US DOT National
University Transportation Center for Safety

Carnegie Mellon University

Penn
University of Pennsylvania

The Ohio State University

CCAC
Community College of Allegheny County

Community College of Philadelphia

The University of Texas Rio Grande Valley

Morgan State University

Safe and Efficient Automated Freeway Traffic Control- Phase 2 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 – July 31, 2025

DISCLAIMER

Technical Report Documentation Page

<table>
<tr><td colspan="2">1. Report No.
Final Report - 500</td><td>2. Government Accession No.</td><td>3. Recipient’s Catalog No.</td></tr>
<tr><td colspan="3">4. Title and Subtitle
Safe and Efficient Automated Freeway Traffic Control- Phase 2 Final Report</td><td>5. Report Date
July 31, 2025
6. Performing Organization Code</td></tr>
<tr><td colspan="3">7. Author(s)
Benjamin Coifman, PhD, https://orcid.org/0000-0002-8201-964X
Yuan Liu, https://orcid.org/0009-0005-4330-8923</td><td>8. Performing Organization Report No.</td></tr>
<tr><td colspan="3">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</td><td>10. Work Unit No.
11. Contract or Grant No.
Federal Grant # 69A3552344811 / 69A3552348316</td></tr>
<tr><td colspan="3">12. Sponsoring Agency Name and Address
Safety21 University Transportation Center
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213</td><td>13. Type of Report and Period Covered
Final Report (July 1, 2024 – June 30, 2025)
14. Sponsoring Agency Code
USDOT</td></tr>
<tr><td colspan="4">15. Supplementary Notes
Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.</td></tr>
</table>

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 and reduce vehicle emissions. 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. This study, 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. Then the methodology is extended further to accommodate entering vehicles from lane change maneuvers.

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	56	

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

Safe and Efficient Automated Freeway Traffic Control

Phase 2 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

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, under [grant number 69A3552344811 / 69A3552348316] from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

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 and reduce vehicle emissions. 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.

This study, 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. Then the methodology is extended further to accommodate entering vehicles from lane change maneuvers.

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 and reduce vehicle emissions because an internal combustion engine is more efficient cruising at a constant speed than when it is accelerating. Connected and autonomous vehicles (**CAV**) 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 shockwaves in human driven vehicles (**HDV**) can also do so in CAV. While CAV 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 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. 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.

The primary focus of this work is the initial transition from unstable stop and go traffic to the first follower with a smooth trajectory, i.e., dissipating large shockwaves after they have formed and begun propagating. As a result, the modeling assumes that all vehicles downstream of a key vehicle are connected-HDV (**cHDV**) that simply communicate their state, but are otherwise HDV; whereas the key vehicle is a CAV that proactively responds to the downstream conditions. After considering the single CAV scenario the method it is extended to multiple CAVs. It is shown 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 would quickly be dissipated. 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 2 of a multi-phase study. 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 (Coifman and Liu, 2024). This study, 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. Then the methodology is extended further to accommodate entering vehicles from lane change maneuvers. To place the Phase 2 work in context for a reader unfamiliar with the Phase 1 results, it is necessary to briefly review that prior work.

1.1. A brief review of Phase 1

The results of Phase 1 are presented in Coifman and Liu, (2024), this section briefly summarizes those findings to provide context and the starting point for the present Phase 2 report. 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 an uninterrupted flow facility, e.g., a freeway. The work assumes that all vehicles downstream of the CAV are connected and report their instantaneous state (location and speed). The CAV integrates the instantaneous state information from all 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. 1A using real human driven vehicle (HDV) trajectories reextracted from the NGSIM I-80 video in lane 2 (Coifman and Li, 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.

Phase 1 modified the method from Coifman (2002) to use conditions at an instant in time (a vertical cutline). Consider Fig. 1B, which shows the complete trajectories up to time t_1, set to 630 sec in this case. Suppose we were interested in establishing the optimal trajectory for the CAV at (t_1, y_1) in the lower left corner of this plot, we do so by first estimating the trajectory for its leader at $(t_1, 100ft)$. For this work we assume that this leader vehicle and all vehicles downstream are connected HDV, or cHDV for short. The cHDV communicate their state but are otherwise HDV. For this work we use the instantaneous location and speed reported by all the downstream cHDV. While the real traffic state varies continuously over time and space, we approximate it with thin bands of constant speed. Specifically, imagine that each cHDV were to continue traveling at the same speed, yielding a chord of constat slope in the time-space plane. Whereas each vehicle passage is taken to represent the start of a new band of constant speed, with the signal between bands propagating upstream at the wave velocity from the triangular fundamental relationship, set to -15 mph in this case. So, each cHDV yields a downstream moving chord at the vehicle's speed and an upstream moving signal at the fixed wave velocity. Then, when a chord reaches the first signal emanated by the vehicle ahead that chord is terminated. The result is illustrated in Fig. 1C. Since the bands are assumed to be at constant speed, we shift the truncated chords end to end, starting at the lead vehicle's current position of $(t_1, 100ft)$, yielding the estimated trajectory of the leader in Fig. 1D. Fig. 1E 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.

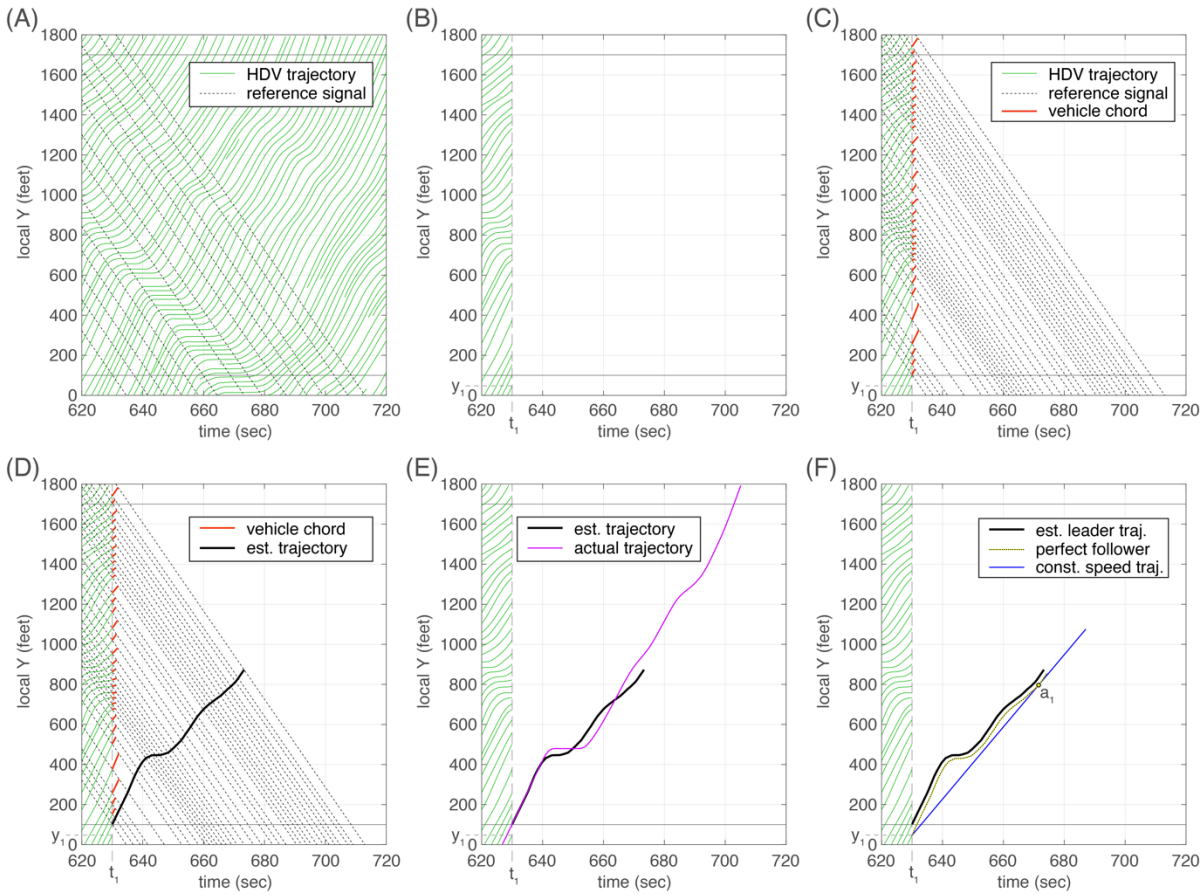

Figure 1, (A) vehicle trajectories from congested traffic and several upstream moving reference signals. (B) Traffic state up to time t_1, (C) roughly steady state bands and their measured vehicle chords, (D) aligning the chords at t_1 end to end to estimate the lead vehicle's trajectory, and (E) comparing the estimate to what will eventually occurs. (F) Establishing the perfect follower trajectory and finding the intersection with the fastest constant speed radial for the CAV at t_1.

The lead vehicle's estimated trajectory bounds where the CAV at (t_1, y_1) can travel. We first project the estimated lead vehicle's trajectory upstream to derive the **perfect follower trajectory**, as shown in Fig. 1F using Newell (2002) with jam spacing, $s_{jam} = 22\ ft$ and reaction time $\tau = 1\ sec$ (with the corresponding wave velocity $w = -15\ mph$). 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. 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. 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.

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

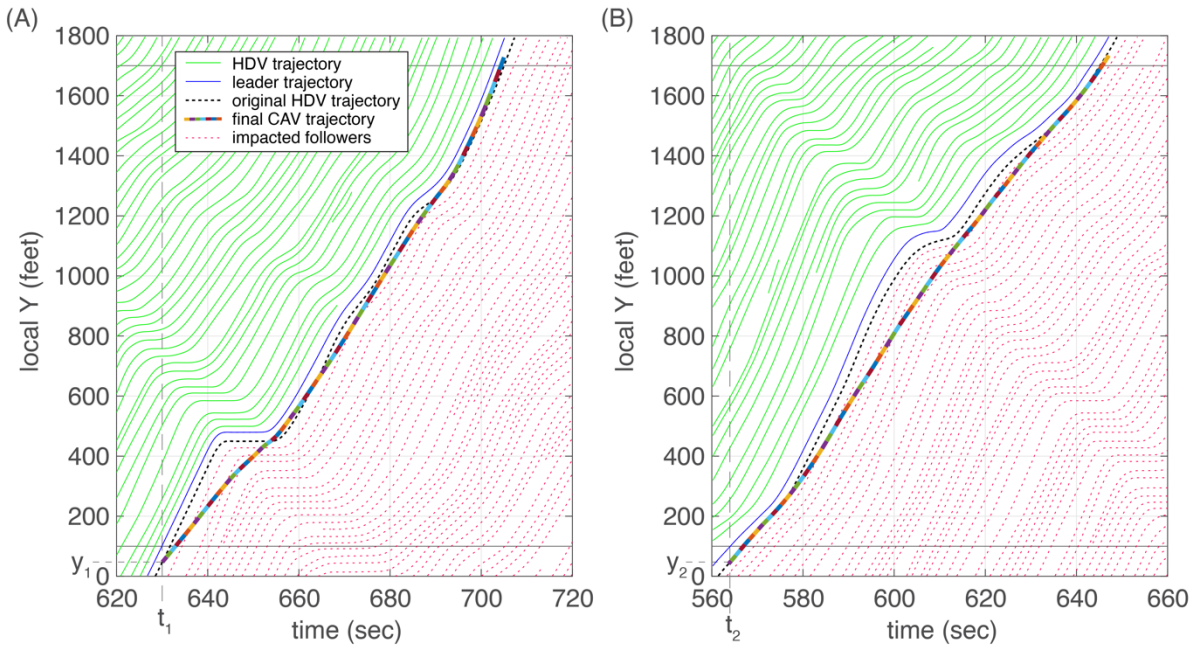

Figure 2, The final CAV trajectory for the CAV in (A) Case 1 starting at (t_1, y_1), and (B) Case 2 starting at (t_2, y_2). In both cases note how the CAV neutralizes the stop wave seen in the leader trajectory and original HDV trajectory.

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. 2A 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, 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. 2B repeats the analysis for another CAV, in this case at (t_2, y_2).

The empirical data are limited to the original surveillance region of NGSIM, which only spans 1800 ft. As a result, when a vehicle approaches the end of the segment the effective range of the look ahead shrinks. So that study also used Coifman (2002) to generate synthetic vehicle trajectories downstream of the segment to extend the look ahead. Leading to three variations of look ahead: strictly empirical data from 0-1800 ft., 0-2500 ft using synthetic data beyond 1800 ft, and a fixed moving window of 1000 ft downstream of the CAV's leader that also uses synthetic data for trajectories beyond 1800 ft.[1] Not surprisingly, the longer the look ahead, the quicker the CAV can respond. By extension, the earlier the response is the smaller the accommodation needs to be.

The Phase 1 study also considered the impacts of limiting the CAV acceleration, a_{CAV}. Where the limits are strictly on acceleration, for safety's sake no limit is placed on the magnitude of deceleration. That study considered three variants: no limit, $a_{CAV} \leq 1\ mphps$, and $a_{CAV} \leq 0.5\ mphps$. Generally, the limits on acceleration did not delay the CAV. When the acceleration is

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

limited, the CAV might momentarily fall behind its leader, but it eventually catches up and once it has, no delay is incurred by the limited acceleration. This, "falling behind" is simply additional smoothing on the traffic state.

1.2. Overview

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

2.1. Extending to multiple CAV followers

As discussed in Section 1.1, the first CAV follower integrates information from the downstream cHDV to estimate the leader's trajectory. So at time step t_o we estimate the first CAV's trajectory as per Fig. 1 and summarized in Fig. 3A for the CAV starting at (t_o, y_1) by taking the fastest constant speed radial that touches the perfect follower trajectory at a single point, a_1.This first CAV seeks to break the pre-existing trends in the preceding HDV trajectories and in the process, its trajectory from (t_o, 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. 3B. This estimated trajectory is then shifted using Newell (2002) to establish the first CAV's perfect follower trajectory, as shown in Fig. 3C. At which point the second CAV's trajectory is then estimated following the same rules used for the first CAV and is 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 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 the signal propagating at the wave velocity from a_1 (as per Fig. 3C) and at or before the signal propagating at the wave velocity from the end of the estimated perfect follower trajectory (as per Fig. 3D), 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.

As we shift to the empirical data, we must work with 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.

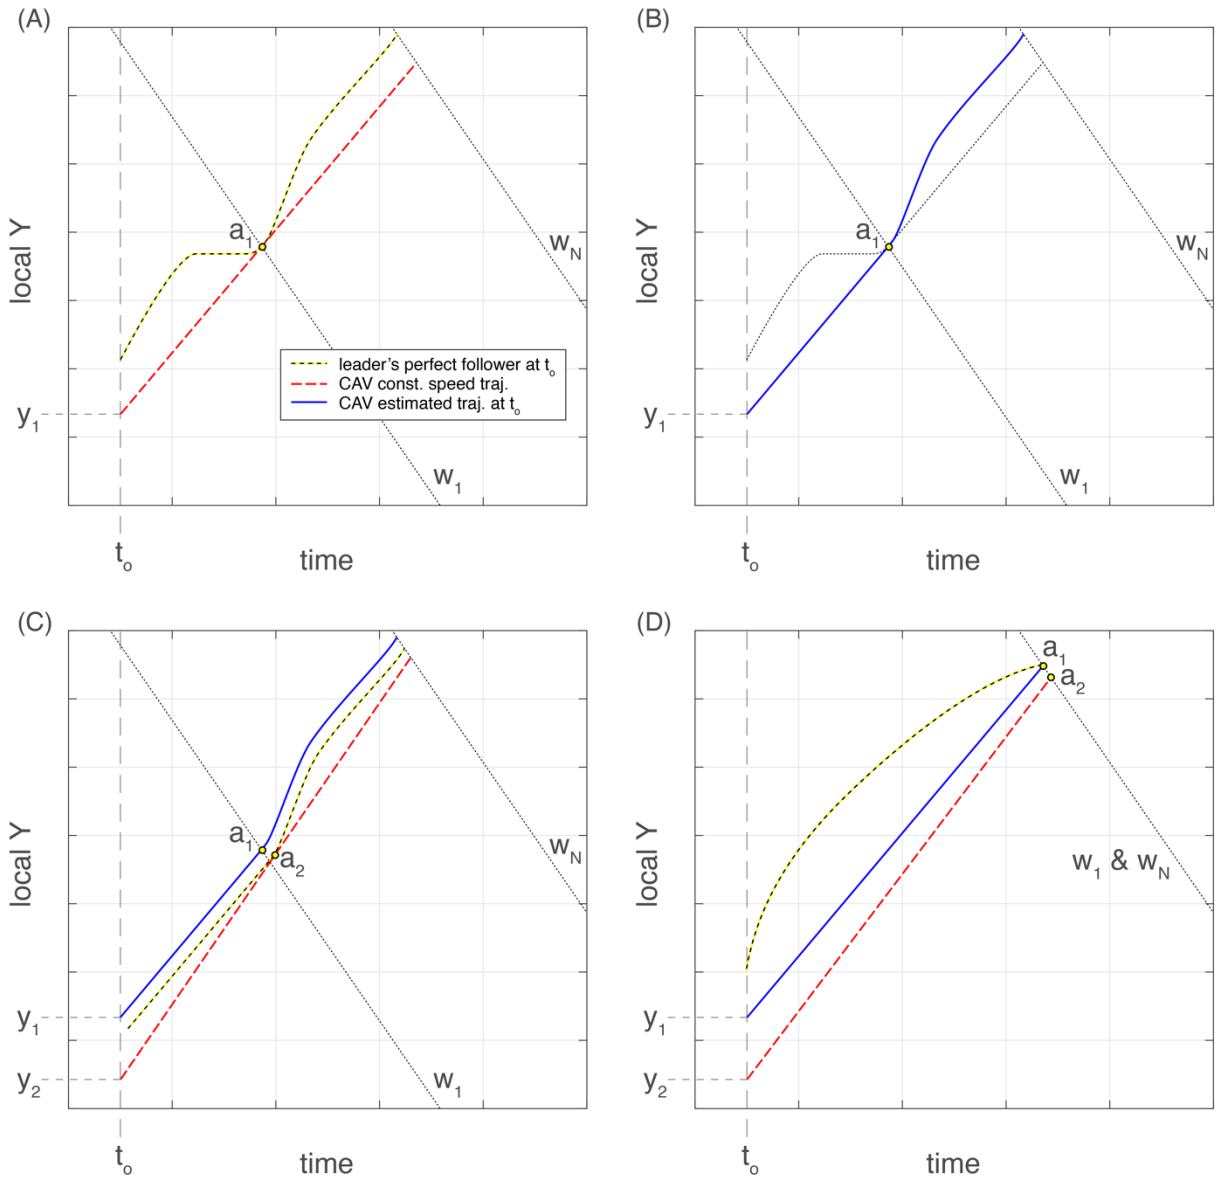

Figure 3, (A) To plan the path of the first CAV follower find the fastest radial from the CAV's location that touches the perfect follower trajectory at a_1, (B) combine the radial to a_1 and the perfect follower trajectory beyond as the estimated trajectory for the first CAV follower, (C) shifting the 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.

The following results use two cases first developed in Phase 1 and presented in the Phase 1 final report. Fig. 2A and 2B show the Phase 1 results for Case 1 and Case 2, respectively, for the first CAV with the 0-1800 ft look ahead and no limit on acceleration. Fig. 4 revisits Case 1 only now with fixed 1000 ft window and no limit on acceleration. More importantly, it now uses 15 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 single CAV trajectory

for this scenario as presented in the Phase 1 report. The next 14 CAVs are all new. 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 constant speed over the same period.

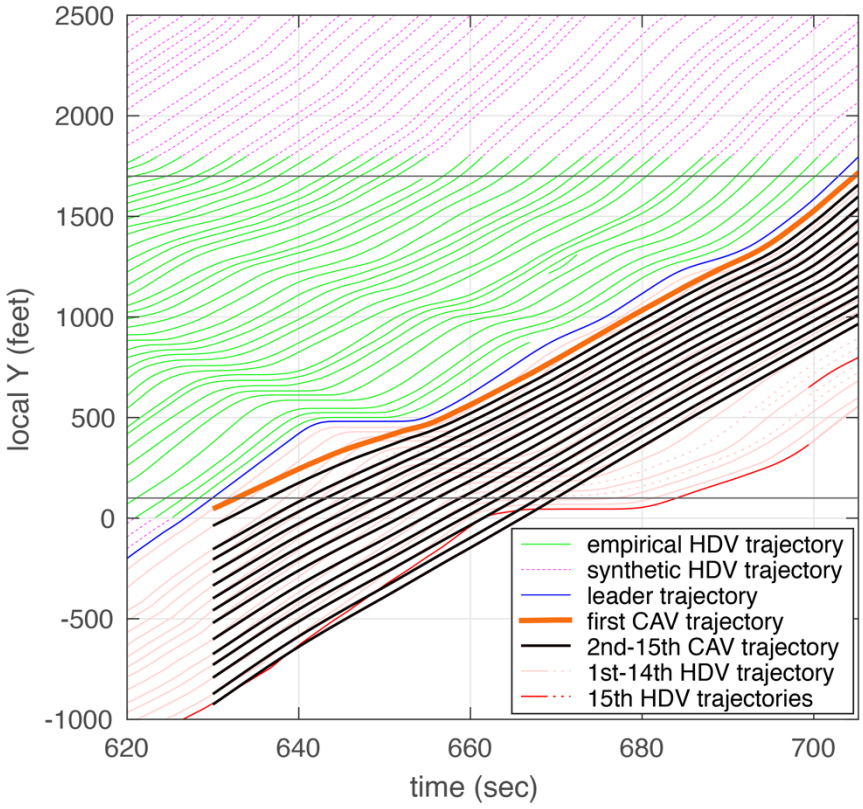

Figure 4, Case 1 with the fixed 1000 ft window and no limit on acceleration using the first follower from before and the next 14 vehicles for a total of 15 successive CAV followers. The CAV trajectories are shown in bold and the original HDV trajectories in a faint color for reference. Note that in the 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.

Fig. 5A 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. 5B 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. 5C-D show the corresponding time series acceleration. The results are consistent, e.g., the 15th CAV shows almost no acceleration before 705 sec.

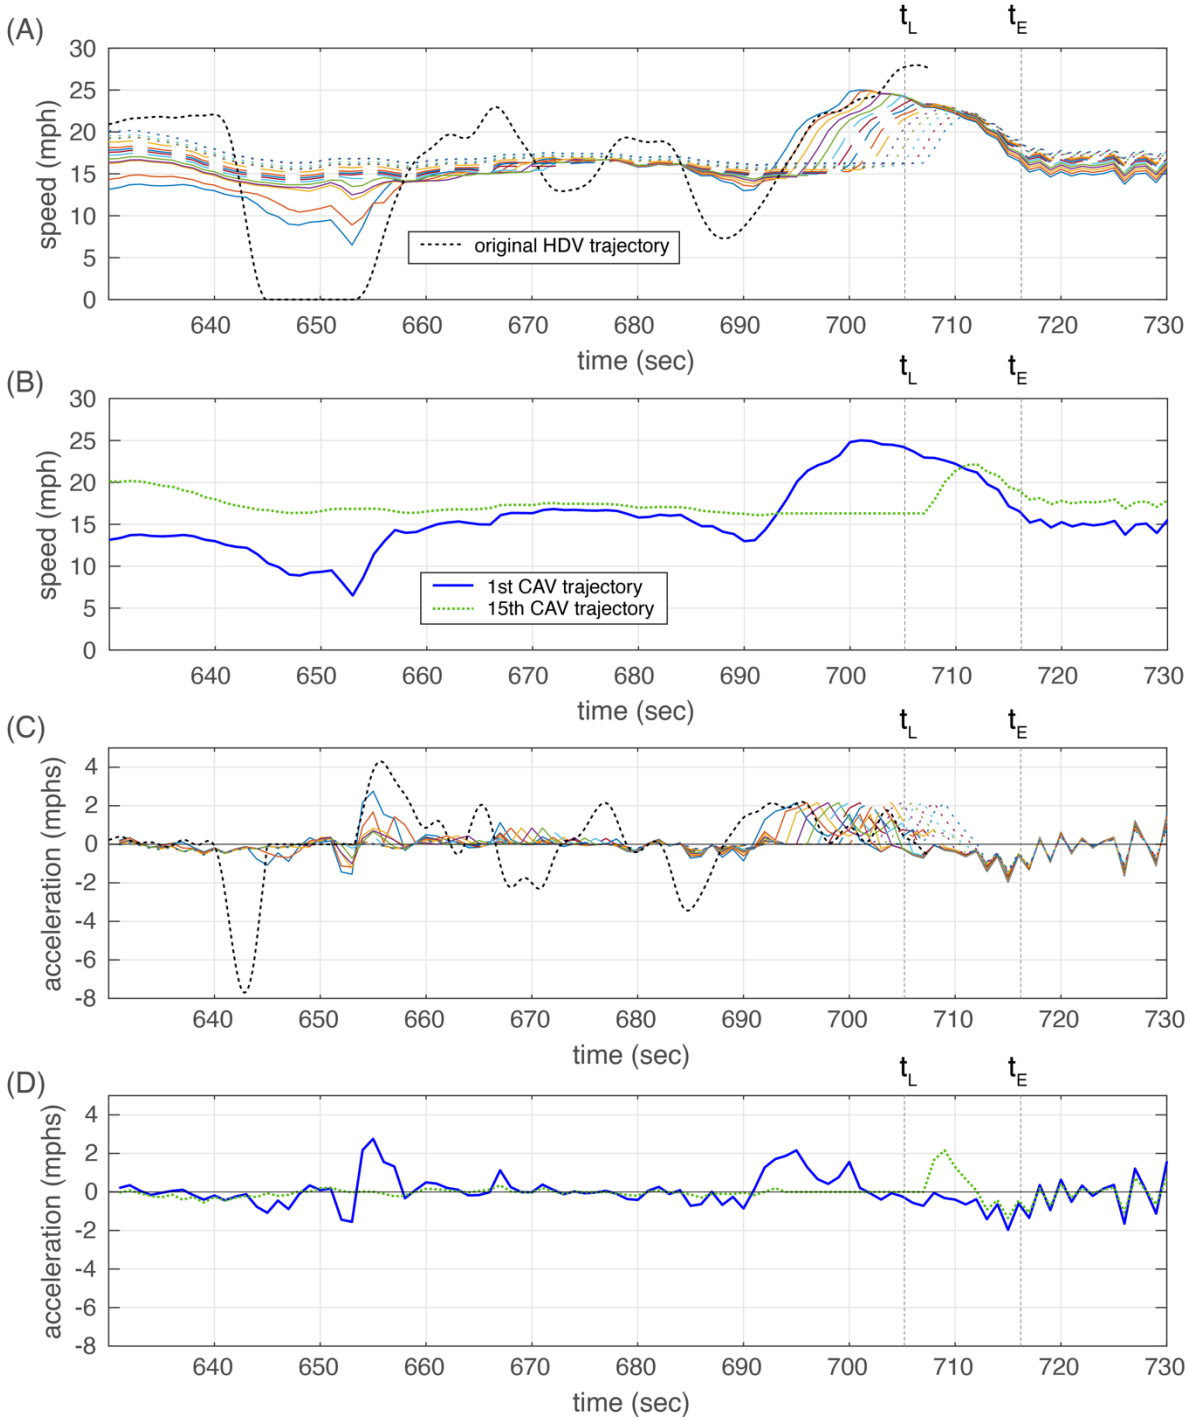

Figure 5, (A) Time series speed from the 15 CAV and the original trajectory for the first HDV follower for Case 1 in Fig. 4. (B) Repeating A only now just showing the first and 15th CAVs. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first and 15th CAVs.

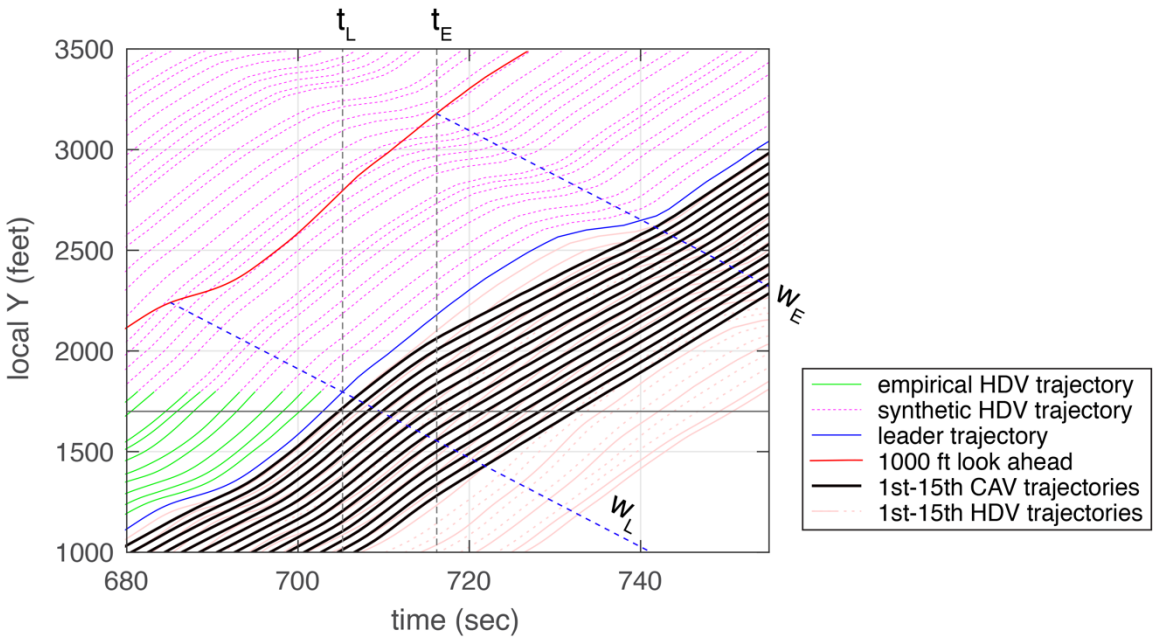

Figure 6, Using the synthetic data to extend the coverage beyond the distance shown in Fig. 4 to show the general dynamics of a major slow wave.

Note that the time range in Fig. 5 extends almost 25 sec past the end of the time range in Fig. 4. Fig. 4 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. 5. Any calculations beyond t_L must be based on synthetic data for the leader. While we are comfortable using the synthetic data to extend 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. 5. First, notice how the different CAVs sometimes exhibit synchronous responses, e.g., the small wobbles in acceleration after 710 sec in Fig. 5C. In this time period the first CAV is slowing, so point a_1 is continually at the end of the perfect follower trajectory, as per Fig. 3D. 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. 5A 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. 5B shows that both the magnitude and duration of this acceleration-deceleration cycle diminishes greatly from the first to the 15th CAV.

Using the synthetic data, Fig. 6 extends the coverage further in time and distance from Fig. 4 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 of the 1000 ft look ahead window), 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. 5A-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. 5A 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. 5A 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. 6. 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. 6 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. 5A. Regardless, as already noted, Fig. 5B shows that both the magnitude and duration of the impact diminishes greatly from the first to the 15th CAV.

Fig. 7-8 repeat the 15 CAV evaluation for Case 2 (previously shown in Fig. 2B for the single CAV with the 0-1800 ft look ahead) with a 1000 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 pursues 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 successfully navigates a gentle transition to the post-stop wave state, as per Fig. 8A, the corresponding HDV exhibits large speed fluctuations and almost comes to a full stop. Fig. 7 shows the stop wave persists to the 15th HDV while using the methodology developed herein, the 15th CAV traverses the former stop wave with nearly constant speed. Although there is no limit placed on acceleration, the 15th CAV keeps acceleration below 1 mphps except for one time step after 650 sec and below 0.5 mphps for most of the time steps in this example.

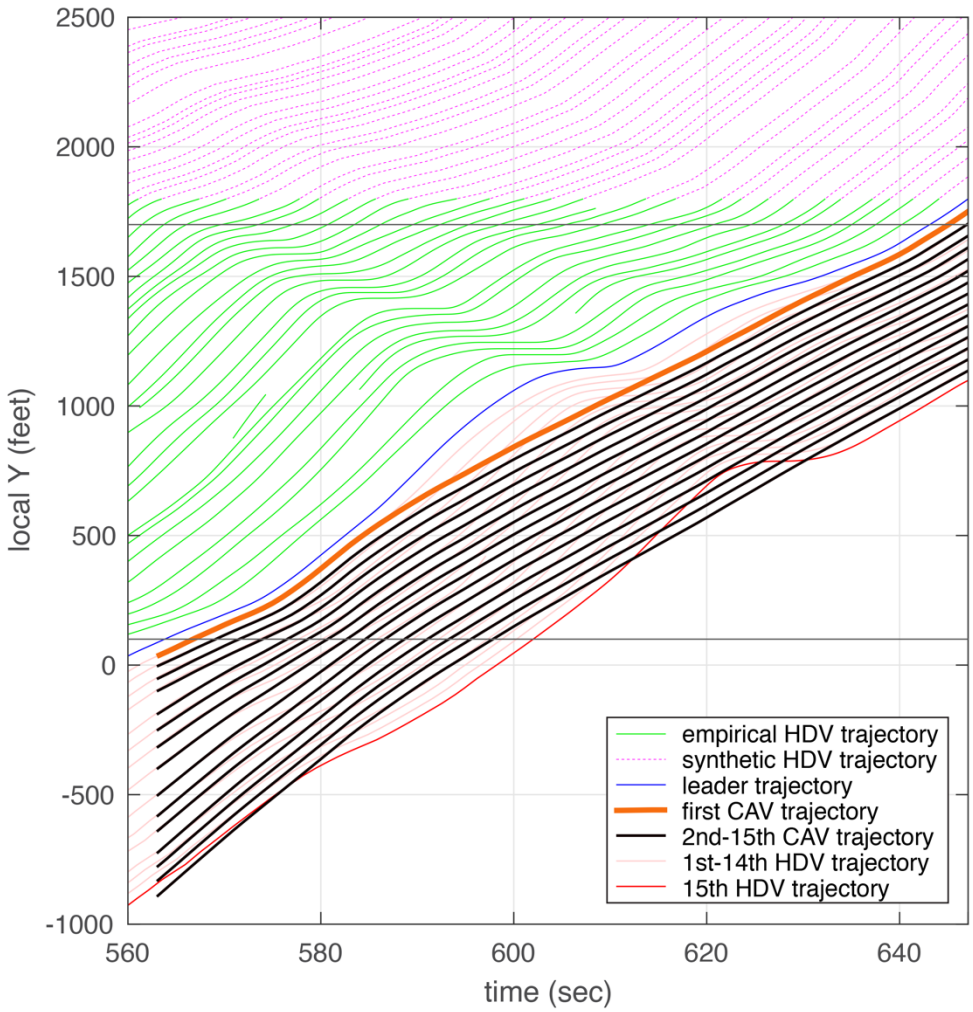

Figure 7, Case 2 with the fixed 1000 ft window and no limit on acceleration using the first follower from before and the next 14 vehicles for a total of 15 successive CAV followers. The CAV trajectories are shown in bold and the original HDV trajectories in a faint color for reference. Note that in the 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.

2.2. Emergent dynamics

Ultimately this anticipatory CAV control deliberately seeks to disrupt normal HDV car following behavior. Fig. 9A shows the progression of 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 ($s_{jam} = 22\ ft$, $\tau = 1\ sec$, and $w = -15\ mph$). Notice how the HDV vehicle exhibits counterclockwise loops, as anticipated by Newell (1962) and shown empirically in subsequent research (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.

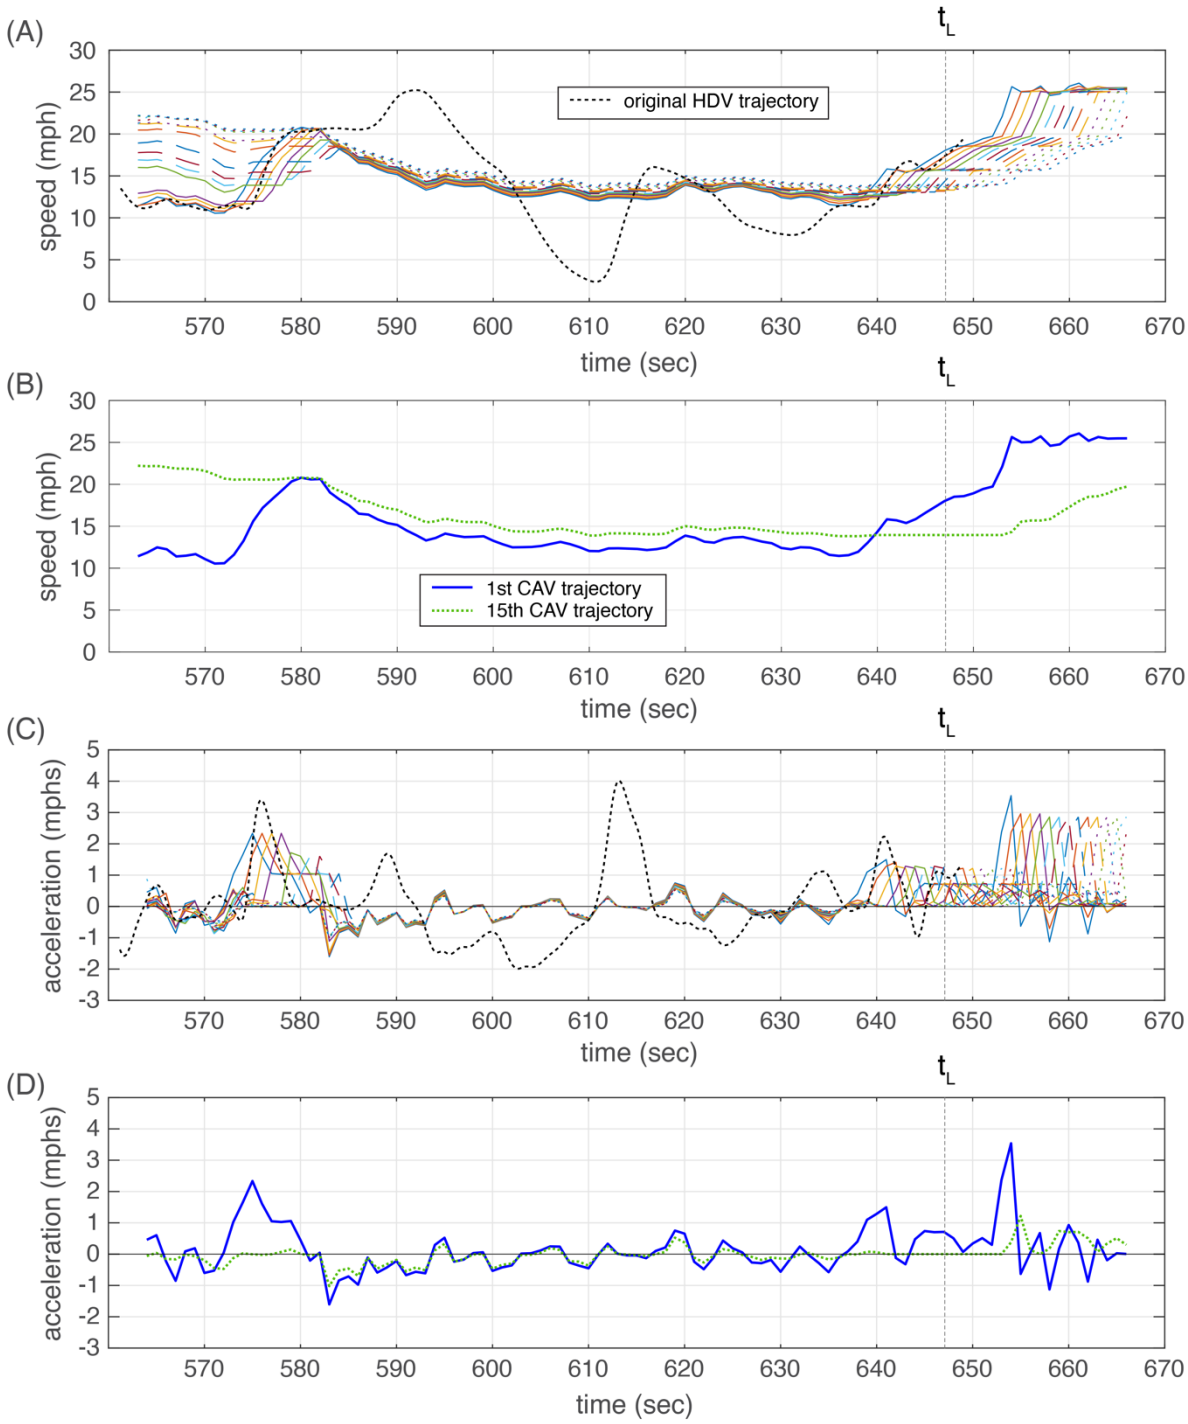

Figure 8, (A) Time series speed from the 15 CAV and the original trajectory for the first HDV follower for Case 2 in Fig. 7. (B) Repeating A only now just showing the first and 15th CAVs. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first and 15th CAVs.

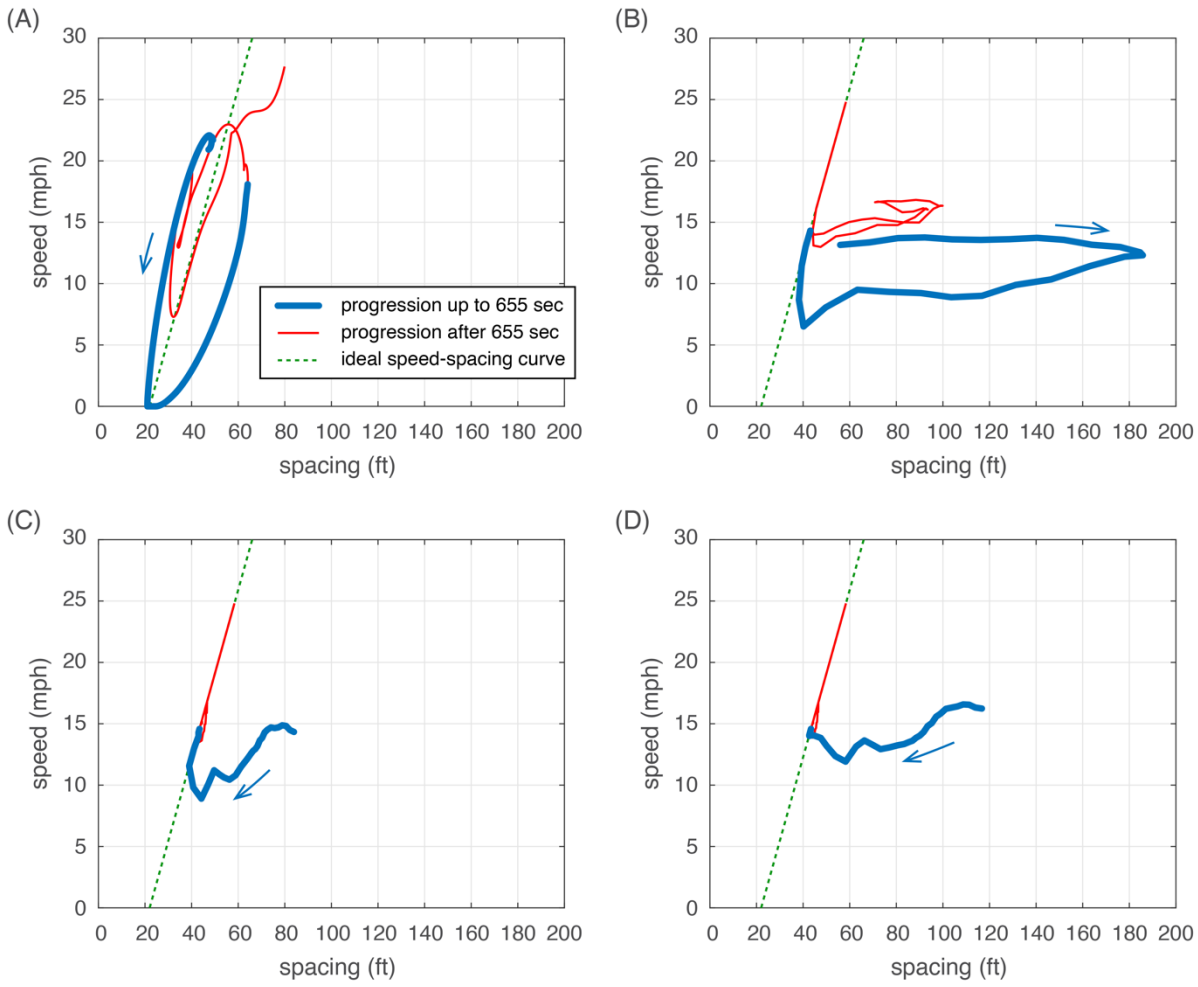

Figure 9, The progression of the speed-spacing for Case 1 underlying the trajectories in Fig. 4 (A) for the original trajectory for the first HDV follower, (B) the first CAV follower, (C) second CAV follower and (D) third CAV follower. For reference the ideal speed-spacing curve is shown with a straight line in all four subplots.

Fig. 5A 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. Fig. 5B shows the first CAV has completed its response to the stop wave by 655 sec, so for reference in Fig. 9 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. 9B shows the corresponding progression when this first HDV is replaced with a CAV (1000 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 forthcoming 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. 1E). The CAV executes a similar loop near 15

mph as the leader passes through two minor disturbances visible in Fig. 4 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. 6.

Fig. 9C-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 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.

Of note is that away from the large disturbances, all of the CAVs stay close to the ideal speed-spacing curve. There is no perceptible looping as was exhibited by the HDV speed-spacing progression in Fig. 9A. Thus, even when the CAVs are simply car following and not proactively maneuvering to absorb disturbances, they are far more stable than the HDVs.

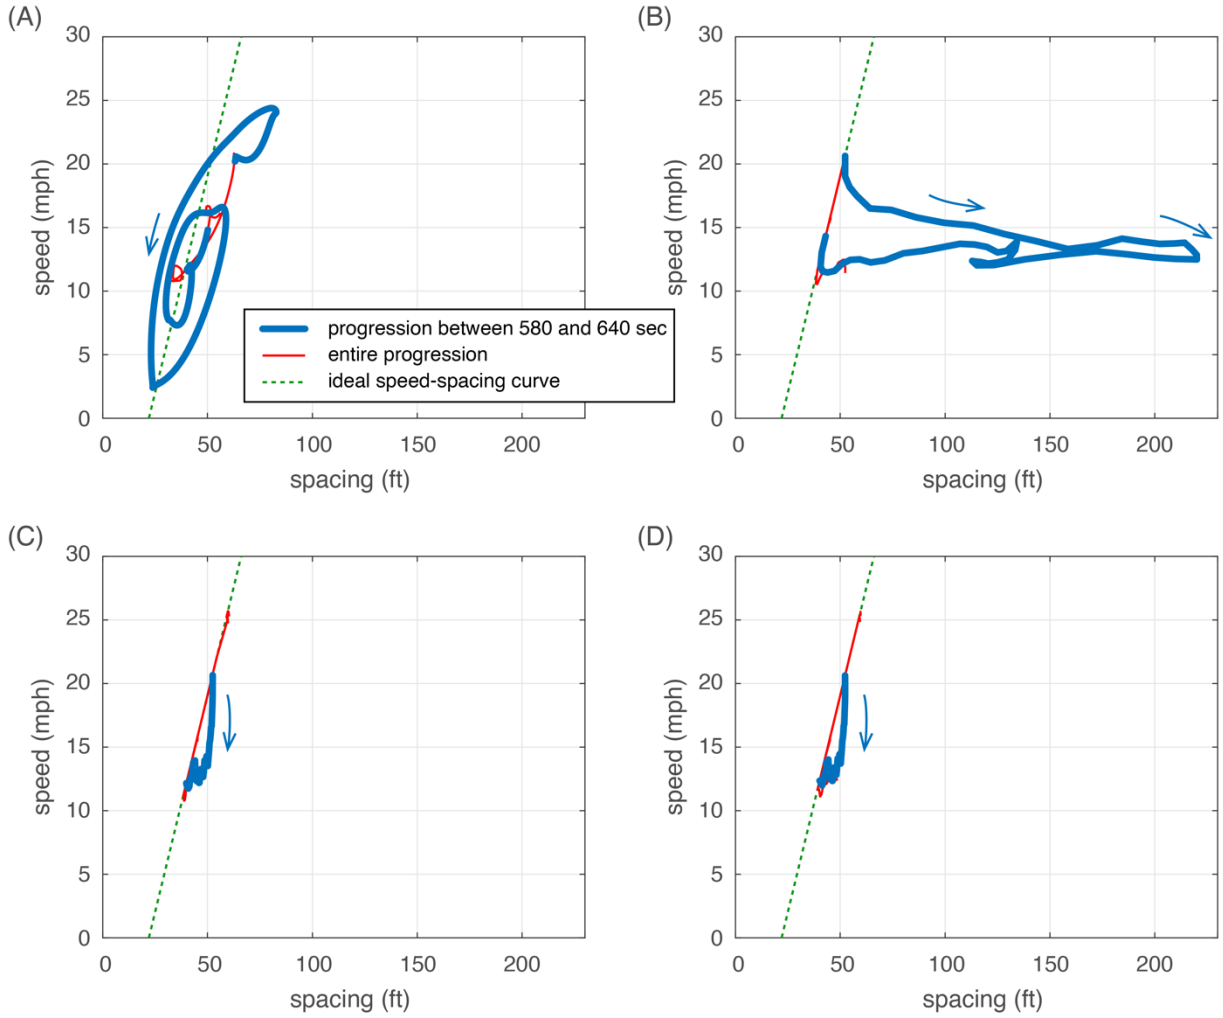

Figure 10, The progression of the speed-spacing for Case 2 underlying the trajectories in Fig. 7 (A) for the original trajectory for the first HDV follower, (B) the first CAV follower (note that the bold curve comes close together but never crosses itself), (C) second CAV follower and (D) third CAV follower. For reference the ideal speed-spacing curve is shown with a straight line in all four subplots.

Fig. 10 repeats this analysis for Case 2. Again, the HDV exhibits the expected counterclockwise looping (Fig. 10A). This time, at the start the CAV have not yet detected the delays of the forthcoming stop wave. Instead, they are simply converging towards car following. Fig. 8A shows that the first CAV starts decelerating just after 580 sec in anticipation of the stop wave, 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. 10B 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. 8 are helpful references when interpreting Fig. 10. 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. 10A initially accelerates from 20 mph into the forthcoming stop wave, while the corresponding CAV decelerates in Fig. 10B. 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 slowing their progression to car following (as per Fig. 9C-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 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 could never exhibit.

2.3. Lane change maneuvers (LCM)

Thus far, the analysis has been confined to single-lane scenarios with no lane change maneuvers (LCMs). Recall that the overall scope is limited to congested traffic, so if one vehicle enters a lane, it does so by delaying all vehicles behind it in the new lane by one headway because it "bumps" everyone in the queue back by one position. As such, this work assumes that LCM will only be taken when needed, i.e., aggressive vehicles will not be allowed to jump back and forth between lanes. Still, vehicles will need to move between lanes in anticipation of exiting the freeway, balancing demand between lanes, and so forth. While the higher-level objective dictating when a vehicle should change lanes is beyond this work, it is necessary to have the enabling building blocks.

If vehicles are platooned and speeds are stable (e.g., Fig. 4 and Fig. 5), the only way another vehicle can enter the lane is by creating a slow wave upstream of the entrance. If traffic is slowing, an entering vehicle will exacerbate the pre-existing slow or stop wave. Both of which run counter to the work in the previous sections, where we sought to smooth out the congested traffic flow by eliminating unnecessary acceleration and deceleration cycles. So, our focus is on LCM that do not require the target lane to decelerate. In other words, we seek a LCM strategy where the vehicle enters on an acceleration cycle. Specifically, as an acceleration wave propagates upstream through a platoon, the new target follower will hold off accelerating and instead maintain their previous speed to create a gap for the entering vehicle to fill.

This policy allows vehicles from adjacent lanes to merge into the target lane while preserving the benefits of smoothing the disturbances enabled by the CAV control strategy. To evaluate this extended approach, Cases 1 and 2 are revisited, now incorporating scenarios in which a vehicle merges into the study lane from an adjacent lane.

This opportunistic approach of admitting LCM only in acceleration waves limits the maximum rate of LCM. If a high LCM rate is needed (e.g., to balance demand across lanes), then it may be necessary to lower the speed in the target lane, but that would be a unilateral speed drop to actively control lateral traffic flow rather than a deceleration-acceleration for individual LCM.

While proscribing target speeds for some region of the time-space plane is an important element for managing CAV traffic, it is beyond the scope of the current work where we simply seek to smooth the flow in a way that does not impact the average speed.

2.3.1. LCM policy

LCMs are a challenging and active topic in transportation engineering, yet comprehensive and satisfactory solutions in the literature are still limited. One of the main difficulties lies in the need for coordinated adjustments in speed and spacing between vehicles across adjacent lanes. This complexity is further amplified under congested conditions, where the additional requirement of sufficient space in the target lane to accommodate the merging vehicle becomes particularly difficult to satisfy. This study focuses on strategies for CAVs in the target lane to create the necessary space for merging vehicles while still maintaining the ability to attenuate stop waves effectively.

Several stipulations are made regarding the entering vehicle during a LCM. First, the entering vehicle is assumed to be a CAV so that its maneuvers can be executed with precision. As in the previous sections, any spacing deliberately created in the current lane is assumed to be reserved and cannot be captured or occupied by other vehicles. It is also assumed that the CAVs in the target lane are informed of a lane change request at a designated time stamp, allowing them to begin making speed and spacing adjustments to accommodate the merging vehicle. However, if a suitable gap is not formed before the following vehicle passes 1800 ft, the test is terminated because the empirical data used in the analysis ends at that point. In other words, although synthetic data is available to extend downstream, the decision to execute a lane change must be made before the end of the empirical data so that it does not depend on a synthetic trajectory for the follower beyond 1800 ft.

For this study, the actual timing and location of the LCM are determined solely by the target lane's decision process after receiving the LCM request. There is no consideration of constraints in the exited lane, i.e., it is assumed that the exited lane can deliver the lane changing vehicle at the time and location dictated by the target lane. In practice this coordination could likely take the form of choosing which vehicle in the target lane will make the gap. So, if the source lane was moving slower than the target lane, the gap would be created such that it would overtake the lane-changing vehicle. While the following examples have the entering vehicle filling the gap as soon as the gap is large enough, it is conceivable that to facilitate synchronizing between lanes one might form a gap and let it persist unfilled for a short time before filling it. Finally, the entering CAV is assumed to follow the same parameters as those in the current lane, including a one second response time and a jam spacing of 22 ft. In practice, if the entering vehicle has a longer response time or larger jam spacing—such as in the case of a truck—the required spacing in the target lane would need to be proportionally larger.

For ease of presentation, this study restricts the entrance to occur between the first CAV follower and its HDV leader, but it is a simple extension to modify this method for any of the CAV followers. At time stamp t_R the CAV in the target lane is notified of the LCM request and it begins responding immediately. In the absence of a LCM request, the CAV would proceed as per Section 2.1 and at each time step it would choose point a_1 in Fig. 1. After a request from a vehicle to enter ahead of the CAV, it will still calculate a_1 in the same manner. Using v_1 to denote the speed of the straight line radial and $v(t)$ for the CAV's current speed, as long as $v_1 < v(t)$ the CAV will need to slow in direct response to its leader, and it will ignore the request to make a gap for the LCM. Whenever $v_1 > v(t)$, instead of adopting v_1 in the next time step, the CAV will maintain its current speed, setting $v(t+1) = v(t)$. Thus, the leader will pull ahead of the CAV, which

enlarges the gap between the two vehicles. Note that the CAV will either maintain its current speed or decelerate in the next time step, it will not accelerate while creating the gap. This process continues until the first time step, t_e, where the estimated trajectory for the CAV falls upstream of the location of the second perfect follower at that instant, $y_2(t_e)$, as given by Fig. 11, where τ is the response time, d is the jam spacing, and w is the wave speed. At this instant the LCM vehicle is allowed to enter. Starting the very next time step, the CAV behind the entrance will behave as the second CAV. If, for some reason the entrance is delayed, then between when the entrance is allowed and when it occurs, the former first CAV will simply shift the estimated leader's trajectory by two positions to estimate the location of the second perfect follower instead of the first perfect follower. It will then calculate a_1 in Fig. 1F from the second perfect follower's trajectory. Meanwhile, throughout this entire process, all other CAVs continue to compute their speeds based on the planned trajectories, following their respective leaders in a one-by-one fashion without change from the normal platooning in Section 2.1.

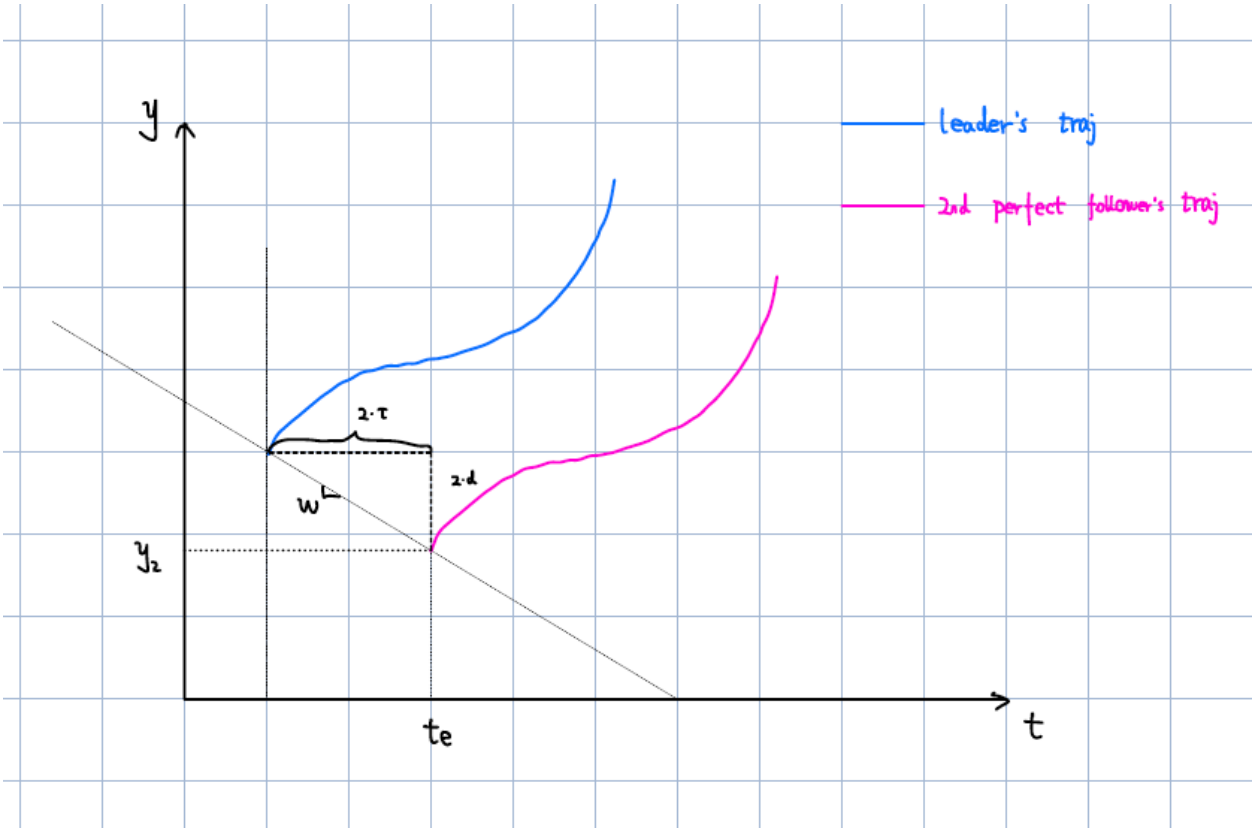

Figure 11, The trajectory of the second perfect follower as calculated from the original leader. At some time, t_e, the location of the second perfect follower can be calculated from the recent history of the leader's trajectory. When the first CAV passes the second perfect follower trajectory it will have created a sufficient gap for an entrance and will now start tracking the second perfect follower trajectory.

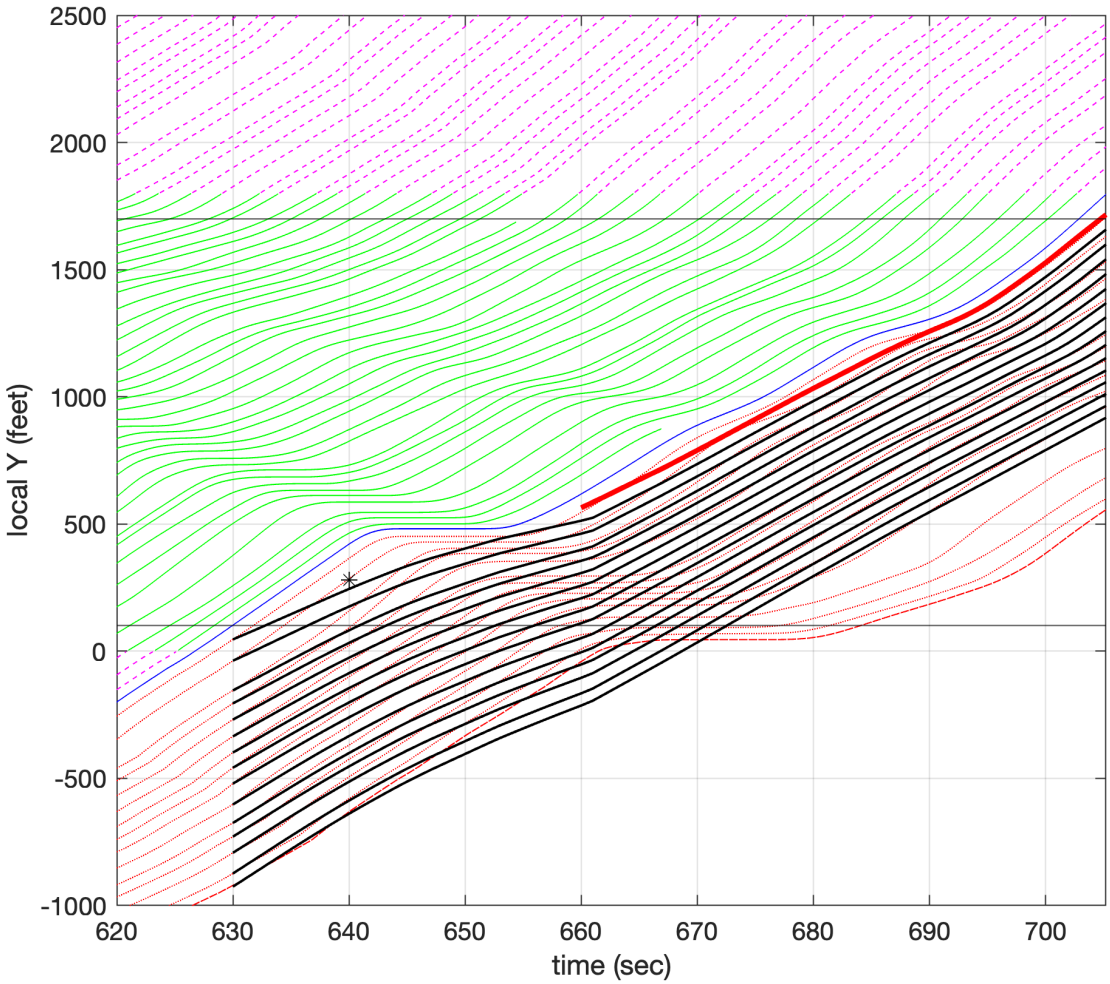

Figure 12, Case 1 with the fixed 1000 ft window and no limit on acceleration, with LCM request at 640 sec in front of the first CAV out of a total of 15 successive CAV followers. The entering vehicle trajectory is plotted in bold red solid curve. The CAV trajectories are shown in bold and the original HDV trajectories in a faint color for reference.

2.3.2. Experimental results

Fig. 12 presents the final trajectories for Case 1 with 15 CAV followers and a lane change request initiated at 640 sec. The look ahead window is consistently set to 1000 ft, in alignment with the settings used in previous analyses. The star marker in the figure denotes the requested time and location of the LCM, which falls within a gap created by the control model to absorb the shockwave. However, in accordance with the LCM policy, the vehicle does not merge at that point. Instead, the entering vehicle waits until the CAVs have passed through the stop wave. Following the policy, instead of accelerating after absorbing the stop wave, the first CAV transitions to a constant-speed trajectory until a sufficiently large gap is created. The entering vehicle then completes the merge at 660 sec. After the vehicle successfully merges, the previous first follower is now in the second follower position and the control reverts to the non-LCM strategy from Section 2.1. Throughout, all vehicles behind the original first CAV simply forecast their trajectory based on their immediate leader without having to directly account for any LCM that occur ahead of that leader.

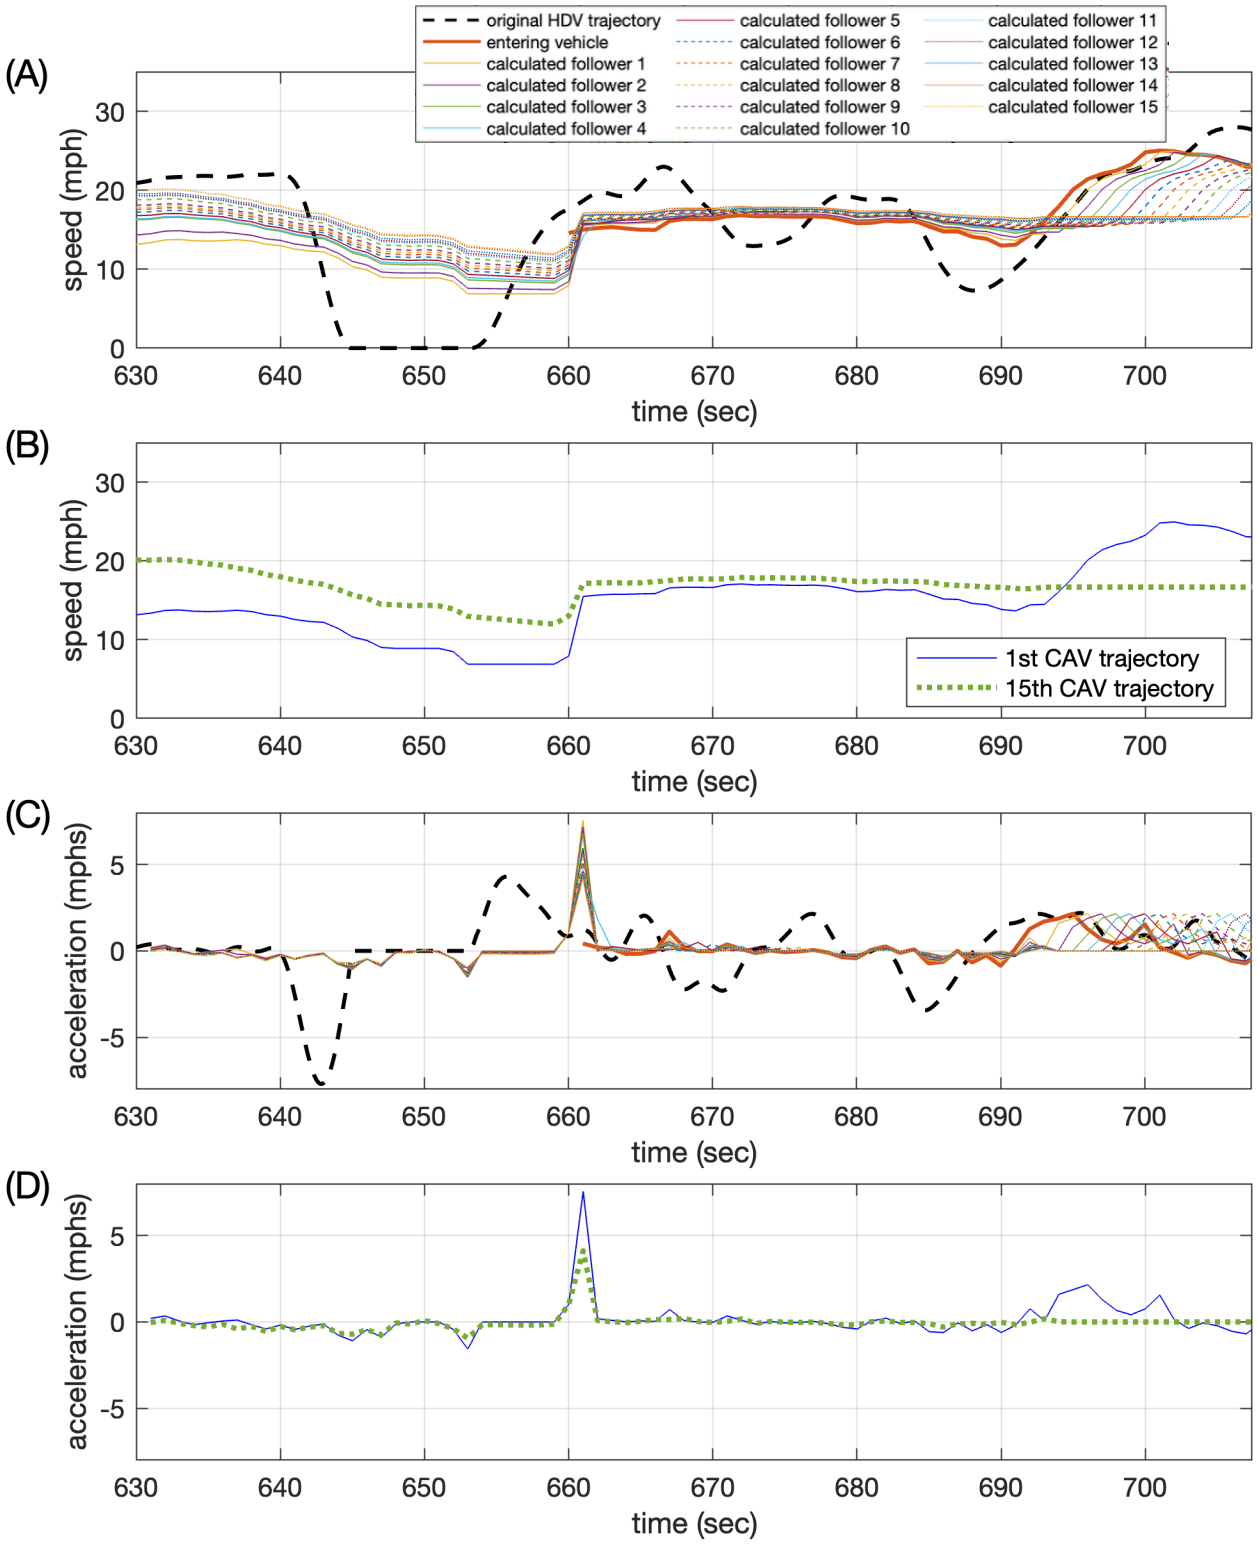

Figure 13, (A) Time series speed from the 15 CAV and the original trajectory for the first HDV follower along with the entering vehicle for Case 1 in Fig. 12. (B) Repeating A only now just showing the first and 15th CAVs. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first and 15th CAVs.

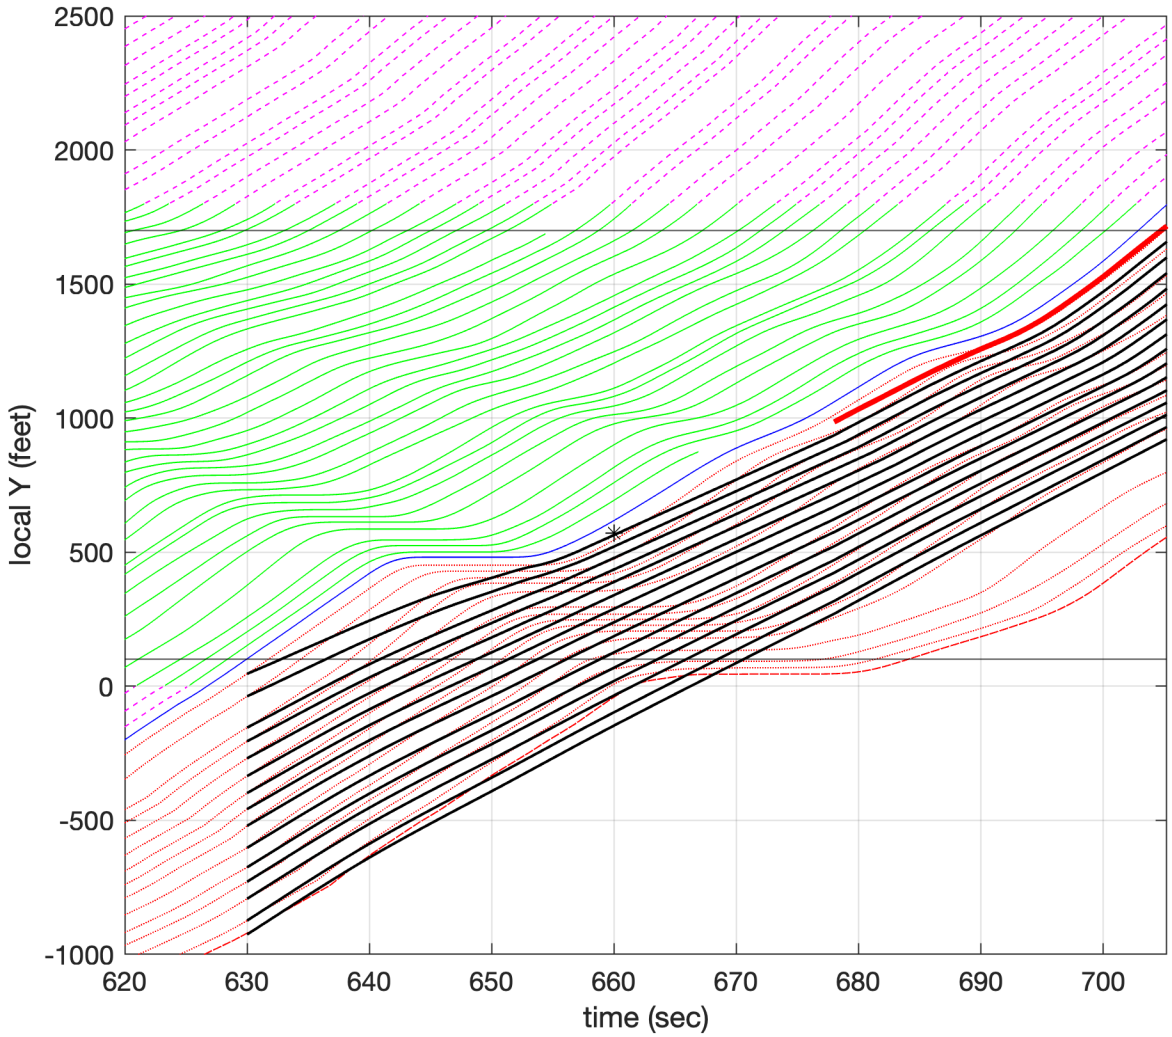

Figure 14, Case 1 with the fixed 1000 ft window and no limit on acceleration, with LCM request at 660 sec in front of the first CAV out of a total of 15 successive CAV followers. The entering vehicle trajectory is plotted in bold red solid curve. The CAV trajectories are shown in bold and the original HDV trajectories in a faint color for reference.

In this implementation the vehicle behind the LCM immediately transitions to its new speed after reaching the second perfect follower trajectory. This change in trajectory is passed to all of the following CAV, as evident in the time series speed, Fig. 13A. When limiting the plot to just the original first and 15th followers, Fig. 13B shows that the speed jump is slowly absorbed in the platoon. Fig. 13C-D show the corresponding acceleration and the transient spike is evident in all of the CAVs, approaching 8 mphps for the first CAV and dropping to 4 mphps for the last CAV in the platoon. We will address this transient in the next subsection.

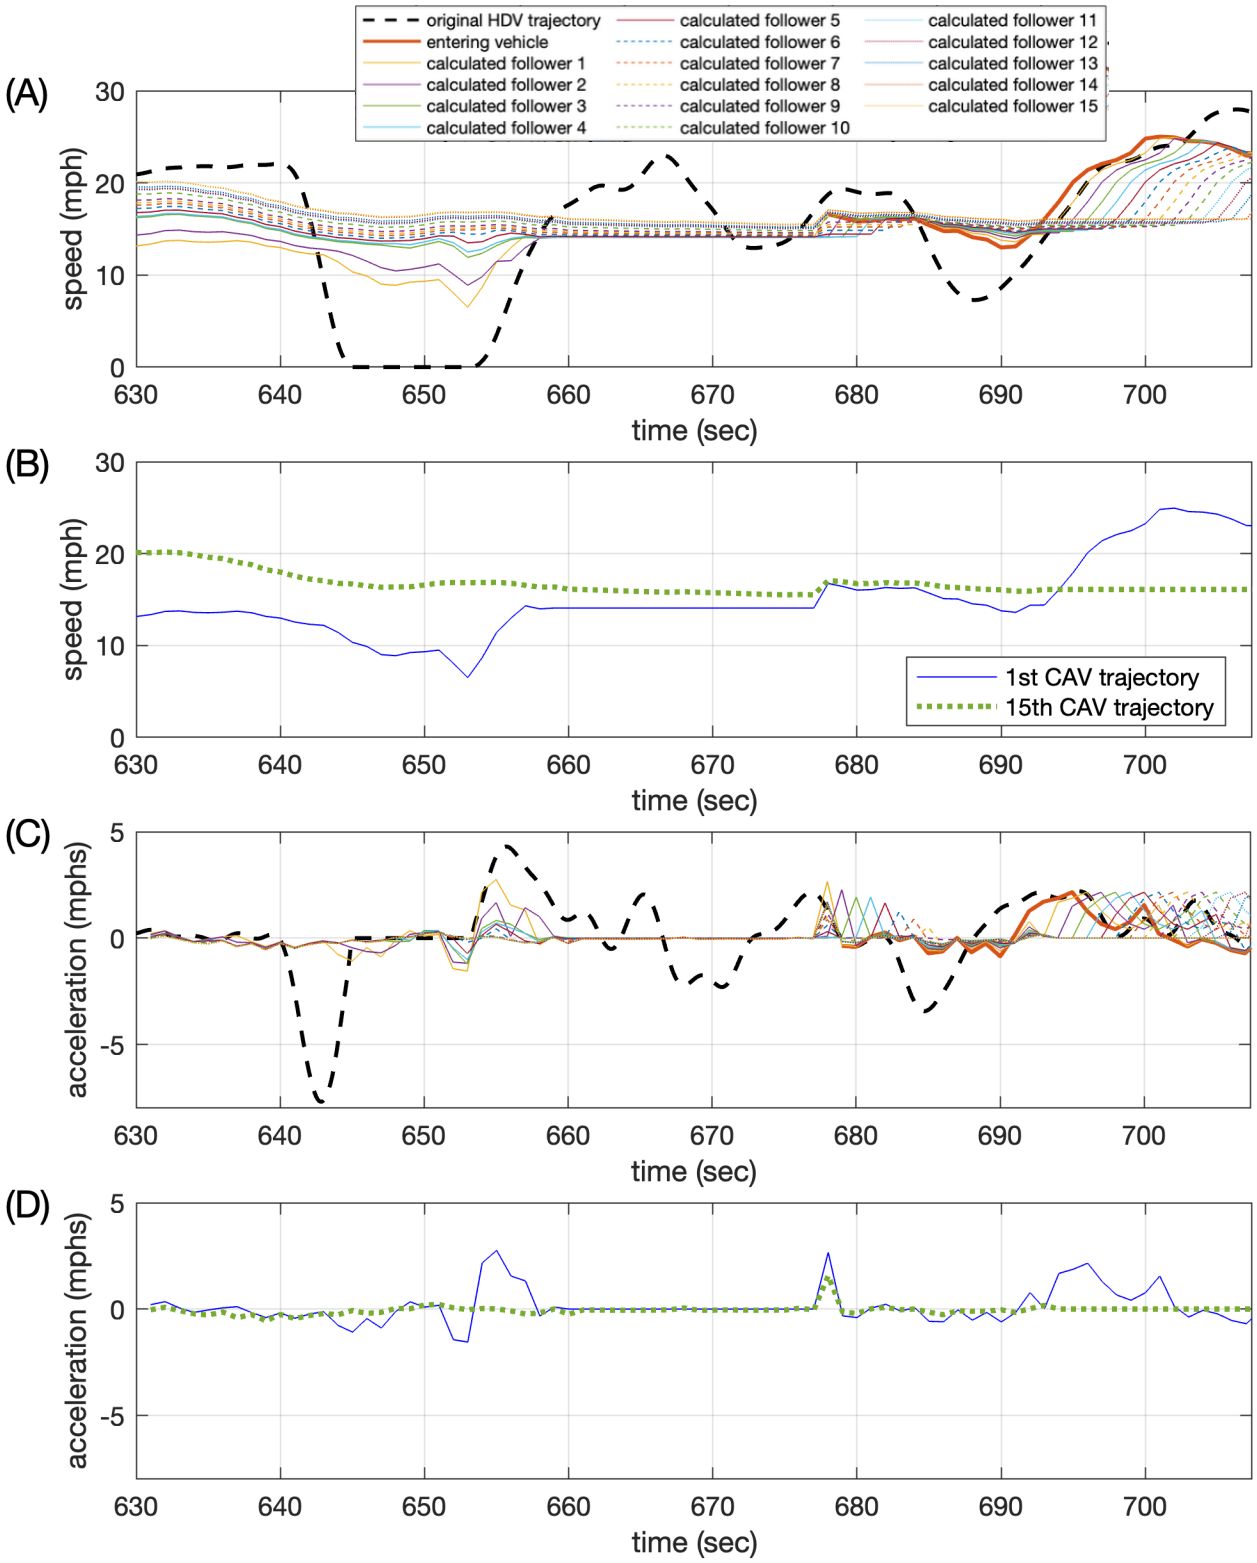

Figure 15, (A) Time series speed from the 15 CAV and the original trajectory for the first HDV follower along with the entering vehicle for Case 1 in Fig. 14. (B) Repeating A only now just showing the first and 15th CAVs. (C) The corresponding acceleration from A, and (D) repeating C only now just showing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

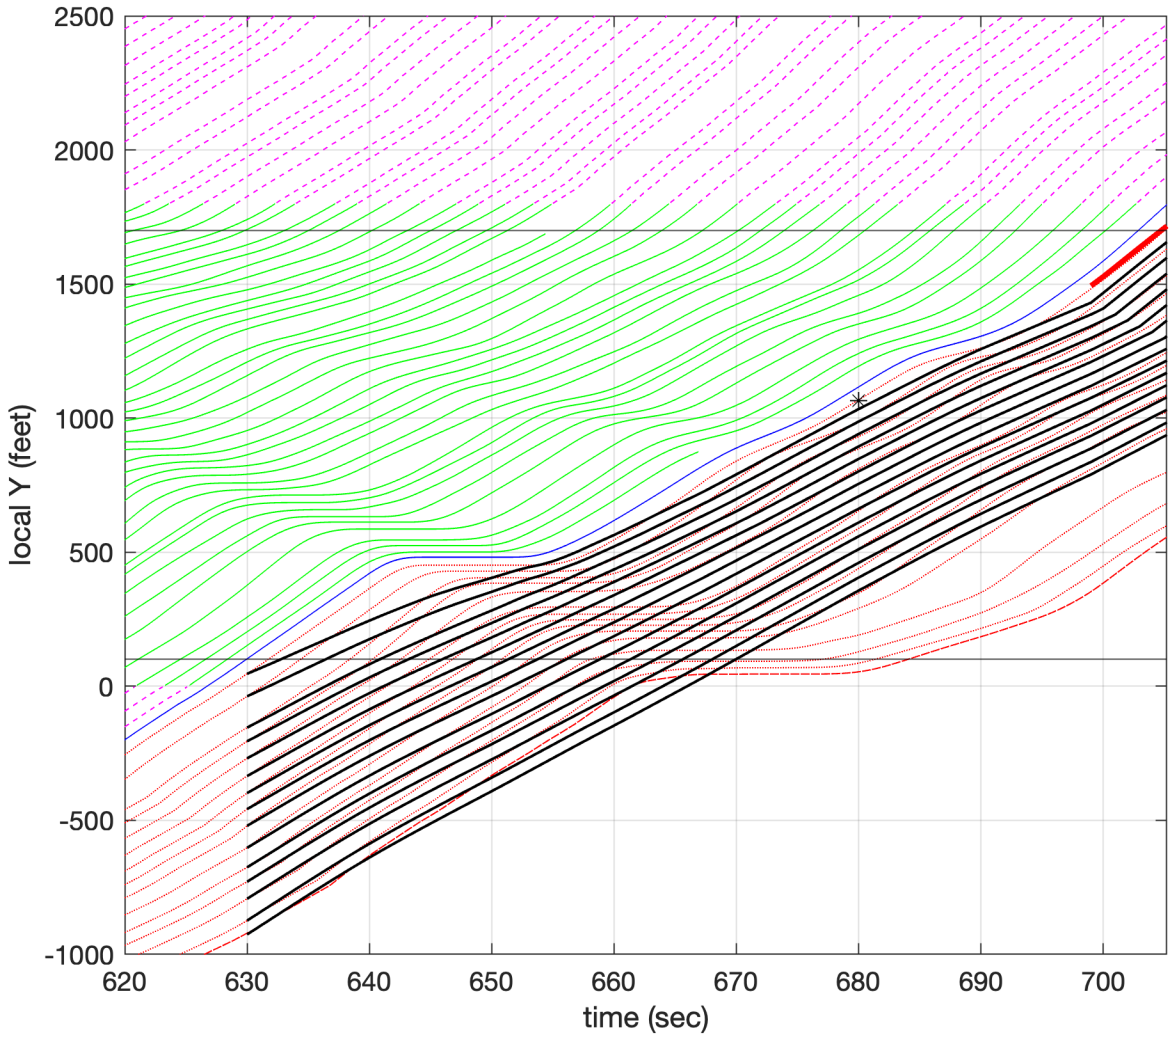

Figure 16, Case 1 with the fixed 1000 ft window and no limit on acceleration, with LCM request at 680 sec in front of the first CAV out of a total of 15 successive CAV followers. The entering vehicle trajectory is plotted in bold red solid curve. The CAV trajectories are shown in bold and the original HDV trajectories in a faint color for reference.

Fig. 14 repeats the experiment for Case 1 only now the lane change request is initiated 20 sec later, at 660 sec. The request location is again marked by a black star in the time–space diagram, occurring after the first CAV has successfully navigated through the stop wave. Upon receiving the request, the first CAV transitions to a constant-speed trajectory. Since the HDV leader is accelerating at this point, a sufficient gap is created relatively quickly, allowing the entering vehicle to merge into the lane at 678 sec. Now acting as the new first follower, the entering vehicle maintains an almost constant speed, aside from minor fluctuations between 680 and 690 sec, before continuing in standard car-following mode. This transition persists until a new stop wave is detected in the 1000 ft look ahead from the downstream cHDV. Fig. 15A-B shows the corresponding time series speed for the CAVs. The flat segment in the speed profile between 660 and 678 sec corresponds to the constant-speed phase of the original first follower prior to the entrance. In this case the speed of the former first follower is just below that of the second perfect follower at the switchover, so there was not a large transient speed jump.

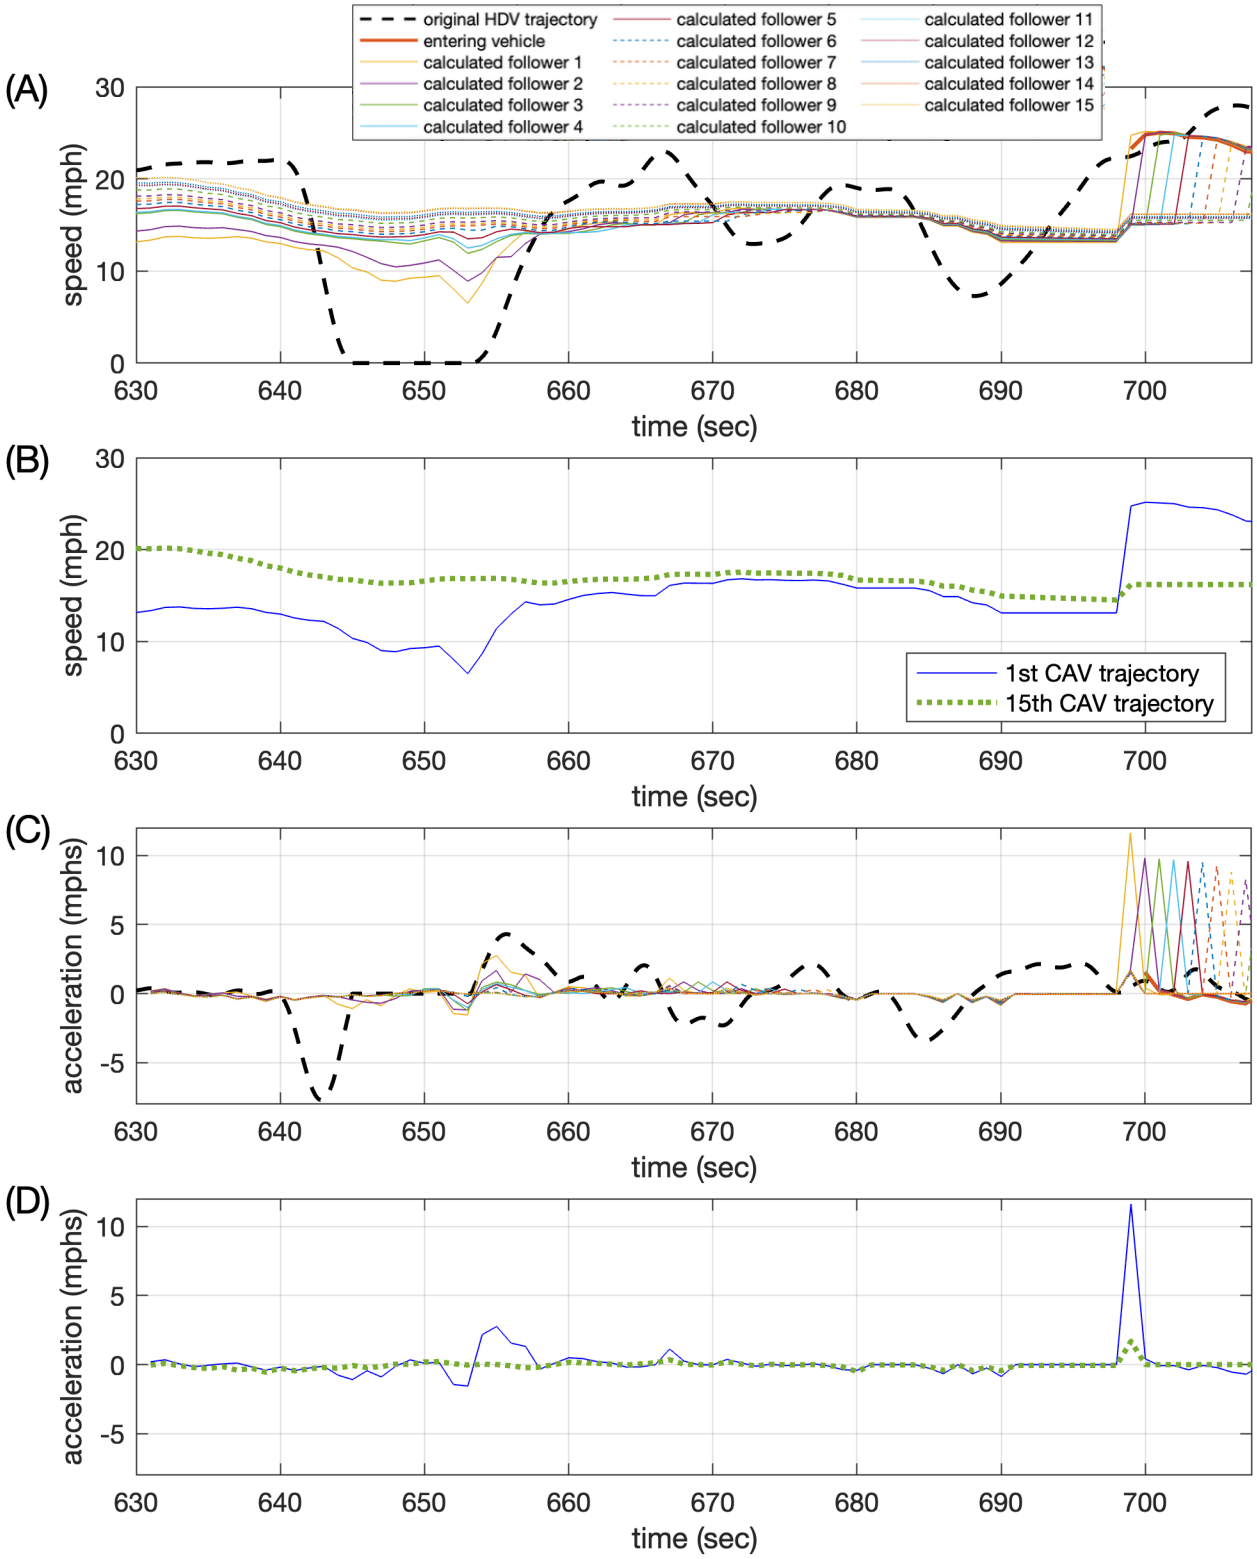

Figure 17, (A) Time series speed from the 15 CAV and the original trajectory for the first HDV follower along with the entering vehicle for Case 1 in Fig. 16. (B) Repeating A only now just showing the first and 15th CAVs. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first and 15th CAVs.

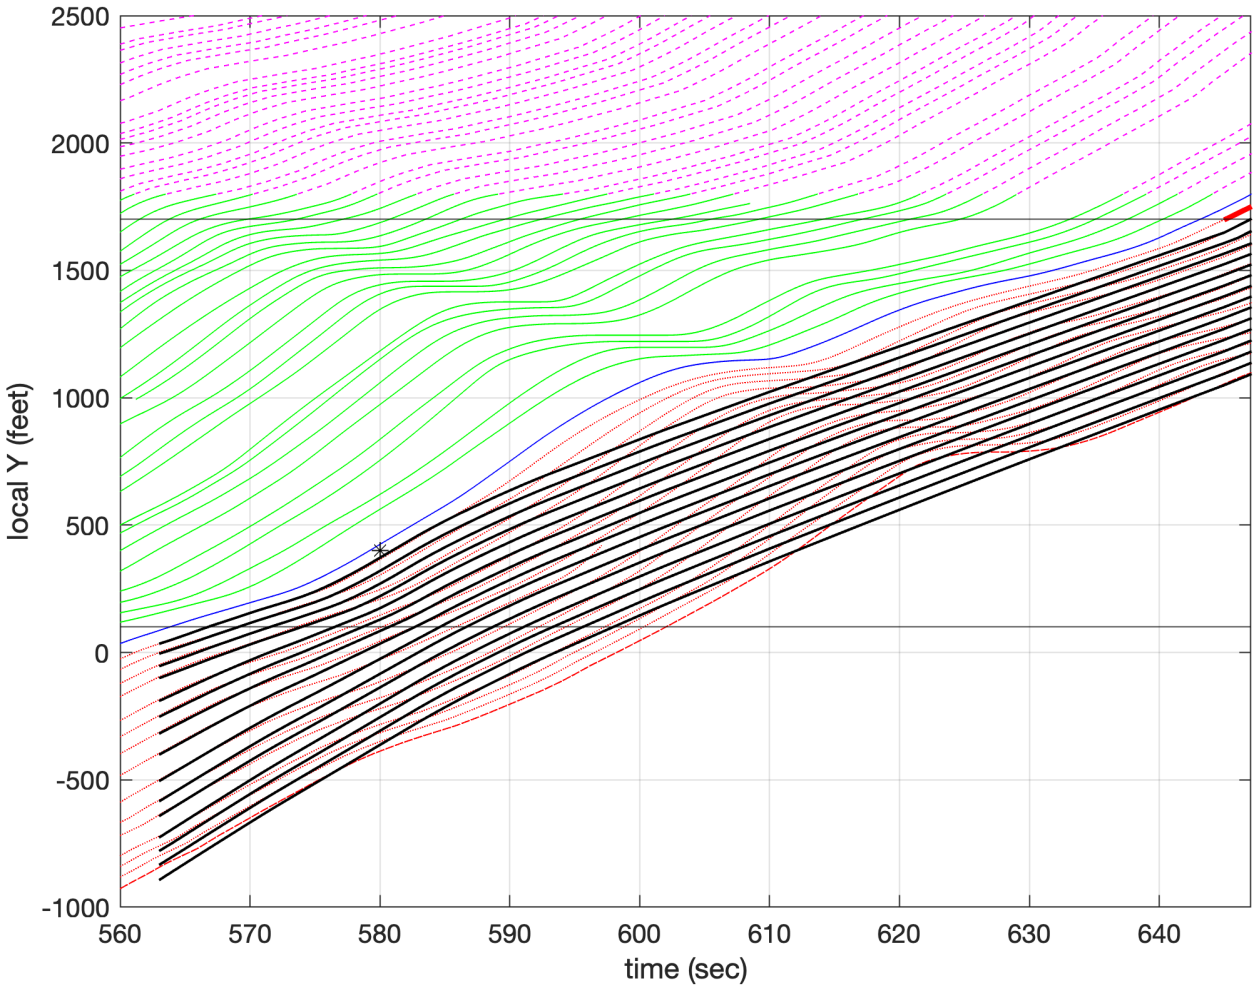

Figure 18, Case 2 with the fixed 1000 ft window and no limit on acceleration, with LCM request at 580 sec in front of the first CAV out of a total of 15 successive CAV followers. The entering vehicle trajectory is plotted in bold red solid curve. The CAV trajectories are shown in bold and the original HDV trajectories in a faint color for references.

Fig. 16 repeats the experiment for Case 1 with the lane change request coming at 680 sec. In this scenario, the entrance is delayed due to the presence of another slow wave occurring between 680 and 690 sec. Consequently, the original first CAV must follow its planned trajectory, as if no LCM is occurring. Once the stop wave subsides, the first CAV switches to a constant-speed trajectory in preparation for the merge. However, because this action takes place immediately after a deceleration phase, the constant speed is significantly lower than that of the leader at the moment of entrance. This mismatch results in a sharp jump in speed for the CAVs behind the entrance, as shown in Fig. 17. Fig. 17D shows the transient acceleration approaching 12 mphps for the first CAV—well beyond the capability of a typical passenger vehicle, but by the 15th follower the transient spike has dropped below 1 mphps.

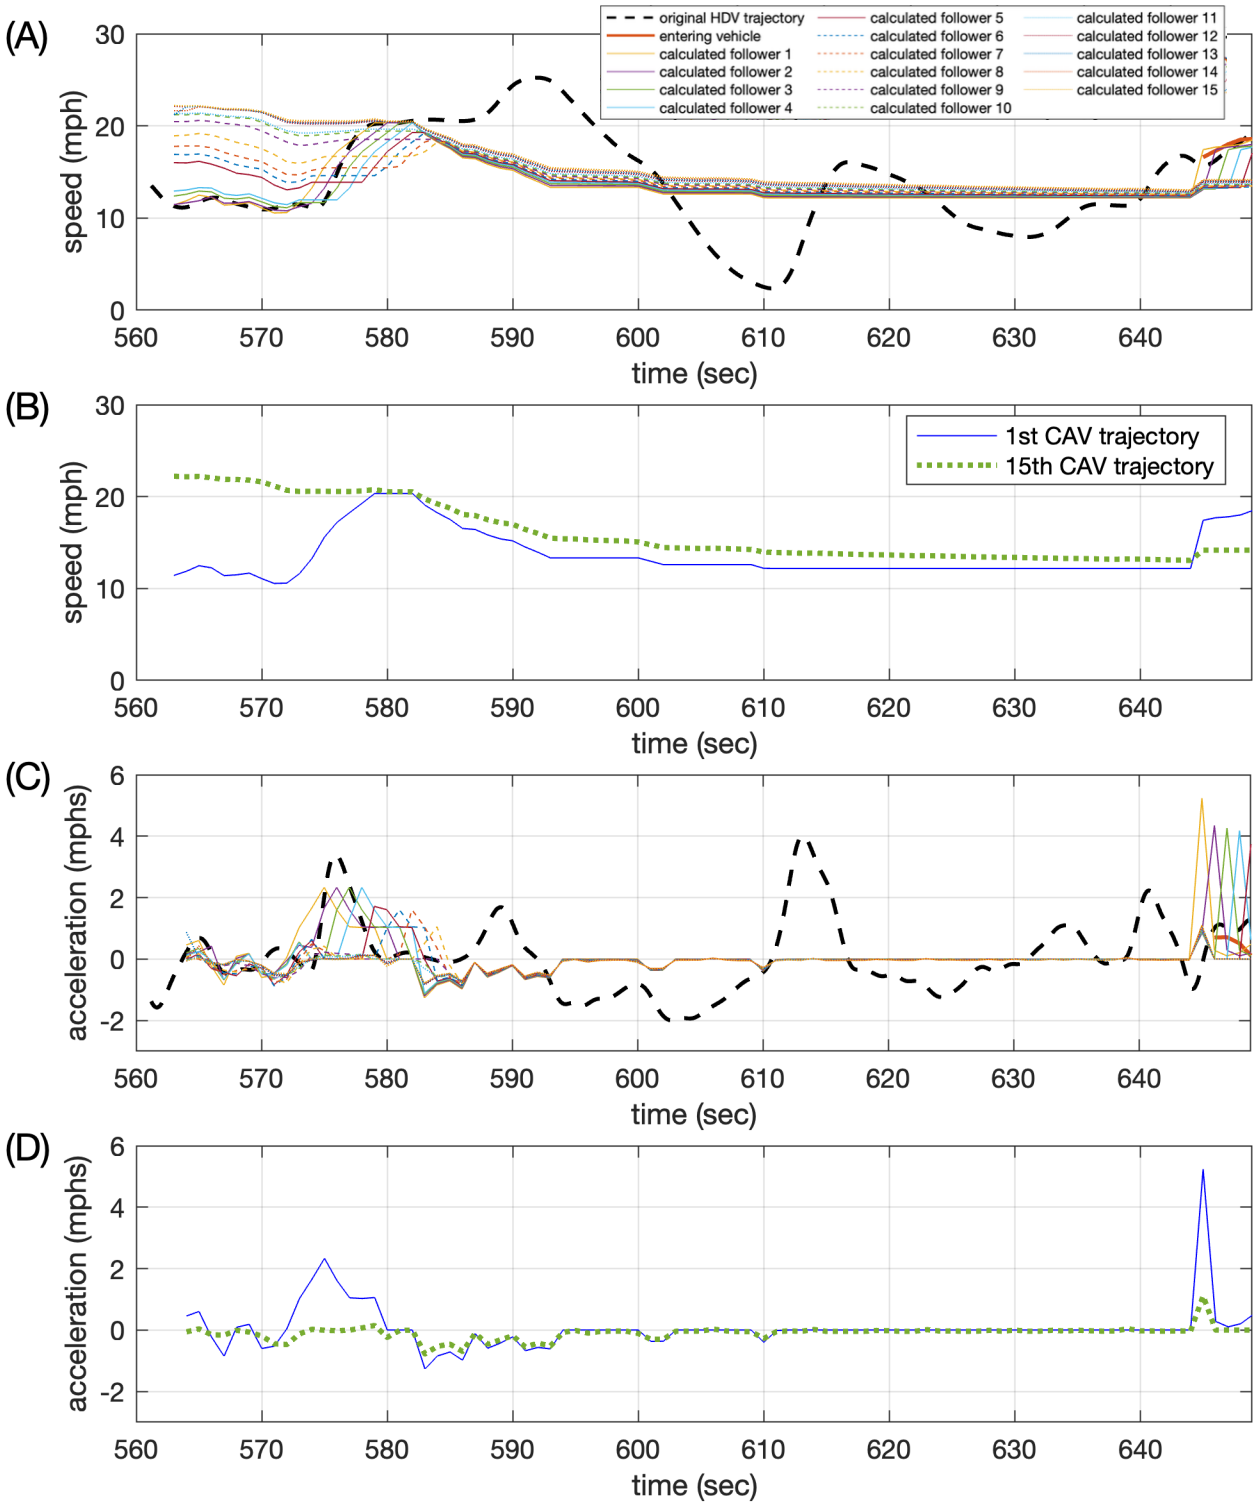

Figure 19, (A) Time series speed from the 15 CAV and the original trajectory for the first HDV follower along with the entering vehicle for Case 2 in Fig. 18. (B) Repeating A only now just showing the first and 15th CAVs. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first and 15th CAVs.

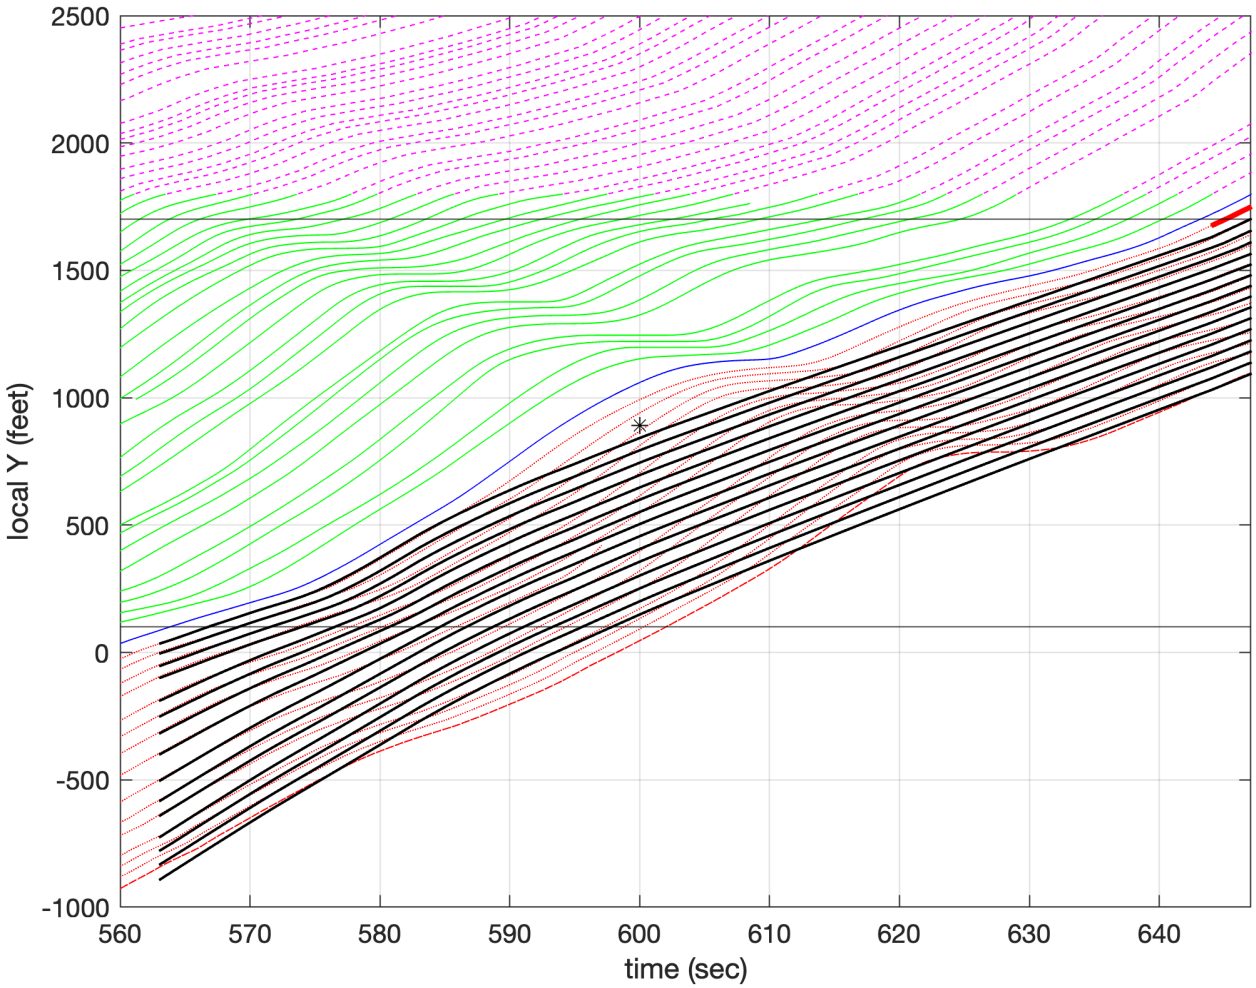

Figure 20, Case 2 with the fixed 1000 ft window and no limit on acceleration, with LCM request at 600 sec in front of the first CAV out of a total of 15 successive CAV followers. The entering vehicle trajectory is plotted in bold red solid curve. The CAV trajectories are shown in bold and the original HDV trajectories in a faint color for reference.

Fig. 18 presents the final trajectories for Case 2 with 15 CAV followers and a lane change request initiated at 580 sec. Unlike Case 1, where the stop wave passes the first CAV around the midpoint of the empirical data—Case 2 exhibits a different dynamic: the stop wave does not reach the first CAV until near the end of the empirical data. In accordance with the established LCM policy, the entering vehicle is not permitted to utilize the gap that was intentionally created to absorb downstream fluctuations. Consequently, the entering vehicle remains in the adjacent lane until after the stop wave has been absorbed, and the LCM vehicle finally enters at 645 sec, just moments before the first CAV exits the supervised region. After the merge, the original first follower becomes the second follower and begins to follow the entering vehicle, adjusting its speed based on the new leader's planned trajectory. Fig. 19 shows a speed jump and acceleration spike across all CAVs in the platoon, similar to what was seen in Case 1. Here too, the transient diminishes through successive CAV, with the acceleration dropping by roughly 70% from the first to last CAV (Fig. 19D).

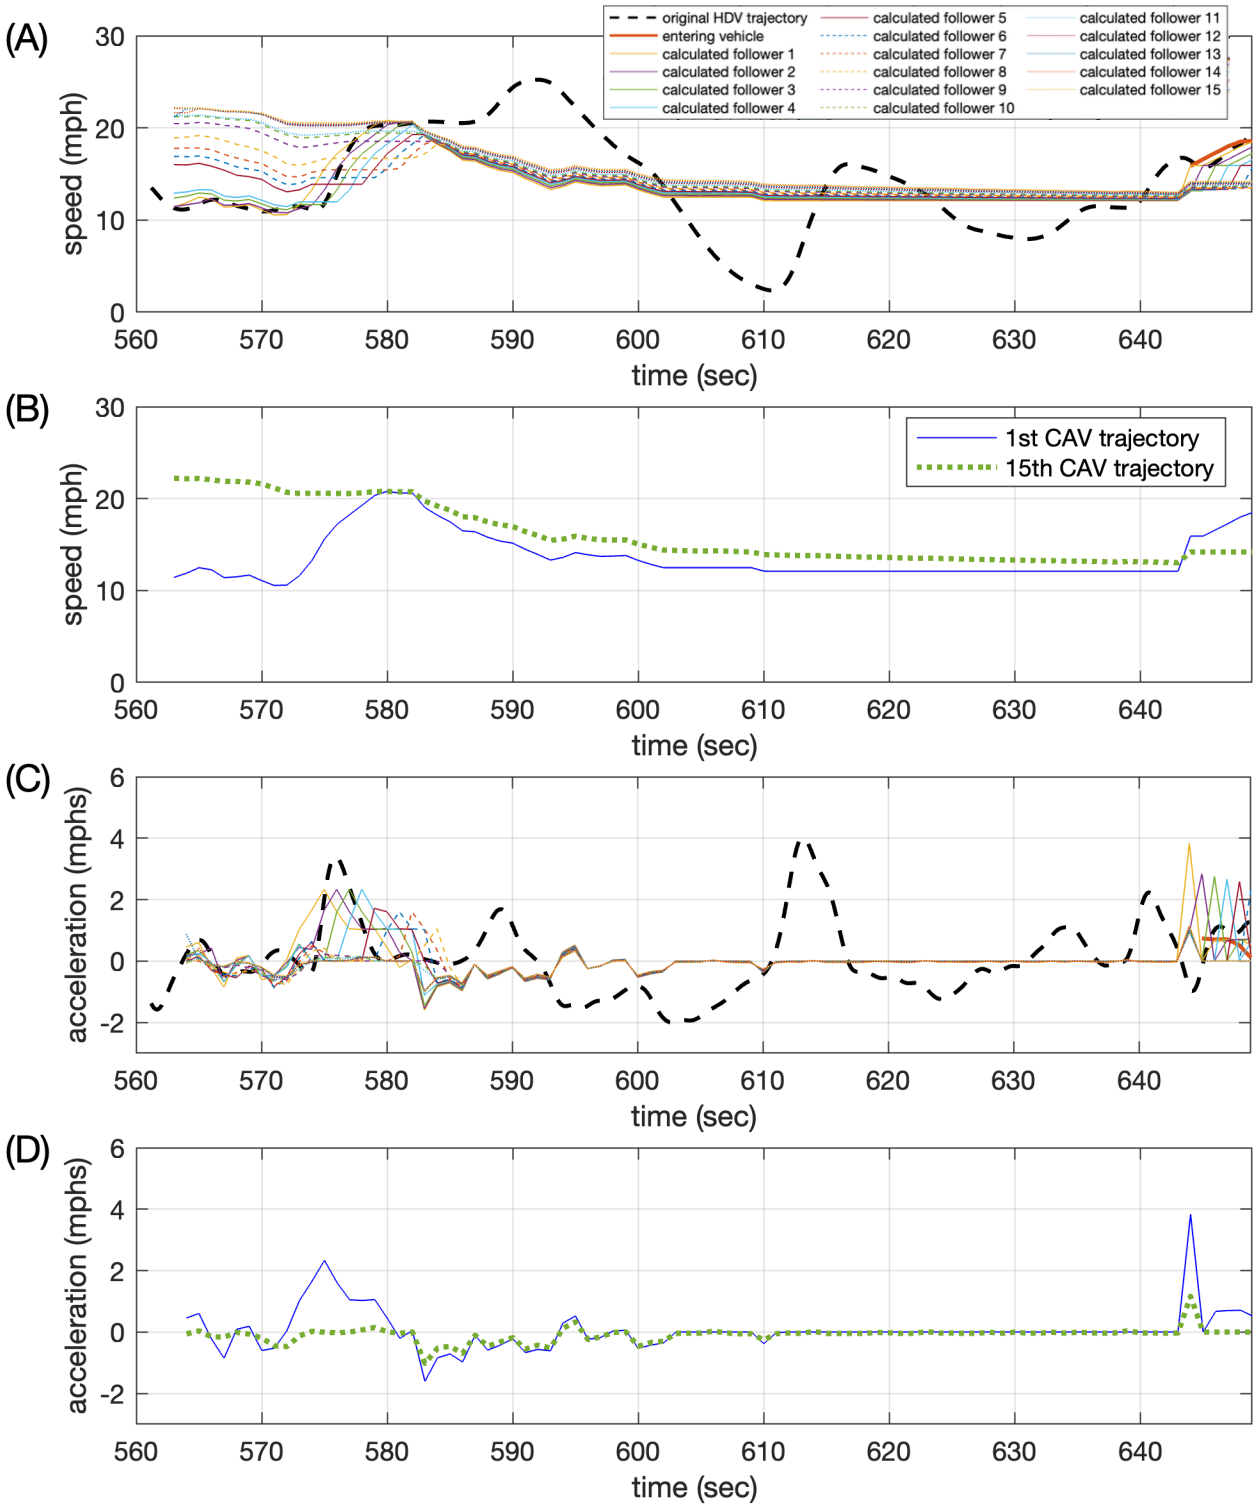

Figure 21, (A) Time series speed from the 15 CAV and the original trajectory for the first HDV follower along with the entering vehicle for Case 2 in Fig. 20. (B) Repeating A only now just showing the first and 15th CAVs. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first and 15th CAVs.

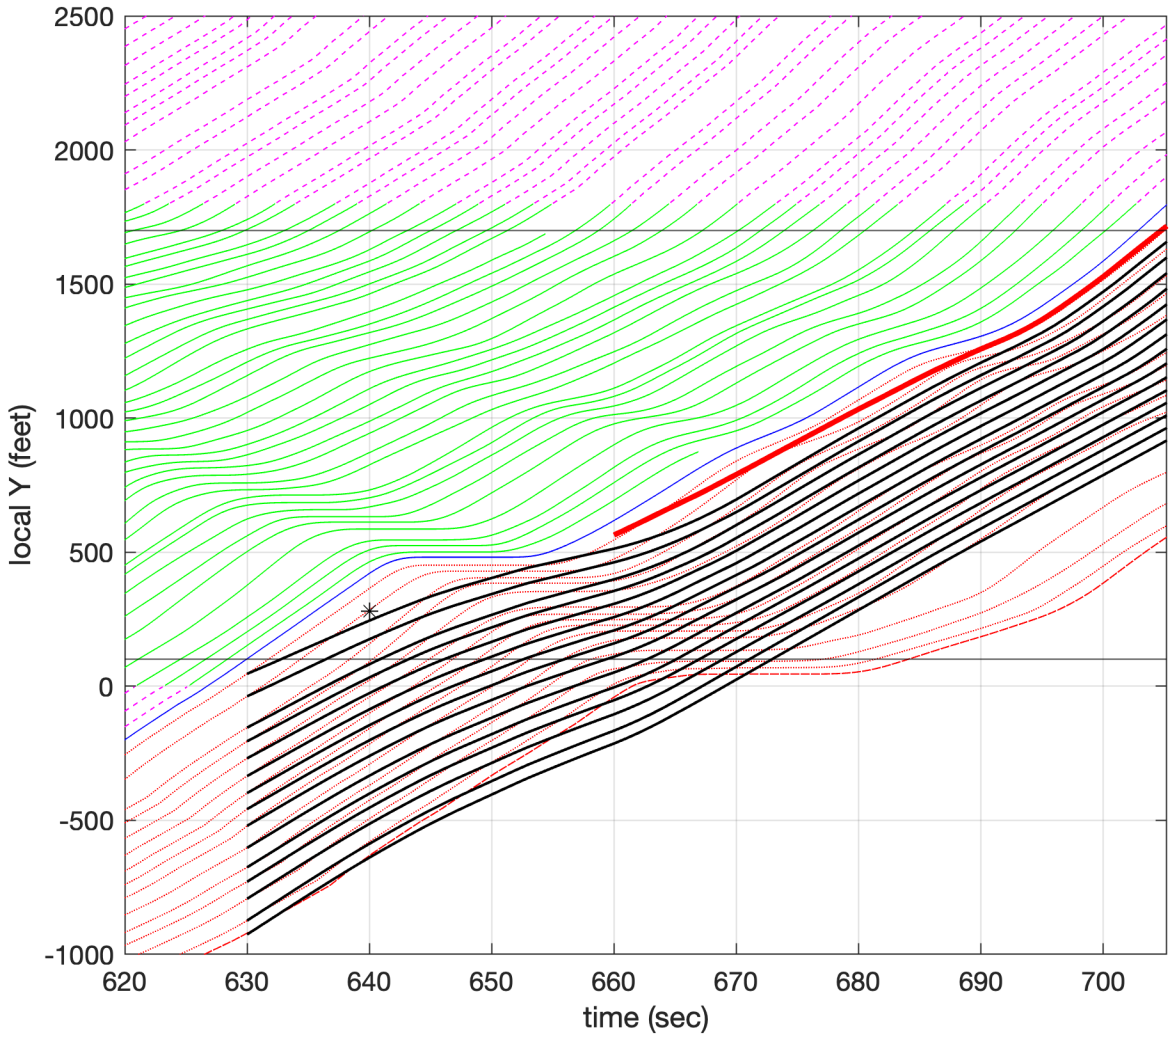

Figure 22, The final CAV trajectory for Case 1 with the fixed 1000 ft window and LCM request at 640 sec with $a_{CAV} \leq 1\ mphps$ only while accommodating the entrance.

Fig. 20 shows the results from Case 2 for a LCM request at 600 sec—timed to fall within the gap deliberately created to absorb the stop wave. As per the LCM policy, again the entering vehicle is held in the adjacent lane until the stop wave has fully passed the first CAV. Fig. 21D shows that after the request, the first CAV slows twice in response to the cHDV leader and then maintains a constant speed from 610 sec until a sufficient gap becomes available for the vehicle to enter. Because the lag before the entering vehicle can affect their maneuver, the results in Fig. 21 are very similar to the previous example. Still, these results demonstrate the platoon's inherent ability to absorb and dissipate localized disturbances, maintaining overall flow stability.

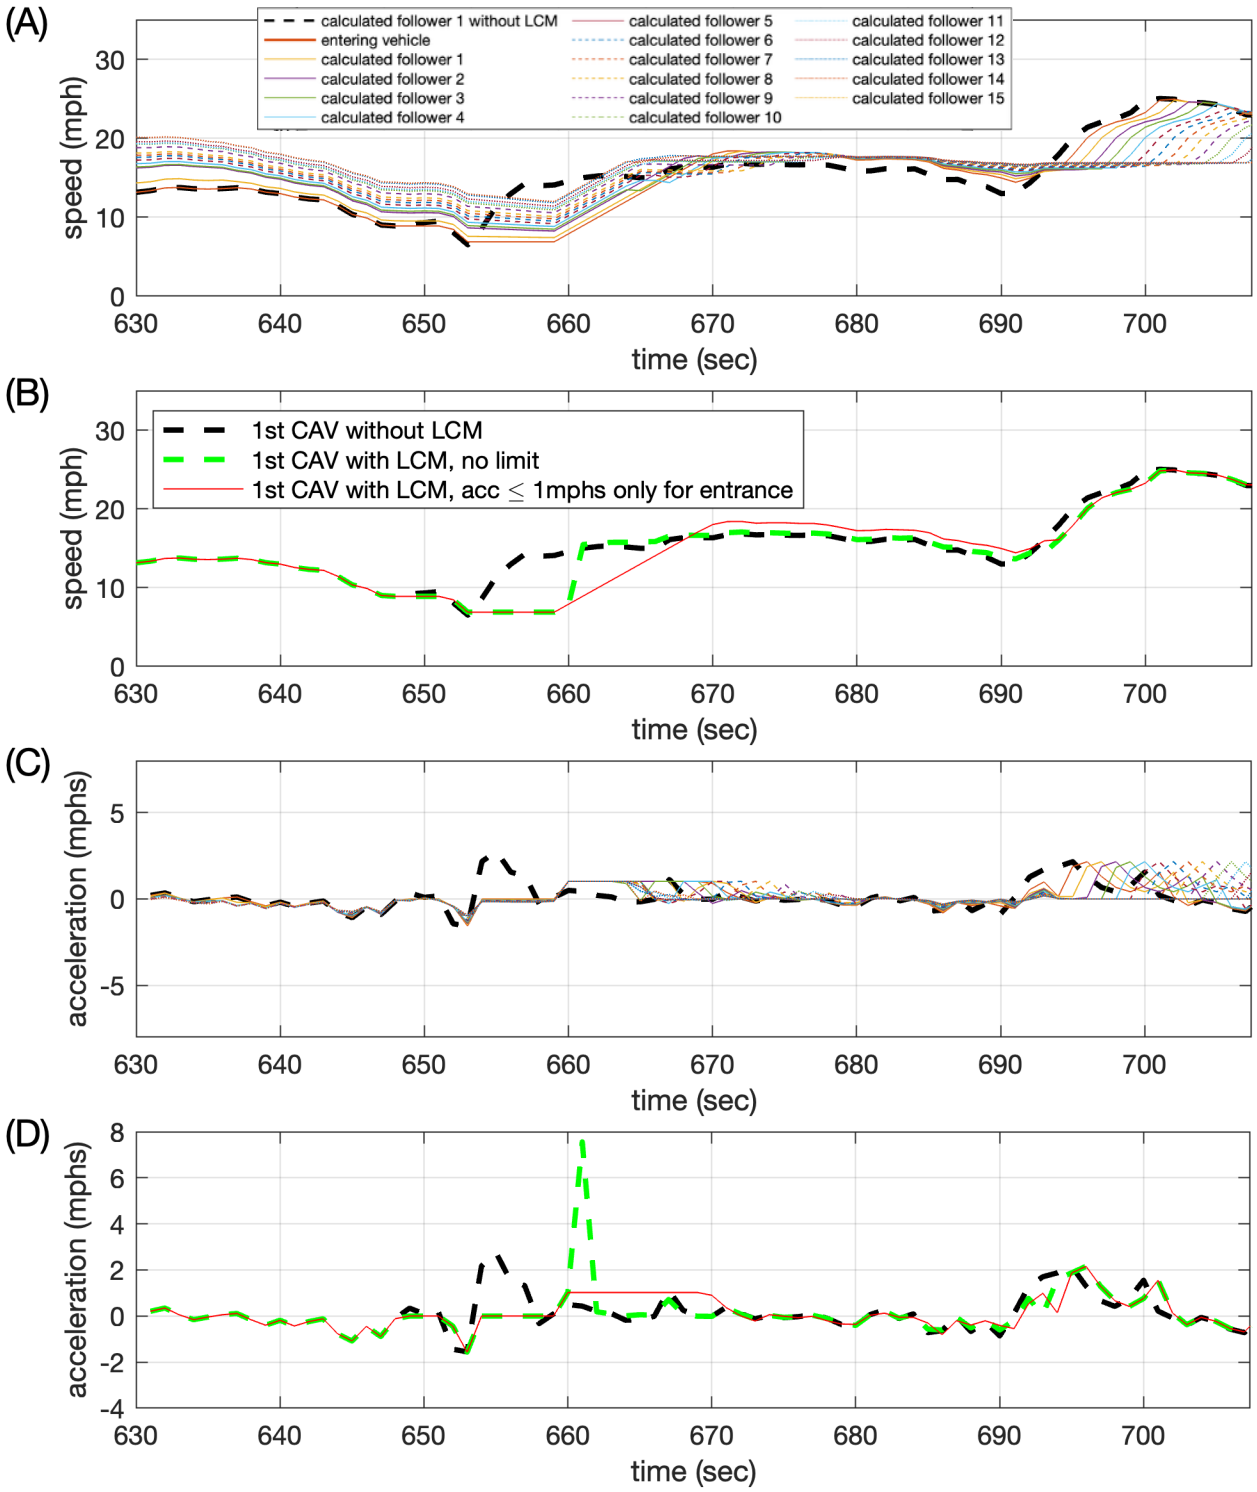

Figure 23, (A) Time series speed for Case 1 with the fixed 1000 ft window and LCM request at 640 sec for the first CAV follower with $a_{CAV} \leq 1\ mphps$ only while accommodating the entrance in Fig. 22. (B) Repeating A only now just showing the first CAV. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first CAV.

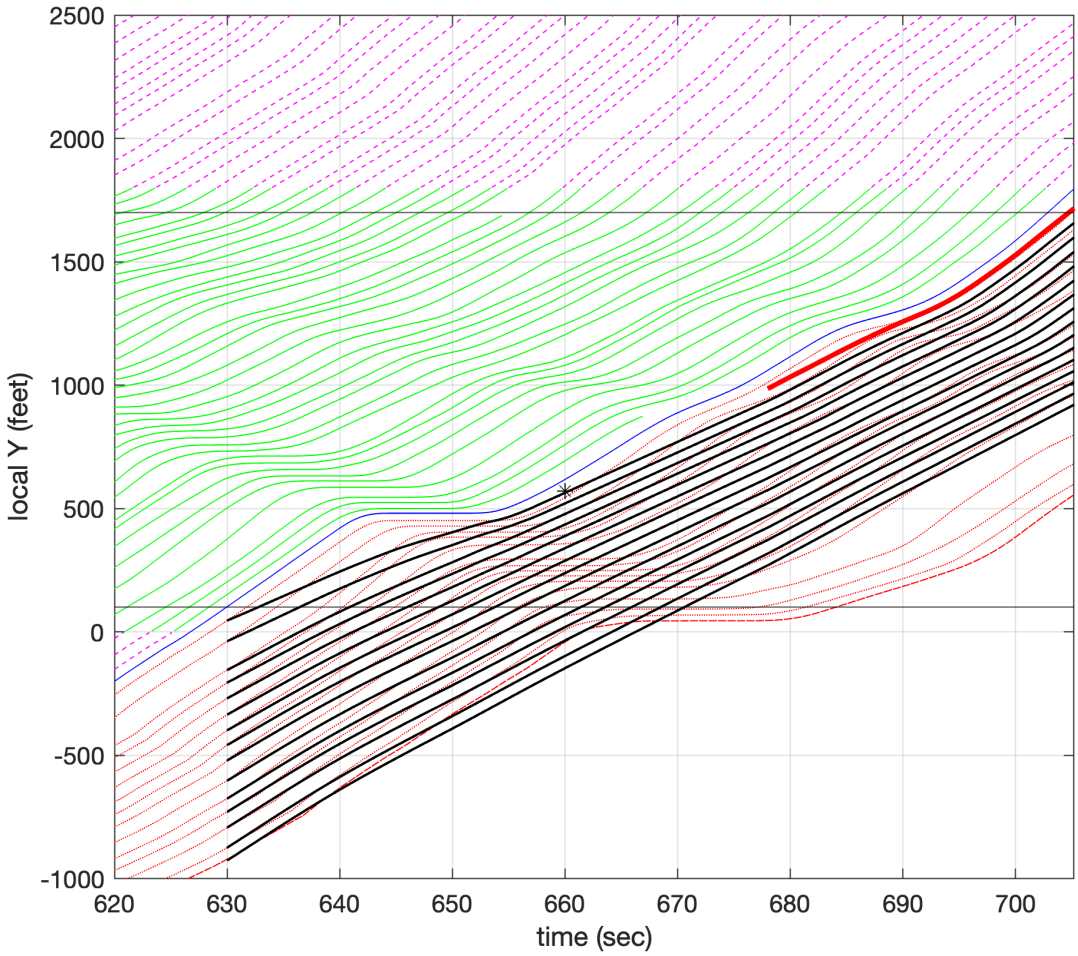

Figure 24, The final CAV trajectory for Case 1 with the fixed 1000 ft window and LCM request at 660 sec with $a_{CAV} \leq 1\ mphps$ only while accommodating the entrance.

2.3.3. Limiting acceleration when accommodating the entrance

As seen in the previous subsection, under the base method the CAV follower behind an entrance abruptly transitions from making a gap to following the new leader. Sometimes the change in speed is large, giving rise to a large transient acceleration, e.g., Fig. 13C-D. This subsection extends the method to constrain the acceleration as the CAV behind an entrance starts to follow its new leader. As a result of the speed differential, a transient gap will form while the follower catches up, which could be misinterpreted as an underutilization of space in the target lane. But like the transient gap behind the cHDV in Fig. 2A, this transient gap simply reflects part of the smoothing process and closes over time. After the gap closes, the fact that it had existed has no impact on the travel time of individual vehicles or the throughput of the lane, i.e., the overall exit time of the vehicles is not significantly impacted. More importantly, the coordinated behavior of the CAV platoon prevents the formation of new stop waves or disturbances among the followers. For the following examples, the maximum acceleration for the CAV behind an entrance is capped at 1 mphps until the vehicle behind has resumed following its leader. In this subsection there is no limit on acceleration away from this maneuver (i.e., all other vehicles, and for the accommodating vehicle both before and after allowing the new leader in). The next subsection will consider the case where acceleration is simply always capped at 1 mphps for all CAVs without any special consideration for LCM.

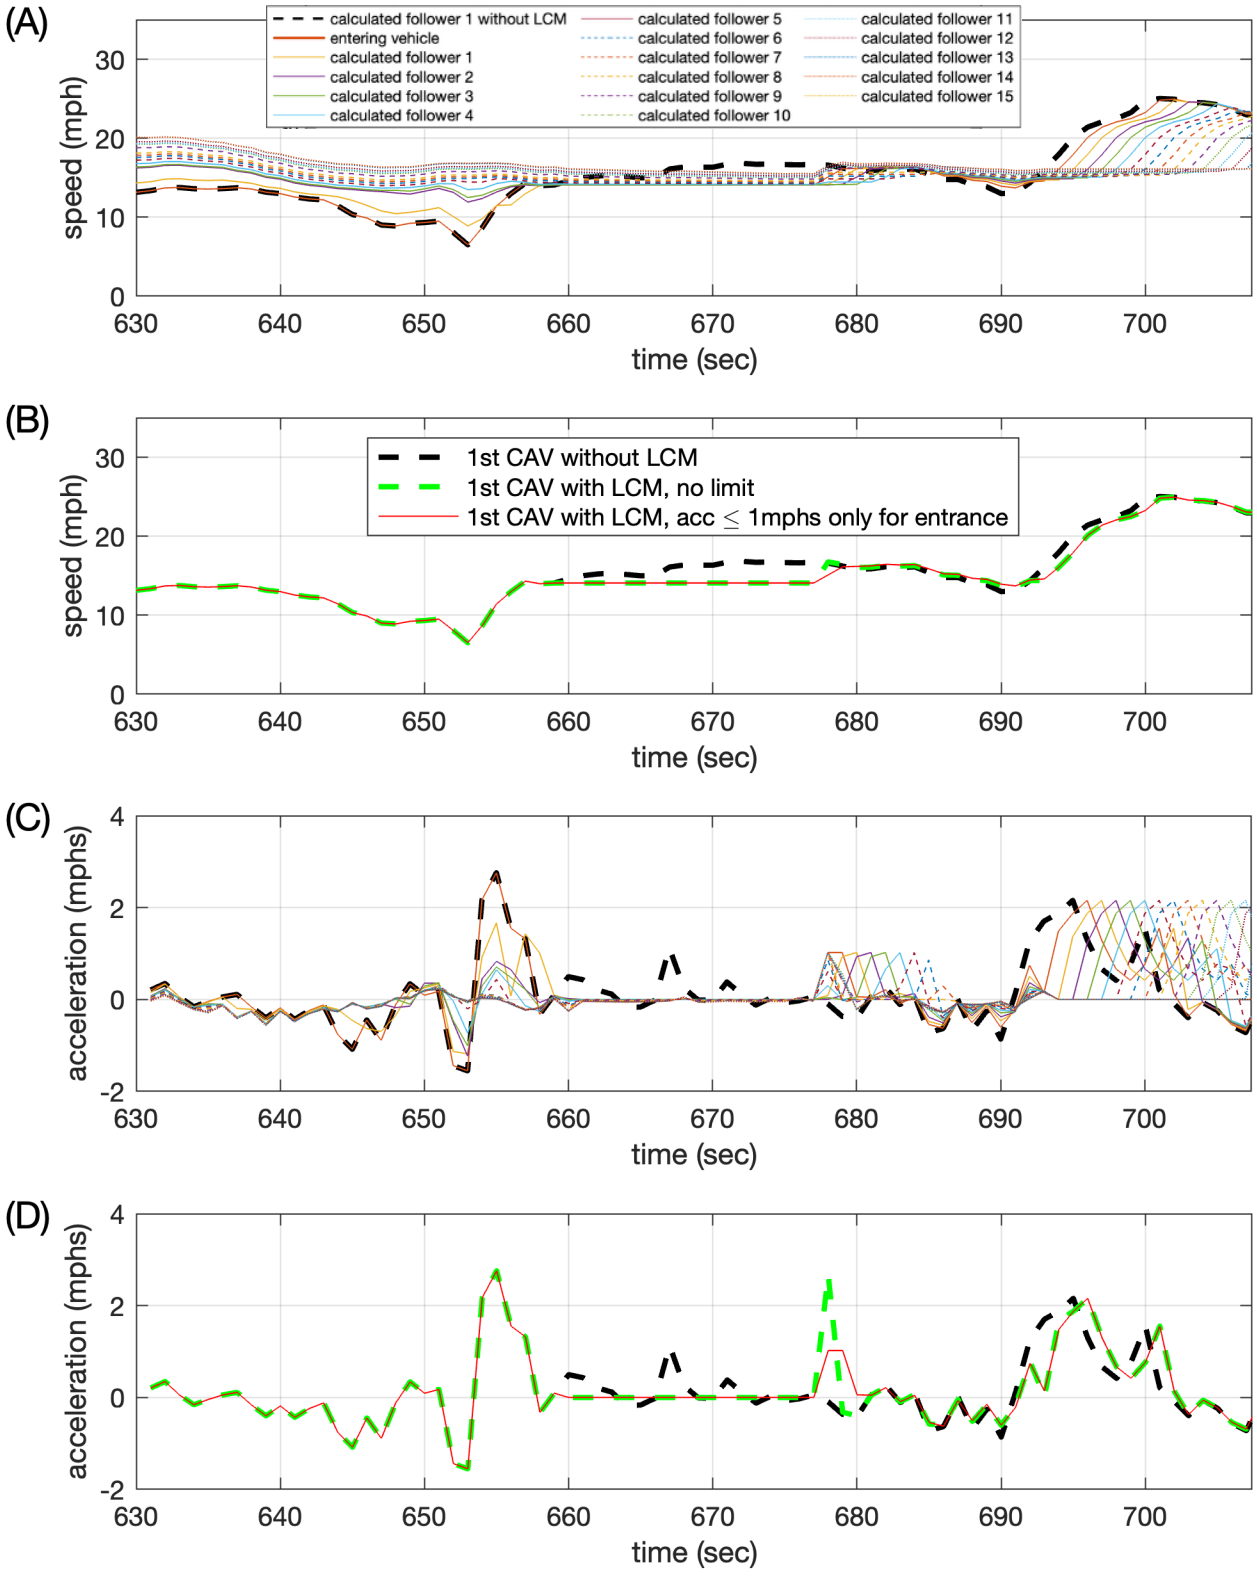

Figure 25, (A) Time series speed for Case 1 with the fixed 1000 ft window and LCM request at 660 sec for the first CAV follower with $a_{CAV} \leq 1\ mphps$ only while accommodating the entrance in Fig. 24. (B) Repeating A only now just showing the first CAV. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first CAV.

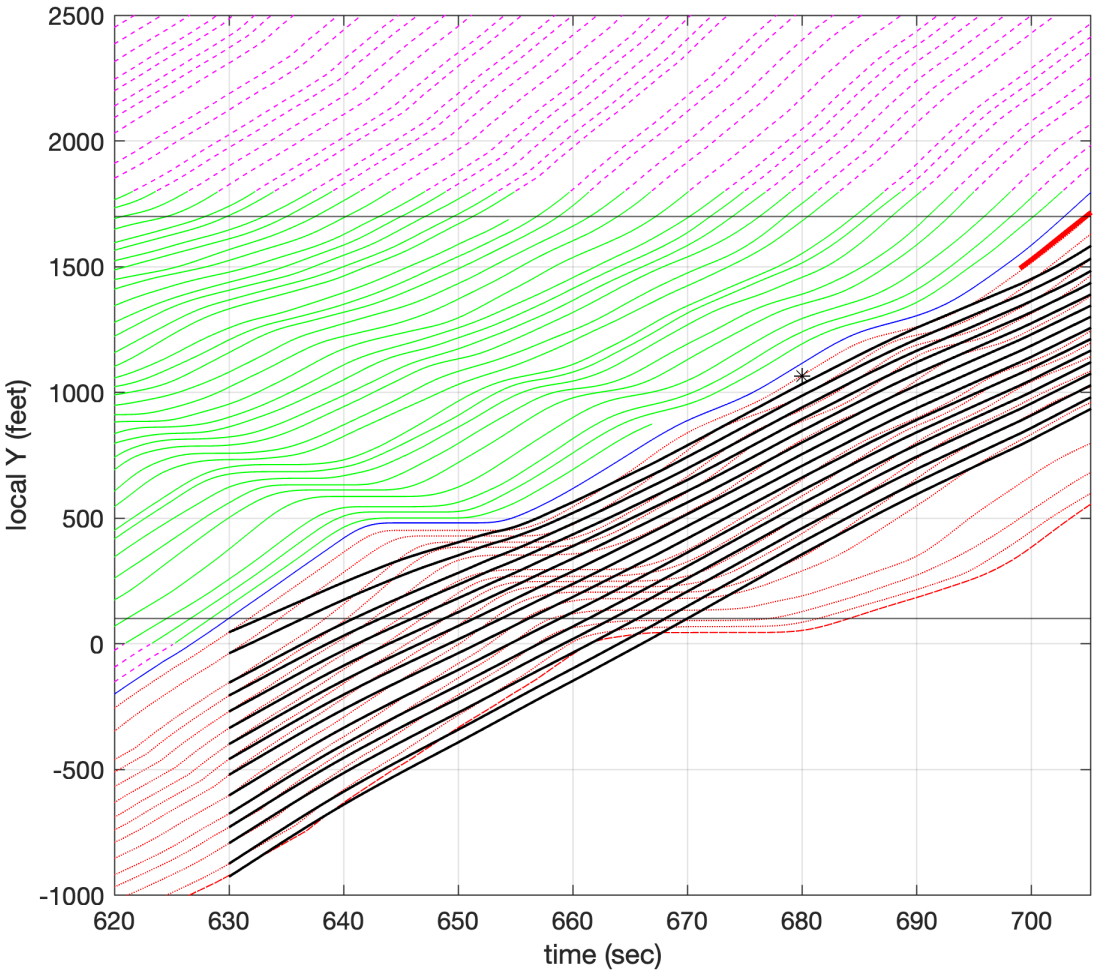

Figure 26, The final CAV trajectory for Case 1 with the fixed 1000 ft window and LCM request at 680 sec with $a_{CAV} \leq 1$ *mphps* only while accommodating the entrance.

Fig. 22 returns to Case 1 with a lane change request at 640 sec (compare to Fig. 12 for the unlimited acceleration version). The look ahead window is maintained at 1000 ft. The acceleration constraint helps the CAV mitigate its acceleration after the entrance at 660 sec. Fig. 23 illustrates the benefits in terms of the time series speed and acceleration plots. In Fig. 23B and D, compare the acceleration "only for entrance" (from Fig. 22) against "no limit" (from Fig. 12). Notice how the speed adjustment starting at 660 sec is now spread over 10 sec instead of a single time step.

Fig. 24 shows the results for Case 1 with a lane change request at 660 sec (compare to Fig. 14 for the unlimited acceleration version). In this more favorable scenario, the difference between constrained and unconstrained trajectories is less pronounced than in the previous example. Fig. 25 shows the corresponding time series speed and acceleration. The difference is most pronounced in Fig. 25D, where the "only for entrance" curve truncates the peak at 678 sec evident in the "no limit" curve. Fig. 26 repeats the analysis for Case 1 with a lane change request at 680 sec. Due to the slow movement of the leader near 690 sec, the entrance is delayed until approximately 698 sec—just before the leader exits the surveillance region. Fig. 27 shows the corresponding time series speed and acceleration. Here the contrast between "only for entrance" and "no limit" curves is much larger, e.g., Fig. 27D shows the large spike has been flattened into a prolonged period of 1 mphps acceleration.

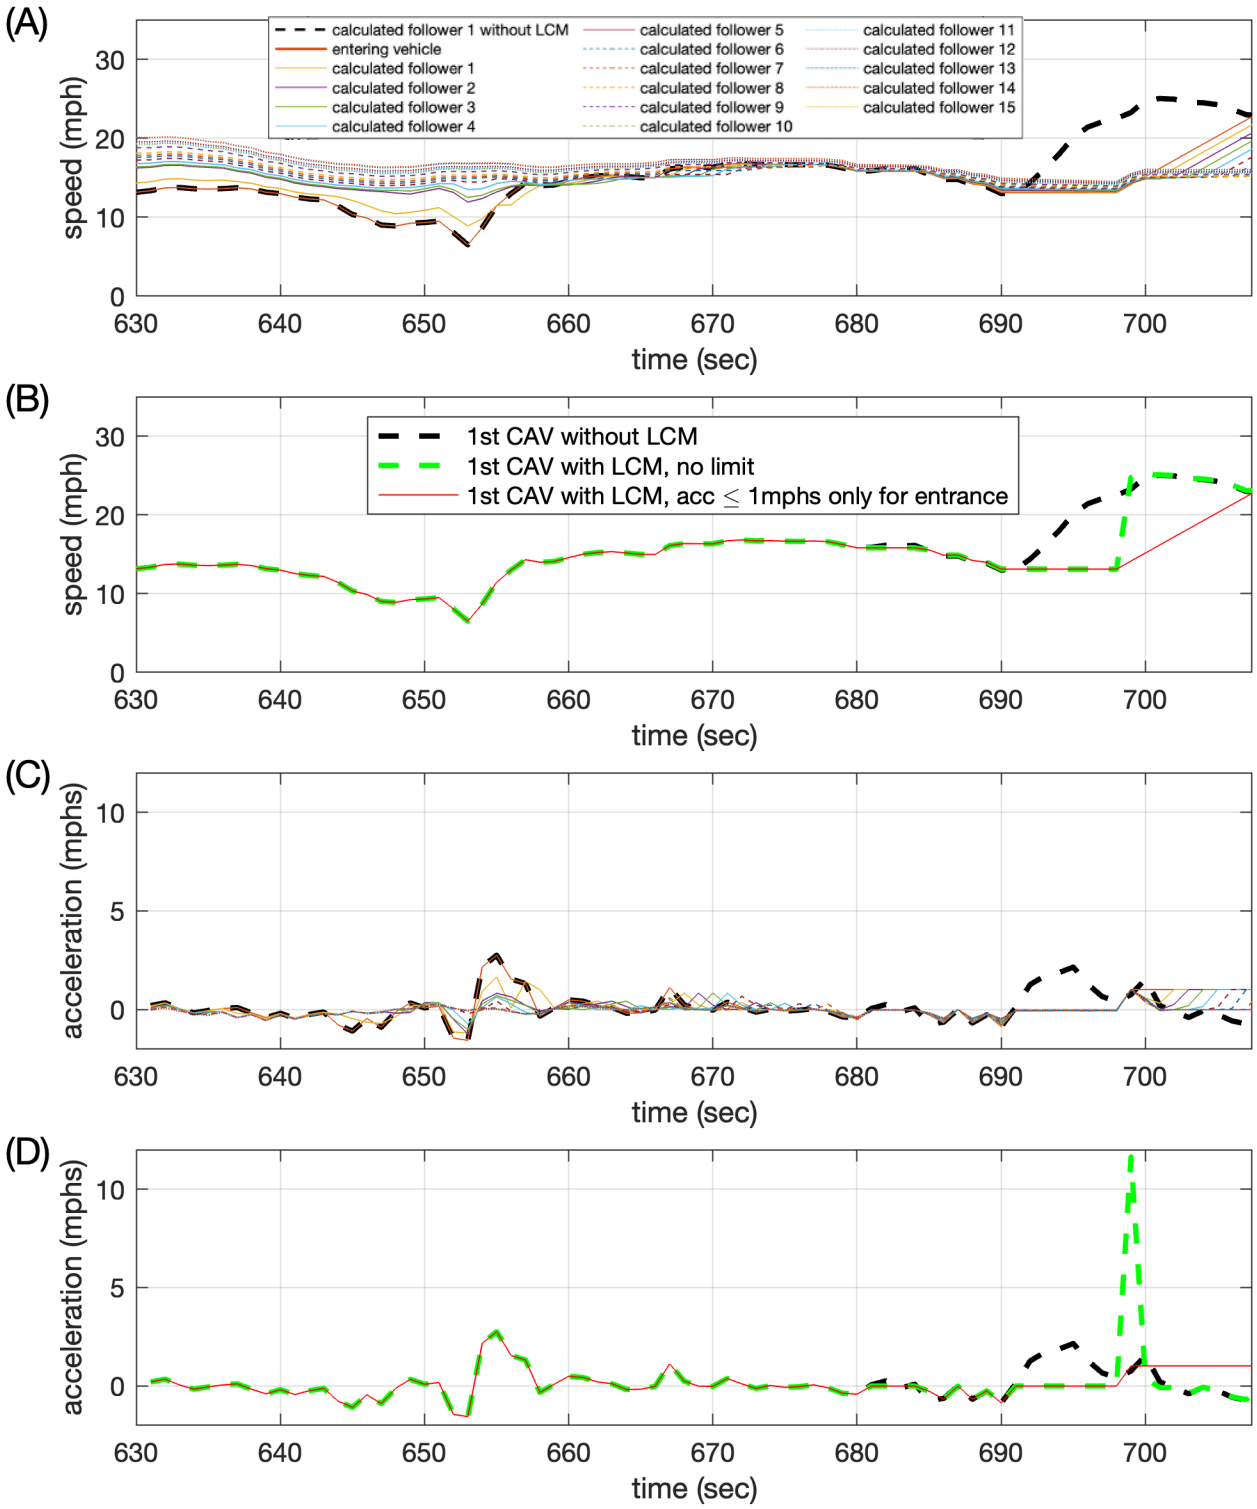

Figure 27, (A) Time series speed for Case 1 with the fixed 1000 ft window and LCM request at 680 sec for the first CAV follower with $a_{CAV} \leq 1\ mphps$ only while accommodating the entrance in Fig. 26. (B) Repeating A only now just showing the first CAV. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first CAV.

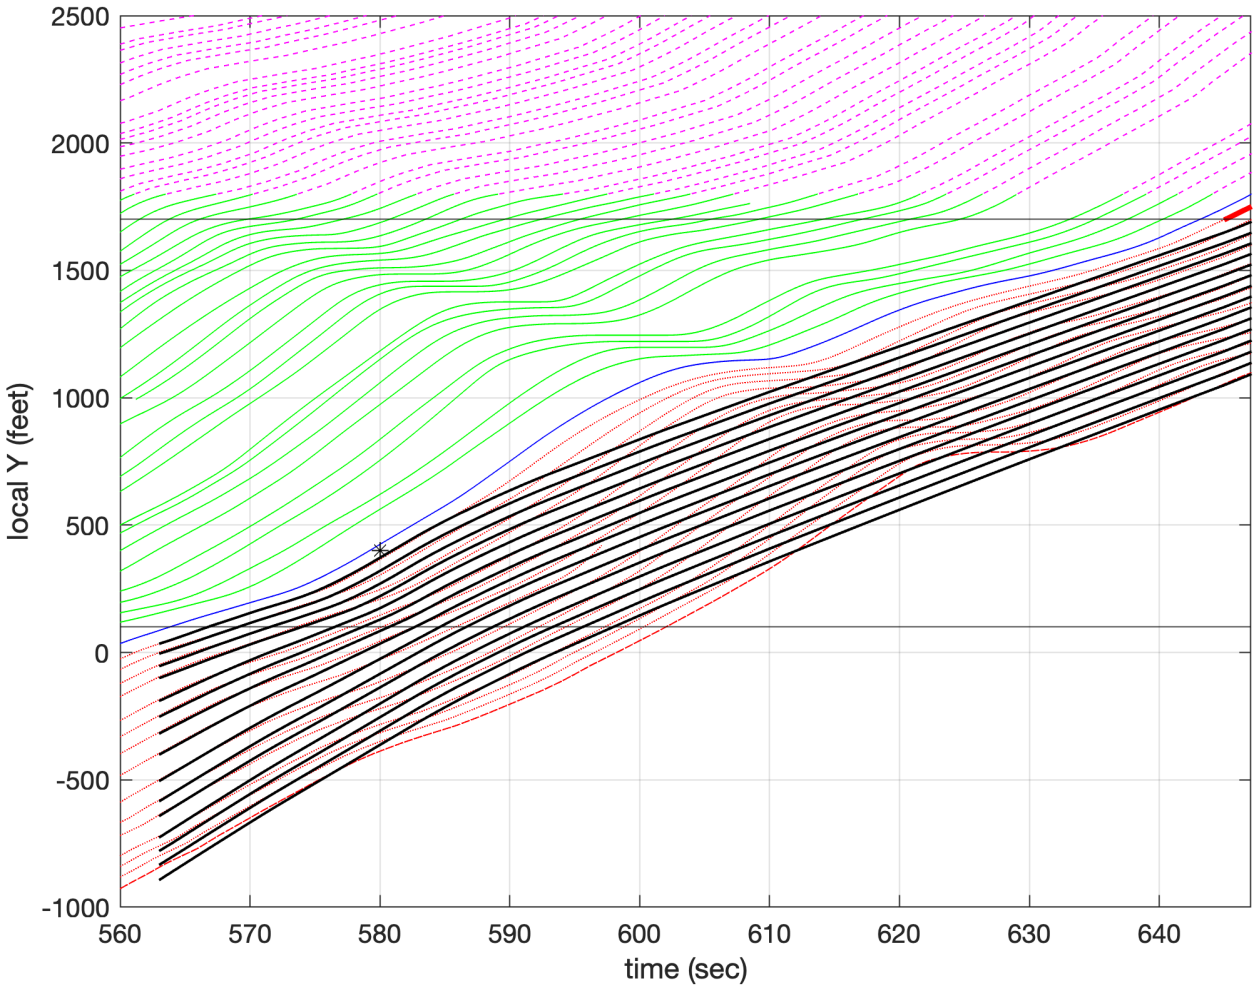

Figure 28, The final CAV trajectory for Case 2 with the fixed 1000 ft window and LCM request at 580 sec with $a_{CAV} \leq 1\ mphps$ only while accommodating the entrance.

Shifting to Case 2 with fixed 1000 ft look ahead window with the 1 mphps constraint on acceleration behind an entrance, Fig. 28 shows the resulting trajectories for the LCM request at 580 sec (compare to Fig. 18). Fig. 29 shows the corresponding time series speed and acceleration. Once more the contrast between "only for entrance" and "no limit" curves is pronounced at the time of the entrance, most notably in Fig. 29D shows that the large acceleration spike has been flattened into a prolonged period of 1 mphps acceleration that extends beyond the right side of the plot. Provided the leader does not continuously accelerate, the follower will eventually catch up and close the gap. Fig. 30 and Fig. 31 repeat the analysis for the LCM request at 600 sec (compare to Fig. 20). The results are very similar to the previous example because in both cases the LCM is delayed until after the same slow wave passes.

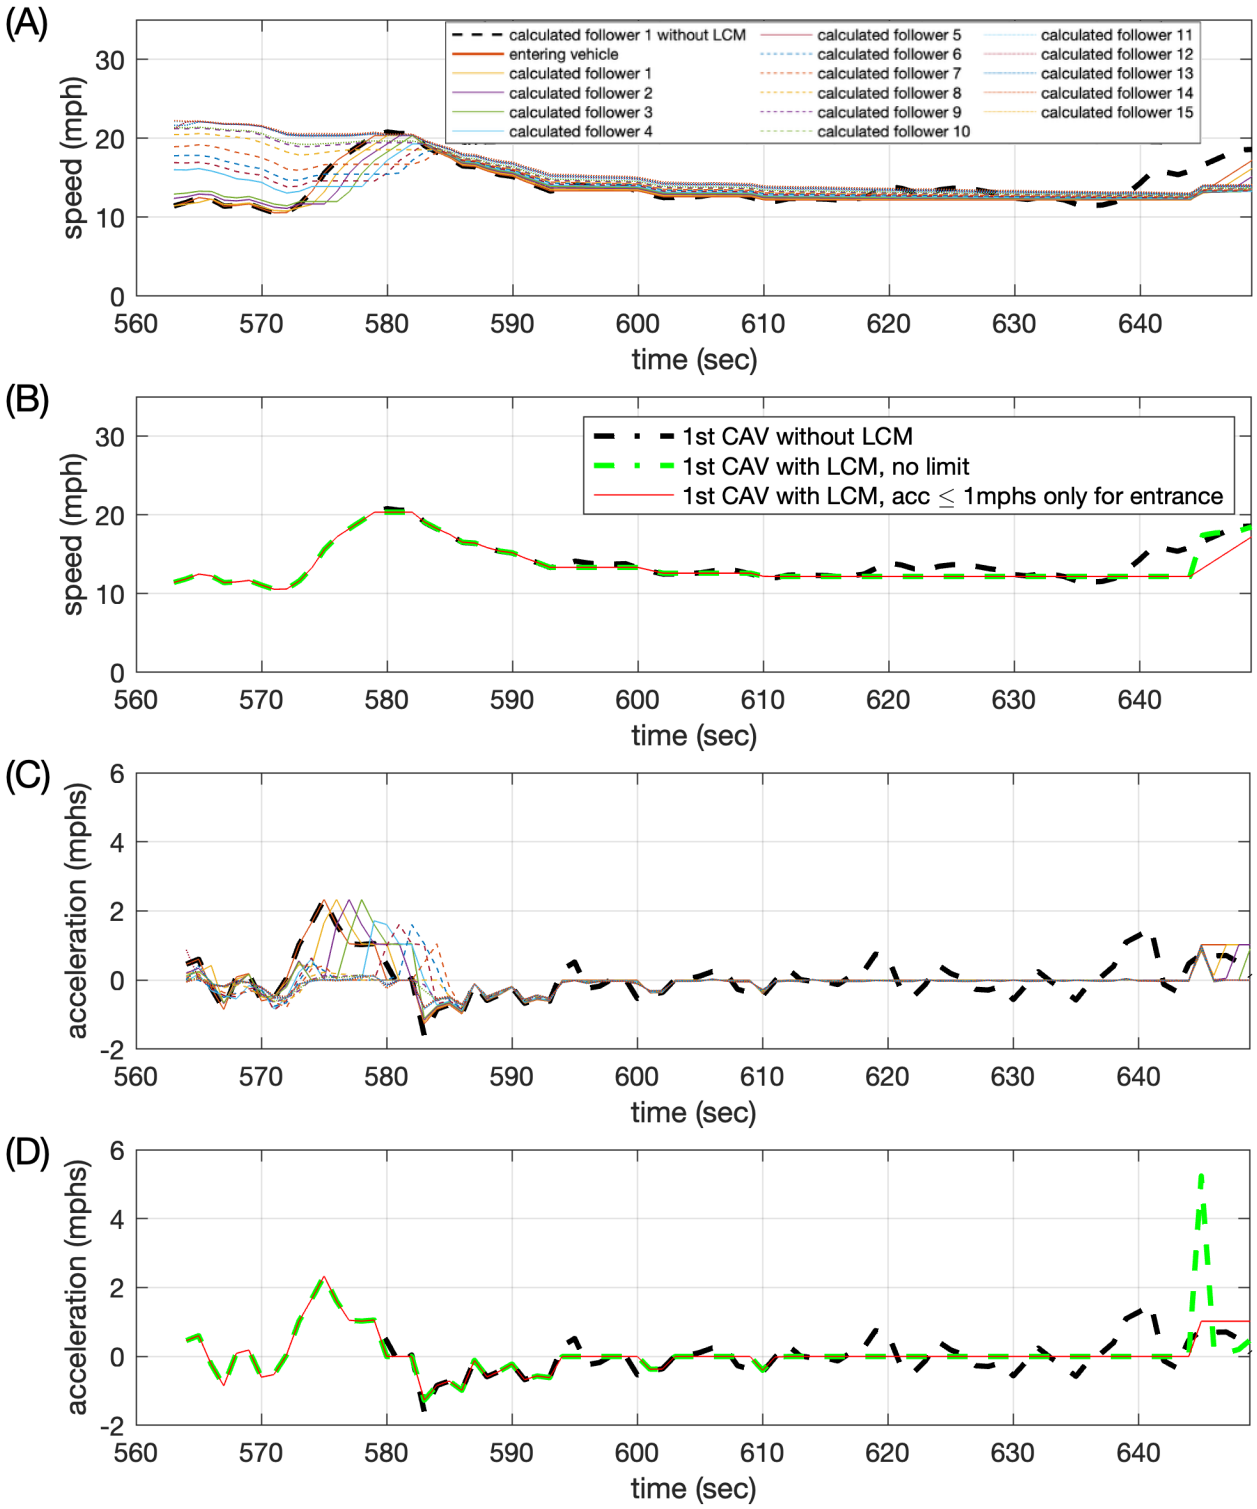

Figure 29, (A) Time series speed for Case 2 with the fixed 1000 ft window and LCM request at 580 sec for the first CAV follower with $a_{CAV} \leq 1\ mphps$ only while accommodating the entrance in Fig. 28. (B) Repeating A only now just showing the first CAV. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first CAV.

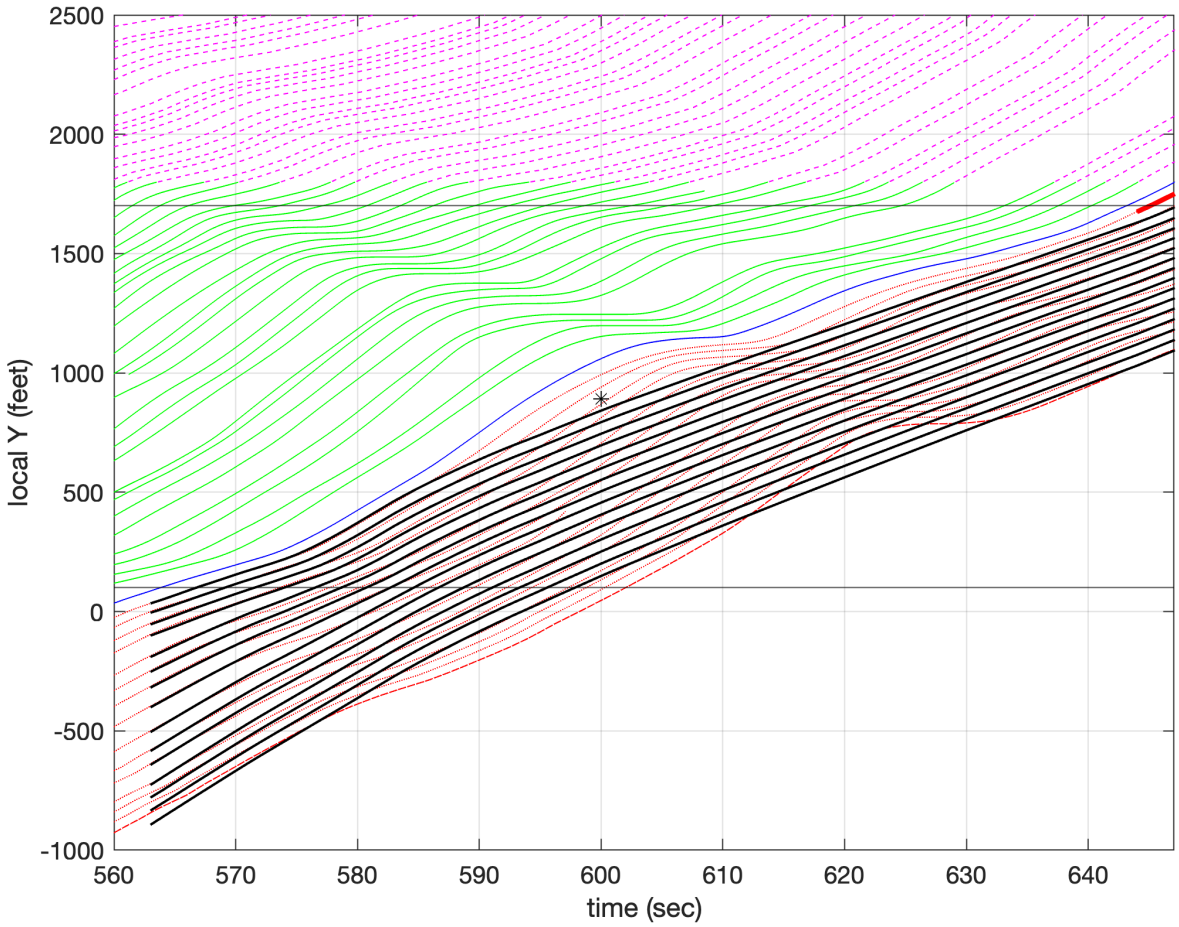

Figure 30, The final CAV trajectory for Case 2 with the fixed 1000 ft window and LCM request at 600 sec with $a_{CAV} \leq 1$ $mphps$ only while accommodating the entrance.

2.3.4. Limiting acceleration always

Section 2.3.2 considered the scenario where the CAV behind an entrance immediately adjusts their speed to follow the new leader, but that led to a large transient acceleration. Section 2.3.3 sought to suppress the transient by limiting acceleration strictly for the immediate follower and strictly during the transition. This subsection now considers the scenario where the acceleration is simply limited at all times for all CAVs. Essentially, this subsection combines the LCM of Section 2.3.1 with the acceleration limitation from Phase I of this study (Coifman and Liu, 2024). This subsection sets the maximum acceleration at 1 mphps at all times.

Fig. 32 returns to Case 1 with a lane change request at 640 sec (compare to Fig. 12 for the unlimited acceleration version and Fig. 22 where the acceleration is only limited for the first CAV after an entrance). The look ahead window is maintained at 1000 ft. Fig. 33 shows the corresponding time series speed and acceleration plots. The results are almost identical to the case where acceleration is only limited for the first CAV after an entrance (Fig. 22 and Fig. 24), except that a separate period of high acceleration has been flattened just after 690 sec. Fig. 34 and Fig. 35 repeat the analysis for the lane change request at 660 sec while Fig. 36 and Fig. 37 repeat the analysis for the lane change request at 680 sec, both with similar results to the preceding example with the request at 640 sec.

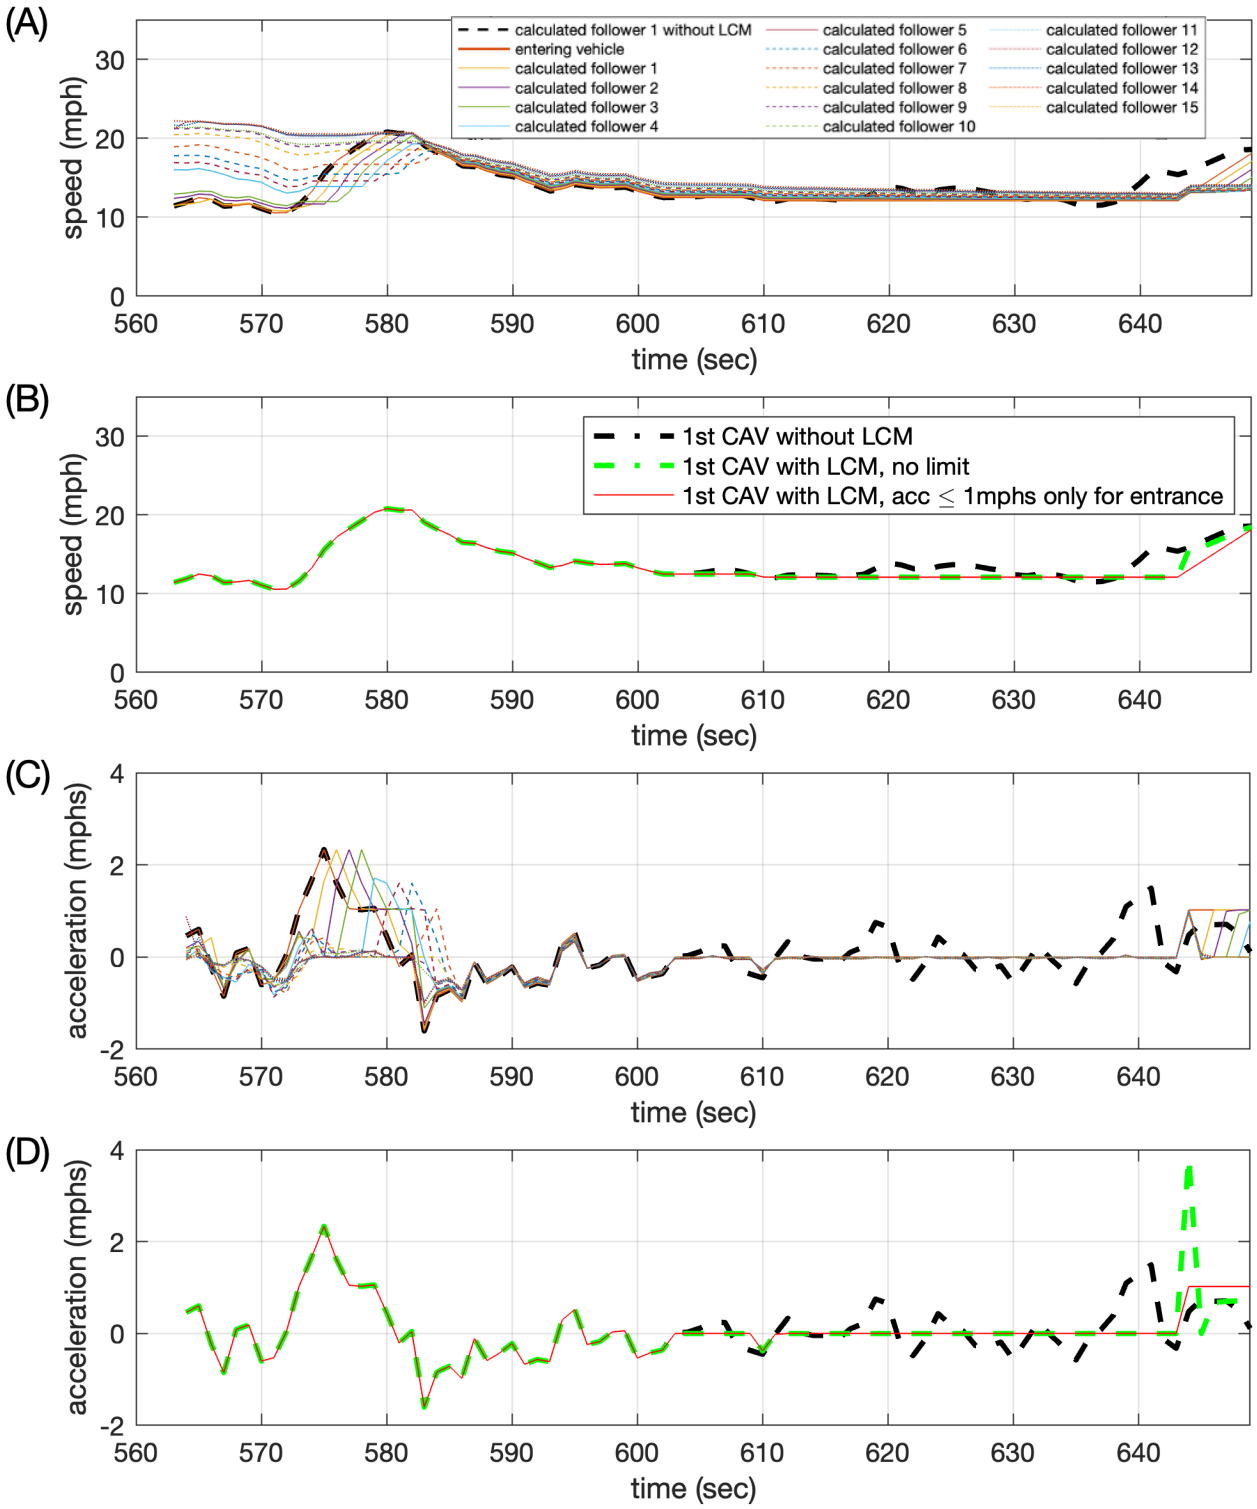

Figure 31, (A) Time series speed for Case 2 with the fixed 1000 ft window and LCM request at 600 sec for the first CAV follower with $a_{CAV} \leq 1\ mphps$ only while accommodating the entrance in Fig. 30. (B) Repeating A only now just showing the first CAV. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first CAV.

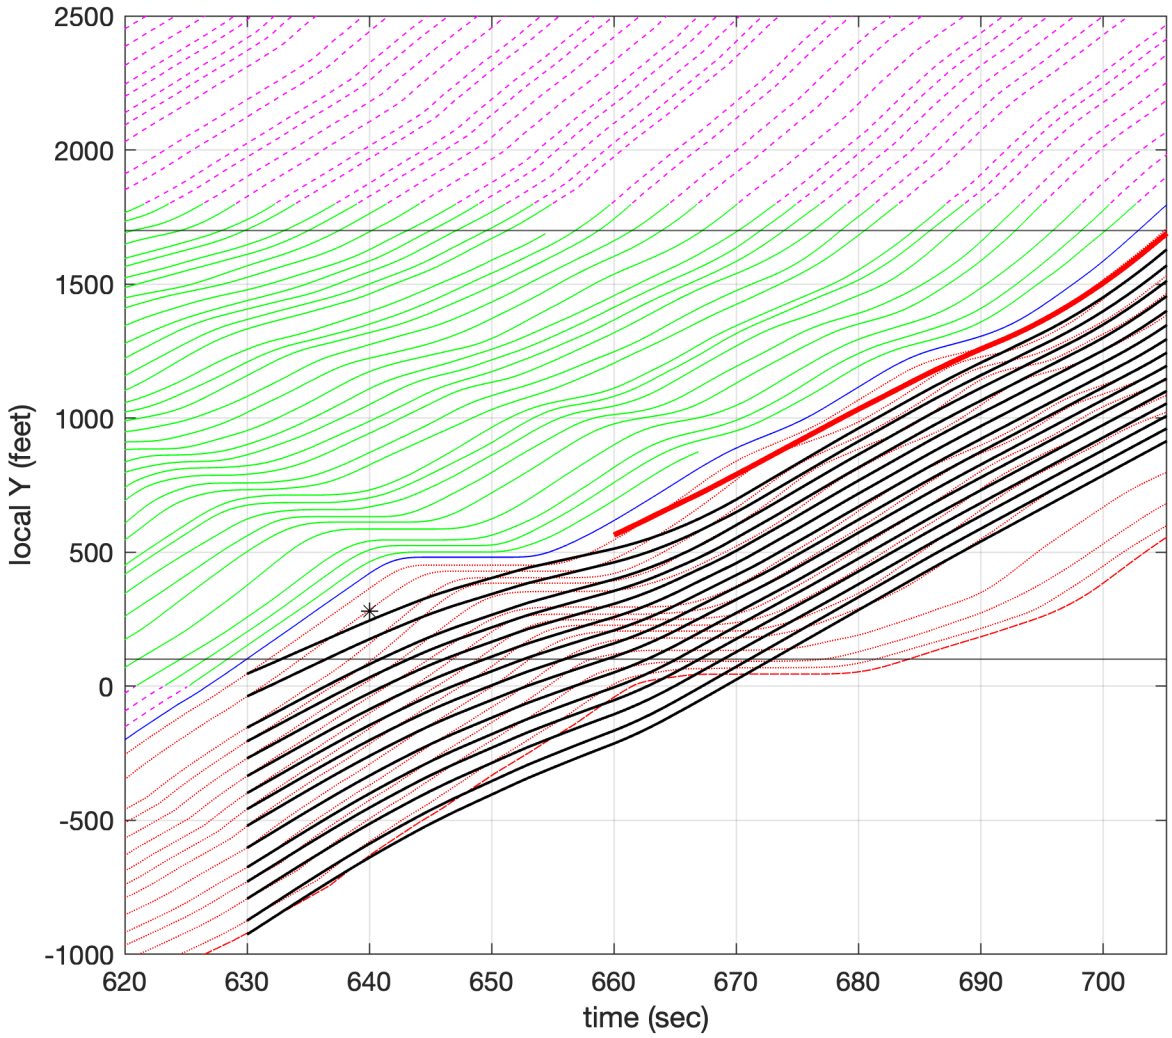

Figure 32, The final CAV trajectory for Case 1 with the fixed 1000 ft window and LCM request at 640 sec with $a_{CAV} \leq 1\ mphps$ at all times.

Fig. 38-41 repeat the analysis for the two lane-change requests in Case 2. In this case the global acceleration limit flattens an acceleration at 573 sec. The LCM policy successfully facilitates the lane change maneuver, and the imposed acceleration constraint further mitigates fluctuations in speed and acceleration.

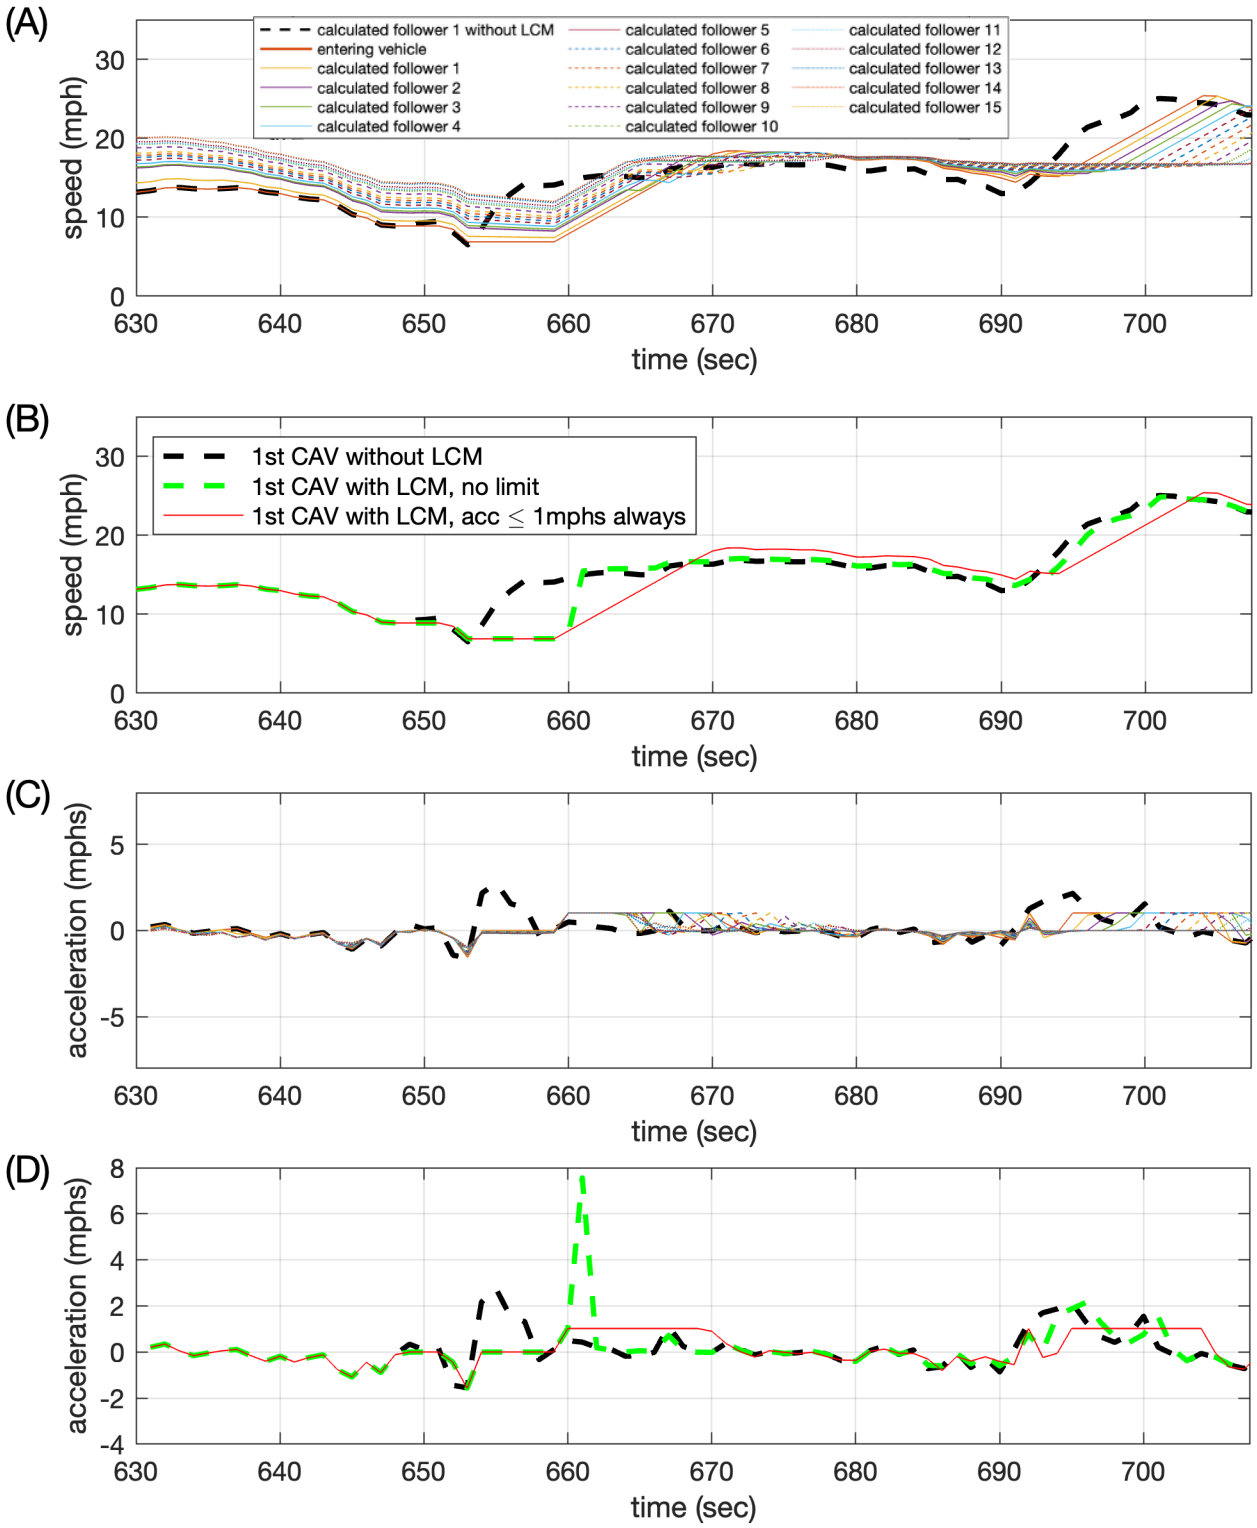

Figure 33, (A) Time series speed for Case 1 with the fixed 1000 ft window and LCM request at 640 sec for the first CAV follower with $a_{CAV} \leq 1\ mphps$ at all times in Fig. 32. (B) Repeating A only now just showing the first CAV. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first CAV.

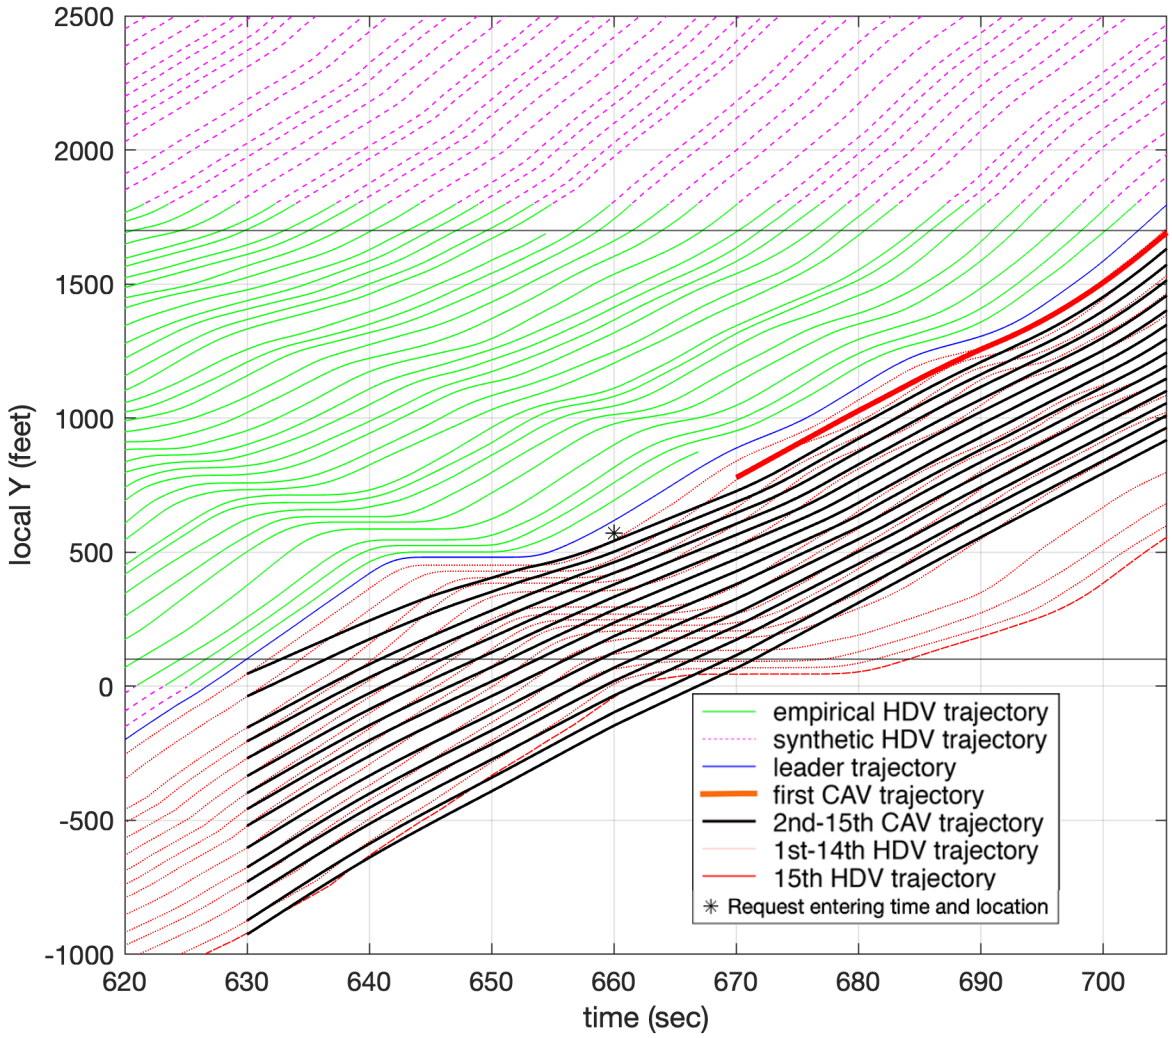

Figure 34, The final CAV trajectory for Case 1 with the fixed 1000 ft window and LCM request at 660 sec with $a_{CAV} \leq 1\ mphps$ at all times.

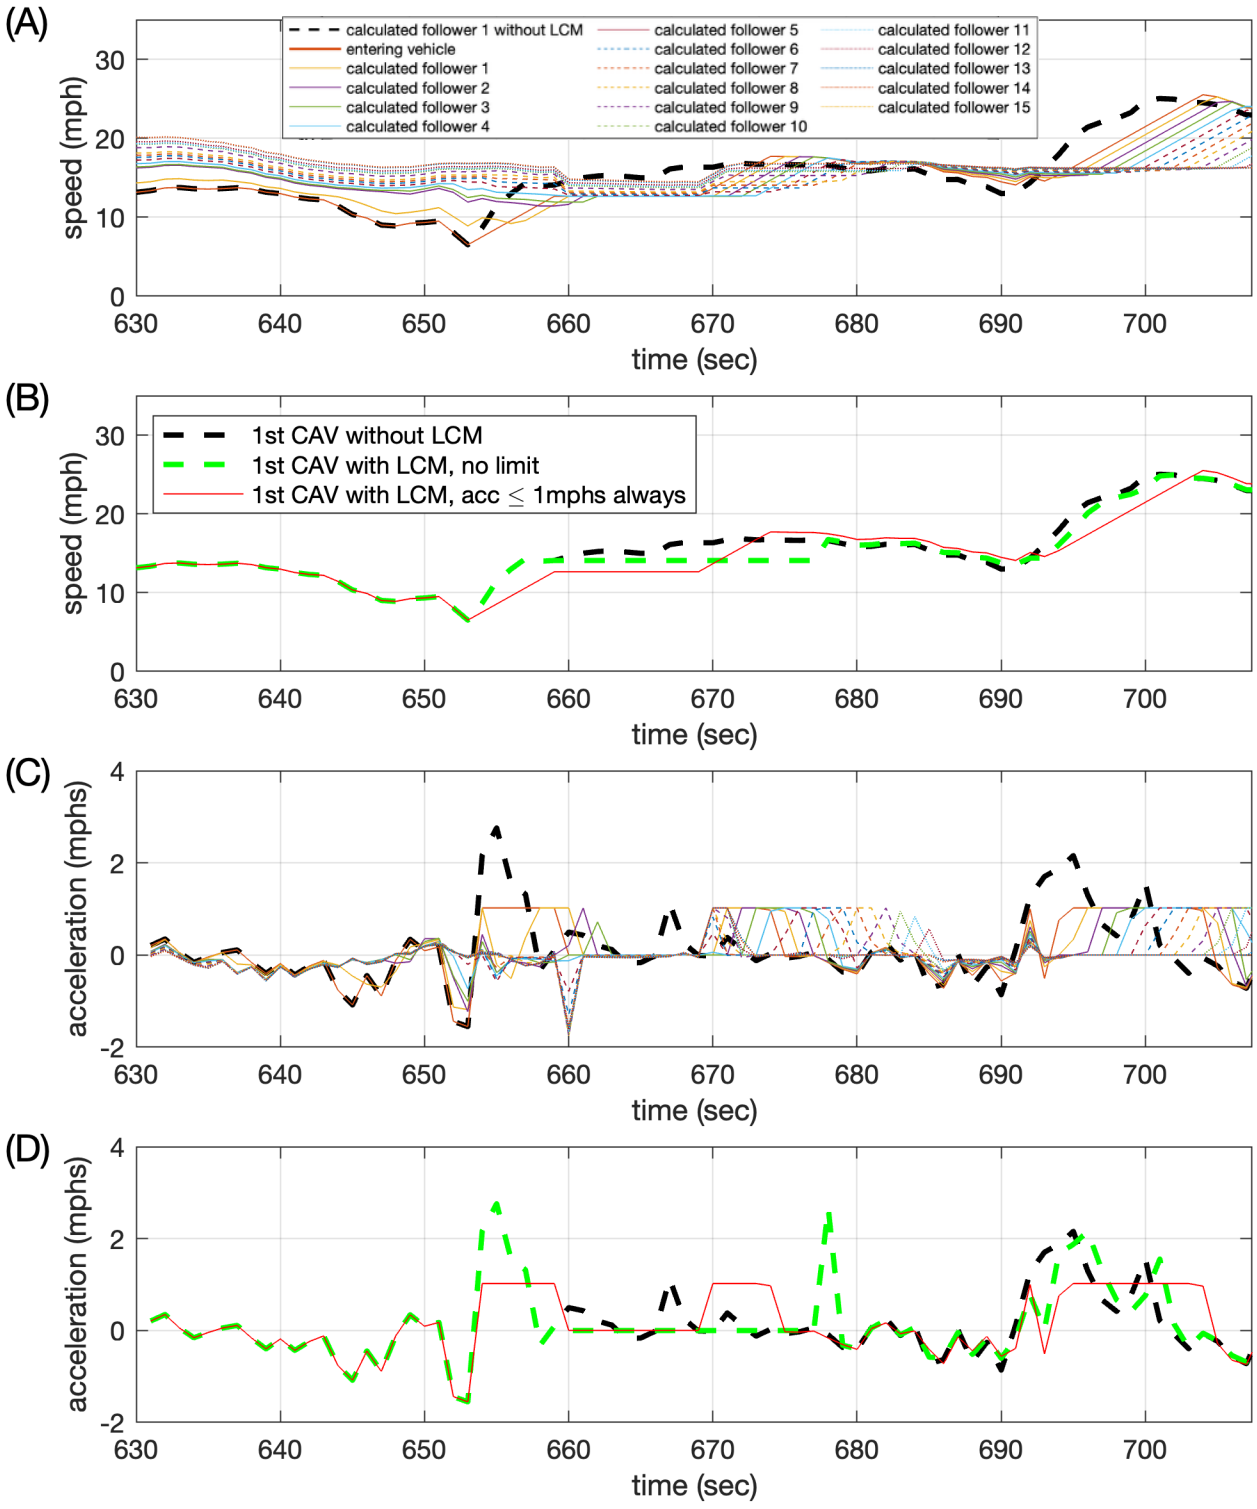

Figure 35, (A) Time series speed for Case 1 with the fixed 1000 ft window and LCM request at 660 sec for the first CAV follower with $a_{CAV} \leq 1\ mphps$ at all times in Fig. 34. (B) Repeating A only now just showing the first CAV. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first CAV.

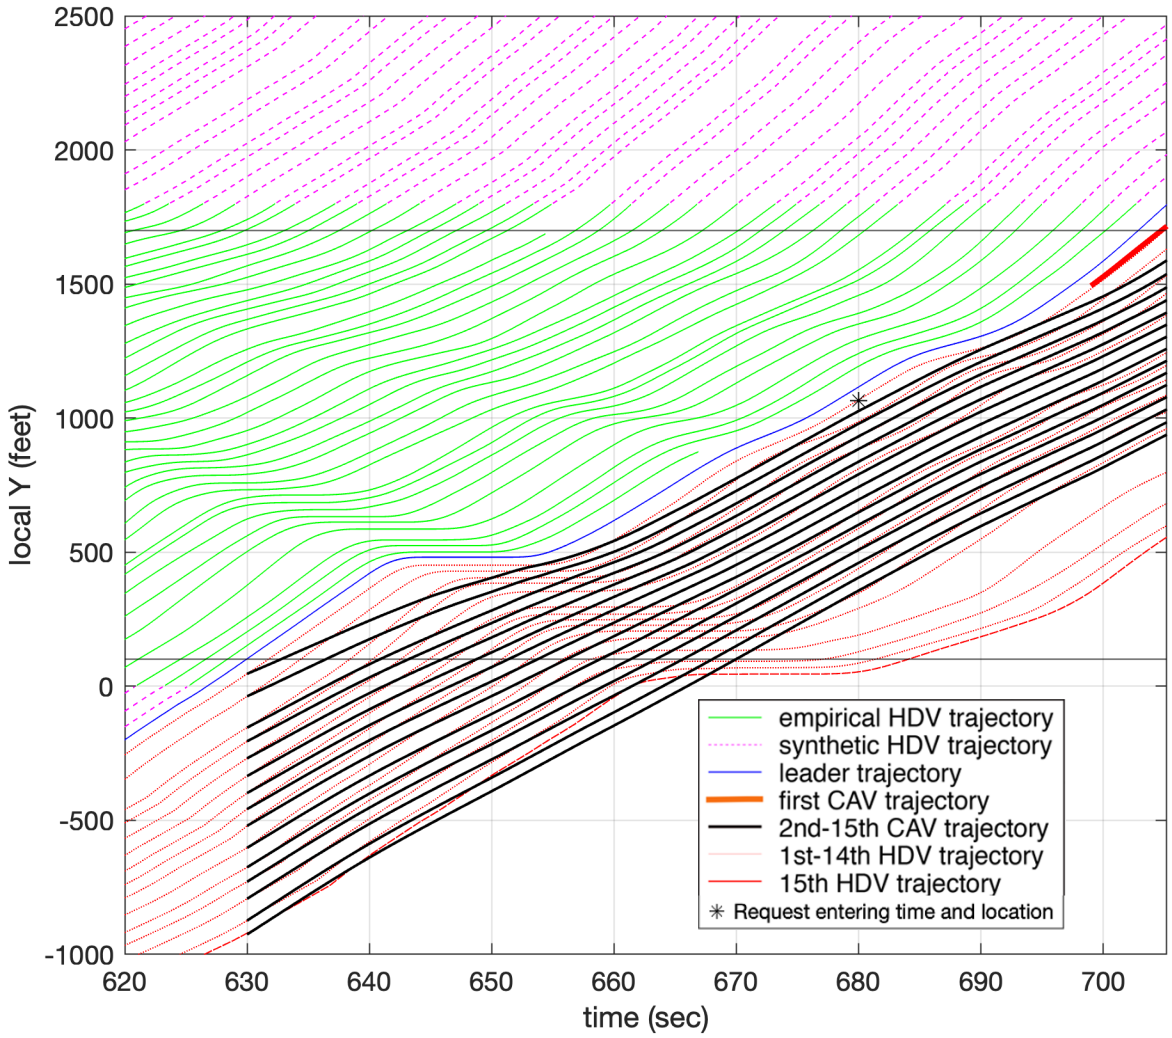

Figure 36, The final CAV trajectory for Case 1 with the fixed 1000 ft window and LCM request at 680 sec with $a_{CAV} \leq 1\ mphps$ at all times.

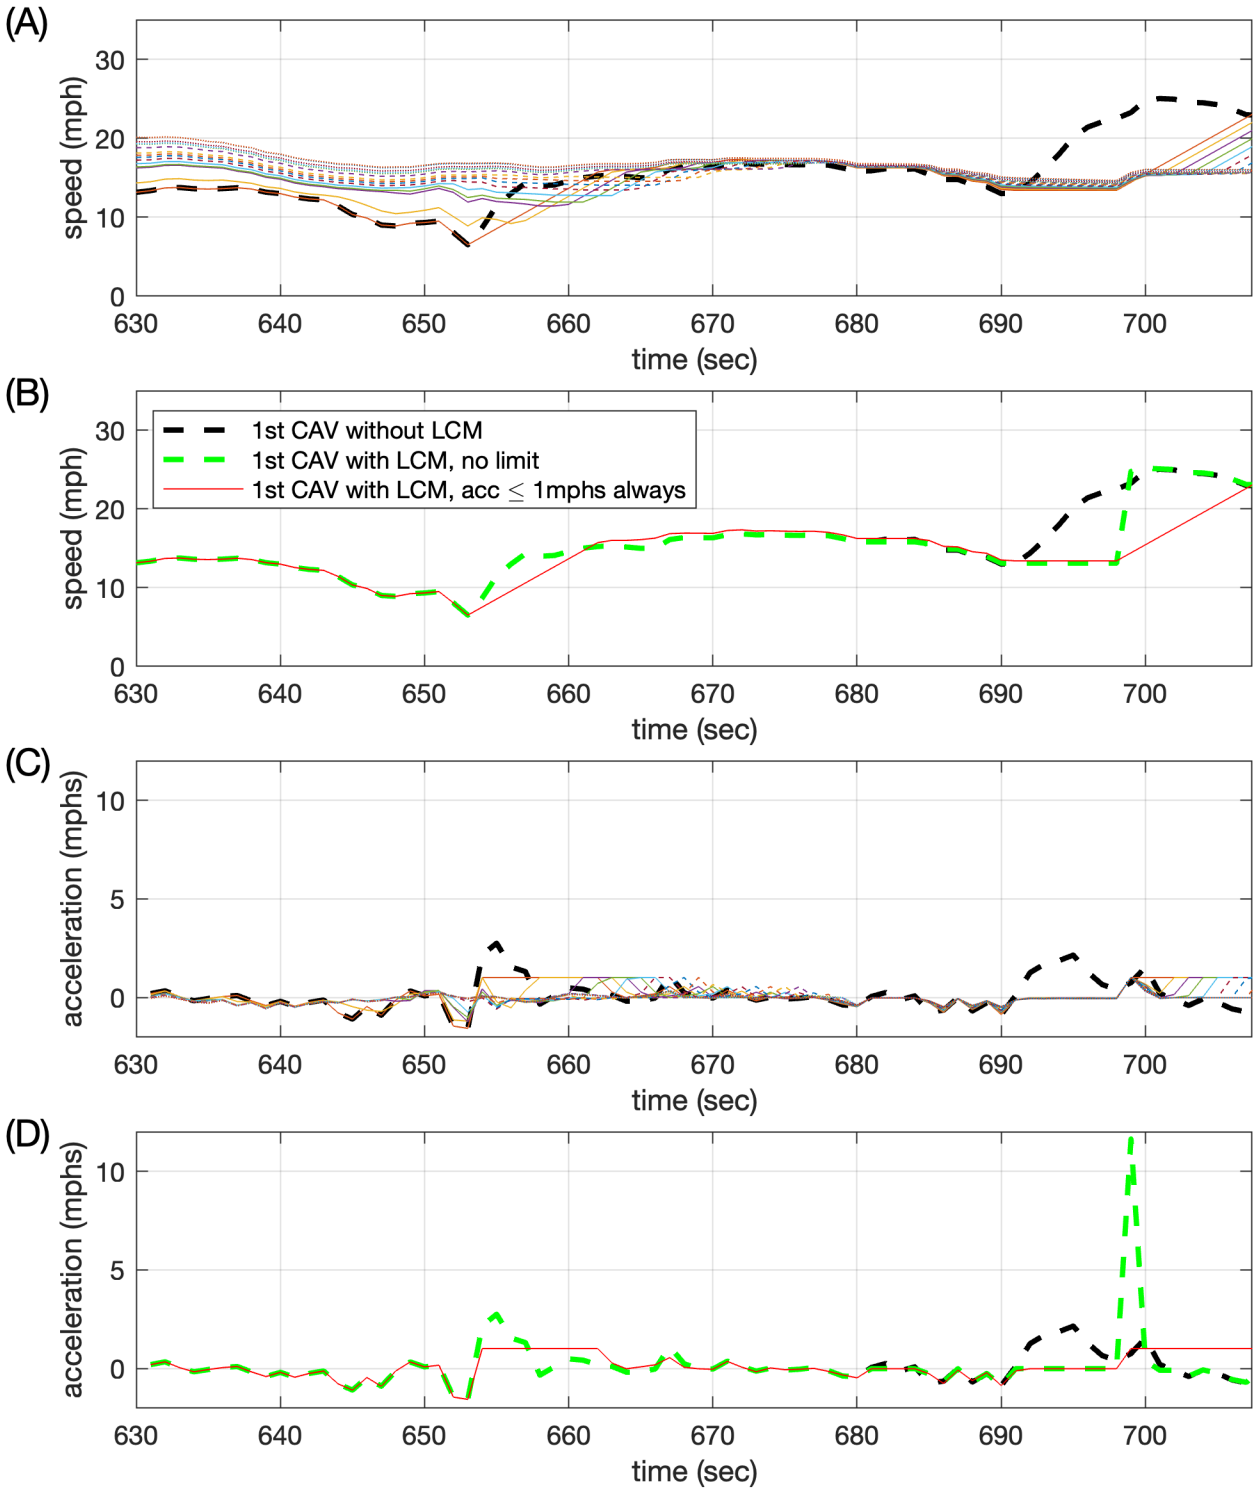

Figure 37, (A) Time series speed for Case 1 with the fixed 1000 ft window and LCM request at 680 sec for the first CAV follower with $a_{CAV} \leq 1\ mphps$ at all times in Fig. 36. (B) Repeating A only now just showing the first CAV. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first CAV.

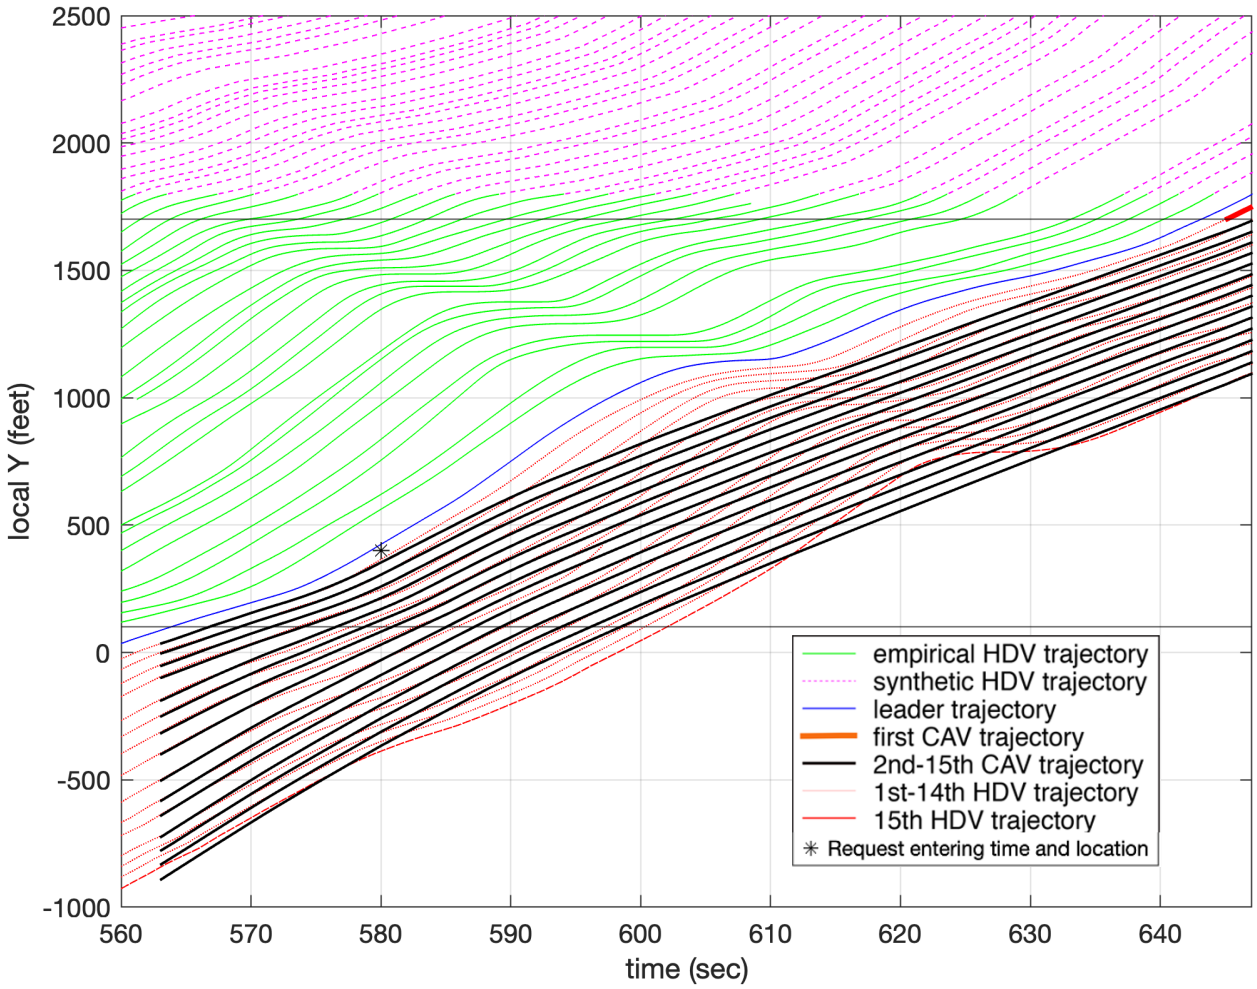

Figure 38, The final CAV trajectory for Case 2 with the fixed 1000 ft window and LCM request at 580 sec with $a_{CAV} \leq 1\ mphps$ at all times.

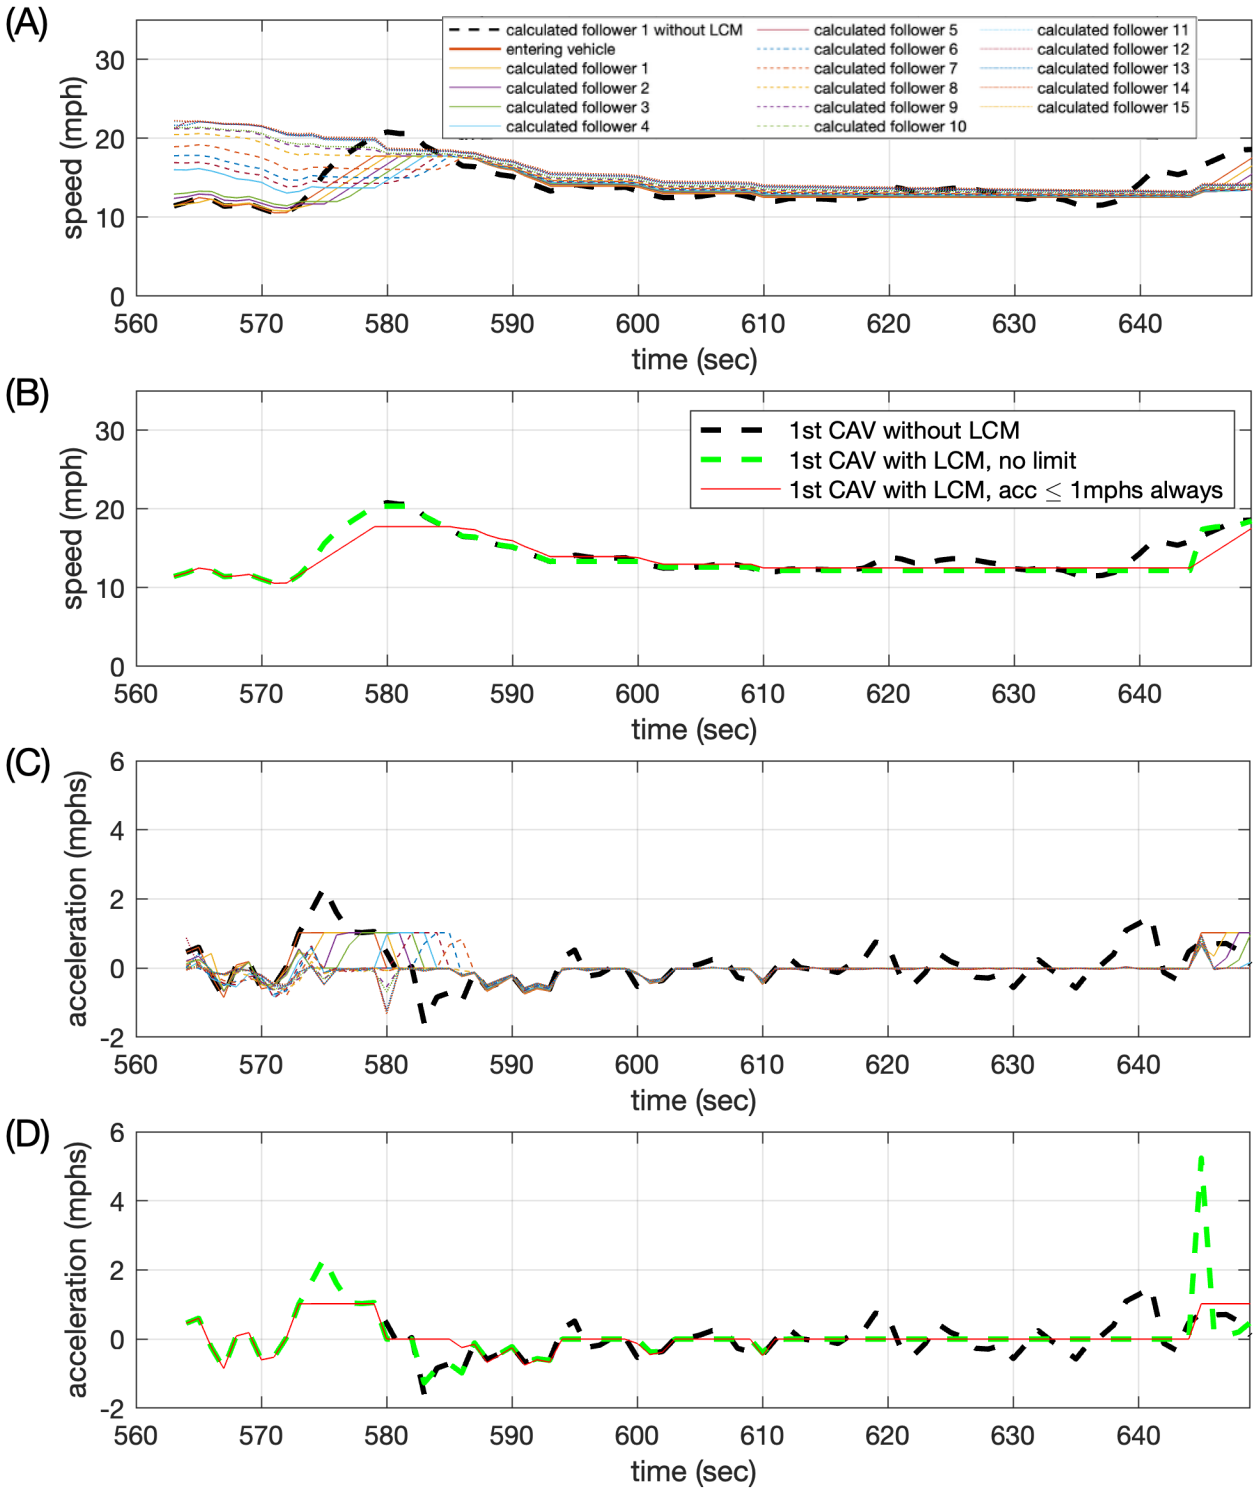

Figure 39, (A) Time series speed for Case 2 with the fixed 1000 ft window and LCM request at 580 sec for the first CAV follower with $a_{CAV} \leq 1\ mphps$ at all times in Fig. 38. (B) Repeating A only now just showing the first CAV. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first CAV.

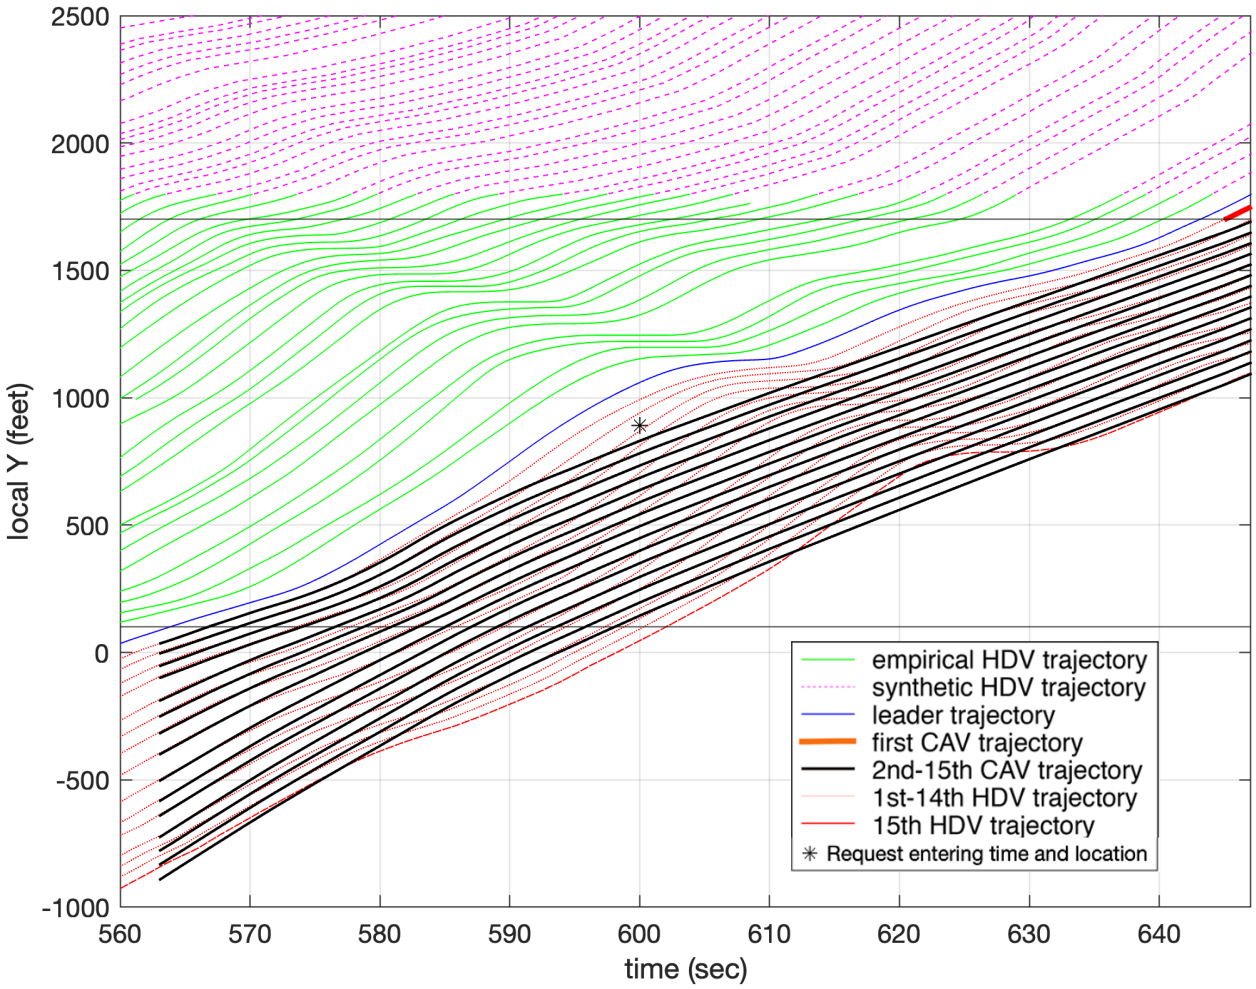

Figure 40, The final CAV trajectory for Case 2 with the fixed 1000 ft window and LCM request at 600 sec with $a_{CAV} \leq 1\ mphps$ at all times.

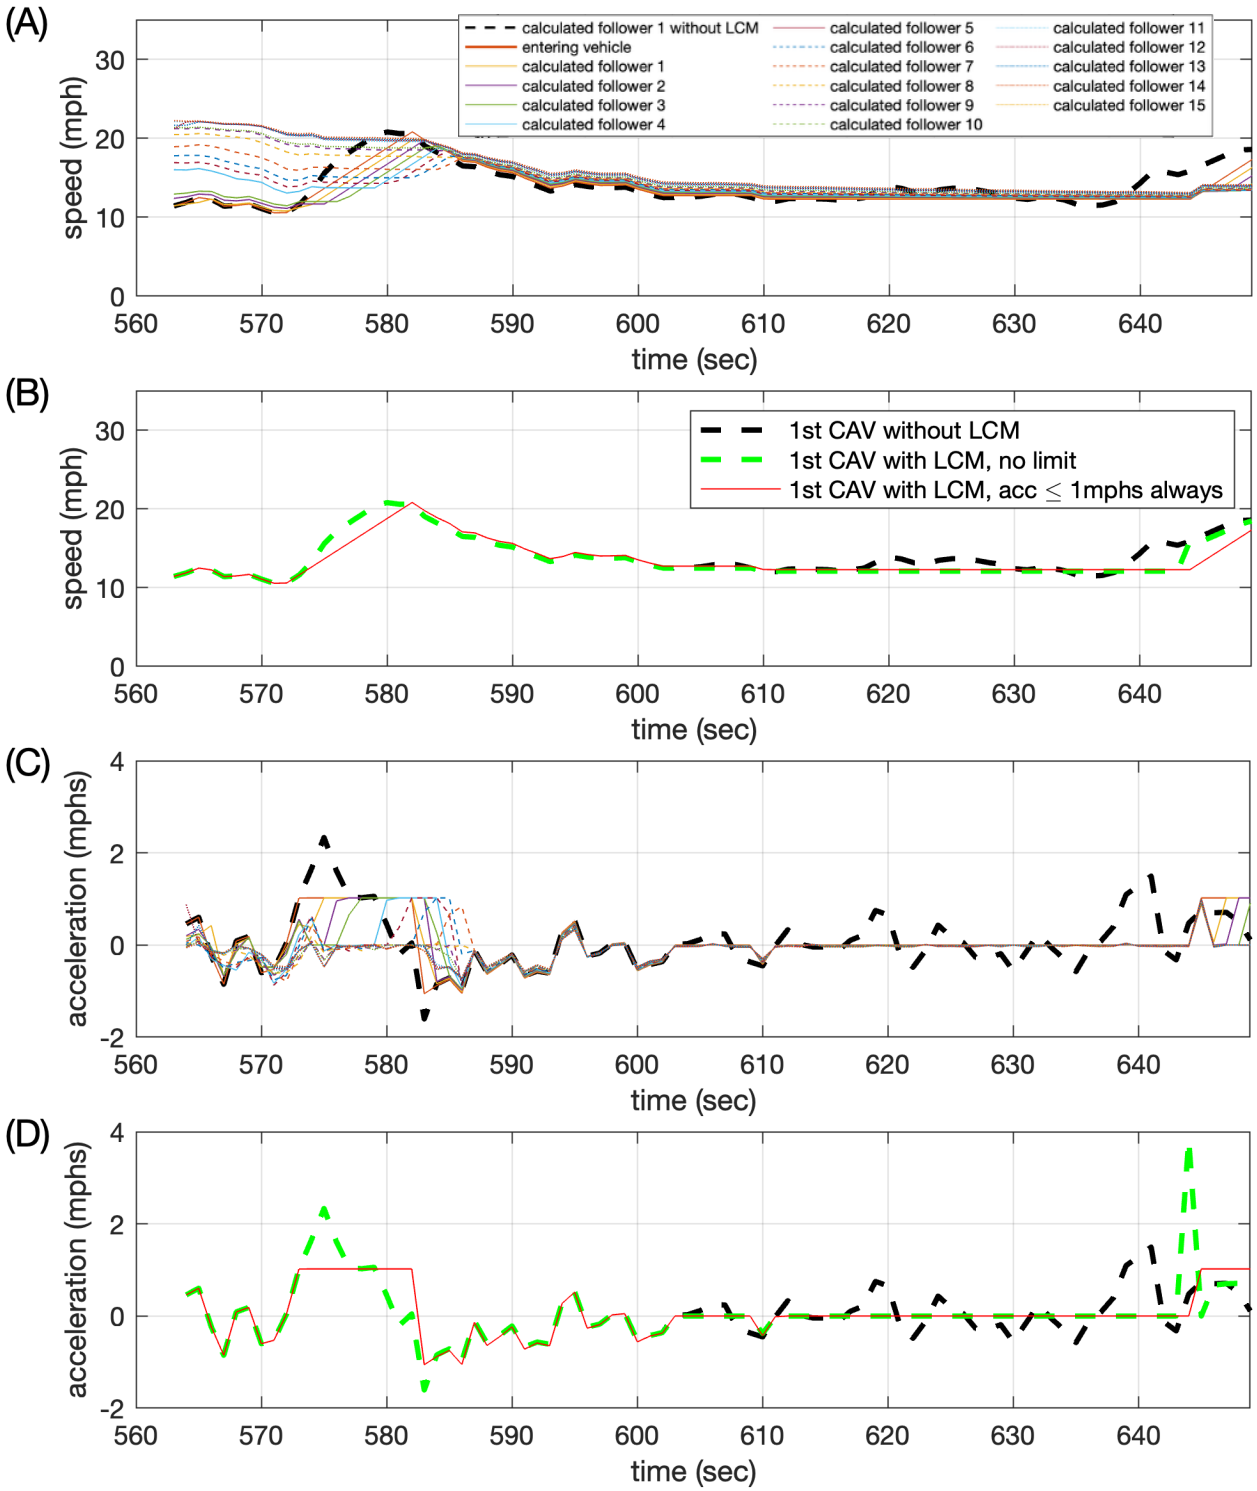

Figure 41, (A) Time series speed for Case 2 with the fixed 1000 ft window and LCM request at 600 sec for the first CAV follower with $a_{CAV} \leq 1\ mphps$ at all times in Fig. 40. (B) Repeating A only now just showing the first CAV. (C) The corresponding acceleration from A, and (D) repeating C only now just showing the first CAV.

3. Conclusions

This study presents a novel anticipatory control method for connected and automated vehicles (CAVs), designed to mitigate and smooth disturbances in congested freeway traffic. By enabling CAVs to anticipate stop waves, they can proactively generate smooth trajectories. The method reduces unnecessary accelerations and decelerations, thereby improving safety, driving comfort, fuel efficiency, and emissions. A key feature of the proposed method is its reliance on vehicle-to-vehicle communication. The first CAV continuously receives instantaneous speed and position data from downstream connected human-driven vehicles (cHDVs) within a specified look ahead range. This information is spatially integrated to forecast the trajectory of the CAV's immediate leader, allowing the CAV to identify and adopt a stable speed trajectory aimed at dissipating disturbances propagating upstream through the traffic. Note that the new trajectory is intended to maintain the same average speed the vehicle would experience but greatly reduce speed variance.

Since the CAV's trajectory is designed to smooth the state from that of its leader, it effectively creates a new traffic state progression. This new progression should benefit any vehicle that follows- HDV or CAV. As a result, even low levels of CAV penetration could yield significant benefits. For example, if there was only a single CAV in the system, its smooth trajectory could guide upstream HDVs toward more stable behavior, effectively attenuating traffic disturbances.

In the period of this report, the methodology first developed in Phase 1 (Coifman and Liu, 2024) was extended to analyze platoons of successive CAVs. Because each CAV effectively overrides the evolving state, it no longer makes sense to assess the state from vehicles further downstream. So instead, each successive CAV follower uses the planned trajectory of their respective leader as the estimated trajectory of their leader. It was shown that combined response of multiple successive CAVs to disturbances is more effectively than a single CAV, which enhances the overall stability of the platoon.

Lane Change Maneuvers (LCMs) are an unavoidable aspect of traffic flow on multi-lane freeways. So, this work developed a LCM policy for the method to allow LCM without disrupting the benefits of the flow smoothing. During congested conditions, each entering vehicle will create delay for all following vehicles in the given lane. This work seeks to do so in a way that avoids deceleration. So, when an entering vehicle notifies the platoon of its intent to change lanes, the CAV in the target lane will wait for an acceleration wave. As that acceleration wave propagates upstream through a platoon, the new target follower will hold off accelerating and instead maintain their previous speed to create a gap for the entering vehicle to fill. This policy allows vehicles from adjacent lanes to merge into the target lane while preserving the benefits of stop wave smoothing enabled by the CAV control strategy. Simulation results show that this policy allows for successful LCMs without sacrificing the method's benefits.

This proof-of-concept study adopts several simplifying assumptions. Notably, it considers all vehicles as either cHDVs or fully automated CAVs and assumes uniform vehicle types (i.e., passenger cars). Future research should address heterogeneous traffic composition, including varying CAV market penetration levels and the presence of larger vehicles such as trucks, which require greater headways and respond more slowly.

Future work will consider how this approach could be used to improve conditions when there are few cHDV and few CAVs. If there is only a small percentage of connected vehicles reporting their state, a CAV could still talk to a leader several hundred ft ahead to provide a similar prediction window of the future state. For example, the actual trajectory in Fig. 1E could be used as an estimate of the state for a CAV that enters the segment afterwards. If that following CAV is

right behind, then the smoothing will be small but still beneficial. Whereas if that following CAV entered at 700 seconds the look ahead might be comparable to a moving 1600 ft window in this work. Of course, challenges arise in terms of addressing any dynamics that arise after the first CAV's passing and preventing HDV from changing lanes into a large void ahead of the CAV that is designed to absorb stop waves.

This study also demonstrates that an unstable congested freeway of HDVs can quickly be stabilized with CAV. From which, whole new control strategies could be implemented. It is anticipated that this study will serve as a jumping-off point for exploring how to control CAV flows in ways that would be impossible for HDV.

From a broader perspective, this 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 environmental 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 efficiency, safety, and sustainability of congested traffic networks.

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., 2002. Estimating Travel Times and Vehicle Trajectories on Freeways Using Dual Loop Detectors, *Transportation Research: Part A*, vol 36, no 4, pp. 351-364.

Coifman, B., Wang, Y., 2005. Average Velocity of Waves Propagating Through Congested Freeway Traffic, *Proc. of The 16th International Symposium on Transportation and Traffic Theory,* July 19-21, 2005, College Park, MD. 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, *Transportation Research Record*, 14p.

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

Newell, G., 1962. Theories of Instability in Dense Highway Traffic. *Journal of Operation Research Society of Japan*, Vol. 5, No. 1, pp. 9–54.

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