



Using connected and automated vehicles to attenuate disturbances in freeway traffic

Benjamin Coifman^{a,*} , Yuan Liu^b 

^a The Ohio State University, Joint Appointment with the Department of Civil, Environmental, and Geodetic Engineering, and the Department of Electrical and Computer Engineering, Hitchcock Hall 470, 2070 Neil Ave, Columbus, OH 43210, United States

^b The Ohio State University, Department of Electrical and Computer Engineering, United States

ARTICLE INFO

Keywords:

Connected and Automated Vehicles
Freeway Traffic
Congestion
Safety
Jam Absorption Driving
Vehicle emissions

ABSTRACT

Congested traffic is characterized by disturbances propagating upstream through the queued vehicles. This work seeks to use connected and automated vehicles (CAVs) to mitigate slow and stop waves in congested freeway traffic with the goals of improving safety by reducing unexpected deceleration waves while improving fuel efficiency and emissions by reducing the need to accelerate after disturbances pass. To smooth out traffic disturbances, the CAV implements a jam absorption strategy that integrates the instantaneous state information from the downstream vehicles to forecast the trajectory of the CAV's leader and proactively respond to changes in state that have not yet reached the lead vehicle. Using empirical vehicle trajectory data, it is shown that the methodology can rapidly nullify stop and slow waves. Examples include the conversion of a 10 sec stop wave to steady and smoothly flowing traffic over a span of 15 CAVs.

One of the key innovations in this work is how the traffic state is forecasted using a model that is much simpler than prior work in this area while still being able to dynamically adapt to changes in the evolving traffic state. This paper also explores the impacts of varying the look ahead range, acceleration limits, and the parameters of the forecasting model. It is shown that the method is robust, still performing well when the forecasting model's calibration is poor. While the different permutations result in small changes, in all the experiments the methodology continues working as intended: the first CAV was able to nullify the stop waves and subsequent CAVs continued to anneal the disturbances.

1. Introduction

This work seeks to use connected and automated vehicles (CAVs) to smooth out traffic disturbances on uninterrupted flow facilities (e.g., freeways), eliminating unexpected slowing and stops to improve safety while reducing accelerations to lessen fuel consumption and vehicle emissions. Specifically, this research seeks to take conventionally unstable queued traffic and have a CAV proactively respond to disturbances propagating upstream through the traffic to stabilize the flow. For this study the queued traffic comes from empirical human driven vehicle (HDV) trajectory data.¹ For most of this paper, the approach assumes all vehicles are connected and

* Corresponding author.

E-mail address: Coifman.1@OSU.edu (B. Coifman).

¹ We use HDV for two reasons, first, we can use empirical data instead of stimulation, and second, presumably, any widely deployed autonomous vehicle (AV) system should be able to meet or exceed the stability of HDV.

report their current state (location and speed). In other words, the initial traffic is strictly connected-HDV (cHDV) that communicate their state but are otherwise HDV (Section 4 relaxes this assumption that all HDV are connected). This work seeks to integrate the state across vehicles so that the CAV can efficiently anticipate and respond to disturbances over large distances. It is this look ahead that allows the system to attenuate shockwaves, resulting in smoother conditions for safety and efficiency. The primary focus of this work is the initial transition from unstable stop-and-go HDV traffic to the first CAV follower as it generates a smoother trajectory. After considering the single CAV scenario, the method is extended to a platoon of multiple CAVs. It is shown that the system can rapidly dissipate a stop wave.

For context, it is important to keep in mind what is feasible with CAV. One can view the roadway as a pipe. If you use CAVs to automate the pipe, it is still subject to the operating limitations of a pipe, e.g., even if capacity increases with CAVs there will still be some maximum feasible throughput. Our focus is on queued traffic, and one must recognize that a queue simply represents the storage of demand that exceeds the downstream capacity. So, the queued vehicles experience delay while they wait to be served. To eliminate a queue requires increasing capacity or decreasing demand. This work envisions CAVs that have speed-spacing constraints and response times that are similar to HDV, as such, the method should not change the maximum possible throughput. Moreover, this work does not seek to achieve freely flowing traffic out of queued traffic. Rather, this work seeks to take conventionally unstable queued traffic and bring it to a stable flow while queued. The queue would travel below the speed limit at near constant speeds while the vehicles wait to be served at the bottleneck. By smoothing out the disturbances this work seeks to reduce the negative impacts of congested traffic: keeping conditions as safe as possible while minimizing fuel consumption and emissions from the queued vehicles.

1.1. Background

The idea of stabilizing queued freeway flow has a long history. Some of the earliest work is focused on variable speed limits (VSL). Within VSL, the initial works typically use fixed location sensing, e.g., loop detectors, to measure the state and changeable speed limit signs to control the flow. Some of the recent works rely on connected vehicles to monitor the traffic state or communicate the target speed to the vehicles (e.g., Alessandri et al., 1999; Hegyi et al., 2005; Hegyi and Hoogendoorn, 2010; Chen et al., 2014; Ma et al., 2016; Wang et al., 2016b; Han et al., 2017a; Han et al., 2017b; Ghiasi et al., 2019; Han et al., 2021a; Du et al., 2023). The dominant themes in these papers are (i) to harmonize speeds by setting the speed limit lower than the speed most drivers are taking, (ii) because flow drops with speed in the congested regime use a lower speed limit to reduce the demand delivered downstream, and (iii) a reduced flow could allow the bottleneck to recover from *capacity drop* and thus, increase bottleneck throughput. Not all of the VSL papers touch on all of the themes. The chief concerns about VSL include (i) high cost of sensing, (ii) high cost of the changeable speed limit signs, (iii) limited resolution of control, e.g., one speed limit at a time for all lanes for a given link, (iv) a long latency between sensing and the control response will degrade performance, and (v) driver compliance. While several of the VSL papers do not directly address any of the concerns, many of the papers specifically address one or more of the concerns.

Jam absorption driving (JAD) is a related body of research that seeks to have one or more vehicles proactively respond to approaching conditions, e.g., having a CAV start slowing down in response to an approaching stop wave and potentially doing so long before their leader encounters the disturbance. While VSL typically has most if not all vehicles adopting some target speed, JAD typically only has one or a few vehicles actively adopting the target speed under the presumption that their followers will then have to match that new speed profile (e.g., Nishi et al., 2013; Taniguchi et al., 2015; He et al., 2016; Stern et al., 2018; Di Vaio et al., 2019; Ibrahim et al., 2019; Zheng et al., 2020; Han et al., 2021; Jiang et al., 2022; Li et al., 2023; Goatin, 2024; Li et al., 2024). While VSL often includes consideration of downstream bottlenecks, usually JAD is focused strictly on smoothing out the speed. The chief concerns about JAD include (i) often the studies do not allow lane change maneuvers, (ii) if lane change maneuvers are allowed, they potentially undermine the effectiveness of JAD because other vehicles could spill around a given JAD controlled vehicle, and (iii) JAD vehicles can introduce new safety issues, e.g., high relative speeds and increased frequency of lane change maneuvers.

Both VSL and JAD typically require a long look ahead distance to provide sufficient time to change the traffic dynamics. Another related body of research focuses on the adjacent vehicles that can be sensed directly from the vehicle making the maneuver. This can be a single vehicle perceiving its immediate leader and possibly one or two vehicles further ahead, e.g., adaptive cruise control (ACC) or autonomous vehicles (AV) without connectivity (e.g., Donà et al., 2022; Zhou et al., 2022; Wang and Jin, 2023; Wang et al., 2023). It can also be a platoon of vehicles that are communicating with each other, but they are not directly communicating with vehicles further downstream, e.g., some variants of cooperative-ACC (CACC) (e.g., Vander Werf et al. 2002; van Arem et al., 2006; Shladover et al., 2012; Wang et al., 2016a; Qin et al., 2021; Tan and Zhang, 2022; Zhou et al., 2024). Typically, these approaches rely on perception sensors to monitor the leader(s) and do not require connectivity (except intra-platoon communication for CACC), but because the look ahead is so short, they typically require many successive active vehicles to dissipate a disturbance.

1.2. Contribution

This paper is most closely aligned with the JAD domain. These approaches require some form of traffic modeling to forecast the trajectory of the CAV and its leader. Typically, JAD studies rely on car following models or machine learning tools, in either case the performance is highly dependent on the quality of the calibration or training. One of the biggest challenges for these systems is that the correct calibration typically differs by location and can rapidly change if the composition of the ambient vehicle fleet or driving styles change.

Our study employs much simpler traffic modeling than the prior work in this area. The traffic model has a single parameter (wave velocity), and the performance is not very sensitive to the chosen value. The initial setting for the wave velocity in this paper is based

on typical values reported in the literature rather than finding an optimal value for the empirical data used in the study. Section 4 shows that our method continues to perform well on the empirical data when the assumed wave velocity in the model is varied by a factor of three (from -7 mph to -23 mph). Despite the simplicity, the iterative nature of sampling and recalculating the target trajectory at each time step makes the approach robust, allowing it to adjust to changes in the evolving traffic state. We also explore limiting the CAV acceleration to increase smoothing.

Another distinction is that most of the JAD studies rely on a single vehicle to do the absorption, while we bring in a second stage that uses multiple CAVs to further smooth out the traffic state, similar to some of the CACC studies. This extension needs a modified approach since the first CAV effectively “overwrites” the information contained in the preceding HDV trajectories. So, instead, each subsequent CAV uses its leader’s estimated trajectory for its forecast.

Finally, upon reviewing the emergent dynamics, it is shown that as the first CAV slows, it creates a large gap before its leader encounters an approaching stop wave. Then, when the HDV leader slows down, the first CAV follower maintains a relatively stable speed as it consumes the excess spacing, thereby starving the stop wave’s progression. These JAD behaviors by the CAV give rise to clockwise progressions in the speed-spacing plane with behavior that HDVs do not exhibit.

1.3. Overview

The remainder of this paper is as follows. Section 2 presents the methodology in the context of a single CAV following the top row of the flowchart in Fig. 1. The presentation starts with the underlying theory. It then empirically demonstrates the approach using microscopic vehicle trajectory data and quantifies the impacts on emissions and fuel consumption. The section also explores the impacts of the look ahead range and limiting acceleration. Next, Section 3 expands the work to 15 consecutive CAVs, adding the complete loop in Fig. 1 for subsequent CAVs and explores how multiple followers can further anneal the smoothed trajectory of the first CAV. Section 4 explores the impact of the assumed wave velocity and the cHDV penetration rate on the results. The paper closes with a brief discussion and conclusions in Section 5.

2. Methodology

Congested traffic is characterized by signals and waves propagating upstream through the queued traffic. This work develops a methodology for a CAV to nullify stop and slow waves from propagating upstream within a queue of vehicles on a freeway. Most of this paper assumes that all vehicles downstream of the CAV are connected and report their instantaneous state (location and speed), but we explore lower penetrations rates in Section 4. In any event, the CAV integrates the instantaneous state information from the downstream vehicles to forecast the trajectory of the CAV’s leader and proactively respond to changes in state that have not yet reached that vehicle.

The method is inspired by Coifman (2002), which used measurements of vehicles passing a dual loop detector station to estimate vehicle trajectories over space for up to one mile away from the station. The basic principle is that many freeways exhibit a triangular fundamental relationship, thus, during congestion the traffic state should propagate upstream through the traffic at roughly constant velocity, corresponding to the “wave velocity,” i.e., the slope of the congested regime of the flow-density curve (see, e.g., Coifman and Wang, 2005, Coifman, 2015). This feature is illustrated in Fig. 2A using real human driven vehicle (HDV) trajectories reextracted from the NGSIM I-80 video in lane 2 (Coifman and Li, 2022, 2024). Thirteen “reference signals” are superimposed on the vehicle trajectories, each signal is traveling at -15 mph. Between any pair of signals one can observe that the trajectories show roughly the same speed over the entire band defined by the pair of signals. Of course, this approximation is not perfect, upon close inspection one can see the traffic state between a given pair of signals slowly change shape over time and space, due to natural human variability and abrupt lane change maneuvers.

While Coifman (2002) derived the trajectory estimation using measurements at a point in space (as would be represented by a horizontal cutline in Fig. 2A), this work re-derives the method using conditions at an instant in time (a vertical cutline). Consider Fig. 2B, which shows the complete trajectories up to time t_1 , set to 630 sec in this case. We illustrate the process of establishing the




Fig. 1. Overview of the processing in each time step. With the long look ahead, the process of sampling the current state each cycle allows the method to adjust to changes in the evolving traffic state.




Fig. 2. (A) Vehicle trajectories from congested traffic and several upstream moving reference signals. (B) Traffic state up to time t_1 , (C) roughly steady state bands and their measured vehicle chords, (D) aligning the chords at t_1 end to end to estimate the lead vehicle's trajectory, and (E) comparing the estimate to what will eventually occur. (F) Establishing the perfect follower trajectory and finding the intersection with the fastest constant speed radial for the CAV at t_1 .

optimal trajectory for a CAV at (t_1, y_1) in the lower left corner of this plot. The first step is estimating the trajectory of its leader at $(t_1, 100\text{ft})$. For this work we assume that this leader vehicle and all vehicles downstream are cHDVs that communicate their state but are otherwise HDVs. While the real traffic state varies continuously over time and space, we approximate it with thin bands of constant speed that are defined by the instantaneous location and speed from the downstream cHDV. Specifically, each cHDV location is taken to represent the start of a new band of constant speed, with the signal between bands propagating upstream at the wave velocity (set to -15 mph in this case). For each band we define a chord – a constant speed trajectory passing through the band at the reported cHDV speed. The result is illustrated in Fig. 2C. Since the bands are assumed to be at constant speed, we shift the chords end to end, starting at the lead vehicle's current position of $(t_1, 100\text{ft})$, yielding the estimated trajectory of the leader in Fig. 2D. Fig. 2E compares the estimated trajectory at t_1 to the actual trajectory that is eventually realized by the same vehicle. While the estimate is not perfect, it does capture the key features like the passage of the stop wave.

The lead vehicle's estimated trajectory bounds where the CAV at (t_1, y_1) can travel. To establish the safe travel zone for the CAV we require a minimum 1 sec reaction time to the leader. This safe travel zone is then bounded by the **perfect follower trajectory**, derived from Newell (2002) using the wave velocity from the lead vehicle trajectory estimation and the 1 sec reaction time from above, i.e., $w = -15\text{mph}$ and $\tau = 1\text{sec}$, with the result for the ongoing example shown in Fig. 2F. Next, we find the fastest constant speed radial from (t_1, y_1) that touches the perfect follower trajectory, in this case doing so at point a_1 . Instead of directly following the leader, the CAV avoids unnecessary acceleration by taking the speed of this radial for the next time step; thus, the CAV effectively “stringlines” the leader's trajectory across an acceleration-deceleration cycle while still passing the most restrictive point, a_1 , without any further delay. Note that the chosen trajectory for the CAV is always checked to ensure spacing, s , is at or above the minimum safe spacing, s_{min} , for the

given speed, v , as per Equation 1. Where jam spacing is the quotient of w and τ , so $s_{jam} = 22\text{ft}$ in this case. Operationally, this means that if the CAV is at or downstream of the perfect follower trajectory, it will take the constant speed trajectory to the point on the perfect follower trajectory one time step away.²

$$s \geq s_{min}(v) = s_{jam} + v^* \tau \quad (1)$$

This process of sampling the current traffic state to estimate the leader's trajectory and then setting the CAV speed to intersect the perfect follower trajectory at the furthest point downstream is repeated at each time step. The result for this example is shown in Fig. 3A using one second time steps. The left-most dashed trajectory shows the path taken by the actual HDV in the reextracted NGSIM data (termed the "original HDV trajectory"), while final trajectory taken by the CAV is shown with a bold, multi-color curve. Each colored segment represents a one second time step. While the speed of the final CAV trajectory slowly varies in response to the evolving forecasts, it anticipates the stop wave and, in this case, neutralizes the stop. The final CAV trajectory is much smoother than the cHDV that preceded it. Fig. 4A-B show the time series speed and acceleration for the CAV and the corresponding original HDV for this case, clearly showing less variability in speed and attenuated accelerations. Meanwhile, Fig. 3B repeats the analysis for another CAV, in this case at (t_2, y_2) . For this second case, while some slowing is evident at the downstream end of the segment at t_2 , the stop wave had not yet formed when we start plotting the CAV trajectory. Still, the CAV is able to avoid an unnecessary acceleration and once more neutralizes the stop wave. Fig. 4C-D show the time series speed and acceleration for the CAV and its corresponding original HDV in this case.

2.1. Extending the look ahead

So far, the cHDV look ahead has been limited to the span of the empirical HDV trajectory data, which ends at 1800 ft. So, at t_1 in Fig. 2D the CAV used a 1700 ft look ahead, which yielded an estimated trajectory to almost 900 ft, or roughly half the distance to the end of the segment. As the CAV travels further downstream, the range of the cHDV look ahead shrinks because it can never "see beyond" the 1800 ft limit of the empirical vehicle trajectory data. To this end, we repeat the method from Coifman (2002) to synthesize the trajectories between 1800 ft and 2800 ft using the last reported speed from each vehicle as they depart the surveillance region. This speed is used to project a chord downstream from the last observation of the vehicle while also generating a signal with a velocity of -15 mph, thus, each of these signals effectively moves backward in time as it travels downstream in space. Each vehicle chord is truncated upon reaching the first signal. Once more, each band is assumed to be constant speed and whenever a vehicle departs the field of view at 1800 ft the corresponding vehicle chords are aligned end to end to estimate that vehicle's trajectory beyond 1800 ft. An example of the resulting synthetic trajectories between 1800 and 2500 ft is shown in Fig. 5.

Thus, we have empirical HDV trajectories to 1800 ft and synthetic trajectories beyond. With the synthetic trajectories we can consider longer look ahead distances for estimating the lead vehicle's trajectory. In addition to the original estimate from 0 to 1800 ft from Fig. 3A and repeated in Fig. 6A for case 1, we consider 0–2500 ft in Fig. 6B and a fixed moving window of 1000 ft downstream of the CAV's leader in Fig. 6C, i.e., the trajectories within the dashed region.³ In this way, the 0–1800 represents the largest feasible window with the purely empirical data, 0–2500 ft shows the benefits of a longer look ahead (at the expense of using synthetic data), and the fixed 1000 ft window represents a practical implementation of the method. Regardless of the look ahead region used, at each time step the lead vehicle's trajectory estimation only uses the cHDV data from that instant. For this case there is very little difference between the final CAV trajectory in the three different scenarios. The 0–1800 CAV trajectory from Fig. 6A is shown below the CAV trajectory in Fig. 6B&C, the other two scenarios. For Fig. 6B the two estimates only differ in the last few samples, around 1700 ft. Whereas the shorter look ahead with the fixed 1000 ft window is more susceptible to the stop wave, pulling a little ahead of the 0–1800 CAV trajectory around 300 ft, and again around the last few samples before 1700 ft. Fig. 7 repeats the comparisons for case 2, seen previously in Fig. 3B. In this case the stop wave passes the CAV further downstream, so by the time the look ahead to 1800 ft captures the stop wave the CAV has started accelerating into the lead up to the wave. Whereas Fig. 7B shows that the look ahead to 2500 ft provides enough forewarning to allow the CAV to maintain a more constant speed as it nullifies the stop wave. On the other hand, Fig. 7C shows that the fixed 1000 ft window is even more sensitive to the acceleration in the lead up to the stop wave as it pulls ahead of the original curve, only to have to slow later. Still, this fixed window look ahead is sufficient to nullify the stop wave.

To consider a longer look ahead, Fig. 8 repeats the analysis from Fig. 3 only this time using the resurrected data set from Treiterer (Coifman et al., 2018). These data were collected from a helicopter traveling with the vehicles, thereby providing a moving window that supports a fixed look ahead of empirical trajectories out to 3000 ft and a variable "full range" look ahead that at times exceeds one mile. In this data set traffic appears to be free flowing until just before 6000 ft, where it joins the end of a queue. We consider three CAV different look ahead ranges: full available range, fixed 3000 ft window, and fixed 1000 ft window. The vertical lines in the figure show the look ahead range for the given CAV at the instant it starts to respond to the downstream delay. One can see the proactive benefit of a longer look ahead, the CAV with full range look ahead is responding to the stop wave from the very start of its trajectory, then at 47 sec the fixed 3000 ft window CAV responds, and finally the fixed 1000 ft window CAV responds at 65 sec. Note that the two fixed window CAVs start responding to the slow traffic even before they are able to see the actual stop. Also note that the two fixed window CAVs lag behind the HDV even before they start responding to the slow traffic, because we set the CAV free flow speed to 60 mph while the HDV travel a little above this speed.

² It is assumed that the CAV has perception sensors to monitor the spacing to the lead vehicle to always ensure a safe spacing. This sensing distinction is not critical for our analysis, it does not matter if a vehicle's position is reported from a cHDV or measured directly from the CAV.

³ Specifically, for the 1000 ft look ahead we use all vehicles within 1000 ft downstream of the leader to estimate the leader's trajectory.




Fig. 3. The final CAV trajectory for the CAV in (A) case 1 starting at (t_1, y_1) , and (B) case 2 starting at (t_2, y_2) . In both cases note how the CAV neutralizes the stop wave seen in the leader trajectory and original HDV trajectory.

Fig. 9C shows the corresponding time series speed for these four follower variants. The HDV does not anticipate the stop wave and rapidly comes to a complete stop for several seconds. The fixed 1000 ft window CAV proactively responds to the forthcoming delay with enough time to keep its speed above 10 mph, while the fixed 3000 ft window CAV responds earlier and remains above 20 mph, and the full range CAV is responding from the first instant and remains above 25 mph throughout. At least in this case, the earlier a CAV starts responding also corresponds to the longer the CAV goes before rejoining the HDV trajectory. So, the longer look ahead provides a higher minimum speed and a longer lasting response both before and after the stop wave. Fig. 9A&B repeat the resulting time series speed for two other leader–follower vehicle pairs that pass earlier in the data set. Note that the longest fixed window look ahead is only 2400 ft in these plots because the vehicles are closer to the start of the data, thereby reducing the number of observed leaders and their corresponding spatial coverage. The general trends are consistent across all three leader–follower pairs.

2.2. Limiting acceleration

Next, we consider the benefits of limiting the CAV acceleration. The limits are strictly on acceleration, for safety's sake no limit is placed on the magnitude of deceleration. We return to the reextracted NGSIM data⁴ using the fixed 1000 ft window for illustration because it has the smallest overall look ahead and thus, also exhibits the largest fluctuations among the three different look ahead scenarios, but the results are consistent with the other scenarios. Fig. 10A repeats the unlimited results from Fig. 6C. Fig. 10B repeats the analysis, only now limiting the CAV acceleration, $a_{CAV} \leq 1 \text{ mph/s}$, while Fig. 10C repeats the process except $a_{CAV} \leq 0.5 \text{ mph/s}$. The differences between the three scenarios only becomes apparent at close inspection, namely, in Fig. 10B the dark curve corresponding to the original unlimited CAV trajectory pulls to the left of the lighter curve for $a_{CAV} \leq 1 \text{ mph/s}$ around $Y = 500 \text{ ft}$ and again around $Y = 1400 \text{ ft}$. Despite the slower response within the segment, the departure time at the end of the segment is nearly unchanged. The differences are easier to see in Fig. 10C where $a_{CAV} \leq 0.5 \text{ mph/s}$. In this case the slower acceleration eliminates one acceleration–deceleration cycle that occurred in the CAV trajectory between 500 and 1300 ft in the other two scenarios. At the end of the link, the $a_{CAV} \leq 0.5 \text{ mph/s}$ trajectory is within 1.5 sec of the unlimited CAV trajectory and is closing the gap. This slower response time is inherent with the limited acceleration and as evident earlier in the progression, primarily serves to smooth out the trajectories. Provided the $a_{CAV} \leq 0.5 \text{ mph/s}$ trajectory eventually rejoins the unlimited trajectory, the lower acceleration threshold does not add any delay. Fig. 11A shows the corresponding speed for these three scenarios along with the original HDV trajectory's speed. All the CAV curves greatly reduce the undulations of the HDV, while the maximum acceleration of the latter two scenarios is evident by the maximum slope of the given speed curve. As one should expect, the $a_{CAV} \leq 0.5 \text{ mph/s}$ shows the least fluctuations as it slowly tracks the changing speed, and this smoother trajectory provides more stable conditions to any vehicle following the CAV. Fig. 11B shows the corresponding time series acceleration for the HDV and three scenarios of the CAV. Note how much the $a_{CAV} \leq 0.5 \text{ mph/s}$ scenario reduces the acceleration spikes throughout the timespan. Although deceleration is not restricted, because of the look ahead, all three of the CAV scenarios keep the magnitude of deceleration below 2 mph/s whereas the HDV deceleration briefly reached a magnitude of almost 8 mph/s.

⁴ While Treiterer's data provides greater spatial coverage, there is only a short distance between the upstream end of the queue at 6000 ft and the stop wave that is strictly seen upstream of 8000 ft. While the previous section shows the method works in that scenario, the interaction of the two features complicates the dynamics and obscures some of the finer points this section demonstrates.




Fig. 4. Case 1 time series (A) speed and (B) acceleration; and case 2 time series (C) speed and (D) acceleration.

Fig. 10D-F repeat the analysis for case 2, with Fig. 10D showing the unlimited results from Fig. 7C. Recall from Fig. 7 that the fixed 1000 ft window was more sensitive to the stop wave than the other two variants that considered a larger look ahead to either 1800 ft or 2500 ft. Specifically, the fixed 1000 ft window pulls ahead of the original curve, only to have to slow later. Fig. 10F shows that the $a_{CAV} \leq 0.5 \text{ mphs}$ trajectory compensates for this aggressive behavior, reducing the impacts of the transient increase in speed on the CAV trajectory. Fig. 11C shows the corresponding speed for these three scenarios along with the original HDV trajectory's speed. As




Fig. 5. An example illustrating how conditions upstream of 1800 ft are used to generate vehicle chords representative of steady state bands. The speed in these bands is combined with the time that vehicles depart the surveillance region at 1800 ft to synthesize the unobserved vehicle trajectories downstream of 1800 ft.




Fig. 6. Estimating the final CAV trajectory for Case 1 (A) using the original look ahead region out to 1800 ft, (B) the extended surveillance region out to 2500 ft, and (C) the fixed 1000 ft moving window.

with case 1, the smoothing properties of the CAV are readily apparent. Fig. 11D shows the corresponding time series acceleration for the HDV and three scenarios of the CAV. Once more, the $a_{CAV} \leq 0.5 \text{ mph/s}$ scenario is notable in how much it flattens the acceleration peaks.

Vehicle emissions and fuel consumption are highly dependent on acceleration. We use the Comprehensive Modal Emission Model, CMEM, (Scora and Barth, 2006) to quantify the impact of the various scenarios under case 1. Although not shown for brevity, case 2 had similar trends. Using the original HDV as the baseline, Fig. 12 shows the relative improvement in the CMEM results for the three




Fig. 7. Estimating the final CAV trajectory for Case 2 (A) using the original look ahead region out to 1800 ft, (B) the extended surveillance region out to 2500 ft, and (C) the fixed 1000 ft moving window.




Fig. 8. Estimating the CAV trajectory using data from lane 1 in Treiterer's data set, which provides a much longer look ahead in the empirical data, thereby avoiding the need for synthetic data. In this case we consider the original HDV follower, and three different CAV look ahead ranges. The vertical lines show the look ahead range for given CAV at the instant it starts to respond to the downstream delay.

different acceleration limits using the fixed 1000 ft window. The bar chart shows that fuel and CO₂ are reduced by almost 20 % (so the CAV only used 80 % of the fuel that the HDV did for the same trip), CO drops by almost 40 % and HC by almost 50 %. Meanwhile, NO_x is particularly sensitive to the maximum acceleration, ranging from roughly 30 % savings with unlimited acceleration to almost 60 % savings with the 0.5 mph/s limit. Of course, the specific values are highly dependent on the prevailing conditions, individual vehicle characteristics, and so forth. We use these metrics in the relative sense, to illustrate that one scenario is better or worse than another scenario. The general trends should be transferable, e.g., A is better than B, but the specific values or exact percent improvement are not likely to transfer. Naturally, the savings from the smoother CAV trajectory should be passed to any following vehicle, even if that follower is a HDV.




Fig. 9. Time series speed from Treiterer's data set for three different leader–follower pairs, each with the original HDV follower and three different CAV followers. Note that parts A and B use a fixed 2400 ft window while part C corresponds to the example in Fig. 8 and uses a fixed 3000 ft window.

3. Extending to multiple CAV followers

While the first CAV follower integrates information from the downstream cHDV to estimate its trajectory, the new trajectory effectively “overwrites” the information contained in the preceding HDV trajectories. So, a different approach is needed to find the trajectory for any subsequent CAV. This process still begins with the first CAV. So, at time step t_0 we estimate the first CAV's trajectory as per Fig. 2 and summarized in Fig. 13A. The first CAV at (t_0, y_1) takes the fastest constant speed radial that touches the perfect follower trajectory at a single point, a_1 . In the process, its trajectory from (t_0, y_1) to a_1 creates new bounds for the following vehicle. But at this point it is not yet clear what the optimal path is beyond a_1 . So, the estimated trajectory for the first CAV is taken to be the straight line trajectory to a_1 and its leader's perfect follower trajectory after a_1 , as shown in Fig. 13B.

To find the safe travel zone for the second CAV, the first CAV's estimated trajectory is shifted using Newell (2002) with the same minimum reaction time and wave velocity as before to find the first CAV's perfect follower trajectory. Fig. 13C shows the result for the ongoing example. Now the second CAV's trajectory is estimated following the same rules used for the first CAV, subject to the following conditions: If the first CAV does not decelerate and the second CAV is car following, the second CAV will continue to car follow. If the




Fig. 10. The final CAV trajectory for case 1 with the fixed 1000 ft window and (A) unlimited acceleration, (B) $a_{CAV} \leq 1\text{mphps}$, and (C) $a_{CAV} \leq 0.5\text{mphps}$. (D)–(F) repeat the analysis for case 2.

first CAV decelerates or the second CAV is not car following, the second CAV will aim for point a_2 to begin car following in the same way the first CAV aims for a_1 . Point a_2 will fall somewhere at or beyond w_1 , the signal propagating at the wave velocity from a_1 (as per Fig. 13C) and at or before w_N , the signal propagating at the wave velocity from the end of the estimated perfect follower trajectory (as per Fig. 13D), which is dictated by the look ahead distance and prevailing traffic conditions. Of course, the CAVs only aim for these target points for one time step before recalculating the a 's in the next time step. If there are more than two CAV in a sequence, this process is repeated for each successive CAV at each time step.

3.1. Empirical results

As we shift to the empirical data, again we run into the spatial limitations of NGSIM, but there are few vehicle trajectory data sets that cover greater distance for a sufficiently large number of vehicles with comparable fidelity as the reextracted NGSIM data. So just as we employed synthetic trajectories to extend the surveillance region beyond 1800 ft, we also use synthetic trajectories to establish the location of the CAV followers upstream of 0 ft. But these locations are only used at the first time step, after which, the CAV evolution is derived using the method presented herein.

Fig. 14 revisits case 1 with the fixed 1000 ft window and no limit on acceleration, now using 15 successive CAVs. All 15 CAV trajectories are shown with bold curves, while the first of these CAVs is shown with a lighter shade than the others. This first CAV is unchanged from the CAV trajectory in Fig. 6C. At each time step, the 2nd to 15th CAV use their respective leader's trajectory to estimate the perfect follower and then find the intercept point, as per above. The resulting trajectory for the 15th CAV is much smoother than that of the first CAV. For comparison, the original HDV trajectories for these 15 vehicles are shown with a faint colored curves. This comparison is limited by the fact that five HDV entered the lane within the 15 original HDVs. The entering vehicles are shown with dashed curves. The 15th HDV at the start is shown with a darker curve until just before 700 sec, after which point, the 15th HDV at the end (after accounting for the entering vehicles) is shown with the dark color. Note how the last HDV still follows the stop wave around 660 sec while the last CAV shows nearly constant speed over the same period.

Fig. 15A shows the time series speed for the first original HDV trajectory as a black dashed curve and all 15 of the CAV trajectories using various color and dash patterns. There are so many CAV curves that it is difficult to keep track of any individual curve. On the left side of the plot each successive CAV takes a speed slightly higher than their leader, so the first CAV is at the bottom, then the second CAV, and so forth to 15th CAV at the top and then above all of the CAVs is the curve from the first empirical HDV. The HDV soon decelerates in response to the stop wave and its speed drops below all the CAVs while all of the CAVs proceed without dropping below 9 mph. A key feature of this plot is how conditions improve across the 15 successive CAVs, to this end Fig. 15B repeats the time series speed for only the first and 15th CAV from part A. Notice how the 15th CAV's speed is nearly constant between 650 and 705 sec. Fig. 15C–D show the corresponding time series acceleration. The results are consistent, e.g., the 15th CAV shows almost no acceleration before 705 sec.




Fig. 11. Time series (A) speed and (B) acceleration for case 1 with the fixed 1000 ft window under the three acceleration limitation scenarios; and the corresponding time series (C) speed and (D) acceleration for case 2.

Note that the time range in Fig. 15 extends almost 25 sec past the end of the time range in Fig. 14. Fig. 14 is deliberately terminated at 705 sec because that is the instant when the HDV leader is last seen in the empirical data and we denote this time with t_l in Fig. 15. Any calculations beyond t_l must be based on synthetic data for the leader. While we are comfortable using the synthetic data to extend




Fig. 12. Fuel and emissions savings relative to the first HDV follower for the CAV in case 1 under the fixed 1000 ft look ahead window for the three different acceleration scenarios.

the look ahead, and we believe the synthetic data are valuable for qualitative insights, we are reluctant to use a synthetic leader trajectory for quantitative comparisons.

With these caveats in mind, there are several features of note near t_L in Fig. 15. First, notice how the different CAVs sometimes exhibit synchronous responses, e.g., the small wobbles in acceleration after 710 sec in Fig. 15C. In this time period the first CAV is slowing. As a result, point a_1 is continually at the end of the perfect follower trajectory, as per Fig. 13D. As the first CAV varies its speed due to changes in the estimated perfect follower trajectory, in response all subsequent CAV will also adjust their speed in that time step because point a_2 moves with a_1 . These synchronized wobbles could be viewed as artifacts of the 10 Hz sampling used in the NGSIM data. With higher temporal resolution data, we could employ a shorter time step and use a low-pass filter to smooth out the high frequency wobbles.

Next, in Fig. 15A notice how successive trajectories accelerate in response to an acceleration wave that passes after 690 sec, with the HDV rising earliest but it has not caught up to the speed of the first CAV when it starts accelerating a few seconds later. Then each successive CAV accelerates one second after its leader (as per the response time used for the CAV) in sequence to the 15th CAV. Shortly thereafter, the CAV followers start to peel away from the leader taking lower speeds in response to an approaching slow wave. Fig. 15B shows that both the magnitude and duration of this acceleration-deceleration cycle diminish greatly from the first to the 15th CAV.

Using the synthetic data, Fig. 16 extends the coverage further in time and distance from Fig. 14 and now shows the downstream end of the fixed 1000 ft window. We see that right at t_L the trajectories at the end of the window have started to slow in response to the approaching slow wave. While this approaching wave is evident in the look ahead, the HDV leader and first HDV follower are unaware and continue accelerating while the first CAV follower starts to decelerate in anticipation of the approaching delay. The first sign of recovery from the slow wave is seen in the look ahead window at t_E (i.e., the accelerating trajectory at the downstream end), and in the plot this signal is projected forward in time along the signal w_E . As per design, at t_E the first CAV ultimately aims to rejoin the perfect follower as it recovers from the slow wave and takes a nearly constant speed from t_E until it passes the recovery along w_E . The resulting impact on the CAV speeds is evident in Fig. 15A-B after t_E . At first glance it might seem counterproductive that each successive CAV accelerates for a short while in the period right before t_L , and that is true when you can see the full time-space plane. But at each time step the look ahead has not yet revealed the magnitude of the delay that is approaching. Conditions within the look ahead indicate that the flow is improving and so the followers should accelerate to keep up with their leaders. Fig. 15A shows that before t_E each successive CAV follower speeds up for a shorter period of time and to a lower maximum speed than its leader, which in turn creates excess spacing between the given pair of vehicles. As a result, in Fig. 15A after t_E each successive CAV takes a speed slightly larger than their leader with the aim to catch up to the leader and resume car following just as they pass w_E in Fig. 16. In this manner, the platoon of CAVs continues to anneal the impacts of the disturbance with successively higher speeds as they smooth the flow. Ultimately, the first CAV in Fig. 16 proactively responded to the slow wave almost 20 sec before the HDV leader does. With a longer look ahead, the CAV could have responded earlier to the major slow wave and avoided the acceleration prior to t_L in Fig. 15A. Regardless, as already noted, Fig. 15B shows that both the magnitude and duration of the impact diminish greatly from the first to the 15th CAV.

Figs. 17 and 18 repeat the 15 CAV evaluation for case 2 with a 1,000 ft look ahead and no limit on acceleration. The general trends are similar to case 1. At the start of this case the stop wave is downstream of the look ahead region, so the first CAV initially adopts the perfect follower trajectory until the delay of the stop wave starts to enter the look ahead region. At which point the first CAV slows and




Fig. 13. (A) To plan the path of the first CAV follower find the fastest radial from the CAV's location that touches the perfect follower trajectory at a_1 , (B) combine the radial to a_1 and the perfect follower trajectory beyond as the estimated trajectory for the first CAV follower, (C) shifting the estimated trajectory for the first CAV follower to establish its perfect follower's trajectory, and then the process is repeated by the second CAV follower to find where the fastest radial touches the perfect follower at a_2 . (D) Sometimes the intersection points a_1 or a_2 can fall at the end of the estimated perfect follower trajectory, as shown here.

successfully navigates a gentle transition to the post-stop wave state, as per Fig. 18A, the corresponding HDV exhibits large speed fluctuations and almost comes to a full stop.⁵ Fig. 17 shows the stop wave persists to the 15th HDV while the 15th CAV traverses the former stop wave with nearly constant speed. Although there is no limit placed on acceleration, Fig. 18D shows the 15th CAV keeps acceleration below 1 mphs except for one time step after 650 sec and below 0.5 mphs for most of the time steps in this example.

Returning to case 1, we extended the CMEM evaluation to all 15 of the CAVs, while also considering all three look ahead scenarios and all three acceleration limits. The results are similar across all nine permutations, so for brevity Fig. 19 shows the improvements relative to the HDV baseline with a 1,000 ft look ahead and no limit on acceleration. This bar chart shows the first CAV follower, 15th

⁵ As discussed in Section 2, technically at this point the stop wave has not fully formed because some vehicles do not come to a complete stop, but after a few more vehicles it turns into a full stop wave.




Fig. 14. Case 1 with the fixed 1000 ft window and no limit on acceleration using the first follower from before and the next 14 vehicles for a total of 15 successive CAV followers. The CAV trajectories are shown in bold and the original HDV trajectories in a faint color for reference. Note that in the original data five HDV entered within the platoon, so the 15th HDV at the start is shown with a darker color up to 700 sec and then the 15th CAV at the end is shown with the darker color.

CAV follower, and average across all 15 CAV followers for case 1. Fuel consumption is reduced by approximately 20 % by the first CAV and 40 % by the 15th CAV. In this case the largest benefits are seen in NO_x, where they dropped 80 % by the 15th CAV. Again, the specific values are highly dependent on the vehicles and prevailing conditions, so these results simply indicate the relative improvements in emissions and fuel consumption that can be achieved with the CAV control, and that the benefits increase with the number of successive CAVs.

3.2. Emergent dynamics

Ultimately this anticipatory CAV control deliberately seeks to disrupt normal HDV car following behavior. Fig. 20A shows the speed-spacing progression from the original HDV trajectory of the first follower in case 1. The progression of the first HDV is consistent with all 15 of the HDV, so for clarity only the single HDV is shown. For reference the ideal speed-spacing curve is shown with a straight line ($s_{jam} = 22\text{ft}$, $\tau = 1\text{sec}$, and $w = -15\text{mph}$). Notice how the HDV vehicle exhibits counterclockwise loops, as anticipated by Newell (1962) and subsequently shown empirically (e.g., Coifman et al., 2018). Essentially, a driver only starts decelerating after they perceive they are closer to their leader than they would prefer at the given speed, and similarly, they will only start accelerating after perceiving the gap ahead has grown. This delayed response is evident with periods of deceleration generally falling to the left of the ideal speed-spacing curve and periods of acceleration generally falling to the right.

Fig. 4A-B shows the HDV begins decelerating just after 640 sec, while the corresponding CAV starts out at a lower speed because it immediately starts responding to the forthcoming stop wave. The CAV has completed its response to the stop wave by 655 sec, so for reference in Fig. 20 the period up to 655 sec is shown with a bold curve and the rest of the time is shown with a thinner line. Fig. 20B shows the speed-spacing progression for the corresponding first CAV (1,000 ft look ahead, no limit on acceleration). Since the CAV is responding to the forthcoming stop wave from the first time step, its initial speed is below that of the HDV. The CAV maintains this lower speed in anticipation of the stop wave, creating a large gap as its faster leader pulls away. Then after achieving the largest spacing, the leader starts moving through the stop wave and the first CAV follower maintains a relatively stable speed as it consumes the excess spacing, thereby starving the stop wave's progression. In this case the looping follows a clockwise progression as the CAV returns to the ideal speed-spacing curve. Unlike the HDV, there is no pre-existing mechanism dictating the looping direction, in this case the follower had to slow because the actual delay proved to be larger than originally estimated (compare the estimated trajectory to the realized HDV trajectory in Fig. 2E).⁶ The CAV executes a similar loop near 15 mph as the leader passes through two minor disturbances visible in Fig. 14 between 660 sec and 690 sec. The CAV reaches 25 mph at the very end of the progression. To keep the presentation as clear as possible, the progression stops at 700 sec, just before the influence of the second stop wave begins, as per Fig. 16.

Fig. 20C&D show the speed-spacing progression for the second and third CAVs, respectively. In both cases the vehicles start with a large spacing that is consumed as their respective leader slows slightly while the HDV stop wave is absorbed. After passing the first stop

⁶ If the original estimate proves too conservative, the CAV will eventually speed up and could exhibit a counterclockwise loop.




Fig. 15. (A) Time series speed from the 15 CAV and the original trajectory for the first HDV follower for case 1 in Fig. 14. (B) Repeating A only now just showing the first and 15th CAVs. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first and 15th CAVs.

wave these vehicles exhibit a mix of car following and near car following behavior where they occasionally take a slightly larger spacing to refine the trajectory of their leader. The remaining 12 CAV simply continue refining these trends.

Fig. 21 repeats this analysis for case 2. Again, the HDV exhibits the expected counterclockwise looping (Fig. 21A). This time, at the start the CAVs have not yet detected the delays of the forthcoming stop wave. Instead, they are simply converging towards car




Fig. 16. Using the synthetic data to extend the coverage beyond the distance shown in Fig. 14 to show the general dynamics of a major slow wave.




Fig. 17. Case 2 with the fixed 1000 ft window and no limit on acceleration using the first follower from before and the next 14 vehicles for a total of 15 successive CAV followers. The CAV trajectories are shown in bold and the original HDV trajectories in a faint color for reference. Note that in the original data one HDV exited from within the platoon, so by the end of this figure the trajectory marked 15th HDV is in the 14th position.




Fig. 18. (A) Time series speed from the 15 CAV and the original trajectory for the first HDV follower for case 2 in Fig. 17. (B) Repeating A only now just showing the first and 15th CAVs. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first and 15th CAVs.

following. The first CAV starts responding to the stop wave just after 580 sec, while both the CAV and HDV complete their response to the stop wave just before 640 sec. As such, in each subplot the period between 580 sec and 640 sec are shown with bold curves while the rest of the progression is shown with a thinner line. Like case 1, Fig. 21B shows the first CAV follower lets its leader pull ahead as it builds up reserve spacing in preparation to arrest the stop wave's progression. The corresponding time series speed and acceleration in




Fig. 19. Fuel and emissions savings relative to the first HDV follower for the respective CAVs in case 1 under the fixed 1000 ft window with unlimited acceleration.

Fig. 18 are helpful references when interpreting Fig. 21. Note how the first HDV and CAV show divergent behavior at the start of the bolded period while the stop wave passes, the HDV in Fig. 21A initially accelerates from 20 mph into the forthcoming stop wave, while the corresponding CAV decelerates in Fig. 21B. Unlike case 1, the second and third CAVs are either car following or close to car following when the first CAV starts the anticipatory maneuvers for the stop wave. So instead of shrinking their spacing to car following (as per Fig. 20C-D), they roughly maintain their spacing as they drop their speed to match their respective leader. In other words, the subsequent CAV will rapidly respond to the first CAV slowing and then after passing the former stop wave, close the gap to gradually return to the ideal speed-spacing curve, following a clockwise progression in the speed-spacing plane with behavior that HDVs do not exhibit.

4. Impacts of the wave velocity and cHDV penetration

Up to this point the wave velocity, w , for our calculations was set to -15 mph. This setting was based on typical values reported in the literature; no effort was made to find an optimal value for the reextracted NGSIM data used in this study. Clearly the shape of the estimated trajectory in Fig. 2D depends on w , if one holds the number of cHDV constant, as w decreases generally the end point of the estimated trajectory moves further away in time and space. This skewing changes the impact of each downstream cHDV's state on the estimated trajectory. Fortunately, our method to estimate the leader's trajectory is robust to the resulting errors. Consider Fig. 2F, typically the distance to target point a_1 is large enough that a change in w translates into a small change in the slope of the line between (t_1, y_1) and a_1 , thus, yielding only a small impact on the first CAV's speed for the next time step. The fact that these estimates are recalculated at every time step, means that a wrong w will only impact the current speed calculation. Of course, if w remains incorrect the next sample will have a similar error, but the errors do not accumulate or compound over time. When going to a platoon of CAV, only the first CAV will encounter the errors directly, the dynamics for the subsequent CAV are defined by the methodology, and so they strictly adhere to the set τ and w .

To explore the impacts of w , Fig. 22A revisits the first CAV in case 1 with the fixed 1000 ft window and no limit on acceleration. Now we consider three values of w : -7 mph, -15 mph and -23 mph. The subplot on the left shows the resulting trajectories from the three scenarios and the subplot on the right shows the corresponding time series speed. The outcome is similar to what we saw when varying the look ahead distance and limits on acceleration: varying w changes the responsiveness and thus, the shape of the speed curve, but the methodology continues working as intended. Generally, $w = -7$ mph results in cautious behavior, slowing early in response to an approaching wave while $w = -23$ mph is aggressive. Fig. 22B shows the corresponding curves for the 15th CAV in case 1. In all three variants the distortions from the first CAV are dampened because the successive CAVs follow the set τ and w , so they progressively smooth the first CAV's trajectory. Regardless, all the trajectories come back together after the disturbance, so the travel time (or average speed, or total delay) does not change. Fig. 22C&D repeat this process for case 2 with similar results.

Next, the assumption of 100 % cHDV participation among the downstream HDV is a strong requirement. To explore the impacts of a lower penetration rate we consider four scenarios: 100 % of HDV are cHDV (as before), 50 % of HDV are cHDV (every other HDV), 33 % of HDV are cHDV (every third HDV), and 25 % of HDV are cHDV (every fourth HDV).⁷ Fig. 23 shows the results for case 1 and case 2, first CAV and 15th CAV with $w = -15$, fixed 1000 ft window look ahead, and no limit on acceleration. For each CAV the subplot on the

⁷ In the three reduced penetration rate scenarios we assume the lead vehicle is always detected by perception sensors onboard the first CAV and then the sequence of non-connected vehicles starts with the first HDV downstream of the CAV's leader.




Fig. 20. The progression of the speed-spacing for case 1 underlying the trajectories in Fig. 14 (A) for the original trajectory for the first HDV follower, (B) the first CAV follower, (C) second CAV follower and (D) third CAV follower. For reference the ideal speed-spacing curve is shown with a straight line in all four subplots.

left shows the resulting trajectories from the four penetration rates, with all four curves falling nearly on top of one another. While the subplot on the right shows the corresponding time series speed. The speed curves are similar, but the four penetration rates become distinct, where decreasing penetration rate corresponds to a small decrease in stability. Once more, the impacts are dampened from the first to the 15th CAV. These scenarios are idealized with the cHDV uniformly sequenced. In practice anything less than 100 % would have random patterns of cHDV arrivals, sometimes with two successive HDV being cHDV and sometimes many successive HDV not connected. We envision a different approach for managing long gaps between cHDV (the subject of ongoing research), so this exercise simply demonstrates that while the method requires a moderate penetration rate, it does not need 100 % of vehicles being connected.

5. Conclusions

This research developed an approach to enable a connected and automated vehicle (CAV) to smooth out traffic disturbances. The first CAV requires the vehicles within the look ahead region to be connected, reporting their current location and speed. These instantaneous individual vehicle states are integrated over space to forecast the trajectory of the CAV's leader. This forecast is then used to derive a "perfect follower" for the leader, which bounds the safe travel zone for the CAV that ensures the CAV will have at least 1 sec reaction time to the leader. At which point the highest constant speed radial is found between the CAV's current location and the perfect follower trajectory. The CAV takes this speed for the current time step and then the process is repeated in the next time step. The smooth CAV trajectory eliminates unexpected stops to improve safety and lessens accelerations which in turn reduces vehicle emissions and fuel consumption. These benefits are shared with any vehicle following the CAV.

The analysis showed that limiting the acceleration can reduce emissions even further without adding additional delay. Although




Fig. 21. The progression of the speed-spacing for case 2 underlying the trajectories in Fig. 17 (A) for the original trajectory for the first HDV follower, (B) the first CAV follower (note that the bold curve comes close together but never crosses itself), (C) second CAV follower and (D) third CAV follower. For reference the ideal speed-spacing curve is shown with a straight line in all four subplots.

deceleration is not limited, the methodology also reduces the magnitude of deceleration because the look ahead provides a smoother final trajectory.

Next, the analysis is expanded to explore how multiple CAVs can provide a coordinated response. Since the first CAV “overwrites” any information from the downstream trajectories, all subsequent CAVs use their leader’s planned trajectory to determine the perfect follower. After this adjustment the rest of the process is unchanged. In so doing, the subsequent CAVs react earlier in the evolution of the developing traffic state and thus, continue to smooth out the traffic state beyond what the first CAV achieved.

It is shown that the smoothness of the final trajectory depends on the look ahead distance, any limits on acceleration, the assumed wave velocity, and the number of CAVs that have conditioned the trajectory. Throughout all of the experiments the first CAV was able to nullify the stop waves. Eliminating stop waves like this should greatly reduce the risk of rear-end collisions. Meanwhile, all the scenarios exhibited a large reduction in fuel consumption and emissions.

When exploring the underlying dynamics of the individual vehicles (HDV and CAV) in the speed-spacing plane, the HDV exhibited typical counterclockwise looping around the steady state speed-spacing relationship. As soon as delays entered the look ahead region of the first CAV, it would slow slightly as its HDV leader pulled ahead, building up a cushion of spacing in anticipation of the forthcoming stop wave. Then as the HDV leader moves through the stop wave the first CAV follower maintains a relatively stable speed as it consumes the excess spacing, thereby starving the stop wave’s progression. When the first CAV starts the anticipatory maneuvers for the stop wave, if there are subsequent CAVs closely following, they roughly maintain their spacing as they drop their speed to match their respective leader. In other words, the subsequent CAV will rapidly respond to the first CAV slowing and then gradually take shorter spacing as they return to the ideal speed-spacing curve. Thus, the subsequent CAV follow a clockwise progression in the speed-spacing plane with behavior that HDVs do not exhibit.




Fig. 22. The results from three different values for the wave velocity showing the trajectories in the time-space plane on the left and corresponding time series speed on the right for, (A) first CAV in case 1, (B) 15th CAV in case 1, (C) first CAV in case 2, (D) 15th CAV in case 2. All plots used the fixed 1000 ft look ahead window and no limit on acceleration.

Ultimately, this work seeks to take conventionally unstable queued traffic and bring it to a stable flow while queued at a near constant speed that is below free speed. This study does not seek to increase capacity nor reduce delay. It is important to recognize that a queue is simply the storage of demand that exceeds the downstream capacity. So, the queued vehicles experience delay while they wait to be served. To eliminate a queue requires increasing capacity or decreasing demand. By smoothing out the disturbances this




Fig. 23. The results from four different penetration rates of cHDV showing the trajectories in the time-space plane on the left and corresponding time series speed on the right for, (A) first CAV in case 1, (B) 15th CAV in case 1, (C) first CAV in case 2, (D) 15th CAV in case 2. All plots used the fixed 1000 ft look ahead window and no limit on acceleration.

work seeks to reduce the negative impacts of congested traffic by keeping conditions as safe as possible while minimizing emissions and fuel consumption from the queued vehicles.

This study has its limits; it relies on an idealized assumption that all vehicles are either connected HDV (cHDV) or CAV. [Section 4](#) shows that the method does not need 100 % of the HDV being connected. Future work will extend the methodology to scenarios that

are more representative of the near term, where CAVs and cHDVs only make up a small fraction of the vehicles. Our ongoing research is exploring how to use the trajectory of one CAV to estimate the leader's trajectory for the next CAV in mixed traffic at a low penetration rate, without knowing anything about the HDVs that are between the two CAVs. Meanwhile, this study only considered a single lane, future work will bring in multi-lane control that includes sequencing and optimizing lane change maneuvers.

The prospect of deploying a variant of this method in traffic that includes HDV will require forethought about how to implement the system with unacquainted humans in the loop. By slowing the CAV in anticipation of a forthcoming stop wave, this approach can give rise to a large but short-lived gap between the CAV and its leader, e.g., Fig. 3A shows a roughly 200 ft gap between these two vehicles at 640 sec. For this work it is assumed that no vehicle will enter that gap from another lane. It is recognized that in real traffic such a large gap might attract HDV from the adjacent lanes, and some means of preventing those maneuvers will need to be implemented. That could be achieved via enforcement, coordination across lanes, or adjusting the CAV control. While it is an important element that needs to be addressed for deployment, it is left as future work at this current proof of concept stage.

CRedit authorship contribution statement

Benjamin Coifman: Writing – review & editing, Writing – original draft, Visualization, Supervision, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Yuan Liu:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This research was sponsored by Safety21 University Transportation Center under Funding agreement: 69A3552344811. The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, by a grant from the U.S. Department of Transportation's University Transportation Centers Program. However, the U.S. Government assumes no liability for the contents or use thereof.

Data availability

Data will be made available on request.

References

- Alessandri, A., di Febraro, A., Ferrara, A., Punta, E., 1999. Nonlinear optimization for freeway control using variable-speed signaling. *IEEE Trans. Veh. Technol.* 48, 2042–2052. <https://doi.org/10.1109/25.806796>.
- Chen, D., Ahn, S., Hegyi, A., 2014. Variable speed limit control for steady and oscillatory queues at fixed freeway bottlenecks. *Transp. Res. B Methodol.* 70, 340–358.
- Coifman, B., 2002. Estimating travel times and vehicle trajectories on freeways using dual loop detectors. *Transp. Res. A* 36 (4), 351–364.
- Coifman, B., 2015. Empirical flow-density and speed-spacing relationships: evidence of vehicle length dependency. *Transportation Research- Part B* 78, 54–65.
- Coifman, B., Li, L., 2022. A new method for validating and generating vehicle trajectories from stationary video cameras. *IEEE Trans. Intell. Transp. Syst.* 16227–16236.
- Coifman, B., Li, L., 2024. Partial trajectory method to align and validate successive video cameras for vehicle tracking. *Transp. Res. C* 158.
- Coifman, B., Wang, Y., 2005. Average velocity of waves propagating through congested freeway traffic. In: *Transportation and Traffic Theory. Flow, Dynamics and Human Interaction*. University of Maryland, College Park, pp. 165–179.
- Coifman, B., Li, L., Xiao, W., 2018. Resurrecting the lost vehicle trajectories of treiterer and myers with new insights into a controversial hysteresis. *Transp. Res. Rec.* 2672 (20), 25–38.
- Di Vaio, M., Fiengo, G., Petrillo, A., Salvi, A., Santini, S., Tufo, M., 2019. Cooperative shock waves mitigation in mixed traffic flow environment. *IEEE Trans. Intell. Transp. Syst.* 20, 4339–4353. <https://doi.org/10.1109/TITS.2018.2883485>.
- Donà, R., Mattas, K., He, Y., Albano, G., Ciuffo, B., 2022. Multianticipation for string stable adaptive cruise control and increased motorway capacity without vehicle-to-vehicle communication. *Transportation research part C: emerging technologies*, 140.
- Du, Y., Makridis, M., Tampère, C., Kouvelas, A., ShangGuan, W., 2023. Adaptive control with moving actuators at motorway bottlenecks with connected and automated vehicles. *Transp. Res. Part C: Emerg. Technol.* 156.
- Ghiassi, A., Li, X., Ma, J., 2019. A mixed traffic speed harmonization model with connected autonomous vehicles. *Transp. Res. Part C: Emerg. Technol.* 104, 210–233.
- Goatin, P., 2024. Dissipation of stop-and-go waves in traffic flows using controlled vehicles: a macroscopic approach. *IEEE Control Syst. Lett.* 8, 628–633.
- Han, Y., Chen, D., Ahn, S., 2017a. Variable speed limit control at fixed freeway bottlenecks using connected vehicles. *Transp. Res. B Methodol.* 98, 113–134.
- Han, Y., Hegyi, A., Yuan, Y., Hoogendoorn, S., Papageorgiou, M., Roncoli, C., 2017b. Resolving freeway jam waves by discrete first-order model-based predictive control of variable speed limits. *Transp. Res. Part C Emerging Technol.* 77, 405–420.
- Han, Y., Yu, H., Li, Z., Xu, C., Ji, Y., Liu, P., 2021. An optimal control-based vehicle speed guidance strategy to improve traffic safety and efficiency against freeway jam waves. *Accid. Anal. Prev.* 163.
- Han, Y., Wang, M., He, Z., Li, Z., Wang, H., Liu, P., 2021. A linear Lagrangian model predictive controller of macro-and micro-variable speed limits to eliminate freeway jam waves. *Transp. Res. Part C Emerging Technol.*, 128.
- He, Z., Zheng, L., Song, L., Zhu, N., 2016. A jam-absorption driving strategy for mitigating traffic oscillations. *IEEE Trans. Intell. Transp. Syst.* 18 (4), 802–813.
- Hegyi, A., Hoogendoorn, S., 2010. Dynamic speed limit control to resolve shock waves on freeways-field test results of the SPECIALIST algorithm. In: *13th International IEEE Conference on Intelligent Transportation Systems*. IEEE, pp. 519–524.

- Hegyi, A., Schutter, B.D., Hellendoorn, J., 2005. Optimal coordination of variable speed limits to suppress shock waves. *IEEE Trans. Intell. Transp. Syst.* 6, 102–112. <https://doi.org/10.1109/TITS.2004.842408>.
- Ibrahim, A., Čičić, M., Goswami, D., Basten, T., Johansson, K.H., 2019. Control of platooned vehicles in presence of traffic shock waves, *Proc. 2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 1727–1734. Doi: 10.1109/ITSC.2019.8917389.
- Jiang, L., Xie, Y., Evans, N., Wen, X., Li, T., Chen, D., 2022. Reinforcement learning based cooperative longitudinal control for reducing traffic oscillations and improving platoon stability. *Transp. Res. Part C Emerging Technol.* 141.
- Li, Z., Levin, M., Qu, X., Stern, R., 2023. A network traffic model for the control of autonomous vehicles acting as moving bottlenecks. *IEEE Trans. Intell. Transp. Syst.* 24 (9), 9004–9015.
- Li, S., Yanagisawa, D., Nishinari, K., 2024. A jam-absorption driving system for reducing multiple moving jams by estimating moving jam propagation. *Transp. Res. Part C: Emerg. Technol.*, 158.
- Ma, J., Li, X., Shladover, S., Rakha, H.A., Lu, X.-Y., Jagannathan, R., Dailey, D.J., 2016. Freeway speed harmonization. *IEEE Trans. Intell. Veh.* 1 (1), 78–89.
- Newell, G., 1962. Theories of instability in dense highway traffic. *Journal of Operation Research Society of Japan* 5 (1), 9–54.
- Newell, G., 2002. A simplified car-following theory: a lower order model. *Transp. Res. B Methodol.* 36 (3), 195–205.
- Nishi, R., Tomoda, A., Shimura, K., Nishinari, K., 2013. Theory of jam-absorption driving. *Transp. Res. B Methodol.* 50, 116–129.
- Qin, Y., Wang, H., Ni, D., 2021. Lighthill-Whitham-Richards model for traffic flow mixed with cooperative adaptive cruise control vehicles. *Transp. Sci.* 55, 883–907. <https://doi.org/10.1287/trsc.2021.1057>.
- Scora, M., Barth, M., 2006. Comprehensive modal emission model (CMEM) Version 3.01 User's Guide. https://www.cert.ucr.edu/sites/g/files/rcwecm1251/files/2019-07/CMEM_User_Guide_v3.01d.pdf [accessed on July 1, 2024].
- Shladover, S.E., Su, D., Lu, X.-Y., 2012. Impacts of cooperative adaptive cruise control on freeway traffic flow. *Transp. Res. Rec.* 2324, 63–70. <https://doi.org/10.3141/2324-08>.
- Stern, R.E., Cui, S., Delle Monache, M.L., Bhadani, R., Bunting, M., Churchill, M., Hamilton, N., Haulcy, R., Pohlmann, H., Wu, F., Piccoli, B., Seibold, B., Sprinkle, J., Work, D.B., 2018. Dissipation of stop-and-go waves via control of autonomous vehicles: field experiments. *Transp. Res. Part C Emerging Technol.* 89, 205–221. <https://doi.org/10.1016/j.trc.2018.02.005>.
- Tan, Y., Zhang, K., 2022. Real-time distributed cooperative adaptive cruise control model considering time delays and actuator lag. *Transp. Res. Record: J. Transp. Res. Board* 2676, 93–111. <https://doi.org/10.1177/03611981221091762>.
- Taniguchi, Y., Nishi, R., Ezaki, T., Nishinari, K., 2015. Jam-absorption driving with a car-following model. *Physica A* 433, 304–315.
- van Arem, B., van Driel, C.J.G., Visser, R., 2006. The impact of cooperative adaptive cruise control on traffic-flow characteristics. *IEEE Trans. Intell. Transp. Syst.* 7, 429–436. <https://doi.org/10.1109/TITS.2006.884615>.
- Vander Werf, J., Shladover, S., Miller, M., Kourjanskaia, N., 2002. Effects of adaptive cruise control systems on highway traffic flow capacity. *Transp. Res. Rec.* 1800 (1), 78–84.
- Wang, Y., Jin, P.J., 2023. Model predictive control policy design, solutions, and stability analysis for longitudinal vehicle control considering shockwave damping. *Transp. Res. Part C Emerging Technol.* 148, 104038. <https://doi.org/10.1016/j.trc.2023.104038>.
- Wang, M., Daamen, W., Hoogendoorn, S.P., van Arem, B., 2016a. Cooperative car-following control: distributed algorithm and impact on moving jam features. *IEEE Trans. Intell. Transp. Syst.* 17, 1459–1471. <https://doi.org/10.1109/TITS.2015.2505674>.
- Wang, M., Daamen, W., Hoogendoorn, S., van Arem, B., 2016b. Connected variable speed limits control and car-following control with vehicle-infrastructure communication to resolve stop-and-go waves. *J. Intell. Transp. Syst.* 20 (6), 559–572.
- Wang, S., Shang, M., Levin, M., Stern, R., 2023. A general approach to smoothing nonlinear mixed traffic via control of autonomous vehicles. *Transp. Res. Part C: Emerg. Technol.* 146.
- Zheng, Y., Zhang, G., Li, Y., Li, Z., 2020. Optimal jam-absorption driving strategy for mitigating rear-end collision risks with oscillations on freeway straight segments. *Accid. Anal. Prev.* 135.
- Zhou, H., Zhou, A., Li, T., Chen, D., Peeta, S., Laval, J., 2022. Congestion-mitigating MPC design for adaptive cruise control based on Newell's car following model: history outperforms prediction. *Transp. Res. Part C Emerging Technol.* 142, 103801. <https://doi.org/10.1016/j.trc.2022.103801>.
- Zhou, Z., Li, L., Qu, X., Ran, B., 2024. PACC: a platoon-based adaptive cruise control strategy based on leader-following information topology to mitigate traffic oscillations under CAV environment. *Physica A* 130117. <https://doi.org/10.1016/j.physa.2024.130117>.