Safety21

INNOVATING FOR SAFETY

US DOT National
University Transportation Center for Safety

Carnegie Mellon University

Penn University of Pennsylvania

The Ohio State University

Morgan State University

CCAC Community College of Allegheny County

The University of Texas Rio Grande Valley

Safe and Efficient Automated Freeway Traffic Control- Phase 3 Final Report

Benjamin Coifman, https://orcid.org/0000-0002-8201-964X

Yuan Liu, https://orcid.org/0009-0005-4330-8923

Ohio State University

Final Report – Aug 31, 2026

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
Final Report - 575		

4. Title and Subtitle Safe and Efficient Automated Freeway Traffic Control- Phase 3 Final Report	**5. Report Date** August 31, 2026 **6. Performing Organization Code**
7. Author(s) Benjamin Coifman, PhD, https://orcid.org/0000-0002-8201-964X Yuan Liu, https://orcid.org/0009-0005-4330-8923	**8. Performing Organization Report No.**
9. Performing Organization Name and Address Ohio State University Department of Civil Environmental and Geodetic Engineering 2070 Neil Ave Hitchcock Hall 470 Columbus, OH 43210	**10. Work Unit No.** **11. Contract or Grant No.** Federal Grant # 69A3552344811 / 69A3552348316
12. Sponsoring Agency Name and Address Safety21 University Transportation Center Carnegie Mellon University 5000 Forbes Avenue Pittsburgh, PA 15213	**13. Type of Report and Period Covered** Final Report (July 1, 2025 – July 31, 2026) **14. Sponsoring Agency Code** USDOT

15. Supplementary Notes Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.

16. Abstract

The goal of this work is to eliminate unexpected stops to improve safety with the added benefit of reducing accelerations to improve fuel efficiency. Congested traffic is characterized by signals and waves propagating upstream through the queued traffic. Freeway drivers do not expect to encounter abrupt drops in speed or stopped traffic, as a result, shockwaves sharply increase the accident rates, particularly in the context of rear end collisions. This work seeks to use connected and automated vehicles (CAV) to smooth out traffic disturbances on a freeway. Prior work in Phase 1 developed a method for a single CAV to integrate 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. Phase 2, extended the methodology to a platoon of 15 CAV's, leading to greater smoothing than the single CAV results.

This study, Phase 3, extended the methodology to address several key aspects. First, to use multiple platoons to buffer against capacity fluctuations and provide finer control. Second, extending the methodology to accommodate lane change maneuvers so that they do not disrupt platoons. Third, develop contingency plans to accommodate situations when the state evolution deviates from forecasts. In this way, the research anticipates and develops the management strategies needed to maintain smoothly flowing queued traffic with CAV's.

17. Key Words	**18. Distribution Statement**
Connected and Automated Vehicles; Freeway Traffic; Congestion; Safety; Stop Waves	No restrictions. This document is available through the National Technical Information Service, Springfield, VA 22161. Enter any other agency mandated distribution statements. Remove NTIS statement if it does not apply.

19. Security Classif. (of this report)	**20. Security Classif. (of this page)**	**21. No. of Pages**	**22. Price**
Unclassified	Unclassified	25	

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Safe and Efficient Automated Freeway Traffic Control

Phase 3 Final Report

Benjamin Coifman, PhD

Professor
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
P: (614) 292-4282
E: Coifman.1@OSU.edu

ORCHID: 0000-0002-8201-964X

Yuan Liu

Graduate Research Associate
The Ohio State University
Department of Electrical and Computer Engineering

ORCHID: 0009-0005-4330-8923

Funding agreement: 69A3552344811/69A3552348316

DISCLAIMER

Abstract

The goal of this work is to eliminate unexpected stops to improve safety with the added benefit of reducing accelerations to improve fuel efficiency. Congested traffic is characterized by signals and waves propagating upstream through the queued traffic. Freeway drivers do not expect to encounter abrupt drops in speed or stopped traffic, as a result, shockwaves sharply increase the accident rates, particularly in the context of rear end collisions. This work seeks to use connected and automated vehicles (CAV) to smooth out traffic disturbances on a freeway. Prior work in Phase 1 developed a method for a single CAV to integrate 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. Phase 2, extended the methodology to a platoon of 15 CAV's, leading to greater smoothing than the single CAV results.

This study, Phase 3, extended the methodology to address several key aspects. First, to use multiple platoons to buffer against capacity fluctuations and provide finer control. Second, extending the methodology to accommodate lane change maneuvers so that they do not disrupt platoons. Third, develop contingency plans to accommodate situations when the state evolution deviates from forecasts. In this way, the research anticipates and develops the management strategies needed to maintain smoothly flowing queued traffic with CAV's.

Keywords

Connected and Automated Vehicles; Freeway Traffic; Congestion; Safety; Stop waves

1. Introduction

This work seeks to use connected and automated vehicles to smooth out traffic disturbances on uninterrupted flow facilities like freeways, eliminating unexpected slowing and stops to improve safety while reducing accelerations to improve fuel efficiency because an internal combustion engine is more efficient cruising at a constant speed than when it is accelerating. Connected and autonomous vehicles (**CAV's**) hold the promise to attenuate and eliminate shockwaves, but only if the system is explicitly designed to do so. The very factors that give rise to shockwaves in human driven vehicles (**HDV's**) can also do so in CAV's. While CAV's offer new ways to manage traffic dynamics, an automated freeway will still be subject to traffic dynamics.

This research essentially seeks to take conventionally unstable queued traffic and have CAV's proactively respond to disturbances propagating upstream through the traffic to stabilize the flow. The approach assumes all vehicles are connected and report their current state (location and speed) to other vehicles. The main objective is to integrate the state across vehicles so that the system can efficiently anticipate and respond to disturbances over large distances. It is this look ahead that will allow the system to detect and attenuate shockwaves, resulting in smoother conditions for safety and efficiency. It is also important to recognize the physical limits, namely, that a queue simply represents the storage of demand in excess of capacity as the excess vehicles wait to be served. To eliminate a queue requires increasing capacity or decreasing demand. 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 at a near constant speed that is below free speed.

This report is for Phase 3 of a multi-phase study. The previous phases of this work focused on the initial transition from unstable HDV stop and go traffic to the first platoon of CAV's with smooth trajectories, i.e., dissipating large shockwaves after they have formed and begun propagating. The methodology is demonstrated using microscopic vehicle trajectory data from real shockwaves in HDV traffic as both the initial conditions and bounding constraints of how the system can respond. Phase 1 developed a method for a single CAV to integrate the instantaneous state information from the downstream vehicles to forecast the trajectory of the CAVs leader and proactively respond to changes in state that have not yet reached the lead vehicle (Coifman and Liu, 2024). Phase 2, extended the methodology to a platoon of 15 CAV's, leading to greater smoothing than the single CAV results of the prior work (Coifman and Liu, 2025). Collectively, the previous phases showed that the system can rapidly nullify a stop wave. If all vehicles were CAV's operating under this approach, then stop waves would rarely arise and when one did form, it could quickly be dissipated.

The Phase 3 work in this report starts with the assumption that CAV's can stabilize flow, as per Phase 1-2, and then seeks to optimize the flow of an indefinite number of CAV's. This work requires 100% penetration of CAV's, i.e., no HDV. Like the previous phases, we use the empirical HDV trajectories from the reextracted I-80 NGSIM data (Coifman and Li, 2024) to determine the underlying demand and capacity. But because the CAV platoons have proven to be very effective in stabilizing the flow, we now use the HDV to determine the cumulative number of CAV that could have departed by the given time step but do not seek to directly follow the precise arrivals or departures of the HDV. In the envisioned full CAV deployment control of a given link would leverage knowledge of conditions far downstream of the link. With the NGSIM dataset we do not have this information. So, instead, we use the HDV departure curve at the downstream end of the link as a proxy for the unobserved downstream conditions. Thus, when a CAV enters the link, we already know the earliest time it can leave the link via the HDV departure curve. But we always have the liberty to hold the vehicle later than its earliest feasible departure time. One downside of using the HDV departure curve is that we preserve the distortions in demand due to the stop waves, but the method can handle those issues, so the results herein are probably more extreme than the expected performance if the downstream links were actually under similar CAV control.

We structure the CAV traffic stream to be a sequence of high-flow platoons and no-flow gaps, as illustrated in Fig. 1B. HDV only maintain states on the fundamental diagram (FD) curve in Fig. 1A, but CAV can also maintain states inside the FD. So, if downstream traffic were restricted to flow q_N HDV would take state N, with speed v_N, equal to the slope of the radial to point N. With the gaps, the platoon can take some higher speed state, e.g., A, at v_A, where the radial to A has a higher slope than the radial to N. But q_A is greater than the downstream capacity at q_N. Which is where the gap comes in. The CAV traffic is essentially engineered to exhibit pulse width modulation (**PWM**), with the high speed and flow pulse at A, balanced by no flow of state O, to yield an average flow of q_N over period "a" shown in Fig. 1B and yielding the average state a in Fig. 1A.

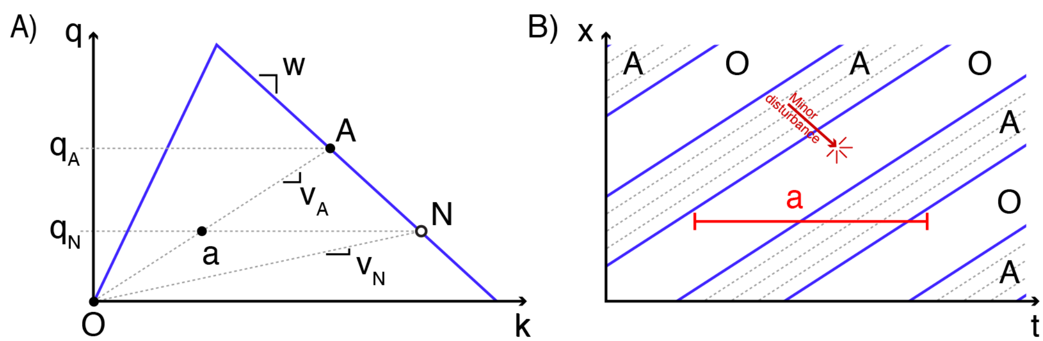

Figure 1, (A) Fundamental diagram of the traffic, (B) example of the platoon/gap structure.

This PWM platooning provides greater flexibility to respond to changing conditions. The gaps naturally absorb disturbances that form within a single platoon, preventing minor disturbances from impacting the other platoons (as illustrated in Fig. 1B). More generally, the gaps provide a buffer where a platoon's dynamics can be changed without impacting any vehicles outside of the platoon. So, if downstream demand exhibits a small unanticipated change, the next platoon could be sped up or slowed down in response without immediately impacting platoons further upstream, and if there is ultimately an impact to the platoons further upstream, the lead time provides the same smoothing benefit that was demonstrated in the Phase 1&2 work. Although we do not demonstrate it in the present study, the gaps between platoons also allow for pass through lane change maneuvers, e.g., from lane 3 to 1, without interacting with the lane 2 vehicles.

In formulating the platoon strategy, we seek to ensure the last vehicle of the platoon departs on time, as per the HDV departure curve. We call this last vehicle the **"key vehicle"** and it departs at time t_f. All preceding vehicles are then scheduled so that they pass one ideal headway at the given state before their follower. For example, suppose a platoon was departing the link operating at state A in Fig. 1A. The second to the last vehicle would be scheduled to depart at $t_f - h_A$, where $h_A = 1/q_A$. The vehicle before that would depart at $t_f - 2h_A$ and so forth. Since the platoon in state A is operating at a higher flow than the HDV at state N, all vehicles in the platoon before the last vehicle are slightly delayed. So, there is a small local cost to create the gap at state O ahead of the platoon, but this local cost is limited to the vehicles in the platoon, it does not accumulate across platoons. In this way, we are using the key vehicle to determine the dynamics of all preceding vehicles in the platoon. Also, it is important to remember that the total flow remains unchanged from the HDV, but the speeds are stabilized.

Next, we adjust the PWM platoons to accommodate lane change maneuvers, **LCM**, so that the inflows and outflows do not disrupt the platoons. To this end, we treat LCM's as an exogenous input that is determined by the HDV trajectory data. Thus, necessitating that we allow the vehicles that enter the lane to change their order relative to other vehicles that started in the lane. This modification is necessary because the new CAV trajectories are completely different than the underlying HDV trajectories used to determine demand and capacity.

In this scheme, the LCM's are effectively random disruptions to a given platoon. We decided to preserve the target lane vehicle and location of each HDV LCM, which typically necessitates changing the time of the maneuver and source lane vehicle. The logic behind this choice is that in a truly efficient CAV traffic stream the synchronization of LCM's would be a complex process that balances origin-destination (OD) flows, lane utilization, and other constraints. While this lateral flow control is a critical element of CAV operations, it is a large research endeavor in itself and requires a move to strictly synthetic data across an extended corridor to establish a meaningful set of OD flows. In the meantime, the methods we develop to accommodate the effectively random LCM's in the present work should be fully compatible with deliberately planned LCM's to accommodate OD flows and lane utilization.

Specifically, any entering vehicles are sequenced at the end of the platoon, with the last entering vehicle becoming the key vehicle. Thus, the last vehicle that started in the lane is moved forward by the number of vehicles that enter the lane. While many past studies focus on platoon splitting, we deliberately avoid such an approach. If the CAV's have been consolidated into a tight platoon, creating a gap will simply add delay to all vehicles upstream of that gap. An equal amount of delay will be incurred if the entering vehicle is forced to join the platoon at the rear rather than in the middle, but the dynamics are far simpler, and the delays are only experienced by that vehicle that changes lanes. In practice, we expect that adjacent lanes would travel at slightly different speeds, in which case, the exact time and location of a LCM might have to change so that the vehicle can leave the source lane just as the rear of a platoon in the target lane passes. But CAV flows should remove the human dynamics of weaving in and out of lanes of congested traffic to get to the locally fastest lane. So, it becomes easier to swap one CAV for another when the objective is simply to balance lane utilization without interfering with the net OD flows. In the worst case, in practice, it might be necessary to occasionally split platoons to create a gap at just the right point to facilitate a critical exchange of vehicles between lanes.

The final element of this work is developing contingency plans to accommodate unforeseen events where the state evolution deviates from forecasts. Forecasts are never perfect, whether it is due to a small drift from the assumed progression or an unanticipated large deviation like a blown tire or an object falling from a truck, the CAV control scheme must be able to handle these unanticipated events. We show how our methods can adjust in the presence of imperfect predictions.

The remainder of this document is as follows. Section 2 presents the methodology, starting with the underlying theory and then empirically demonstrating the approach using microscopic vehicle trajectory data. Then this report closes with a brief discussion and conclusions in Section 3.

2. Methodology

The methodology is split between platooning the CAV in Section 2.1 and addressing contingency plans in Section 2.2.

2.1. Platooning

This work is focused on platooning the CAV's, with the goal of having the last vehicle in each platoon depart the link precisely on time. To this end, the vehicles for each platoon are first grouped based on the number of vehicles departing, which comprise the through vehicles and those vehicles that entered the lane via LCM's. We chose to aim for a platoon length of 15 simply because it was large enough to form meaningful gaps between platoons while being small enough that large disturbances are less likely to form within a given platoon. Obviously, disturbance formation is a function of the CAV dynamics, which are mostly beyond the scope of this proof of concept study. Ultimately, a dynamic platoon size would likely be a robust option.

Because we are designing to target specific departures, the number of vehicles that exit the link via LCM's does not factor into the platoon size constraint for the given lane. Furthermore, vehicles that pass through the lane (i.e., entering and exiting via LCM) within the link are excluded, they are to be routed between platoons.

Fig. 2A shows an example of successive HDV's from the NGSIM data that will eventually be grouped into a single CAV platoon. We start with the first vehicle after the last completed platoon, and count off a total of 15 departures, i.e., vehicles crossing the cutline at 1,700 ft near the end of the segment. If the 15th vehicle is not a through vehicle, we continue adding vehicles to the group until we get to a through vehicle. This restriction on using a through vehicle keeps the subsequent steps simple, but obviously it could be relaxed without many complications.

In any event, the last through vehicle is taken to define the key vehicle trajectory, which is done by finding the straight line trajectory from its point of entry to its point of departure, as per Fig. 2B. Each vehicle that enters the lane via a LCM is given a virtual trajectory from the start of the link to the point of entry.[1] All of the entering vehicles are then moved to the end of the platoon and sorted by the location of their respective LCM, where the first of these vehicles is the one that entered most upstream in the link, and each successive vehicle entered further downstream of its leader.

The rationale for sorting the entering vehicles this way is the fact that a LCM takes more space to execute than simple car following. With each LCM at the rear of the platoon the new vehicle makes their entrance in the gap, slightly upstream of the optimal location, then they undertake the necessary speed and spacing adjustments to correct for the speed differential between lanes and achieve the tight spacing of the platoon. Doing so in the gap at the end of the platoon minimizes the potential disruption to the following vehicles.

The last of the entering vehicles is placed in the key vehicle slot, and the previous vehicle is placed in the "perfect leader" position (i.e., the key vehicle's trajectory shifted earlier in time and space via Newell, 2002). Each subsequent leader is then placed in the perfect leader position. After all of the entering vehicles have been allocated their slots (Fig. 2C), all through and exiting vehicles are added at the front of the platoon, retaining their relative arrival order from the start of the link.

[1] As discussed in Section 1, the location of the HDV LCM is preserved, but the time is allowed to change. The presumption is that in centrally controlled CAV flows, LCM's will primarily be for balancing lane utilization. So, in this context, the number of vehicles that leave the lane is important, but which vehicle changes lanes is not. When dealing with adjacent lanes like this, the entering vehicle would simply be whatever vehicle was immediately upstream of the entering location at the time of the LCM.

Like the vehicles that entered via LCM, each added vehicle is done so with a perfect leader spacing. The result is shown in Fig. 2D.

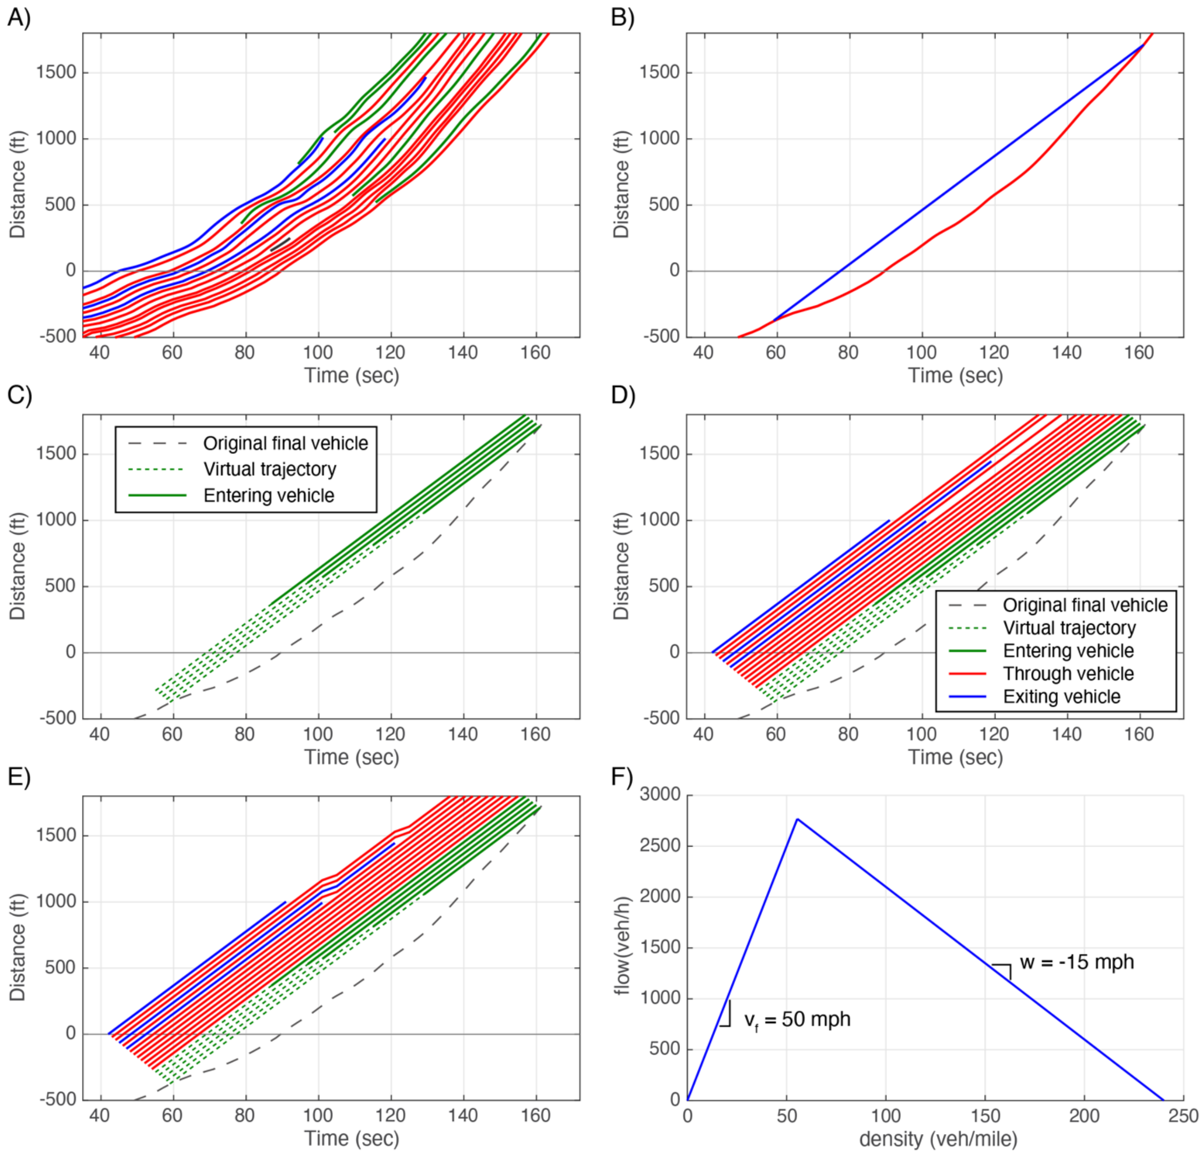

Figure 2, Platoon formation: (A) Counting out all vehicles until reaching the 15th through/entering vehicle. (B) Finding the straight line between the start and end points for the final vehicle, establishing the trajectory for the key vehicle. (C) Put all entering vehicles at the rear of the platoon, sorted by the location where they enter. (D) Then add all through and exiting vehicles that passed the start of the link, without reordering them. (E) Close the gaps left behind each vehicle exit. (F) The FD used throughout this example.

Next, the gap downstream of each exiting vehicle is closed using the following technique. Leveraging the fact that the CAV's do not behave like HDV's, we have all of the vehicles in the platoon downstream of the exit slowdown in unison, taking state B in Fig. 3A, and then return to the original state A once the gap is closed. This maneuver is possible because the CAV can exhibit inhuman behavior and take states inside the FD curve. From a traffic flow theory perspective, the vehicles downstream of the exit will drop from the speed of state A to the speed of state B together, maintaining the density, k, and giving rise to an interface that travels instantaneously, as shown in Fig. 3B. The vehicle immediately behind the exit sees its headway double and then while the downstream vehicles are moving slower, they close the gap and return to the original state A. Since this vehicle maintains their speed throughout, its state A' evolves on the radial at the speed of state A. There is no problem if the response to two (or more) exits overlap, the accommodation period simply lasts longer to close the two gaps. The net result for the ongoing example of Fig. 2 is the final platoon shown in Fig. 2E.

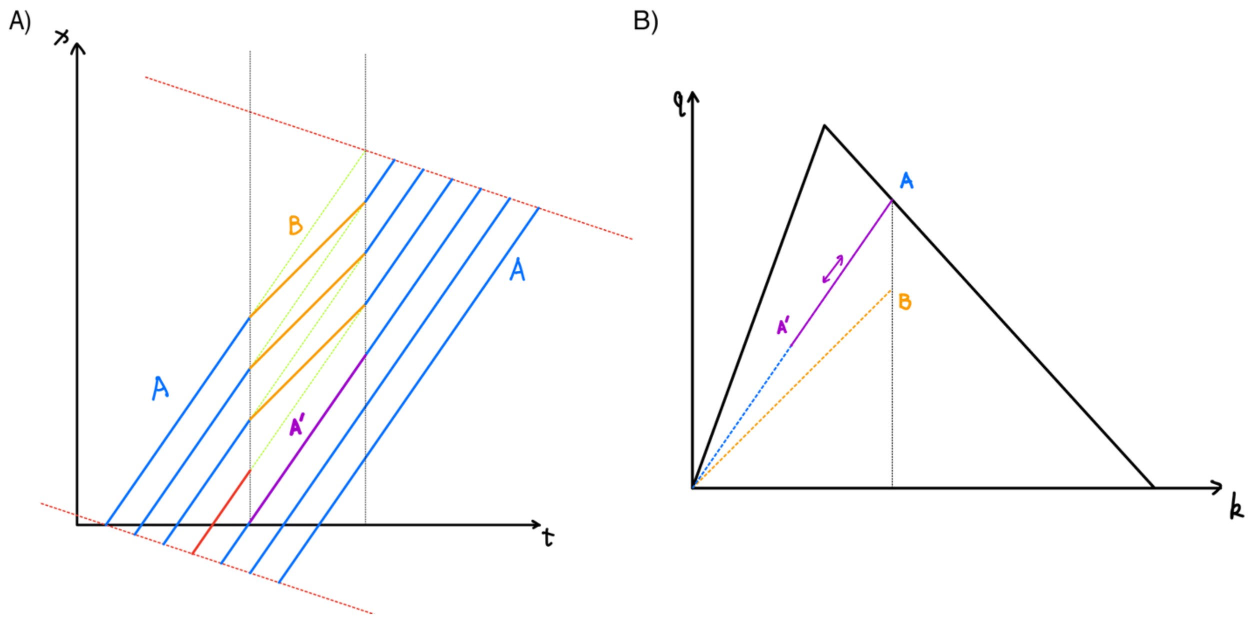

Figure 3, (A) Schematic showing the state changes used to close the gap left after an exiting vehicle left the lane. (B) The corresponding states in the flow density plane.

After the platoon is formed, it is checked against the previous platoon to ensure it never overtakes the safety boundary of its predecessor. As long as there are no disruptions to the departure rate and each platoon is able to leave on time, the following platoon should not overtake the preceding platoon. This safety check remains necessary because sometimes forecasts are wrong. If the following platoon catches up to its leader, the following platoon will slow and temporarily join the leading platoon. The platoons will separate when the downstream departure rates (as seen in downstream links) will allow the lead platoon to pull ahead. Once the platoon's path has been finalized, we find the safety boundary for this platoon. At any given location the boundary is set to be the "perfect follower" spacing (again, shifting the trajectory via Newell, 2002). Note that this step ignores the virtual trajectories of any vehicle that has not yet entered the lane.

Now consider the final results for the reextracted NGSIM data. Fig. 4 shows an example of the first six platoons in lane 2. The top plot shows the original HDV trajectories, with black curves denoting the last vehicle allocated to each successive platoon. The bottom plot shows the corresponding CAV platoons that emerge from the HDV arrival and departure pattern. Note that throughout this report, all of the CAV platoons derived from the NGSIM data are based on the FD in Fig. 2F.

Fig. 5 shows the resulting platoons in lane 2 during the entire available time period, and Fig. 6 repeats the presentation for lane 3. The varying gaps between successive platoons reflect artifacts from the HDV arrivals and departures. This unevenness arises from when and where delays impact the HDV's. The downstream departure pattern maintains the variable rates from the HDV traffic (including stop waves). For example, the key vehicle in one platoon would exhibit a slow speed if the original HDV passed through a stop wave, meanwhile, the key vehicle in the next platoon will take a much higher speed if the stop wave exited the link before that key vehicle entered. This outcome is one of the liabilities of using the empirical HDV trajectories, but this work shows the method is robust to the fact that fluctuations in HDV demand modulate the CAV traffic dynamics. If we had CAV control of the downstream link the departures would be more uniform since the uneven delays from stop waves would have been dissipated from the traffic stream and the more uniform departures would be even easier to control.

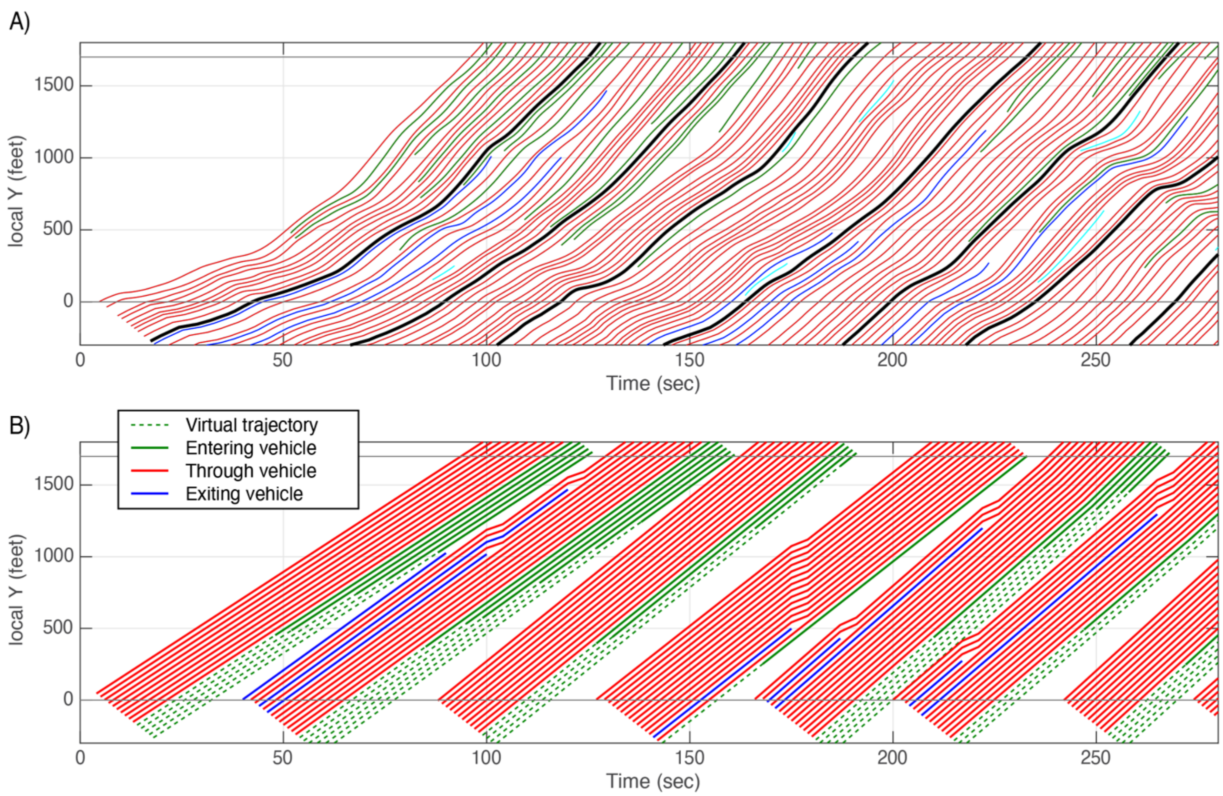

Figure 4, Using lane 2 data, (A) the original HDV trajectories from the reextracted NGSIM data, and (B) the corresponding CAV platoons from our method. Note that the black lines in the top plot correspond to the last vehicle in the corresponding CAV platoon.

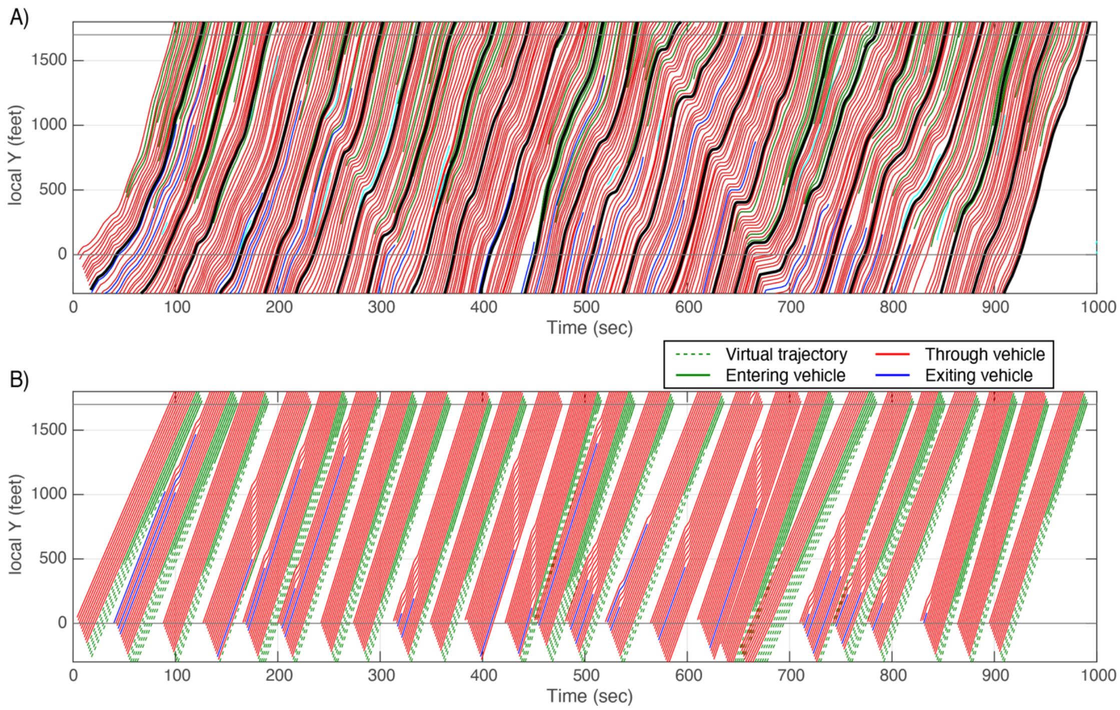

Figure 5, Using lane 2 data and zooming out from the previous figure, (A) the original HDV trajectories from the reextracted NGSIM data, and (B) the corresponding CAV platoons from our method for the entire time period.

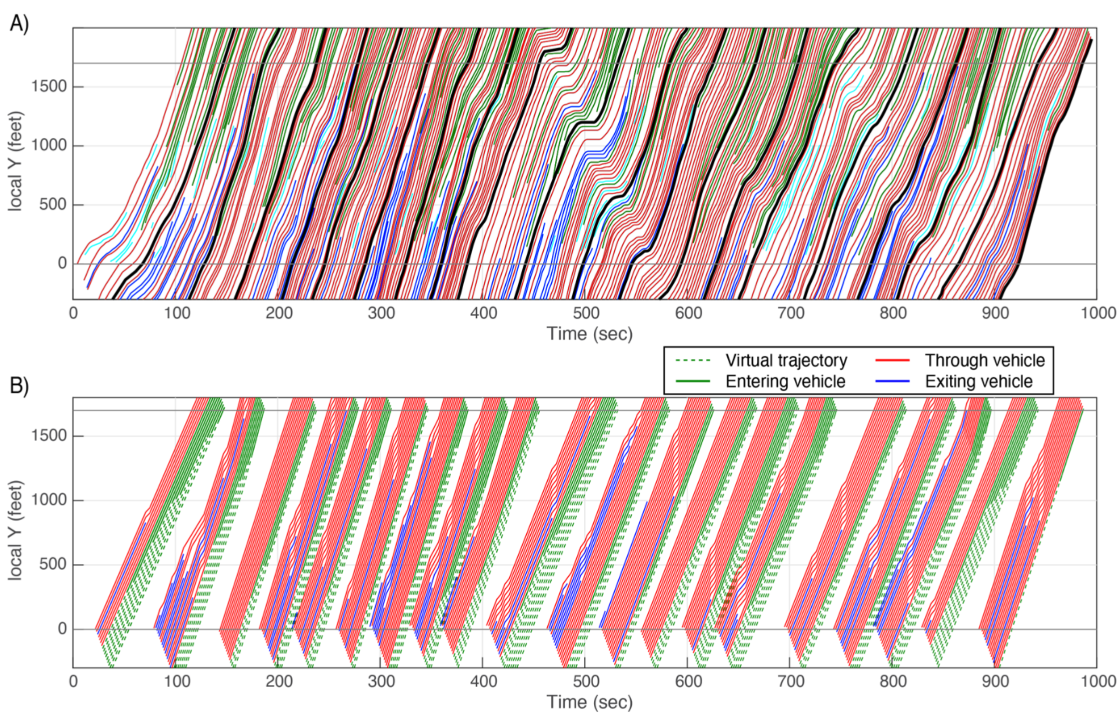

Figure 6, Using lane 3 data, (A) the original HDV trajectories from the reextracted NGSIM data, and (B) the corresponding CAV platoons from our method for the entire time period.

2.2. Contingency Plans

The platoon formation in the previous section relies upon the forecasted departure curve. Under normal conditions this assumption is reasonable for the envisioned system where an extended corridor is under central coordination, and all vehicles follow prescribed trajectories. But even the best forecasts are imperfect. For example, the forecasts would fail if a CAV had to undertake an unexpected maneuver, e.g., in response to debris in the road, a flat tire, or sensor failure. This section develops a suite of contingency plans for when the state progression diverges from the forecasts. We specifically focus on situations where vehicles need to move slower than originally forecast, i.e., unanticipated delays, because these events represent a safety hazard.

We start with the base case, at some time, t_s, during the period that a well-formed platoon (e.g., Fig. 2E) is traversing the segment and an updated forecast postpones the departure time. In particular, the new forecast reports that the first vehicle in the platoon can now depart at some time, t_r, that is later than its original departure time. In the ***base case***, this t_r proves to be accurate. Next, in the ***extended case***, the forecasted t_r proves to be too early and the platoon must be further delayed. Then in the ***truncated case*** the forecasted t_r proves to be too restrictive and the platoon can be released earlier than anticipated. Next, in the ***slowed case*** the forecasted t_r proves to be correct but the speed after the departure is slower than anticipated. Each of these cases are discussed in their own section below. Some of these cases include the following platoons, but not all of them do. So, we also include a separate section ***impacts on the following platoon*** that explicitly considers the following platoons. The discussion of contingency plans closes with a section providing overall comments that apply to the various cases.

2.2.1. Base Case

At t_s the revised forecast indicates that the first vehicle in the platoon can now depart at t_r and in the base case, ultimately this revised forecast proves to be accurate. However, the final outcome is not known at t_s, only the revised forecast. So, at t_s all vehicles in the platoon are slowed, with each CAV after the first taking a speed slightly faster than their leader. Thereby giving rise to an instantaneous transition at t_s and a controlled transition to the final state beginning at t_r, as shown in Fig. 7. Specifically, the first vehicle in the platoon is slowed from v_A to v_{B1} so that it will depart the link at exactly t_r and simultaneously take on the new speed v_C after leaving the link. Initially, all vehicles in the platoon are uniformly spaced corresponding to the density of state A. The second vehicle in the platoon aims to become a perfect follower (as per Newell, 2002) in state C as it crosses the signal chord emanating from the lead vehicle at t_r that travels with the **wave velocity, w**, corresponding to the slope of the congested regime. To do so, the second vehicle must move faster than its leader to close the gap, taking v_{B2}. In terms of the FD, Fig. 7B shows that the second CAV transitions instantaneously from state A to state B_2, and then maintaining that speed, slowly closes the gap and progresses to state $B_2{}'$ that falls on the FD. The progression is timed so that at the instant the second vehicle reaches the FD it instantaneously transitions to the final state C. This process continues for each successive vehicle in the platoon, with $v_{B1} < v_{B2} < v_{B3} < \cdots$. In this nomenclature, one could conceive of states B_1 and $B_1{}'$, but because the first vehicle has a large gap by design, that vehicle's state would generally fall somewhere on the radial at v_{B1} far inside of the FD.

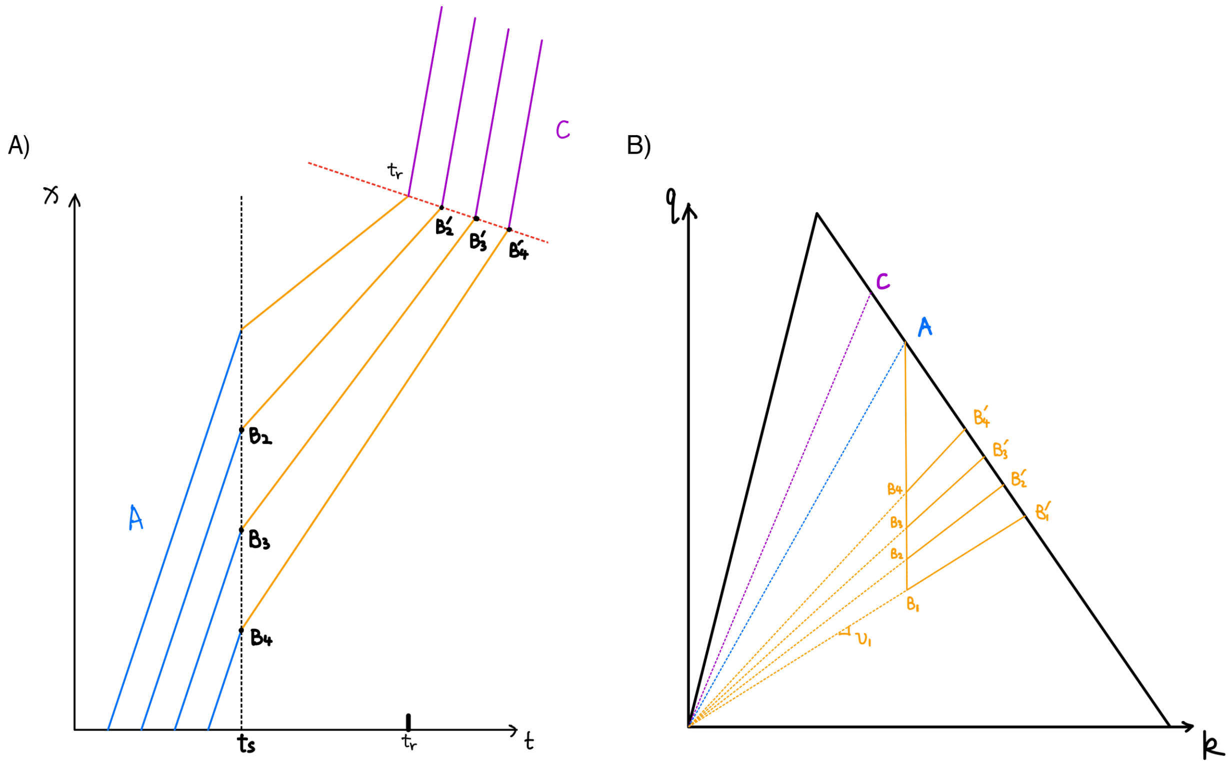

Figure 7, A hypothetical example of the unanticipated delay for the base case. (A) Vehicle trajectories in the time-space plane. (B) The corresponding states taken by the vehicles. Note that the transition from state A to the family of states B_i is instantaneous while the transition from B_i to $B_i{}'$ is a continuous evolution.

With the unanticipated delay, the rerouted platoon might impinge on the planned path of the following platoon. So immediately after calculating the revised path for the current platoon it must be checked against the existing plans for the following platoon. If there is an impingement, the corrections will need to continue into the subsequent platoon. It might be that the following platoon simply has to maneuver around the tail of the first platoon but can still depart on time (effectively turning the corner of the lead platoon using the techniques of Coifman and Liu, 2026), or it may have to temporarily join the end of the previous platoon until such point that the previous platoon can pull away. Fig. 8 shows an example where the third platoon in lane 2 with $t_S = 140\ sec$ and $t_r = 197\ sec$. This disruption impacts the paths of the next two platoons, but both of the following platoons were able to turn the corner from the first platoon and depart on time. Each gap and subsequent platoon absorbs the disturbance until it completely disappears. In this case the sixth platoon and all subsequent platoons are unaffected by the disturbance.

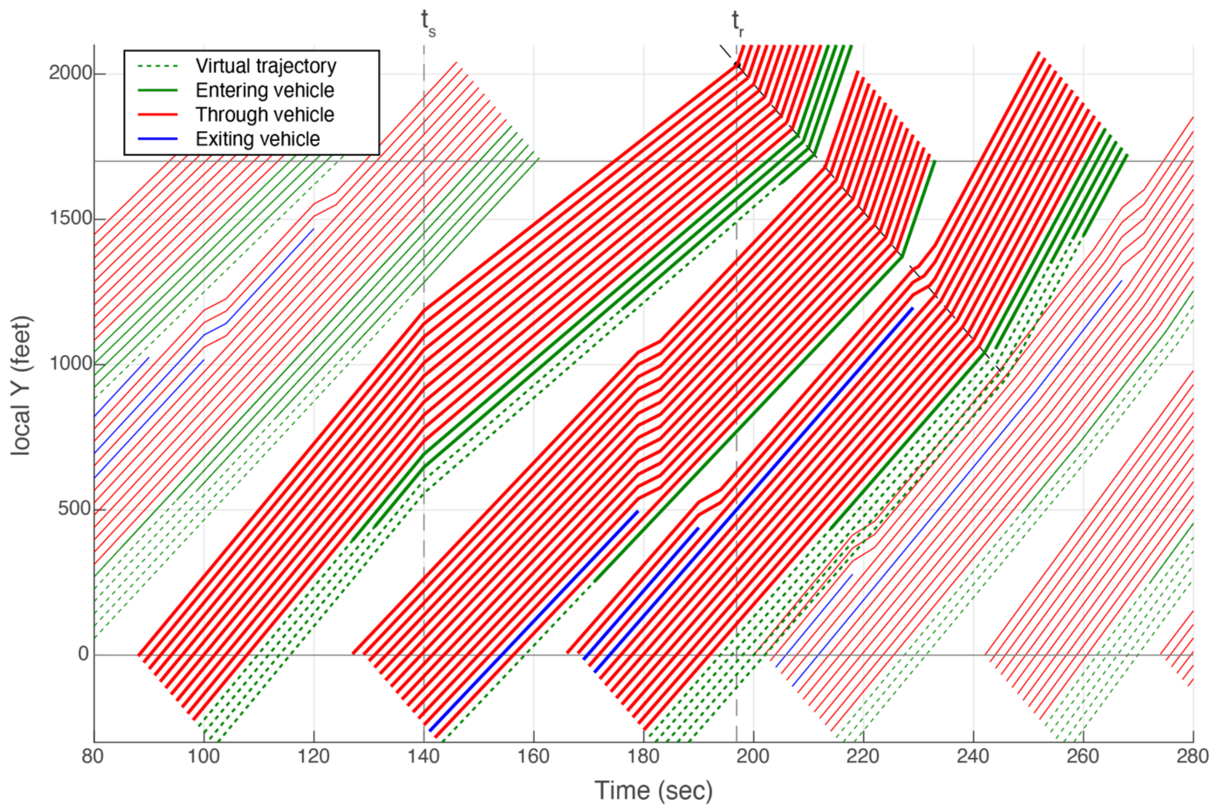

Figure 8, A hypothetical example of the response to unanticipated delay for the base case. Bold curves show the three impacted. Note that all three platoons respond at t_s but the speed change is small enough that it is not perceptible for the middle platoon and the final impacted platoon enters with the speed change already in place.

2.2.2. Extended Case

Fig. 9 shows that the extended case begins with the base case. At some point in time, this case diverges when it becomes clear that the platoon cannot start its recovery at t_r. This discovery does not need to be known in advance, it could happen as late as t_r. Instead of recovering as planned, the lead vehicle in the platoon must maintain v_{B1} for some time beyond t_r. The end point is denoted t_r^*, which may or may not be known at the time of the extension. Since the first vehicle does not speed up as anticipated, each subsequent vehicle must slow down and take state B_1' upon crossing the signal cord from the lead vehicle at t_r as it becomes a perfect follower. So now all vehicles will pass through state $B_1{}'$. Eventually the first vehicle can transition to v_C and all subsequent vehicles in the platoon transition to state C as perfect followers.

Clearly, this case is very flexible in terms of timing, but it relies on v_{B1} continuing to be viable for the first vehicle. If a lower speed is necessary, then one needs to reinvoke the Base Case with that lower speed, as per the Truncated Case below with $t_r^* > t_r$.

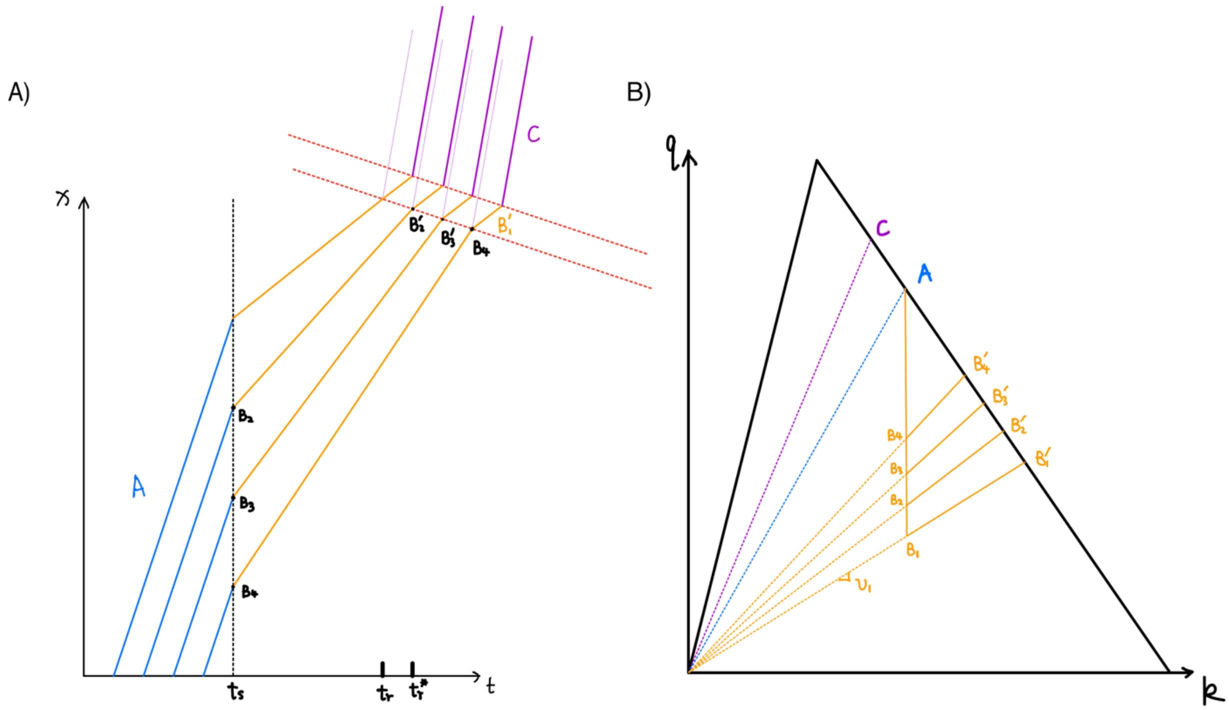

Figure 9, A hypothetical example of the unanticipated delay for the extended case. (A) Vehicle trajectories in the time-space plane, (B) the corresponding states taken by the vehicles.

2.2.3. Truncated Case

Fig. 10 shows that the truncated case also begins with the base case, however, at time t_2 a revised departure time, t_r^*, is determined. As originally envisioned, $t_r^* < t_r$, hence the word "truncated" in the title, but the method also works with $t_r^* > t_r$. Essentially, at t_2 the base case is reinvoked with t_r^*, except that the traffic immediately before t_r^* is not in a steady state. The first vehicle continues at v_{B1}. Then for each subsequent vehicle, as per Fig. 10, the ith vehicle is in its own state B_i'' along the progression from B_i to B_i' while it is traveling at v_{Bi}. Like the base case, each vehicle instantaneously changes speed from v_{Bi} to v_{Di} at t_r^* while preserving their spacing. The vehicles then continue to state D_i' on the FD. Once more, they become a perfect follower in state C upon reaching D_i' and crossing the signal chord from the lead vehicle at t_r traveling at w. Note that Fig. 10B shows $v_{Di} < v_{B(i+1)}$, but that is simply for clarity of presentation, the method also works when $v_{Di} \geq v_{B(i+1)}$. It is a simple extension of this case to have the first vehicle change speeds at t_2 and all subsequent vehicles adjust accordingly.

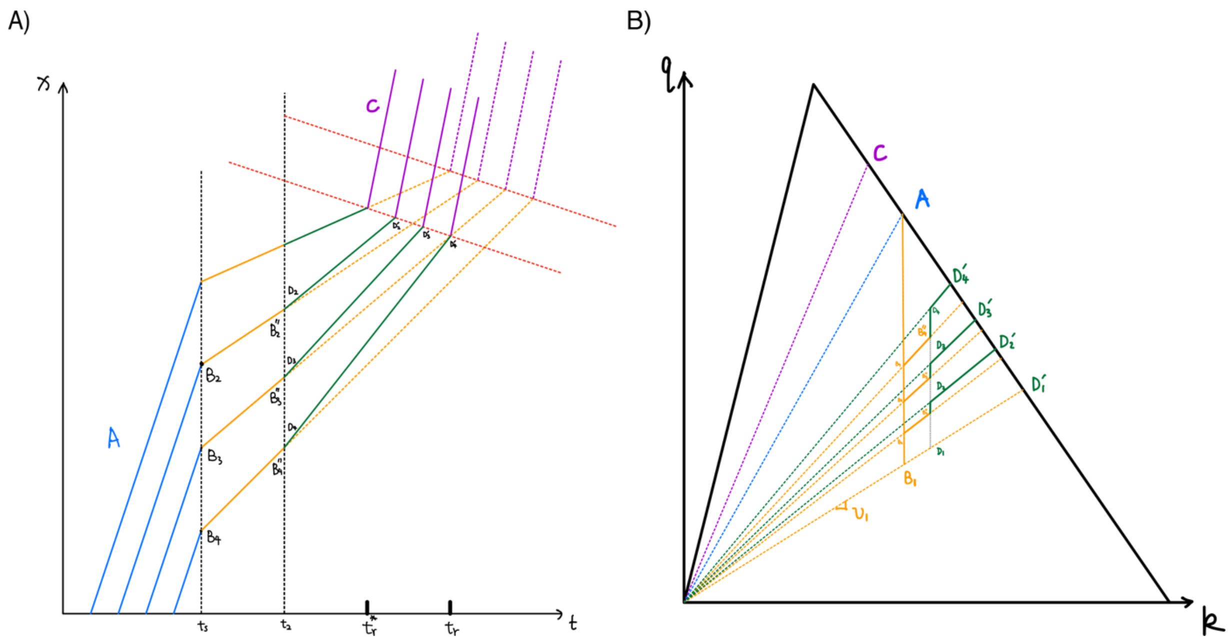

Figure 10, A hypothetical example of the unanticipated delay for the truncated case. (A) Vehicle trajectories in the time-space plane, (B) the corresponding states taken by the vehicles.

2.2.4. Slowed Case

Fig. 11 shows that the slowed case follows the base case up to t_r. At t_r instead of transitioning to the forecast state C at v_C, the state is only allowed to reach state C' at $v_{C'}$, where $v_{C'} < v_c$. As such, the traffic is delayed, but all of the delay occurs in the final state.

2.2.5. Impacts on the Following Platoon

Both the extended case and the slowed case can impact the following platoons, this section explains how we address those cases. If there is enough warning and either t_r^* is known for the extended case or $v_{C'}$ is known for the slowed case, then the following platoon can adjust using the process from the base case. Otherwise, if the final states are not known, we let the following platoon catch up to the current platoon and temporarily join the earlier platoon. Fig. 12A illustrates this process when the following platoon catches up after the slowed case. Here, the duration that the lead platoon will stay at $v_{C'}$ is not known. So initially the following platoon responds as per the base case, speeding up to the intended final state when the vehicles reach the signal chord from the lead vehicle at t_r traveling at w. However, the lead platoon remains at the slower $v_{C'}$, so the second platoon is forced to slow and join the lead platoon. Naturally, this strategy can also be applied to the extended case too.

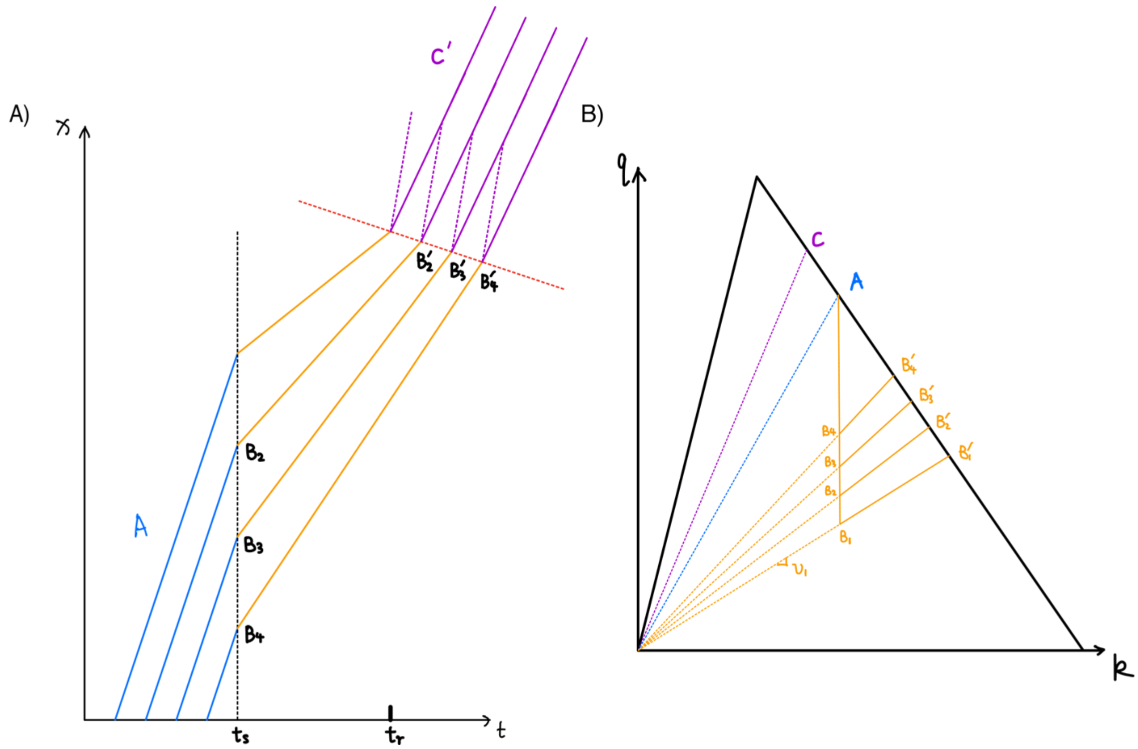

Figure 11, A hypothetical example of the unanticipated delay for the slowed case. (A) Vehicle trajectories in the time-space plane, (B) the corresponding states taken by the vehicles.

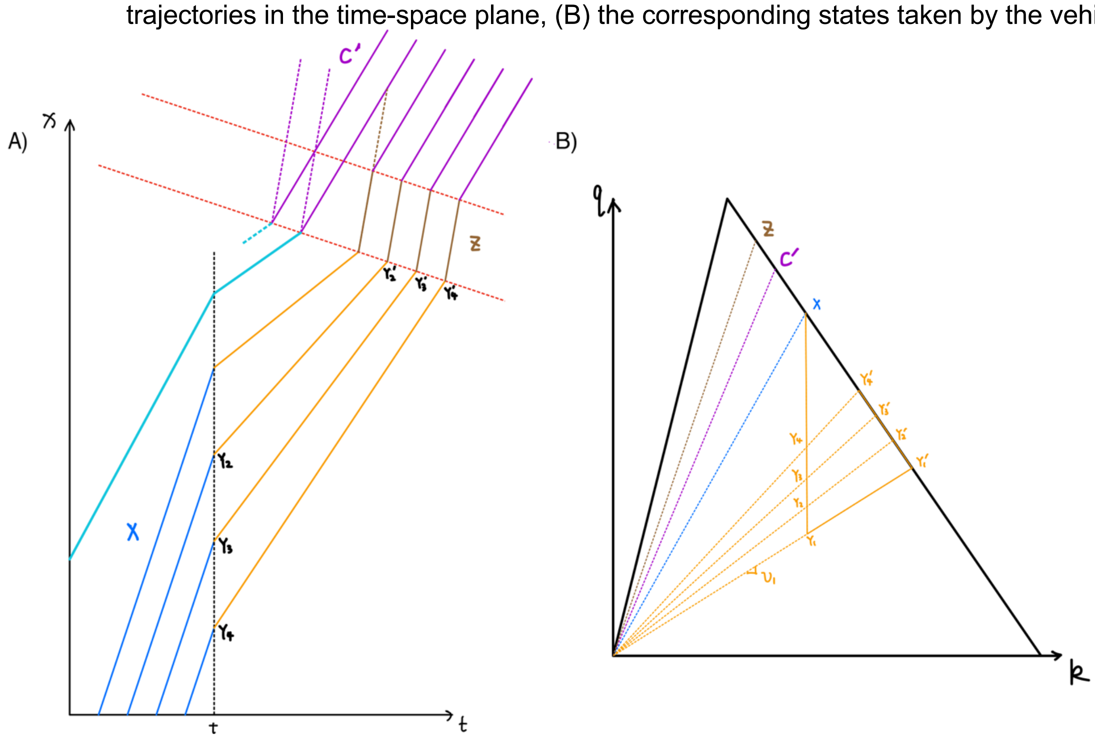

Figure 12, A hypothetical example of how addressing the unanticipated delay in the first platoon can impact the following platoon(s). This figure illustrates the scenario of the slowed case, with the key vehicle from the first platoon shown on the far left and the first vehicles in this platoon entering in state X. To distinguish the second platoon from the first platoon we use states X, Y, and Z in place of A, B and C. (A) Vehicle trajectories in the time-space plane, (B) the corresponding states taken by the vehicles.

2.2.6. Overall Comments

The following comments apply to all of the various cases. If the departure time is continually evolving, a new case can be invoked before the prior one is resolved. Next, throughout the examples we used non-zero speed, but these methods are compatible with a full stop for the final speed in the case of a major event or complete closure.

To fully address these unanticipated delay scenarios, we need downstream information on vehicles beyond 1,800 ft that we do not have. When a delayed platoon exits, it needs to do something, whereas its following platoon is targeting the 'corner' point where the last vehicle speeds up. To avoid having to synthesize the unobserved downstream data, we set the last CAV so that the leading platoon takes the final speed of the first CAV in the following platoon. If the downstream link's dynamics were known (as would be the case in the envisioned system), the departure speeds would be adjusted accordingly. Finally, there are a myriad of anticipated and unanticipated delays and early recovery scenarios, our coverage is not meant to be exhaustive, rather, it is illustrative of representative examples that could be used to address most delays while also recognizing that other solutions are possible.

3. Conclusions

This report presents the results from Phase 3 of a larger study. The overall study seeks to develop a novel anticipatory control method for CAV's, designed to mitigate and smooth disturbances in congested freeway traffic. By enabling CAV's to anticipate stop waves, they can proactively generate smooth trajectories. The method reduces unnecessary accelerations and decelerations, thereby improving safety, driving comfort, and fuel efficiency. This research essentially seeks to take conventionally unstable queued traffic and have CAV's proactively respond to disturbances propagating upstream through the traffic to stabilize the flow. From a broader perspective, the overall study recognizes the limits of automation in resolving fundamental capacity constraints. While the proposed CAV control strategy significantly smooths traffic flow and reduces the negative impacts, it does not increase maximum roadway throughput. Congestion ultimately stems from demand exceeding capacity, and true resolution may require infrastructure expansion or demand management. Nonetheless, by stabilizing flow and minimizing stop-and-go dynamics, this approach represents a substantial step forward in improving the safety and efficiency of congested traffic networks.

The previous phases of this work focused on the initial transition from unstable HDV stop and go traffic to the first platoon of CAV's with smooth trajectories, i.e., dissipating large shockwaves after they have formed and begun propagating. Collectively, the previous phases showed that the system can rapidly nullify a stop wave. If all vehicles were CAV operating under this approach, then stop waves would rarely arise and when one did form, it could be quickly dissipated.

The Phase 3 work in this report starts with the assumption that CAV's can stabilize flow, as per Phase 1-2, and then seeks to optimize the flow of an indefinite number of CAV's. This work assumes that all vehicles are CAV's. Like the previous phases, we use the empirical HDV trajectories from the reextracted I-80 NGSIM data (Coifman and Li, 2024) to determine the underlying demand and capacity. In the envisioned full CAV deployment, control of a given link would leverage knowledge of conditions far downstream of the link. With the NGSIM dataset we do not have this information. So, instead, we use the HDV departure curve at the downstream end of the link as a proxy for the unobserved downstream conditions. Thus, when a CAV enters the link, we already know the earliest time it can leave the link via the HDV departure curve. But we always have the liberty to hold the vehicle later than its earliest feasible departure time.

This study begins with the platoon formation based on the forecasted departure curve. We structure the CAV traffic stream to be a sequence of high-flow platoons and no-flow gaps. This platooning provides greater flexibility to respond to changing conditions. The gaps naturally absorb disturbances that form within a single platoon and provide a buffer where a platoon's dynamics can be changed without impacting any vehicles outside of the platoon.

The platoon formation seeks to collect the vehicles into platoons of roughly the same length as they depart the link. In formulating the platoon strategy, we seek to ensure the last vehicle of the platoon departs on time, as per the HDV departure curve. We use this key vehicle to find the constant speed trajectory from its entrance point to exit point. Then we re-sort the vehicles from their HDV order so that all entrances occur at the rear of the platoon. Leveraging the gap at the end of the platoon, it is far easier to use this space for the vehicle to enter the lane and match the speed, rather than attempting to split a platoon to make room for an entrance. Then, for exiting vehicles we exploit the fact that CAV can exhibit inhuman dynamics and have the platoon's vehicles ahead of the exit briefly slow to close the gap. Thus, keeping the platoon at the highest possible flow for the given speed. By closing the gap in this way, each exiting vehicle effectively removes the first vehicle position in the platoon.

Taken together, exiting vehicles shrink the platoon from the front and entering vehicles grow the platoon from the rear. Furthermore, the vehicles within the platoon exhibit an average speed higher than the platoon itself.

After the platoon is formed, it is checked against the previous platoon to ensure it never overtakes the safety boundary of its predecessor. Once the platoon's path has been finalized, we find the safety boundary for this platoon.

We are not presently assessing optimal platoon length nor how to maintain platoon lengths in the presence of LCM. There are some opportunistic strategies that can be employed, e.g., if several vehicles leave the lane, that might also be a good time to split a platoon. In practice, CAV's at the front of a platoon could accelerate to join the platoon ahead or vice versa for vehicles at the rear of a platoon. These dynamics could be realized with modifications to the controls we use to coordinate departures from the link. However, the greater issue is lane balancing, and as noted below, that is a topic for future research related to managing an extended corridor.

Of course, reordering vehicles like we do assumes that the LCM's are transferable, i.e., we can swap any two CAV's. The present study treats the LCM as an exogenous input determined by the HDV trajectory data. The logic behind this choice is that in a truly efficient CAV traffic stream the synchronization of LCM's would be a complex process that balances OD flows, lane utilization, and other constraints. While this lateral flow control is a critical element of CAV operations, it is a large research endeavor in itself and requires a move to strictly synthetic data across an extended corridor to establish a meaningful set of OD flows. In this context, although there will be some vehicles that need to remain in their given lane (e.g., vehicles approaching their freeway exit or trucks that should stay in the slow lanes), most vehicles do not have a lane dependence and thus, swapping which one of these vehicles changes lanes will not impact the individual OD requirements.

More generally, the human tendency of drivers jockeying to find the fastest lane in queued traffic will be negated by centrally coordinated CAV flows that balance speeds across lanes, so the total number of LCM's for CAV flows should be lower than HDV. We anticipate that an optimal CAV LCM strategy would seek to optimize lane utilization and minimize LCM, e.g., vehicles traveling a long distance would be moved to the left lane while vehicles traveling only a short distance would be kept in the right lanes. While a few vehicles might be shuffled back and forth between lanes to optimize lane utilization, most vehicles would exhibit monotonic LCM behavior. After entering the freeway, they would either maintain their lane or move to the left. Eventually that would switch and they would either maintain their lane or move to the right. But this LCM optimization requires a much longer corridor study and is left to future research.

Using the HDV to determine demands and capacities of the CAV's inadvertently retains artifacts from the HDV dynamics. This unevenness arises because of when and where delays impact the HDV's, e.g., the downstream departure pattern maintains the variable rates from the HDV traffic (including stop waves). This outcome is one of the liabilities of using the empirical HDV trajectories, but this work shows the method is robust to the fact that fluctuations in HDV demand do not correspond to the CAV traffic dynamics. If we had CAV control of the downstream link the departures would be more uniform since the uneven delays from stop waves would have been dissipated from the traffic stream before reaching the study link and the more uniform departures would be easier to control.

Fig. 6 shows many large gaps between platoons in lane 3. We suspect many of these are due to long vehicles being treated as if they are cars. One could modify the code to accommodate multi-class traffic (i.e., different vehicle lengths). When contemplating multiple classes using Newell (2002), one must be careful to recognize the model implicitly constrains the relationship between reaction time, jam spacing, and wave velocity. As per Coifman (2015), these parameters exhibit a strong dependency on vehicle length. Of course, CAV can be engineered to deviate from HDV behavior, but care must be taken to ensure the multi-class model interacts with the platoon dynamics as desired. Such multi-class modeling is left to future work. For the current study the fact that we do not account for the impact of vehicle length simply lowers the entering demand and departure capacity of the link when we use the single class model.

The platoon formation relies upon the forecasted departure curve. Under normal conditions this assumption is reasonable for the envisioned system where an extended corridor is under central coordination, and all vehicles follow prescribed trajectories. But even the best forecasts are imperfect. For example, the forecasts would fail if a CAV had to undertake an unexpected maneuver, e.g., in response to debris in the road, a flat tire, or sensor failure. So, we develop a suite of contingency plans to employ when the state progression diverges from the forecasts. We specifically focus on safety hazards, where vehicles move slower than originally forecast, i.e., unanticipated delays.

Acknowledgements

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.

References

Coifman, B. (2015). Empirical Flow-Density and Speed-Spacing Relationships: Evidence of Vehicle Length Dependency, *Transportation Research- Part B*, Volume 78, pp 54-65.

Coifman, B., Li, L., (2024), Partial Trajectory Method to Align and Validate Successive Video Cameras for Vehicle Tracking, *Transportation Research Part C*. Vol. 158.

Coifman, B., Liu, Y., (2024). *Safe and Efficient Automated Freeway Traffic Control- Phase 1 Final Report*, Safety21 University Transportation Center, USDOT, 18p.

Coifman, B., Liu, Y., (2025). *Safe and Efficient Automated Freeway Traffic Control- Phase 2, Final Report*, Safety 21 A USDOT National University Transportation Center, 56p.

Coifman, B., Liu, Y., (2026). Using Connected and Automated Vehicles to Attenuate Disturbances in Freeway Traffic, *Transportation Research Part C*, Vol 183.

Newell, G., (2002). A simplified car-following theory: a lower order model. *Transportation Research Part B: Methodological*, 36(3), pp. 195-205.