment. Traffic waves appear in all experiments, and the unsteady traffic is allowed to persist for at least 45 seconds before a controller is activated. For some controllers, a desired average velocity is communicated from an external observer.

Precisely, each experiment consists of the following phases: (i) setup; (ii) evacuation; (iii) initialize; (iv) drive; (v) stop; and (vi) conclusion, summarized below.

- i. Setup: Vehicles are distributed equally according to the spacing of their front-left tire. Drivers are individually instructed to turn on their in-vehicle data recorders. Additional driver instructions (if any) are delivered to individual drivers through the window.
- ii. Evacuation: The central camera is switched on. All research team personnel evacuate the track.
- iii. Initialize: An air horn sounds to instruct all drivers to switch gears from Park to Drive, without moving.
- iv. Drive: An air horn sounds, to instruct all drivers to begin driving.
- v. Stop: An air horn sounds, instructing drivers to come to a safe stop and switch gears into Park.
- vi. Conclusion: Experiment personnel enter the track after all vehicles have stopped. Drivers are individually instructed to turn off their in-vehicle recorders. The central camera is switched off.

The following instructions are provided to each driver prior to the start of the experiments. “Drive as if you were in rush hour traffic. Follow the vehicle ahead without falling behind. Do not pass the car ahead. Do not hit the car ahead. Drive safely at all times. Do not tailgate. But put an emphasis on catching up to the vehicle ahead, if a gap starts opening up.” The purpose of these instructions is to explicitly prevent the human drivers from intentionally smoothing out traffic waves themselves. This is important, so that the wave-smoothing effect solely caused by the CAT Vehicle can be studied.

In the event of an unsafe scenario, the drivers are instructed to steer out of the circle, at which the experiment coordinator will sound the air horn and the experiment will stop. All drivers are instructed that the CAT Vehicle would be switching back and forth between autonomous and manual mode, and that they should focus on their driving rather than attempting to guess what mode the vehicle is in at any given time. The vehicle directly following the CAT Vehicle is told to drive as if the CAT Vehicle were in the same lane, and that the CAT Vehicle will be driving at a larger radius (1/2 vehicle width) to facilitate evasive emergency maneuvers.

3. Description of controllers of the autonomous vehicle

This section presents the velocity controllers implemented on the CAT Vehicle to dampen traffic waves on the ring track. The controllers are generally motivated by the fact that mathematical models of vehicular traffic can be stabilized via the control of a small number of vehicles, see for example [49, 50, 51, 52, 63]. One important notion that can help stabilize the overall flow is to have a subset of vehicles drive with a smooth driving profile relative to the traffic conditions, which is the basis of the works [55, 56, 57].

One possible way to create a smooth driving profile is to follow the average speed of the vehicles ahead, which drive faster than the average speed before a stop-and-go wave, and slower than the average speed during the wave. By simply driving at the average speed, a vehicle covers the same distance in the same amount of time, but with less acceleration and braking. Additional logic is necessary to prevent extremely small gap (unsafe) situations from appearing, or large gaps that may induce lane changing on multi-lane roadways. With these ideas in mind, we propose two possible control laws and show experimentally that they are able to stabilize the flow on the ring.

The general structure of the controllers is as follows. The CAT Vehicle continuously tracks its velocity v^{AV} , and measures (at a sampling rate of 30Hz) the *gap* Δx , defined as the distance from its front bumper to the rear bumper of its lead vehicle ahead. This signal, suitably smoothed, is used to calculate the *velocity difference* $\Delta v = \frac{d}{dt} \Delta x$ between the lead vehicle and the AV. The lead vehicle’s (i.e., the car ahead of the AV) velocity is estimated on-board the CAT Vehicle as $v^{\text{lead}} = v^{AV} + \Delta v$. Moreover, a *desired velocity* U is defined (obtained in various ways, see below), which, when chosen correctly, can stabilize the traffic flow. From the desired velocity, the gap, and the velocities of the CAT Vehicle and lead vehicle, a *commanded velocity* v^{cmd} is determined. The commanded velocity is then passed to a low-level controller on the CAT Vehicle that translates it into an actuation of the accelerator or brake. Note that all of these quantities are functions of time; but the time argument is frequently omitted for notational efficiency. Below, we describe: strategies to define a desired velocity, the *FollowerStopper* controller and a PI controller with saturation, the low level controls, and a control law implemented by a trained human driver.

3.1. The FollowerStopper controller

The premise of this controller is to command exactly the desired velocity U whenever safe (i.e., as in a standard cruise controller), but to command a suitable lower velocity $v^{\text{cmd}} < U$ whenever safety requires, possibly based on the lead vehicle’s velocity.

Using the gap Δx and the velocity difference $\Delta v = \frac{d}{dt} \Delta x = v^{\text{lead}} - v^{AV}$, the Δx - Δv phase space is divided into regions (see also Figure 2):

- i. a safe region, where $v^{\text{cmd}} = U$;
- ii. a stopping region, where a zero velocity is commanded; and an
- iii. an adaptation region (two parts), where some average of desired and lead vehicle velocity is commanded.

The boundaries between the regions are parabolas in the Δx - Δv phase space (trajectories that the AV/lead vehicle pair would traverse when decelerating at constant rates), defined as

$$\Delta x_k = \Delta x_k^0 + \frac{1}{2d_k} (\Delta v_-)^2, \quad \text{for } k = 1, 2, 3. \quad (1)$$




Figure 2: Regions defined in FollowerStopper controller.

Here Δx_k^0 is a parameter that defines the intercept in the Δx - Δv phase space, and d_k controls the curvature and is interpreted as a deceleration rate. Moreover, $\Delta v_- = \min(\Delta v, 0)$ is the negative arm of velocity difference, i.e., the case of the CAT Vehicle falling behind is treated just like the case $v^{AV} = v^{\text{lead}}$.

Using the region boundaries defined in (1), the commanded velocity is

$$v^{\text{cmd}} = \begin{cases} 0 & \text{if } \Delta x \leq \Delta x_1 \\ v \frac{\Delta x - \Delta x_1}{\Delta x_2 - \Delta x_1} & \text{if } \Delta x_1 < \Delta x \leq \Delta x_2 \\ v + (U - v) \frac{\Delta x - \Delta x_2}{\Delta x_3 - \Delta x_2} & \text{if } \Delta x_2 < \Delta x \leq \Delta x_3 \\ U & \text{if } \Delta x_3 < \Delta x \end{cases} \quad (2)$$

In (2), the velocity $v = \min(\max(v^{\text{lead}}, 0), U)$ is the lead vehicle velocity (if positive) or the desired velocity, whichever is smaller. In the adaptation regions ($\Delta x_1 < \Delta x \leq \Delta x_3$), the commanded velocity transitions continuously from stopping ($v^{\text{cmd}} = 0$, for short gaps) to safe driving ($v^{\text{cmd}} = U$, for large gaps), via a transition involving the lead vehicle's velocity.

As implemented, in (1) and (2) we set $\Delta x_1^0 = 4.5$ m, $\Delta x_2^0 = 5.25$ m, and $\Delta x_3^0 = 6.0$ m, and the deceleration rates are $d_1 = 1.5 \frac{\text{m}}{\text{s}^2}$, $d_2 = 1.0 \frac{\text{m}}{\text{s}^2}$, and $d_3 = 0.5 \frac{\text{m}}{\text{s}^2}$. Note that the Δx_k boundaries of the regions depend strongly on the velocity difference between the CAT Vehicle and the lead vehicle. For instance, if the CAT Vehicle is catching up rapidly at $\Delta v = -3 \frac{\text{m}}{\text{s}}$, then $\Delta x_1 = 7.5$ m, $\Delta x_2 = 9.75$ m, and $\Delta x_3 = 15$ m.

3.2. The PI with saturation controller

The idea behind this controller is that the CAT Vehicle may estimate the average speed of the vehicles in front, and then drive according to the average speed. When stop-and-go waves are present, it allows a gap to open up in front of the CAT Vehicle when the lead vehicle accelerates, which is then closed when the lead vehicle decelerates. An estimate of the average speed required by the controller is obtained by measuring the CAT Vehicle speed over a large enough time horizon.

The controller determines a command velocity v^{cmd} following a standard proportional integral control logic [64], where the deviation from the average speed is treated as the error signal in the PI controller. This simple idea needs to be paired with saturation: for small gaps the CAT Vehicle should follow the lead vehicle speed to avoid dangerous situations, while for large gaps, the CAT Vehicle should catch up to the lead vehicle.

More precisely, this controller estimates the desired velocity, U , as a temporal average of the CAT Vehicle's own velocity over an interval. Letting $v_1^{\text{AV}}, \dots, v_m^{\text{AV}}$ denote the CAT Vehicle velocities over the last m measurements, the desired velocity is computed as the temporal average $U = \frac{1}{m} \sum_{j=1}^m v_j^{\text{AV}}$. In practice, we choose m corresponding to a 38 second interval, which is approximately the time required to travel one lap around the ring.

The desired average velocity is then translated into a target velocity depending on the current gap between the

CAT Vehicle and lead vehicle:

$$v^{\text{target}} = U + v^{\text{catch}} \times \min(\max(\frac{\Delta x - g_l}{g_u - g_l}, 0), 1), \quad (3)$$

which is up to v^{catch} above U . This allows the CAT Vehicle to drive faster than the average velocity and catch up to the lead vehicle, should it face a gap above the lower threshold g_l , while at lower gaps the target velocity reduces to the average U . At gaps above the upper gap limit g_l , CAT Vehicle vehicle should close the gap by traveling v^{catch} above U .

The commanded velocity sent to the low level CAT Vehicle controller is updated via the rule

$$v_{j+1}^{\text{cmd}} = \beta_j(\alpha_j v_j^{\text{target}} + (1 - \alpha_j)v_j^{\text{lead}}) + (1 - \beta_j)v_j^{\text{cmd}}, \quad (4)$$

where the subscript j denotes the time step. This rule (4) chooses the new commanded velocity as a weighted average of the prior commanded velocity, the target velocity, and the lead vehicle's velocity. The weights α_j and β_j depend on the gap as follows: $\beta = 1 - \frac{1}{2}\alpha$, and

$$\alpha = \min(\max(\frac{\Delta x - \Delta x^s}{\gamma}, 0), 1). \quad (5)$$

In (5), the distance Δx^s is a safety distance. We have $\alpha = 0$ if $\Delta x \leq \Delta x^s$ and $\alpha = 1$ if $\Delta x \geq \Delta x^s + \gamma$, meaning that for relatively short gaps, only the lead vehicle's velocity matters, while for relatively large gaps, only the target velocity is averaged with the commanded velocity. The parameter γ controls the the rate at which α transitions from 0 to 1, and is set to $\gamma = 2$ m in the current implementation. This means that when the gap is short, the CAT Vehicle has the same speed of the lead vehicle, while when the gap is larger the CAT Vehicle speed tends towards the target vehicle, which allows the CAT Vehicle to reduce the gap with the lead vehicle. The parameter β determines how rapidly the controller adjusts to new situations (with more rapid adjustments occurring in more safety-critical situations). At its core, this is a PI controller, but with a saturation at small gaps (for safety purposes), and a saturation at large gaps (so that the CAT Vehicle closes gaps).

The model parameters for both controllers were determined via testing in a simulation environment (with a data-fitted human-driver model), as well as via car-following field tests (before the actual experiments). As a result we set the lower gap limit $g_l = 7$ m, the upper gap limit as $g_u = 30$ m and $v^{\text{catch}} = 1$ m/s. The safety distance is implemented as $\Delta x^s = \max(2 \text{ s} \times \Delta v, 4 \text{ m})$. The term $2 \text{ s} \times \Delta v$ represents the recommended following distance according to the 2 seconds rule, with a lower bound of 4 m.

3.3. Low level vehicle controls

The commanded velocity produced by the two controllers described above is translated to the actual vehicle controls (i.e., gas and brake signals) via a multi-mode controller. The use of a multi-mode controller permits different gains to be used for acceleration or for braking, to enable faster braking when needed while avoiding chattering at steady-state velocities.

Each mode is a PID controller, with gains determined through system identification of the CAT Vehicle at constant velocities representative of those recorded in the experiment, using a similar structure as in [65]. The CAT Vehicle plant is simplified as a first-order model based on constant accelerator inputs. The controller's design is thus:

$$a_{j+1} = \begin{cases} h_1(v_j, v_j^{\text{cmd}}) & \text{if } v_j^{\text{cmd}} - v_j > -0.25 \frac{\text{m}}{\text{s}} \\ h_2(v_j, v_j^{\text{cmd}}) & \text{if } v_j^{\text{cmd}} - v_j \leq -0.25 \frac{\text{m}}{\text{s}} \\ 0 & \text{otherwise,} \end{cases}$$

in which $a_{j+1} \in [-100, 100]$ represents the next commanded "acceleration" value, where 100 is the maximum depression of the accelerator, and -100 is the maximum depression of the brake. Moreover, v_j is the current speed of the CAT Vehicle, and v_j^{cmd} is the desired speed. When $a < 0$ the brake is depressed, and when $a > 0$ the accelerator is depressed. Controller h_1 is designed to accelerate to the desired reference speed and maintain that desired speed primarily through control of the vehicle's accelerator, and controller h_2 is designed to effect more rapid speed reduction via the brake. Thus when the desired speed is less than $0.25 \frac{\text{m}}{\text{s}}$ of the current speed, the brake is used,

and otherwise the accelerator is used to control speed (as in normal driving when releasing the accelerator reduces speed). These controllers are provided sampled data at 20 Hz and are permitted to send new updates to the CAT Vehicle at 20Hz.

The performance characteristics of each controller are provided for a change of input as a step function of 1 m/s (-1 m/s for braking). The controller h_1 has a rise time of approximately 1.6 s with an overshoot of 5% and a settling time of approximately 5.52 s. The controller h_2 has a rise time of approximately 0.8 s, with an 11% overshoot and settling time of approximately 1.94 s.

Given the dynamics of the CAT Vehicle, a tradeoff must be performed on the comfort of the ride and the physical dynamics for a change in reference speed. The switching nature of the controller provides robustness to noise in sampled speed, since the accelerator is primarily used to control speed at steady state. Finally, the PID controllers are reset at 0 velocity, and standard approaches for windup avoidance are used to prevent unsafe acceleration [66].

3.4. Human driver controller

The driver who implements human control of the CAT Vehicle (coauthor M. Bunting) is instructed to attempt to maintain a desired velocity, but to slow down to avoid collision with the vehicle ahead. This is similar to the control law used in Experiment A, with the notable exception that the desired velocity is given in miles/hour (the primary readout of the speedometer in the CAT Vehicle). The driver received training from the University of Arizona on safe driving of high-occupancy vehicles, and had extensive practice to drive in this way, before the actual experiment is performed.

4. Experimental results: Dampening traffic waves with a single vehicle

The experimental results are presented in this section. To be able to effectively compare the results of the experiments, it is important to define metrics, which are consistent across the experiments. To this end, we present the metrics used to describe the traffic flow in Section 4.1. With the metrics fully defined, we present the results of the three experiments conducted in Section 4.2.

4.1. Definition and calculation of metrics

At each time step, we have the following data. From the image processing, we have the position x , the velocity v , and the acceleration a of each vehicle. From the OBD-II sensors, we have the instantaneous fuel consumption c of each vehicle. Let f_j^i denote the sample of a quantity f , corresponding to vehicle i , at time t_j . A temporal average of a quantity over an interval $t \in [t_{\text{start}}, t_{\text{end}}]$ is calculated as $\bar{f}^i = \frac{1}{m} \sum_{j=1}^m f_j^i$, where f_1^i, \dots, f_m^i are the samples of vehicle i in that time interval. Likewise, an average over all n vehicles at an instant t_j is given by $\bar{f}_j = \frac{1}{n} \sum_{i=1}^n f_j^i$, where j denotes the time step. Finally, spatio-temporal averages are given by $\bar{f} = \frac{1}{mn} \sum_{i=1}^n \sum_{j=1}^m f_j^i$. The precise quantities of interest are defined next.

At each time instant we compute the spatially-averaged instantaneous velocity is computed by summing the velocity of each vehicle $i = 1, \dots, n$ at a given time indexed by j , and dividing by the number of vehicles as:

$$\bar{v}_j = \frac{1}{n} \sum_{i=1}^n v_j^i.$$

Over a given time interval with m velocity samples per vehicle, we compute the average (over all vehicles and over the time interval) as:

$$\bar{v} = \frac{1}{mn} \sum_{j=1}^m \sum_{i=1}^n v_j^i.$$

Similarly, we compute the velocity standard deviation of all vehicles and over the interval as:

$$\sigma = \left(\frac{1}{mn - 1} \sum_{j=1}^m \sum_{i=1}^n (v_j^i - \bar{v})^2 \right)^{\frac{1}{2}}.$$

Interval	Time (s)	Velocity st. dev. (m/s)	Fuel consumption ($\ell/100\text{km}$)	Braking (events/vehicle/km)	Throughput (veh/hr)
Exp. start	0	1.87	18.8	1.66	1809
Waves start	79	3.31	24.6	8.58	1827
Autonomy 6.50m/s	126	1.69	18.0	3.45	1780
Autonomy 7.00m/s	222	0.67	15.0	0.21	1915
Autonomy 7.50m/s	292	0.64	14.1	0.12	2085
Autonomy 8.00m/s	347	1.56	17.7	2.50	1952
Autonomy 7.50m/s	415	1.14	16.7	0.31	1938
Disable Autonomy	463	1.44	17.4	2.95	2133
Exp. end	567	-	-	-	-

Table 1: Summary metrics over all vehicles by interval with corresponding start time, for Experiment A.

Given the fuel consumption of vehicle i at time t_j , the average (over all vehicles and over a time interval) consumption is computed as:

$$\bar{c} = \frac{1}{mn} \sum_{j=1}^m \sum_{i=1}^n c_j^i.$$

The throughput of traffic is computed as the product of the average velocity and the density (obtained from the number of vehicles and the length of the track $L = 260$ m), and is given as as:

$$q = \frac{n}{L} \bar{v}.$$

We also quantify braking events. Given a threshold deceleration τ , a brake event (deceleration peak) is defined as a contiguous region in time when $-a_j^i > \tau$ that has the additional property that the signal $-a_j^i$ must drop by more than τ on either side of a peak. This deceleration peak count is encoded in the function ρ_τ , which takes the signal a_1^i, \dots, a_m^i as an input, and outputs the number of peaks in that interval for vehicle i . The final calculation of κ , the rate of braking events, normalizes by the number of vehicles and total distance traveled (in kilometers). The threshold τ is chosen as the average standard deviation of deceleration, taken over all vehicles in the uncontrolled interval when waves are active. This is computed as:

$$\kappa = \frac{1}{n} \sum_{i=1}^n \frac{1}{x_m^i - x_1^i} \rho_\tau(a_1^i, \dots, a_m^i),$$

where τ is always calculated over the interval when waves are active. The interval is determined based on the standard deviation of the velocity as shown in Appendix 5.2.

4.2. Experimental results

Experiment A contains 21 vehicles, including the CAT Vehicle. The CAT Vehicle is initially under human control, and the first traffic wave is observed 79 seconds into the experiment. The *FollowerStopper* wave-dampening controller is activated 126 seconds into the experiment, and set with a desired velocity of $U = 6.50$ m/s. Over the next several minutes, the desired velocity is varied step by step to test the dependence of the traffic conditions on the set point. It is changed to 7.00 m/s (222 seconds into the experiment), 7.50 m/s (292 seconds into the experiment), and finally 8.00 m/s at 347 seconds into the experiment. At 415 seconds, the desired velocity is reduced to 7.50 m/s, where it remains for 48 seconds. At 463 seconds into the experiment, the human driver resumes control of the CAT Vehicle speed. The experiment is ended at 567 seconds.

Experiment B also involves 21 vehicles and follows a similar design as Experiment A. The main difference is that in Experiment B, after the wave initially appears, a trained human driver implements the control strategy described

in Experiment A but without the aid of automation. The human-executed control strategy is to maintain a desired velocity, calculated by an external observer as the average velocity of the previous lap, without colliding with the vehicle ahead. The CAT Vehicle is always under human control, but the driver switches from initially following the instructions given to all human drivers to mimicking the control strategy in Experiment A after a wave appears. A traffic wave is first observed 55 seconds from the start of Experiment B, and the active control to dampen the wave begins after 112 seconds at a desired velocity of 6.25 m/s (14 mph in the units displayed in the CAT Vehicle dashboard) with the command to “drive with an average speed of 14 miles per hour, unless safety requires slower speeds.” After 202 seconds, the CAT Vehicle operator is instructed to increase the desired speed “to 16 miles per hour,” (7.15 m/s) which is maintained for 98 seconds before reverting to typical human driving behavior. The experiment is ended after 409 seconds.

Experiment C is conducted with 22 vehicles. Note that one vehicle was added for this experiment, to demonstrate that instabilities and wave damping are not specific to having exactly 21 vehicles on the ring. At 161 seconds into the experiment, a traffic wave appears. At 218 seconds, the *PI controller with saturation* wave damping controller is activated, and remains active until the end of the experiment at 413 seconds. Because the controller directly determines the desired velocity as part of the control algorithm, there are no external parameters that are changed during the experiment.

After conclusion of the experiments, the data gathered through a 360° camera placed at the center of track, in-vehicle devices, and CAT Vehicle control computers are analyzed to characterize the performance of each experiment with respect to dampening traffic waves.


Displacement data describe each vehicle’s distance traveled from their initial position on the ring at the beginning of the experiment. These data are extracted from the central camera using standard computer vision techniques to track vehicle positions along the ring. Velocity data are derived through discrete differentiation of displacement data (likewise, acceleration from velocity data) [61], which are validated against the CAT Vehicle control computers and the in-vehicle devices. The full data set and code used to generate the plots can be found in [67].

Notable events in each experiment define time intervals over which to evaluate the state of traffic. In each analysis we quantify the traffic state according to (a) the standard deviation of traffic velocity; (b) the fuel consumption [68]; (c) the rate of excessive braking events; and (d) the traffic throughput. A detailed description of how these calculations are performed, as well as the specific definitions of common metrics of fuel consumption, velocity (average and standard deviation), and excessive braking are available in Section 4.1.

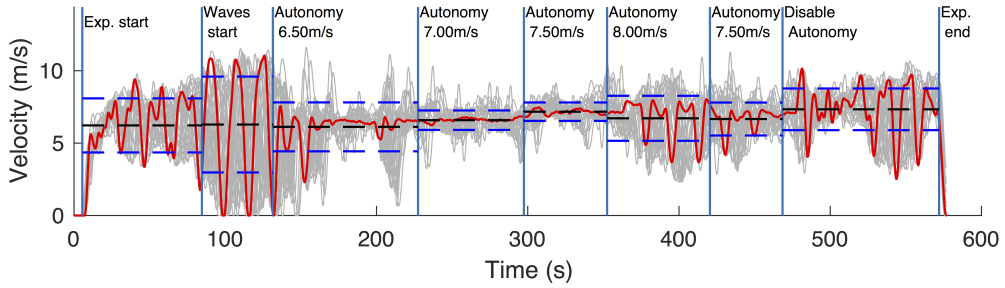
In Experiment A, the traffic state is quantified over eight time intervals throughout the experiment. In each interval, the fuel consumption, velocity standard deviation, excessive braking, and throughput are reported in Table 1. The position of each vehicle’s center over time is shown in Figure 3a, with the CAT Vehicle is shown in red and all other vehicles are shown in gray. In Figure 3b, the velocity profile of the CAT Vehicle is shown in red, and all other vehicle velocities are plotted in gray. For each interval, the black dashed line denotes the average velocity of traffic over that interval, and the blue dashed lines denote the average velocity plus/minus one standard deviation.

When the CAT Vehicle is initially under human control, a wave begins to appear. After the controller is activated the wave is noticeably affected, but not fully removed, through the interval when the CAT Vehicle control is activated at a desired velocity of 6.50 m/s. The wave-dampening effect of the controller is observed in the velocity profiles (Figure 3b), which exhibit a lower magnitude of oscillation from the mean after control begins. When the CAT Vehicle desired speed set in the *FollowerStopper* is increased to 7.00 m/s 222 seconds into the experiment, further wave dampening is observed, and at a desired speed of 7.50 m/s, the best performance of the controller is achieved. Compared to the initial period where a wave was present under human control, the velocity standard deviation is reduced by 80.8%, the fuel consumption is reduced by 42.5%, and the excessive braking events are reduced from 8.58 events per vehicle per kilometer to 0.12 events per vehicle per kilometer. Because the average velocity of traffic on the ring is also increased, the throughput on the roadway increases by 14.1%.

At 7.50 m/s, the CAT Vehicle’s speed matches the average traffic speed almost precisely, and consequently it does not need to slow down. However, at 347 seconds, the CAT Vehicle desired velocity is increased further to 8.00 m/s and the CAT Vehicle is, on average, faster than the flow of traffic—which inevitably induces a wave again. The reappearance of a wave has the effect of increasing fuel consumption relative to the slower desired velocities, but still represents a benefit of the control compared to the uncontrolled traffic. At 415 seconds the CAT Vehicle’s desired velocity is reduced to 7.50 m/s, and the wave is once again dampened. The *FollowerStopper* controller is deactivated after 463 seconds, and the traffic wave reappears.



(a) Trajectories of all vehicles in Experiment A, CAT Vehicle shown in red.



(b) Velocity profiles of all vehicles (gray) and the CAT Vehicle (red) in Experiment A. Horizontal blue dashed lines are one standard deviation above and below the mean speed of traffic in the interval.

Figure 3: Trajectories and standard deviation in velocity for Experiment A.


The control of the CAT Vehicle has the impact of reducing the total fuel consumption of the traffic, as shown in Table 1. The lowest fuel consumption of $14.13 \ell/100\text{km}$ is observed when the *FollowerStopper* is operated with the set point of 7.50 m/s . This is also the lowest fuel consumption amongst all experiments conducted. A video of Experiment A is provided in Movie S1 in the supplementary materials.

In Experiment B, the CAT Vehicle operator initially drives according to the same instructions as the other vehicle operators and a traffic wave appears at 55 seconds. At 112 seconds into the experiment, the CAT Vehicle operator begins to drive at a desired velocity of 6.26 m/s without colliding with the vehicle in front. Later the desired velocity is increased before returning to the gap closing instructions followed by all other human drivers. The trajectories are shown in Figure 4 and the traffic state is quantified in each interval in Table 2.


The dampening effect of the human-implemented controller is quantified by the standard deviation of the velocities, which is reduced by 49.5% when the 6.26 m/s control is active compared to when it is not. Similarly, excessive

Interval	Time (s)	Velocity st. dev. (m/s)	Fuel consumption ($\ell/100\text{km}$)	Braking (events/vehicle/km)	Throughput (veh/hr)
Exp. start	0	2.11	20.6	3.88	1665
Waves start	55	2.36	21.8	9.50	1828
Control 6.26m/s	112	1.58	17.9	4.22	1822
Control 7.15m/s	202	1.19	17.0	2.27	2008
Disable control	300	2.25	21.6	9.43	1908
Exp. end	409	-	-	-	-

Table 2: Summary metrics over all vehicles by interval with corresponding start time, for Experiment B.



(a) Trajectories of all vehicles in Experiment B, CAT Vehicle shown in red.



(b) Velocity profiles of all vehicles (gray) and the CAT Vehicle (red) in Experiment B. Horizontal blue dashed lines are one standard deviation above and below the mean speed of traffic in the interval.

Figure 4: Trajectories and standard deviation in velocity for Experiment B.

Exp.	Velocity st. dev. (m/s)			Fuel consumption ($\ell/100\text{km}$)			Braking (events/veh/km)			Throughput (veh/hr)		
	WS	CA	%	WS	CA	%	WS	CA	%	WS	CA	%
A	3.31	0.64	-80.8	24.6	14.1	-42.5	8.58	0.12	-98.6	1827	2085	+14.1
B	2.36	1.19	-49.5	21.8	17.0	-22.1	9.50	2.27	-76.2	1828	2008	+9.8
C	3.85	1.74	-54.7	29.0	20.9	-28.1	9.66	2.47	-74.4	1755	1711	-2.5

Table 3: Summary metrics of the flow in each experiment under the first interval when the traffic wave starts (WS) without wave dampening control, and under the best interval when control is active (CA). The percent change from WS to CA in each experiment is also reported.

braking is reduced by 76.2% when control is applied.

The desired speed given to the CAT Vehicle driver also influences the reduction in the velocity variability. When the desired speed of the CAT Vehicle is increased to 16 mph (7.15 m/s), the standard deviation reaches the minimum values for the experiment. The throughput is also higher during this period compared to when the traffic is uncontrolled (2008 veh/hr vs. 1828 veh/hr). Once the CAT Vehicle returns to a typical gap closing behavior (i.e., no longer under human control to maintain a desired speed), the traffic wave reappears, and the velocity standard deviation increases.

When the CAT Vehicle is under wave dampening control by a human, the control reduces the overall fuel consumption when compared to the uncontrolled case (Table 2), where average fuel consumption ($\ell/100\text{km}$) is recorded for each period of the experiment. The result is a decrease in fuel consumption when the CAT Vehicle begins to dampen the traffic wave (a decrease from 21.8 $\ell/100\text{km}$ to 17.9 $\ell/100\text{km}$, (17.9%)). A further decrease in fuel consumption (to 17.0 $\ell/100\text{km}$, a decrease of 22.0% compared to when a wave is present initially) is observed when the desired velocity is increased from 14 mph (6.25 m/s) to 16 mph (7.15 m/s). Finally, when the CAT Vehicle returns to human-driven behavior and stops actively dampening the traffic wave, fuel consumption increases again and returns to a level similar to the pre-control fuel consumption (21.8 $\ell/100\text{km}$) before control starts compared to (21.6 $\ell/100\text{km}$) after control ends and the traffic wave reappears).

Finally, in Experiment C, an additional vehicle is added to the track, bringing the total number of vehicles to 22.

Interval	Time (s)	Velocity st. dev. (m/s)	Fuel consumption ($\ell/100\text{km}$)	Braking (events/vehicle/km)	Throughput (veh/hr)
Exp. start	0	1.62	17.7	1.11	2160
Waves start	161	3.85	29.0	9.66	1755
Autonomy	218	1.74	20.9	2.47	1711
Exp. end	413	-	-	-	-

Table 4: Summary metrics over all vehicles by interval with corresponding start time, for Experiment C.

As in other experiments, the CAT Vehicle begins the experiment in human control. The traffic is relatively smooth until the first strong wave occurs at 161 seconds. Comparing the traffic conditions under no control before the strong wave appears to when it is observed, the traffic wave results in a 37.7% increase in the average velocity standard deviation, a 63.8% increase in fuel consumption, and a 18.8% reduction in throughput.

At 218 seconds, the *PI controller with saturation* wave damping traffic controller is activated, and the wave is substantially reduced. Compared to the interval when the wave is present, the controller results in a reduction of the speed variability (54.7% reduction in standard deviation see Table 4). It also reduces the fuel consumption by 27.9%, and the rate of excessive braking is reduced by 74.4%. The throughput is also slightly reduced by 2.5%. Note that this controller only uses information directly measured by the CAT Vehicle itself, and it does not require any external information. One consequence of this more local nature is that the wave dampening is not as perfect as in Experiment A. In particular, the controller slightly reduces the average velocity. Nevertheless, the velocity standard deviation, excessive braking, and fuel consumption are substantially reduced compared to when uncontrolled traffic waves are present.

By examining the trajectories (Figure 5), it is apparent that the control law is able to eliminate the initially present wave, but another wave is generated during the control period. The second wave is also damped by the control law, but its presence for a period of the active control accounts for the difference relative to the period where no waves occurred in the uncontrolled period at the start of the experiment.

To conclude, these experiments demonstrate that traffic flow control via low penetration rate automated vehicles is in fact possible. Moreover, the data collected quantify the benefits of conducting control via the AV. Specifically, under proper control, (a) the velocity standard deviation reduces noticeably; (b) the fuel consumption is reduced by a significant margin; (c) braking events are substantially reduced; and (d) in some experiments, even the average velocity (and thus the throughput) is increased (Table 3).

Velocity standard deviation: The velocity standard deviation is reduced in all experiments, ranging between 49.5% for Experiment B, 54.7% for Experiment C, and a high of 80.8% for Experiment A.

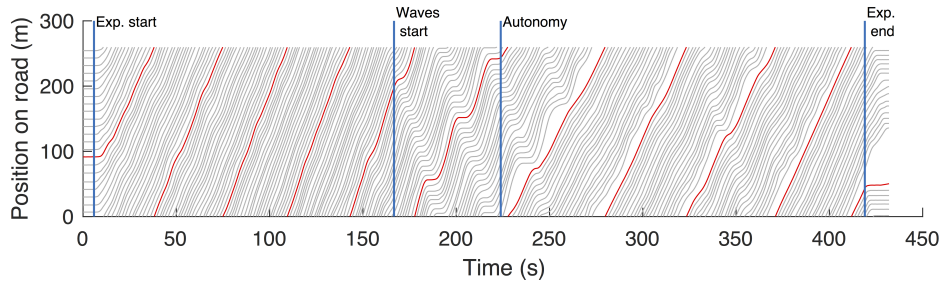
Fuel consumption: The fuel consumption in all experiments is reduced from when waves are present and under human control compared to when the autonomous car is active. The improvements include a reduction of 42.5%, 22.1%, and 28.1% in Experiments A, B, and C, respectively.

Excessive braking: The number of excessive braking events is also substantially reduced, from 8.58 to 9.66 excessive braking events/veh/km down to 2.47 events/veh/km in the worst performing controller and nearly complete elimination (0.12 events/veh/km) in the best performing controller.

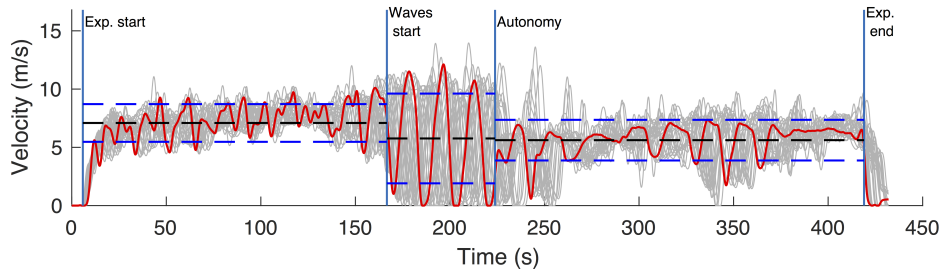
Throughput: Changes in throughput for each experiment are +14.1% for Experiment A, and +9.8% for Experiment B, with a -2.5% change in Experiment C.

5. Conclusions

AVs can revolutionize the control of traffic flow. They offer the potential to shift from localized control measures, like ramp metering, and centralized ones, as variable speed limit gantries, to Lagrangian actuators immersed in the traffic stream. Strikingly, it is not necessary for all vehicles to be automated in order to benefit from mobile actuation. A single autonomous vehicle can control the flow of at least 20 human-controlled vehicles around it, with substantial reductions in velocity standard deviation, excessive braking, and fuel consumption.



(a) Trajectories of all vehicles in Experiment C, CAT Vehicle shown in red.



(b) Velocity profiles of all vehicles (gray) and the CAT Vehicle (red) in Experiment C. Horizontal blue dashed lines are one standard deviation above and below the mean speed of traffic in the interval.

Figure 5: Trajectories and standard deviation in velocity for Experiment C.

Moreover, this study demonstrates that these benefits can be achieved via structurally very simple control strategies, based only on the AV's velocity, its spatial gap between the vehicle immediately in front, and some estimate of the average velocity of traffic flow. The *PI with saturation* controller (Experiment C) is a fully automatic control, while the *FollowerStopper* (Experiment A) and the human-implemented control (Experiment B) have an external input (dependent on observed traffic conditions). This simple structure implies that a noticeable impact on congested traffic flow can in principle be achieved by means of adaptive cruise control systems that are already in place in certain new vehicles, and the use of intelligent infrastructure and/or connected vehicles to provide the required external inputs.

Most contemporary traffic control strategies (implemented in practice, and/or proposed in the literature) are based on centralized interventions, such as ramp-metering, variable speed limits, and traffic light controls. For example, a successful variable speed limit control may yield a 5% increase in capacity [28]. However, those traditional control approaches will always have limited effect on the traffic dynamics that emerge between the fixed control points. In contrast, the control of traffic flow via a sparse set of Lagrangian actuators (AVs or trained human drivers) enables new opportunities for control, with a direct positive effect on the dynamics of traffic flow, and without the need of a dedicated actuation infrastructure.

The presented ring experiments represent a stretch of single-lane roadway. However, the theory extends also to multi-lane freeways, on which lane changing can serve as an additional trigger of stop-and-go waves. The lane changes can also open up gaps in the vacated lane, which can serve to dampen waves in that lane. The central challenge in the multi-lane setting is to have controllers that dynamically dampen waves, but without leaving too large gaps, because large gaps may trigger additional lane changing, which may reduce the effectiveness of the strategy. The results of Experiment A demonstrate that the controller does in fact not leave a large gap once the waves have been damped. To fully quantify the benefits of Lagrangian actuators on urban freeways, future multi-lane experiments are needed.

The control of complex multi-agent systems has impact beyond vehicular traffic flow, including coordinated robots [69], social networks [70], animal swarming [71, 72], and many other applications. However, in contrast to many applications in robotics or fleet control, the human agents play a crucial role in traffic flow dynamics. Moreover, in contrast to other human-in-the-loop cyber-physical systems, the automated controller and the human agents are spatially separated, and they do not work cooperatively. Rather, the AVs counteract the humans' tendency to

Veh. Num.	Year	Make	Model	Length (m)	Consumption City ($\ell/100\text{km}$)	Consumption Hwy. ($\ell/100\text{km}$)
1	2013	Chevrolet	Silverado	5.22	15.67	10.68
2	2013	Dodge	Grand Caravan	5.15	13.83	9.47
3	2015	Chevrolet	Malibu	4.86	9.47	6.53
4	2012	Chevrolet	Malibu	4.87	10.69	7.13
5	2012	Dodge	Grand Caravan	5.15	13.83	9.47
6	2013	Dodge	Grand Caravan	5.15	13.83	9.41
7	2014	Chevrolet	Malibu	4.86	9.41	6.53
8	2016	Chevrolet	Malibu	4.92	8.71	6.37
9	2013	Chevrolet	Impala	5.09	13.07	7.84
10	2014	Chevrolet	Malibu	4.86	9.41	6.53
11	2016	Chevrolet	Malibu Limited	4.86	9.80	6.92
12	2015	Chevrolet	Suburban	5.69	14.71	10.22
13	2014	Chevrolet	Silverado	5.21	13.07	9.80
14	2014	Dodge	Grand Caravan	5.15	13.83	9.41
15	2012	Chevrolet	Malibu	4.87	10.70	7.13
16	2016	Dodge	Grand Caravan	5.15	13.83	9.41
17	2014	Chevrolet	Malibu	4.86	9.41	6.53
18	2012	Chevrolet	Malibu	4.87	10.70	7.13
19	2012	Dodge	Grand Caravan	5.15	13.83	9.41
20	2016	Chevrolet	Suburban	5.70	14.71	10.22
21	2009	Ford	Escape Hybrid	4.44	6.92	7.84
22	2012	Dodge	Grand Caravan	5.15	13.83	9.41

Table 5: Vehicle properties of all vehicles used in the experiments.

produce unstable traffic situations. The results shown here imply that this concept is not a far future but instead could be, in principle, implemented with already existing technology.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. CNS-1446715 (B.P.), CNS-1446690 (B.S.), CNS-1446435 (J.S.), and CNS-1446702 (D.W.). The authors thank the University of Arizona Motor Pool in providing the vehicle fleet. They offer additional special thanks for the services of N. Emptage in carrying out the experiment logistics.

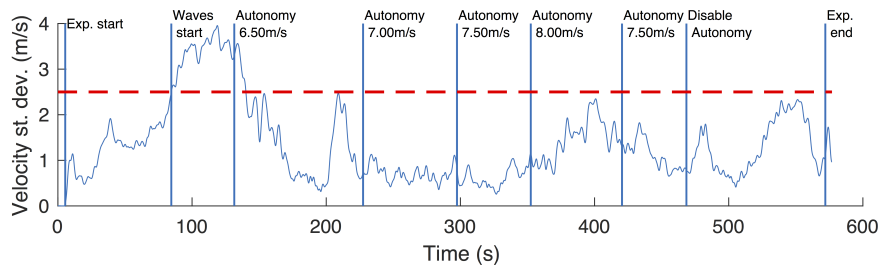
Appendix

5.1. Vehicle specifications


A full description of the year, make, model, length, and EPA rated fuel consumption of each vehicle used in the experiments is presented in Table 5. For Experiments A and B, vehicles 1 through 21 are used. In Experiment C, vehicle 22 is also used.

5.2. Identifying onset of oscillatory traffic behavior


We briefly describe the method used to identify the onset of oscillatory traffic behavior. The mean and standard deviation of the instantaneous vehicle velocities are computed at each timestep (i.e., for each timestep we compute the average of all vehicle speeds at that timestep, and then compute the sample standard deviation of all vehicle speeds at that timestep). When the standard deviation of the instantaneous vehicle velocities exceeds 2.5 m/s, the traffic



(a) Instantaneous standard deviation of vehicle velocities for Experiment A with dashed line indicating threshold for traffic waves in flow.



(b) Instantaneous standard deviation of vehicle velocities for Experiment B with dashed line indicating threshold for traffic waves in flow.



(c) Instantaneous standard deviation of vehicle velocities for Experiment C with dashed line indicating threshold for traffic waves in flow.

Figure 6: Instantaneous velocity standard deviation for all three experiments showing onset of traffic wave.

is considered to contain a traffic wave. This threshold is used to define the time at which waves first appear in the traffic experiments. Figures 6a, 6b, and 6c show a timeseries of the instantaneous velocity standard deviation for Experiments A, B, and respectively with the threshold plotted in a red dashed line.

References

- [1] J. Treiterer, J. Myers, The hysteresis phenomenon in traffic flow, *Transportation and traffic theory* 6 (1974) 13–38.
- [2] Y. Sugiyama, M. Fukui, M. Kikuchi, K. Hasebe, A. Nakayama, K. Nishinari, S. ichi Tadaki, S. Yukawa, Traffic jams without bottlenecks – experimental evidence for the physical mechanism of the formation of a jam, *New Journal of Physics* 10 (3) (2008) 033001.
- [3] M. R. Flynn, A. R. Kasimov, J.-C. Nave, R. R. Rosales, B. Seibold, Self-sustained nonlinear waves in traffic flow, *Phys. Rev. E* 79 (5) (2009) 056113. doi:10.1103/PhysRevE.79.056113.
- [4] B. S. Kerner, *The physics of traffic: empirical freeway pattern features, engineering applications, and theory*, Springer, 2012.
- [5] S. Tadaki, M. Kikuchi, M. Fukui, A. Nakayama, K. Nishinari, A. Shibata, Y. Sugiyama, T. Yosida, S. Yukawa, Phase transition in traffic jam experiment on a circuit, *New Journal of Physics* 15 (2013) 103034.
- [6] R. Jiang, M.-B. Hu, H. Zhang, Z.-Y. Gao, B. Jia, Q.-S. Wu, B. Wang, M. Yang, Traffic experiment reveals the nature of car-following, *PLoS one* 9 (4) (2014) e94351.
- [7] J. A. Laval, C. F. Daganzo, Lane-changing in traffic streams, *Transportation Research Part B: Methodological* 40 (3) (2006) 251–264.

- [8] J. A. Laval, Stochastic processes of moving bottlenecks: Approximate formulas for highway capacity, *Transportation Research Record: Journal of the Transportation Research Board* 1988 (2006) 86–91.
- [9] Z. Zheng, S. Ahn, D. Chen, J. Laval, Applications of wavelet transform for analysis of freeway traffic: Bottlenecks, transient traffic, and traffic oscillations, *Transportation Research Part B: Methodological* 45 (2) (2011) 372–384.
- [10] M. Bando, H. K., A. Nakayama, A. Shibata, Y. Sugiyama, Dynamical model of traffic congestion and numerical simulation, *Phys. Rev. E* 51 (2) (1995) 1035–1042.
- [11] K. Nagel, M. Schreckenberg, A cellular automaton model for freeway traffic, *J. Phys. I France* 2 (1992) 2221–2229.
- [12] M. Garavello, K. Han, B. Piccoli, *Models for Vehicular Traffic on Networks*, American Institute of Mathematical Sciences, 2016.
- [13] M. Brackstone, M. McDonald, Car-following: a historical review, *Transportation Research Part F: Traffic Psychology and Behaviour* 2 (4) (1999) 181–196.
- [14] D. Chowdhury, L. Santen, A. Schadschneider, Statistical physics of vehicular traffic and some related systems, *Physics Reports* 329 (4) (2000) 199–329.
- [15] D. Helbing, Traffic and related self-driven many-particle systems, *Rev. Mod. Phys.* 73 (2001) 1067–1141. doi:10.1103/RevModPhys.73.1067. URL <http://link.aps.org/doi/10.1103/RevModPhys.73.1067>
- [16] H. J. Payne, Models of freeway traffic and control, *Proc. Simulation Council* 1 (1971) 51–61.
- [17] G. B. Whitham, *Linear and nonlinear waves*, John Wiley and Sons, New York, 1974.
- [18] A. Aw, M. Rascle, Resurrection of second order models of traffic flow, *SIAM J. Appl. Math.* 60 (2000) 916–944. doi:10.1137/S0036139997332099.
- [19] H. M. Zhang, A non-equilibrium traffic model devoid of gas-like behavior, *Transp. Res. B* 36 (2002) 275–290.
- [20] J. M. Greenberg, Congestion redux, *SIAM J. Appl. Math.* 64 (2004) 1175–1185. doi:10.1137/S0036139903431737.
- [21] D. C. Gazis, C. H. Knapp, On-line estimation of traffic densities from time-series of flow and speed data, *Transportation Science* 5 (3) (1971) 283–301.
- [22] Y. Wang, M. Papageorgiou, Real-time freeway traffic state estimation based on extended Kalman filter: a general approach, *Transportation Research Part B: Methodological* 39 (2) (2005) 141–167.
- [23] S. Blandin, A. Couque, A. Bayen, D. B. Work, On sequential data assimilation for scalar macroscopic traffic flow models, *Physica D: Nonlinear Phenomena* 241 (17) (2012) 1421–1440.
- [24] J. C. Herrera, D. Work, R. Herring, X. Ban, Q. Jacobson, A. Bayen, Evaluation of traffic data obtained via GPS-enabled mobile phones: The mobile century field experiment, *Transportation Research Part C: Emerging Technologies* 18 (4) (2010) 568–583.
- [25] D. B. Work, S. Blandin, O.-P. Tossavainen, B. Piccoli, A. M. Bayen, A traffic model for velocity data assimilation, *Applied Mathematics Research eXpress* 2010 (1) (2010) 1–35.
- [26] C. De Fabritiis, R. Ragona, G. Valenti, Traffic estimation and prediction based on real time floating car data, in: *Proceedings of the IEEE Conference on Intelligent Transportation Systems*, 2008, pp. 197–203.
- [27] A. Hofleitner, R. Herring, P. Abbeel, A. Bayen, Learning the dynamics of arterial traffic from probe data using a dynamic bayesian network, *IEEE Transactions on Intelligent Transportation Systems* 13 (4) (2012) 1679–1693.
- [28] A. Nissan, H. N. Koutsopoulos, Evaluation of the impact of advisory variable speed limits on motorway capacity and level of service, *Procedia-Social and Behavioral Sciences* 16 (2011) 100–109.
- [29] A. Hegyi, B. De Schutter, J. Hellendoorn, Optimal coordination of variable speed limits to suppress shock waves, *IEEE Transactions on Intelligent Transportation Systems* 6 (1) (2005) 102–112.
- [30] S. Smulders, Control of freeway traffic flow by variable speed signs, *Transportation Research Part B: Methodological* 24 (2) (1990) 111–132.
- [31] A. Hegyi, S. Hoogendoorn, M. Schreuder, H. Stoelhorst, F. Viti, SPECIALIST: A dynamic speed limit control algorithm based on shock wave theory, in: *Proceedings of the IEEE Conference on Intelligent Transportation Systems*, IEEE, 2008, pp. 827–832.
- [32] A. Popov, A. Hegyi, R. Babuška, H. Werner, Distributed controller design approach to dynamic speed limit control against shockwaves on freeways, *Transportation Research Record: Journal of the Transportation Research Board* 2086 (2008) 93–99.
- [33] M. Papageorgiou, A. Kotsialos, Freeway ramp metering: An overview, *IEEE Transactions on Intelligent Transportation Systems* 3 (4) (2002) 271–281.
- [34] G. Gomes, R. Horowitz, Optimal freeway ramp metering using the asymmetric cell transmission model, *Transportation Research Part C: Emerging Technologies* 14 (4) (2006) 244–262.
- [35] M. Papageorgiou, H. Hadj-Salem, J.-M. Blosseville, ALINEA: A local feedback control law for on-ramp metering, *Transportation Research Record: Journal of the Transportation Research Board* 1320 (1991) 58–67.
- [36] A. Hegyi, B. De Schutter, H. Hellendoorn, Model predictive control for optimal coordination of ramp metering and variable speed limits, *Transportation Research Part C: Emerging Technologies* 13 (3) (2005) 185–209.
- [37] I. Papamichail, K. Kampitaki, M. Papageorgiou, A. Messmer, Integrated ramp metering and variable speed limit control of motorway traffic flow, *IFAC Proceedings Volumes* 41 (2) (2008) 14084–14089.
- [38] X.-Y. Lu, Z. Qiu, P. Varaiya, R. Horowitz, S. E. Shladover, Combining variable speed limits with ramp metering for freeway traffic control, in: *Proceedings of the American Control Conference*, 2010, pp. 2266–2271.
- [39] Y. Han, A. Hegyi, Y. Yuan, S. Hoogendoorn, M. Papageorgiou, C. Roncoli, Resolving freeway jam waves by discrete first-order model-based predictive control of variable speed limits, *Transportation Research Part C: Emerging Technologies* 77 (2017) 405–420.
- [40] W. Levine, M. Athans, On the optimal error regulation of a string of moving vehicles, *IEEE Transactions on Automatic Control* 11 (3) (1966) 355–361.
- [41] D. Swaroop, J. Hedrick, String stability of interconnected systems, *IEEE transactions on automatic control* 41 (3) (1996) 349–357.
- [42] S. E. Shladover, Review of the state of development of advanced vehicle control systems (avcs), *Vehicle System Dynamics* 24 (6-7) (1995) 551–595.
- [43] R. E. Fenton, R. J. Mayhan, Automated highway studies at the Ohio State University-an overview, *IEEE transactions on Vehicular Technology* 40 (1) (1991) 100–113.
- [44] S. Darbha, K. Rajagopal, Intelligent cruise control systems and traffic flow stability, *Transportation Research Part C: Emerging Technologies*

- 7 (6) (1999) 329 – 352.
- [45] B. Besselink, K. H. Johansson, String stability and a delay-based spacing policy for vehicle platoons subject to disturbances, arXiv preprint arXiv:1702.01031.
- [46] P. Ioannou, Z. Xu, S. Eckert, D. Clemons, T. Sieja, Intelligent cruise control: theory and experiment, in: Proceedings of the 32nd IEEE Conference on Decision and Control, IEEE, 1993, pp. 1885–1890.
- [47] M. Buehler, K. Iagnemma, S. Singh, The DARPA urban challenge: autonomous vehicles in city traffic, Vol. 56, Springer, 2009.
- [48] R. Rajamani, S. B. Choi, J. K. Law, B. K. and Hedrick, R. Prohaska, P. Kretz, Design and experimental implementation of control for a platoon of automated vehicles, AMSE Journal of Dynamic Systems, Measurement, and Control 122 (3) (1998) 470–476.
- [49] L. Davis, Effect of adaptive cruise control systems on traffic flow, Physical Review E 69 (6) (2004) 066110.
- [50] A. Talebpour, H. S. Mahmassani, Influence of connected and autonomous vehicles on traffic flow stability and throughput, Transportation Research Part C: Emerging Technologies 71 (2016) 143–163.
- [51] M. Guériau, R. Billot, N.-E. El Faouzi, J. Monteil, F. Armetta, S. 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 (2016) 266–279.
- [52] M. Wang, W. Daamen, S. P. Hoogendoorn, B. van Arem, Cooperative car-following control: Distributed algorithm and impact on moving jam features, IEEE Transactions on Intelligent Transportation Systems 17 (5) (2016) 1459–1471.
- [53] M. Wang, W. Daamen, S. P. Hoogendoorn, B. 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) (2016) 559–572.
- [54] Y. Han, D. Chen, S. Ahn, Variable speed limit control at fixed freeway bottlenecks using connected vehicles, Transportation Research Part B: Methodological 98 (2017) 113 – 134.
- [55] G. van de Weg, A. Hegyi, H. Hellendoorn, S. E. Shladover, Cooperative systems based control for integrating ramp metering and variable speed limits, in: Transportation Research Board 93rd Annual Meeting, no. 14-1432, 2014.
- [56] R. Nishi, A. Tomoeda, K. Shimura, K. Nishinari, Theory of jam-absorption driving, Transportation Research Part B: Methodological 50 (2013) 116–129.
- [57] Z. He, L. Zheng, L. Song, N. Zhu, A jam-absorption driving strategy for mitigating traffic oscillations, IEEE Transactions on Intelligent Transportation Systems.
- [58] Y. Taniguchi, R. Nishi, A. Tomoeda, K. Shimura, T. Ezaki, K. Nishinari, A Demonstration Experiment of a Theory of Jam-Absorption Driving, Springer International Publishing, 2015, pp. 479–483.
- [59] J. Ma, X. Li, S. E. Shladover, H. A. Rakha, X.-Y. Lu, R. Jagannathan, D. J. Dailey, Freeway speed harmonization, IEEE Transactions on Intelligent Vehicles 1 (1) (2016) 78–89.
- [60] X.-Y. Lu, S. E. Shladover, I. Jawad, R. Jagannathan, T. Phillips, Novel algorithm for variable speed limits and advisories for a freeway corridor with multiple bottlenecks, Transportation Research Record: Journal of the Transportation Research Board 2489 (2015) 86–96.
- [61] F. Wu, R. E. Stern, M. Churchill, M. L. D. Monache, K. Han, B. Piccoli, D. Work, Measuring trajectories and fuel consumption in oscillatory traffic: Experimental results, in: Proceedings of the 96th Transportation Research Board Annual Meeting, 2017 (to appear).
- [62] E. P. Agency, Control of air pollution from new motor vehicles and new motor vehicle engines; regulations requiring availability of information for use of on-board diagnostic systems and emission-related repairs on 1994 and later model year light-duty vehicles and light-duty trucks, Federal Register 60 (153) (1995) 40474–40498.
- [63] S. Cui, B. Seibold, R. E. Stern, D. Work, Stabilizing traffic flow via a single autonomous vehicle: Possibilities and limitations, in: Proceedings of the IEEE Intelligent Vehicles Symposium, 2017.
- [64] K. J. Åström, R. M. Murray, Feedback systems: An introduction for scientists and engineers (2008).
- [65] J. E. A. Dias, G. A. S. Pereira, R. M. Palhares, Longitudinal model identification and velocity control of an autonomous car, IEEE Transactions on Intelligent Transportation Systems 16 (2) (2015) 776–786. doi:10.1109/TITS.2014.2341491.
- [66] K. J. Åström, T. Hägglund, Advanced PID Control, ISA, Research Triangle Park, NC, 2005.
- [67] R. E. Stern, S. Cui, M. L. Delle Monache, R. Bhadani, M. Bunting, M. Churchill, N. Hamilton, R. Haulcy, H. Pohlmann, F. Wu, B. Piccoli, B. Seibold, J. Sprinkle, D. B. Work, Data and source code, <https://uofi.box.com/v/trajectoryPaperData> (2016).
- [68] R. P. Larrick, J. B. Soll, The MPG illusion, Science 320 (5883) (2008) 1593–1594. arXiv:<http://science.sciencemag.org/content/320/5883/1593.full.pdf>, doi:10.1126/science.1154983.
URL <http://science.sciencemag.org/content/320/5883/1593>
- [69] R. Olfati-Saber, J. Fax, R. Murray, Consensus and cooperation in networked multi-agent systems, Proceedings of the IEEE 95 (1) (2007) 215–233. doi:10.1109/JPROC.2006.887293.
- [70] M. Mesbahi, M. Egerstedt, Graph theoretic methods in multiagent networks, Princeton University Press, 2010.
- [71] C. Hönicke, P. Bliss, R. F. A. Moritz, Effect of density on traffic and velocity on trunk trails of *formica pratensis*, The Science of Nature 102 (3) (2015) 17. doi:10.1007/s00114-015-1267-6.
URL <http://dx.doi.org/10.1007/s00114-015-1267-6>
- [72] A. Berdahl, C. J. Torney, C. C. Ioannou, J. J. Faria, I. D. Couzin, Emergent sensing of complex environments by mobile animal groups, Science 339 (6119) (2013) 574–576.