A Distributed Framework for Coordinated Heavy-Duty Vehicle Platooning

Jeffrey Larson, Kuo-Yun Liang, and Karl H. Johansson, *Fellow, IEEE*

***Abstract*—Heavy-duty vehicles (HDVs) traveling in single file with small intervehicle distances experience reduced aerodynamic drag and, therefore, have improved fuel economy. In this paper, we attempt to maximize the amount of fuel saved by coordinating platoon formation using a distributed network of controllers. These virtual controllers, placed at major intersections in a road network, help coordinate the velocity of approaching vehicles so they arrive at the junction simultaneously and can therefore platoon. This control is initiated only if the cost of forming the platoon is smaller than the savings incurred from platooning. In a large-scale simulation of the German Autobahn network, we observe that savings surpassing 5% when only a few thousand vehicles participate in the system. These results are corroborated by an analysis of real-world HDV data that show significant platooning opportunities currently exist, suggesting that a slightly invasive network of distributed controllers, such as the one proposed in this paper, can yield considerable savings.**

***Index Terms*—Automated highways, intelligent systems, road transportation.**

I. Introduction

HEAVY-DUTY vehicles (HDVs) driving in a platoon use significantly less fuel than when driving separately. By traveling in single file with small intervehicle distances, trailing vehicles experience reduced aerodynamic drag. This reduces the total fuel consumed and, therefore, significantly reduces the cost to vehicle owners. For example, transport in Europe uses approximately 180 million tons of diesel fuel [1] every year; therefore, even modest decreases in fuel use can yield dramatic savings. In addition, since HDVs are a significant contributor to greenhouse gas emissions (road transport generates 16% of the CO_2 pollution in Europe [2]), platooning vehicles can also have a notable environmental impact. A photograph of platooning vehicles can be seen in Fig. 1(a).

Manuscript received December 19, 2013; revised February 25, 2014; accepted April 10, 2014. Date of publication July 1, 2014; date of current version January 30, 2015. The work of J. Larson was performed while he was with KTH. This work was supported in part by VINNOVA within the FFI program under Contract 2011-01144, by the EU FP7 Project COMPANION, by the Swedish Research Council, by the Swedish Foundation for Strategic Research, by the Knut and Alice Wallenberg Foundation, and by the U.S. Department of Energy, Office of Science, under Contract DE-AC02-06CH11357. The Associate Editor for this paper was B. De Schutter.

J. Larson is with the Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439 USA (e-mail: jmlarson@anl.gov).

K.-Y. Liang is with the ACCESS Linnaeus Centre and the Department of Automatic Control, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden. He is also with Scania CV AB, 151 87 Södertälje, Sweden (e-mail: kyliang@kth.se; kuo-yun.liang@scania.com).

K. H. Johansson is with the ACCESS Linnaeus Centre and the Department of Automatic Control, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden (e-mail: kallej@kth.se).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TITS.2014.2320133

While research into platooning vehicles has been ongoing for the past 50 years (e.g., [3], [4]), most of the work has focused on the control of vehicles already in a platoon. For example, methods for communicating within platooning vehicles [5], paradigms for visually recognizing obstacles for platooning vehicles [6], and procedures for coordinated maneuvering of platooning vehicles [7] have been studied. For some recent implementations and demonstrations of HDV platooning, see the PATH [8], SARTRE [9], Scoop [10], CHAUFFEUR [11], KONVOI [12], or Energy ITS [13] projects.

The majority of research in platooning has addressed single platoons or individual vehicles within a platoon. The coordination and optimization of HDVs and platoons of HDVs throughout a large-scale real-world road network have largely been neglected until now. The architecture for automated highway systems described in [14] captures the necessary information between layers needed in any such complex transport system based on vehicle platooning. There are several reasons why few studies on large-scale optimization of coordinated HDV platooning exist today. First, there is no central location to find every HDV's current location and eventual destination, information that would apparently be necessary to route vehicles in a coordinated manner. Second, there is no global coordinator with the authority to suggest routes that provide platooning opportunities. Third, even if complete knowledge of every HDV was available and a global coordinator could direct vehicles along any route, the general problem of finding fuel-optimal routing using platooning is known to be NP-hard, even for simple road networks [15].

In this paper we develop a system of distributed controllers that coordinate platoon formation by slightly adjusting the speeds of HDVs as they approach an intersection in the road network. By coordinating the arrival of multiple vehicles approaching a junction in the road, our controllers facilitate platoon formation and, therefore, decrease the total fuel consumed. We believe this distributed control framework is much more appropriate for real-world coordination of vehicles than any global approach. Practical road networks see HDVs entering and leaving the network haphazardly; it is more appropriate to coordinate platoon formation in real time as vehicles approach intersections, by slightly adjusting their speed, as opposed to planning routes days beforehand.

Of course, altering the speed of vehicles can be detrimental to their fuel economy. By knowing the destination of the approaching vehicles, our controllers determine the possible fuel savings incurred from platooning (assuming that the vehicles coordinate their arrival at the intersection). The controller advises a velocity adjustment only if the savings from platooning

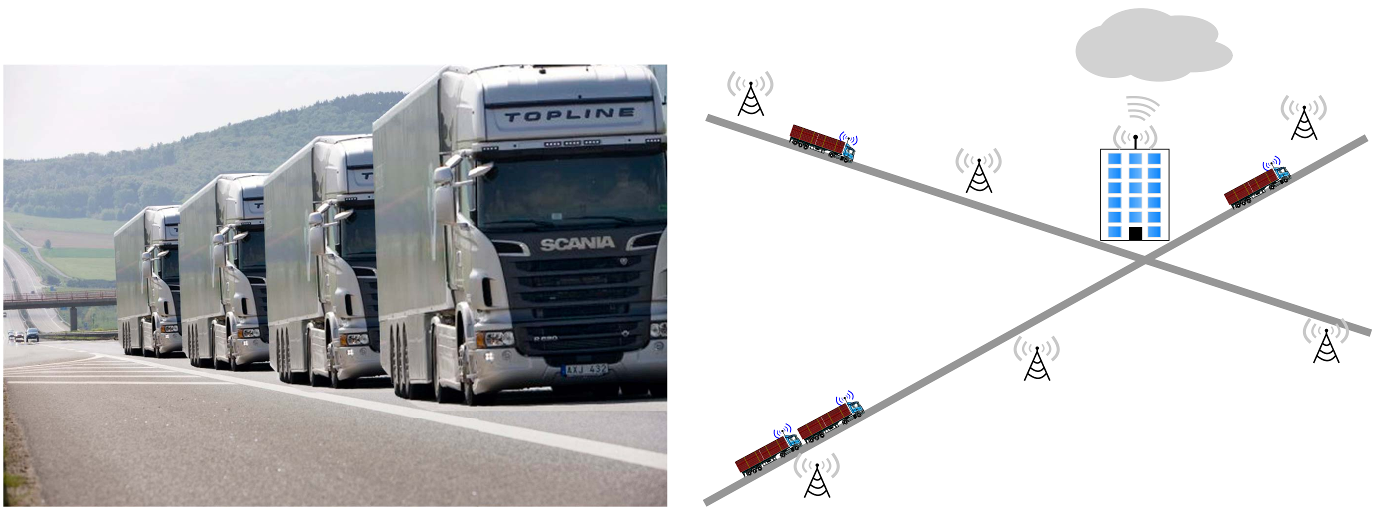

(a) (b)

Fig. 1. (a) Four platooning HDVs. (b) A single HDV and a platoon approaching a local virtual controller at a highway road junction.

are more than the cost of adjusting speeds to meet at the intersection. The controller naturally can be constrained to ensure that advised velocities do not violate speed limits or cause a vehicle to arrive late at its final destination.

We demonstrate the performance of our distributed control architecture in a large-scale simulation of the German Autobahn network. Although the amount of fuel saved clearly depends on the number of vehicles moving through the network (more participating vehicles provide more opportunities to platoon and, therefore, save more fuel), our simulation shows that 2000 participating vehicles can reduce their fuel consumption by over 5% when using our approach. Since Germany has over 400 000 registered HDVs [16], this is a relatively small number of vehicles. We believe that similar savings can be achieved on a road network with a similar geography, for example, the U.S. Interstate Highway System.

Implementations of HDV platooning have been limited to small case studies, in part because of legality concerns surrounding driving HDVs at small intervehicle distances. Although many jurisdictions are actively amending regulations to allow for platooning, we currently cannot validate our simulation results on a real road network. Nevertheless, we support the results observed in our simulations by analyzing real-world road data from over 7500 vehicles in a section of Europe during a 24-h period. By observing each vehicle's longitude and latitude at 5–10-min intervals, we can infer the number of platooning opportunities available from a small group of HDVs. Our data show that significant platooning opportunities exist in Europe today, and we believe that only slight changes in speed (such as those suggested by our distributed network of controllers) can facilitate appreciable savings.

We consider our work to significantly extend the work in [17], which addresses a single HDV increasing its speed to catch other vehicles or platoons. In this case, it only benefits the trailing HDV to speed up (and temporarily use more fuel) if the platoon is maintained long enough to justify the extra expended fuel. The proposed controller could be viewed as solving this single-HDV case whenever a vehicle approaches an intersection. Furthermore, our work is a natural extension of the eco-routing concept [18], [19] in which minimum-fuel-use routes are desired over minimum-distance routes. The routes suggested by the controller may slightly differ from the individual HDV's shortest path if considerable platooning savings can be achieved. We also mention the work of Meisen *et al.* in [20], which attempts to increase platooning throughout a network by using data-mining techniques to identify common routes in which platoons should be formed.

Our distributed control algorithm can be applied to a variety of platooning paradigms. The virtual controllers could physically reside at major intersections, they could be integrated into individual HDVs, or they could be simulated in a fleet management dispatch center. The only information required of the HDVs is their location, speed, and destination; such information is likely available for any vehicle using a GPS navigator.

An outline of this paper is as follows. Section II formally describes the model and introduces the local controller algorithm at the heart of our distributed control architecture. In Section II-C, we develop an efficient algorithm for solving this local control problem for two HDVs, and in Section II-D, we present one possible method for the multiple-HDV case. In Section III-A, we simulate our control paradigm on a large-scale simulation of the German Autobahn network. Section III-B discusses the relationship between the number of HDVs in the simulation and the possible savings in fuel consumption. Section III-C relaxes the constraint that HDVs only travel longer than the time required to traverse their shortest path. We substantiate the savings observed in our simulation by analyzing real-world data from a single day in a region of the European road network in Section IV. Section V concludes this paper with a brief summary and an outline of future work.

II. Coordinated HDV Platooning

In the following, we present a framework for modeling vehicles traveling on a road network. We then outline our local controller system to coordinate platoon formation.

A. Road and Vehicle Platoon Models

We can model a given road network as a graph $G = (V, E)$, where the edges E represent the road segments in the network, and the vertices V are nodes connecting the road segments. Furthermore, we define a vertex $\nu \in V$ with more than two connecting road segments as a junction. Without loss of generality, we can assume that the edges of E all have unit length. Any longer edges in an initial graph can be subdivided to satisfy this assumption. With these conditions, the vertices V represent possible HDV locations when traveling through the network. Of course, in reality, the vehicles are continuously traversing the road segments. Considering HDVs at vertices is equally valid, is much easier to handle computationally, and provides a natural platooning indicator: If two HDVs are at vertex ν_i at time t and an adjacent vertex ν_j at the next time step $t + 1$, then we consider them to have platooned over edge e_{ij}. Notice that in our network, all roads are considered flat, and each HDV is represented as a massless point location.

For each HDV T in the road network, we can assign a current starting location $s_T \in V$ and destination $d_T \in V$. Trailing vehicles in a platoon save a fraction $\eta \in (0,1)$ of the fuel used when traveling alone. For very small intervehicle distances η can be up to 21% [11]; we assume a more reasonable savings of 10% throughout this paper. We note that the leading vehicle in a platoon will also consume fuel at a slower rate because of reduced air drag [11]; these savings are considerably less than for the trailing vehicles, so we assume that only $n - 1$ vehicles in an n-vehicle platoon obtain the 10% reduction in fuel use.

We naturally want to use this 10% reduction in fuel use to minimize the total fuel expended, while ensuring that all HDVs reach their destinations on time. To accurately describe reality, additional constraints on the solution space are needed. Real-world HDV drivers will not go significantly out of their way to facilitate platooning formation. If $\mathcal{T}(s_T, d_T)$ is the time required to travel the shortest path from s_T to d_T, we must ensure that T does not travel more time than $\mathcal{T}(s_T, d_T) + m_T$. For the majority of this work, we consider $m_T = 0$, since most drivers will not drive any more time than necessary.

For further simplicity, we assume that the edges of the graph have been subdivided so that it takes one unit of time to travel between adjacent vertices. Therefore, vehicles also travel at a unit speed, except when the controller advises a vehicle to speed up to form a platoon.

B. Local Vehicle Coordination for Platooning Models

A global controller attempting to coordinate the timing and routes of every HDV in a real-world scenario is beyond current capabilities, not only because no such controller currently exists but also because coordinating every vehicle in a network centrally is computationally intractable. We therefore simplify the problem considerably by distributing controllers at junctions in the road network (similar to the hierarchical control system of [21], [22]).

For example, consider the scenario in Fig. 1(b). A single HDV and a platoon of two HDVs are approaching a location where they could possibly form a larger platoon. Knowing only the HDVs' current location, speed, and final destination, the controller can decide whether the single HDV should adjust its speed to form a platoon at the intersection or keep traveling alone. We define this problem as the "local controller problem." By placing local controllers at junctions, our method can coordinate fuel-efficient platoons in a distributed fashion while only slightly altering an HDV's route.

We make some assumptions to simplify the presentation. We only consider speeding up as an option for platoon formation, which means that the HDVs are not traveling at their maximum speed. Allowing vehicles to slow down provides more opportunities for platoon formation. Moreover, we only consider single vehicles increasing their speed to form platoons; the cost of increasing speed of multiple vehicles would rarely be beneficial. If platoons did speed up to catch another vehicle, all vehicles would incur additional fuel costs, but only one would eventually receive aerodynamic savings. For the majority of this paper (aside from Section III-C) we only consider vehicles taking routes that are of the same time as the time required to traverse their shortest path. We also do not consider the possibility of splitting existing platoons so that a given vehicle can increase its speed to form a platoon in the future. Although we show significant savings with these limiting assumptions, relaxing them can only increase the number of platooning opportunities and, therefore, the possible fuel savings.

A flowchart for the local controller's logic is given in Fig. 2(a). As a vehicle T_0 approaches an intersection, the controller must know the current speed, location, and destination of T_0 and any other approaching vehicles T_i. If any of the vehicles T_i can feasibly adjust its speed to meet (e.g., without violating posted speed limits) and if the savings from platooning after the intersection is larger than the cost of adjusting speeds to meet at the intersection, then the controller informs the HDV to modify its speed. If the approaching vehicle cannot platoon with any other vehicle, the controller naturally adds T_0 to the set of approaching vehicles in case viable platooning opportunities exist later. Of course, once a vehicle has passed the controller, it can be removed from consideration.

Algorithm 1: Pseudocode for the savings calculation in Fig. 2a for two HDVs.

Input A starting node s and two destinations d_1, d_2, and the matrix $\mathcal{F}(i, j)$ with entries corresponding to the fuel required to go from i to j;
Output The node where the platoon should split N_s and the savings SV;
Start:
$N_s \leftarrow s$; $Best \leftarrow \mathcal{F}(s, d_1) + \mathcal{F}(s, d_2)$;
$m_i \leftarrow 0\ \forall i$;
for ν *in* V **do**
  if $((2-\eta)\mathcal{F}(s,\nu) + \mathcal{F}(\nu, d_1) + \mathcal{F}(\nu, d_2) < Best)$ &
  $(\mathcal{F}(s,\nu) + \mathcal{F}(\nu, d_1) \leq \mathcal{F}(s, d_1) + m_1)$ &
  $(\mathcal{F}(s,\nu) + \mathcal{F}(\nu, d_2) \leq \mathcal{F}(s, d_2) + m_2)$ **then**
    $N_s \leftarrow \nu$;
    $Best \leftarrow (2-\eta)\mathcal{F}(s,\nu) + \mathcal{F}(\nu, d_1) + \mathcal{F}(\nu, d_2)$;
    Update m_1 or m_2 if needed;
  end
end
$SV = \mathcal{F}(s, d_1) + \mathcal{F}(s, d_2) - Best$;

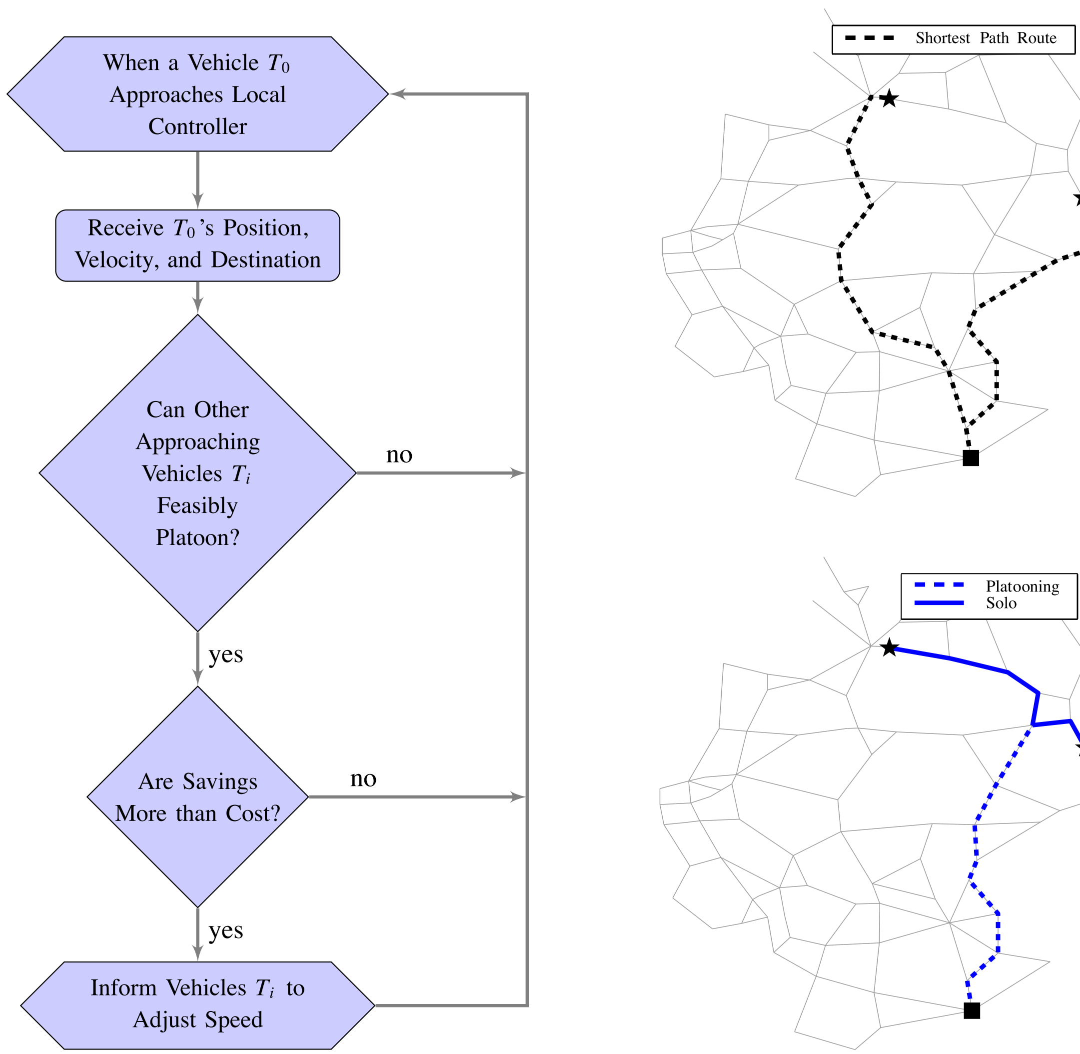

(a) (b)

Fig. 2. (a) Controller's logical flowchart. (b) Two HDVs taking (top) their shortest paths versus (bottom) the fuel-optimal route using platooning.

C. *Shortest Paths versus Fuel-Optimal Routes*

We now present a single instance of a local controller facilitating HDV platooning. Consider the map of Germany in Fig. 2(b) where two HDVs are approaching a junction (denoted by the square), and each HDV has a different destination (denoted by the stars). We assume that if the HDVs were independent, they would each take their respective shortest paths, shown in Fig. 2(b) (top). However, if one HDV slightly adjusts its speed so that both arrive at the junction simultaneously, they could form a platoon. The local controller, located at the square, must decide whether the additional fuel required to form the platoon will be offset by the savings from platooning.

Assuming that the controller has access to the all-pairs shortest path matrix $\mathcal{D}(i, j)$, it can quickly determine the most fuel-efficient route for the HDVs, comparing at most $|V|$ values in Algorithm 1. Notice that in this example, we only consider alternative paths with length equivalent to the shortest paths for each HDV (since $m_T = 0$ for all HDVs). That is, neither HDV must increase its travel time to follow the recommendations from the local controller. We consider increasing m_i slightly in Section III-C. We see in Fig. 2(b) (bottom) that the fuel-optimal routes returned from Algorithm 1 can allow for considerable platooning savings.

If the platooning savings are more than the cost of forming the platoon, the controller can advise a platoon to be formed. To calculate the platoon formation cost as two HDVs approach a node s, let v_1, v_2 be their respective velocities and let D_1, D_2 be their respective distances from s. Assume HDV 1 must speed up for the platoon to be formed (i.e., $D_1/v_1 > D_2/v_2$). Let us consider a cost for vehicle i defined as $F_i = D_i v_i^2$, which represents approximately how the additional fuel consumed due to air drag depends on distance and velocity.[1] For the platoon to be formed

$$\Delta v_1 = \frac{D_1}{D_2} v_2 - v_1 \text{ and } \Delta F_1 = D_1 \left(\left(\frac{D_1}{D_2} v_2 \right)^2 - v_1^2 \right) \quad (1)$$

[1]We assume no detailed model of the vehicles; hence, we are not able to estimate the absolute fuel costs. However, since our study focuses on the relative benefits from platooning, it suffices to consider the additional fuel consumed due to air drag.

where Δv_1 is the change in velocity, and ΔF_1 is the increased fuel cost for HDV 1 over D_1. If ΔF_1 is less than the platooning savings SV from Algorithm 1, the platoon should be formed. Algorithm 1 assumes all vehicles arrive at the controller simultaneously. The controller therefore compares whether the savings SV from Algorithm 1 is more than the cost ΔF_1 from (1). If $SV > \Delta F_1$, the controller informs the HDVs that a platoon should be formed.

We adopt the convention that $\mathcal{D}(i,j)$, $\mathcal{T}(i,j)$, and $\mathcal{F}(i,j)$ are given in per unit, since we have assumed that edges have unit length, traveling across an edge takes unit time, and the corresponding fuel cost is one unit. Hence, these parameters can be used interchangeably.

D. Control for More Than Two HDVs

Before showing the benefit of multiple controllers distributed throughout a large-scale network, we must first generalize Algorithm 1 for more than two HDVs. As we developed above, when two HDVs are approaching an intersection at s, at most $|V|$ quantities must be examined to find the most fuel-efficient route. If g HDVs are approaching the same controller, computing the optimal platooning route requires exponentially more comparisons. This computational growth prevents finding the exact optimum for $g \geq 4$; however, we highlight a fast heuristic that closely approximates the optimal routes.

If g HDVs are approaching a controller, finding a fuel-efficient route can be broken down into $\binom{g}{2}$ pairwise decision problems, which can be quickly solved by Algorithm 1. If no pair of HDVs has platooning savings that outweigh the cost of formation, no controller action is taken. Otherwise, the pair of HDVs that incurs the largest savings considered fixed, that platoon is formed and considered one unit. This process is repeated with $\binom{g-1}{2}$ pairs of HDVs and continues until every vehicle is assigned to a platoon, or none of the pairwise savings from Algorithm 1 outweigh the cost of platoon formation.

Although the exact calculation of the optimal routes for general g is too time consuming in practice, we can examine how the proposed heuristic compares for moderately sized problems. Since the optimal solution for four HDVs is attainable within reasonable time and since a situation with five or more HDVs rarely occurred in our simulation, we chose to evaluate the heuristic for this number. We place four HDVs at a random node s in the Germany network [shown in Fig. 4(c)] and assign each a random destination. We can then compare the amount of fuel saved by the controller with the amount of fuel saved by the optimal solution. We repeat this random experiment 1000 times for four HDVs and show the relative difference between the optimal and pairwise solutions in Fig. 3. The figure shows the cumulative distribution of savings. We observe that the routes returned from Algorithm 1 using our proposed heuristic were less than 80% of the optimal fuel savings for only 2% of the experiments. The data point (90%, 3.5%) means our heuristic returned a solution that was worse than 90% of the optimal savings in only 3.5% of the cases. We see that for 96% of the experiments, our proposed heuristic computes the path with optimal savings.

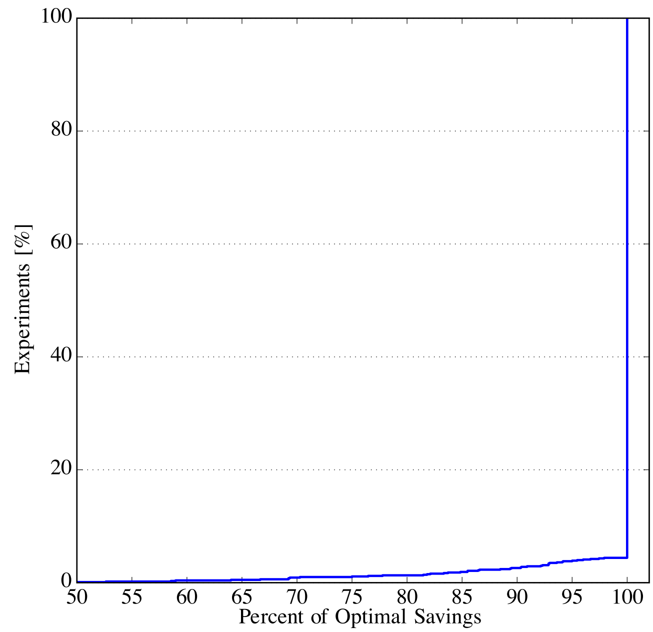

Fig. 3. Cumulative distribution function of the percentage of the amount of fuel saved by the local controller compared with the optimal solution for four HDVs, 1000 random experiments.

III. Simulation of the German Road Network

The following results show the strengths and savings of our platooning control methodology.[2]

A. Experimental Framework

To evaluate the performance of our algorithm on a large scale, we use the simplified graph of the German Autobahn network with 647 nodes, 695 edges, and 12 destinations shown in Fig. 4(c). For this simulation, we consider a static network with fixed number of HDVs and do not consider the speed to be influenced by traffic. Furthermore, an HDV can speed up to catch another vehicle only if it is trailing by, at most, one edge length. In the German road network, each edge length corresponds to roughly 13 km, which we consider to be a reasonable distance for an HDV to catch up. Of course, this catch-up of one edge length must be spread out over a stretch of road long enough to prevent illegal speed. For example, if two HDVs driving at the same speed are approaching a local controller, respectively 10 and 11 edge lengths away, the latter HDV could increase its speed by 10% to facilitate platoon formation. On the other hand, if the respective distances were only 1 and 2 edge lengths, the latter HDV could not double its speed to form a platoon.

Vehicles are placed at random points throughout the network and started toward their random destination (following their shortest path). They move from node to node in the network along their shortest path, moving one node in each time step. Whenever a vehicle leaves an intersection, it broadcasts its route to the next intersection on its path. At every time step, the controller attempts to find a routing using the pairwise

[2] MATLAB code to recreate all simulation data and duplicate all figures can be found here: `http://people.kth.se/~jeffreyl/Platooning/`.

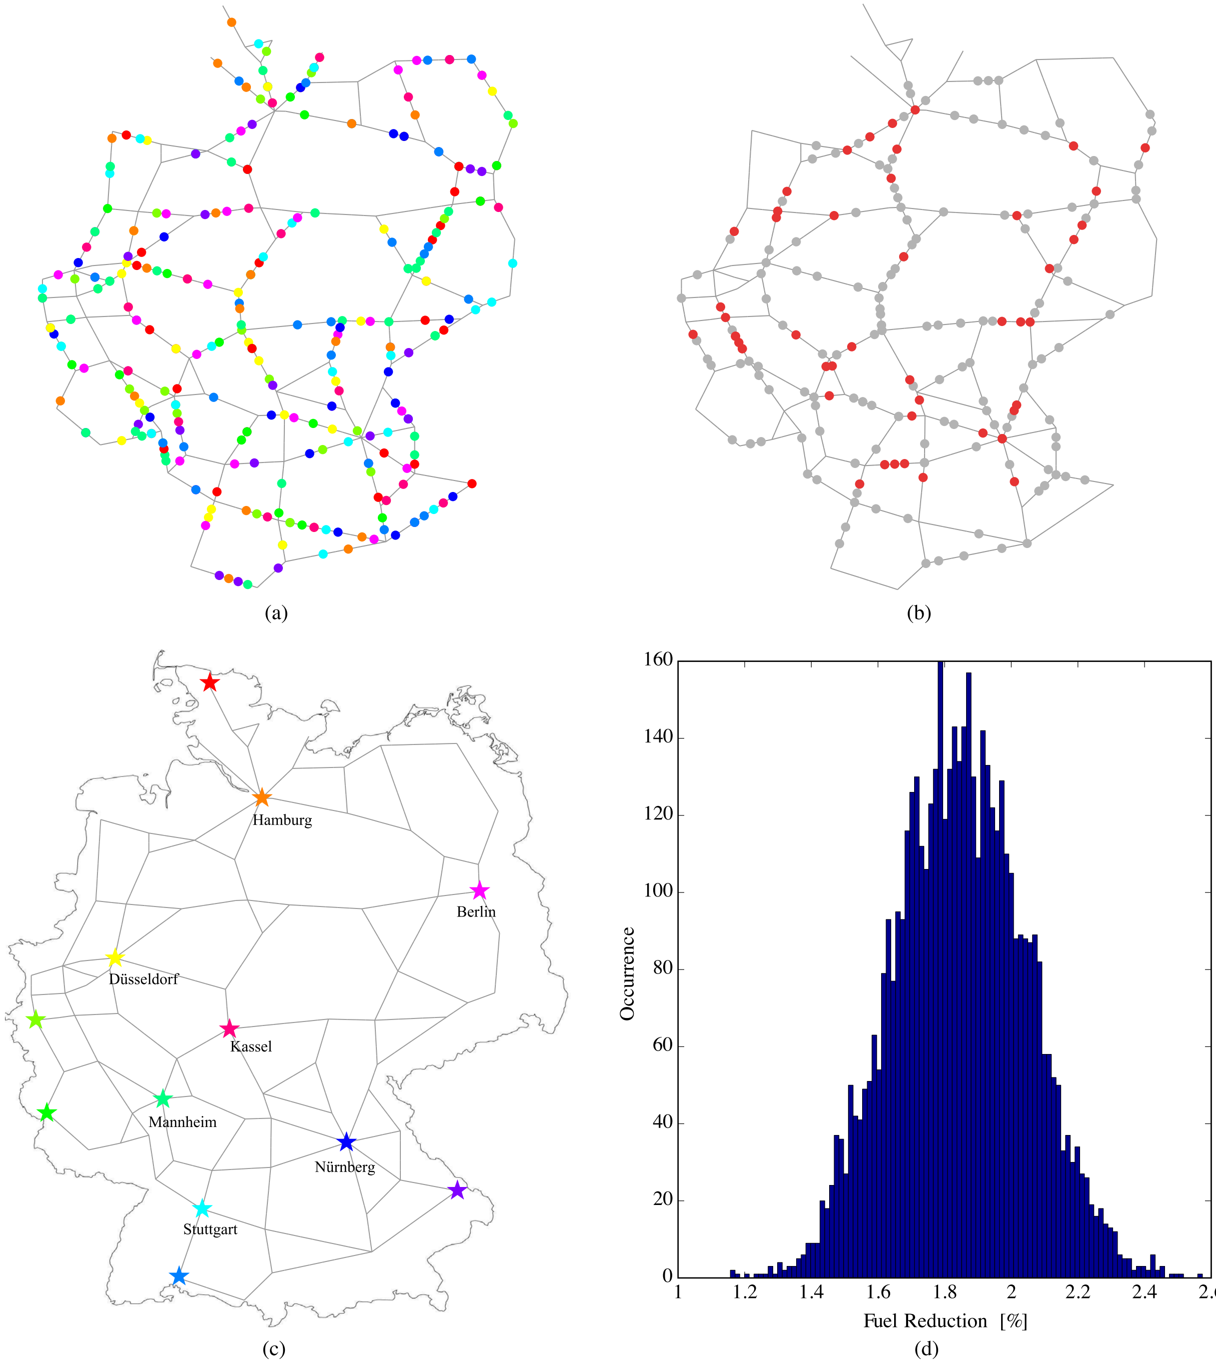

Fig. 4. Large-scale simulation results. (a) Snapshot of an initial state of our system, with dots representing HDVs. The color of an HDV designates its destination according to image (c). (b) Snapshot of the same system as in image (a) after ten time steps. Platoons are depicted as red dots, whereas single vehicles appear in gray. (c) Visualization of the simplified German Autobahn network used in our simulations (nodes are not shown). Possible HDV destinations are represented by stars. (d) Percentage of fuel saved by our approach, compared with every HDV taking its shortest path, for 300 HDVs, repeated 5000 times with random starting points and destination.

algorithm presented in Section II-D. If platooning savings are larger than the cost of adjusting speeds (so the relevant vehicles arrive at the intersection simultaneously), the vehicles' speeds are adjusted. This is represented in the simulation by allowing the vehicles to travel t nodes on the path to the controller in $t-1$ time steps. The simulation stops when all vehicles have reached their destination.

One example initial state of 300 HDVs is depicted in Fig. 4(a), with HDVs represented by dots, in which the color matches the destination's color in Fig. 4(c). After starting the simulation using the initial state shown in Fig. 4(a), we pause after ten time steps and observe the network in Fig. 4(b). The platooning vehicles are in red, and the solo HDVs in gray. Approximately 30% of the vehicles have formed platoons at this stage.

We repeat the simulation 5000 times for 300 HDVs with random starting points and destinations. As might be expected, some random configurations allow for more platooning opportunities than others. In Fig. 4(d), we see a histogram of the percentage of total fuel saved by our approach in each simulation

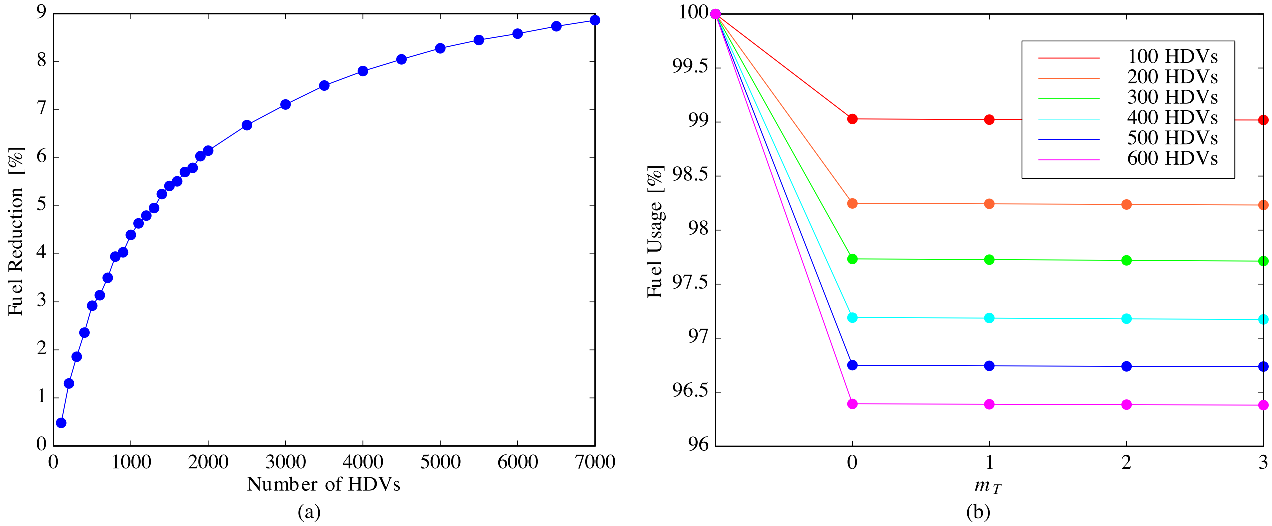

Fig. 5. (a) Percentage of fuel saved by our method compared with every HDV taking its shortest path and not platooning. Each data point is an average of the savings from five random instances. The point (300, 1.82) mirrors the results in Fig. 4(d). (b) Number of additional edges that an HDV can travel versus the average fuel savings for various numbers of HDVs on the Germany network. Longer detours result in lower average fuel consumption, but the benefits are very small for this network.

(compared with every HDV taking its shortest path). We see that even for this relatively low number of HDVs, the average total fuel consumption has been decreased by almost 2%.

B. Benefit of Increasing the Number of HDVs

Here, we analyze how the total fuel use changes as the number of HDVs in the network increases. Intuitively, one would assume that if the density of vehicles in the network is low, there are few opportunities for platooning; few HDVs will take a route other than their shortest path. As the HDV density in the network increases, more HDVs will avail themselves of platooning options; hence, more savings will be exploited. Eventually, once the number of HDVs in the network grows large, all opportunities for savings are extracted from the network topology; adding more vehicles will not decrease the average fuel use considerably.

We see that this intuition is true (at least for the Autobahn) in Fig. 5(a). Average fuel savings rapidly increase between 0 and 2000 vehicles. As the network becomes "saturated" with vehicles, nearly every edge can be traveled in a platoon; hence, nearly every HDV uses 10% less fuel (compared with the fuel use when driving its shortest path alone), and adding more HDVs will result in only marginal savings.

C. Increasing Allowable Detours

All previous results assume that vehicle routes are of the same length as an HDV's shortest path from start to destination. We now examine the possibility of allowing routes for an HDV that are longer than the length of its shortest path from start to destination.

Intuitively, by adding extra edges, a vehicle that could travel would have quickly diminishing returns. Allowing an HDV to travel 10 or 20 km extra will help improve the average fuel use because more platooning options will be available. However, allowing an HDV 60 km of additional travel is unlikely to provide much additional savings; a vehicle that travels 60 km extra must be platooned an exceptionally long time to offset the costs of platoon formation.

We find this intuition to hold in Fig. 5(b), in which we partially resimulate our German road network with a slight modification to Algorithm 1. Instead of defining $m_T = 0$ for all vehicles, we assign each HDV an upper bound m_T and ensure that the controller never returns a route for k_T that will result in a total travel time more than $\mathcal{T}(s_T, d_T) + m_T$. For example, if $m_T = 3$ for some vehicle T that has already traveled two additional edges before approaching an intersection, the local controller at $\nu \in V$ can only look for routes with one or zero edges more than $\mathcal{T}(\nu, d_T)$. In Algorithm 1, m_i would be nonzero, and m_T is updated at "Update m_1 or m_2 if needed."

The results of increasing m_T uniformly for all T are seen in Fig. 5(b). A given experiment starts by randomly assigning HDVs starting points and then observing the savings produced by our approach. The possible savings are then recomputed using the same starting points and a larger value of m_T. Each point in Fig. 5(b) is then the average of five such experiments. We see that the majority of the savings produced by our local controllers arise from synchronization. Since increasing m_T from 0 to 1 results in almost imperceptible savings, we conclude that the local controllers are rarely routing HDVs off their shortest path routes, at least for the network in question. (It may also be that there are few paths from a given point to a destination that are only a few edges longer than the shortest path.) We consider the fact that most of the savings arise from coordination as favoring our system's possible adoption: More HDV drivers are willing to participate in a system that does not significantly modify routes they are already traveling.

It might seem surprising that allowing every HDV to possibly detour 10 km does not allow for a relatively larger savings than when no detour is allowed. Although the small detour does allow for significantly more platooning options, those options are rarely long enough to justify driving 10 km out of the way. If the savings for platooning are only 10%, then the platoon must stay together for 100 km to warrant the detour.

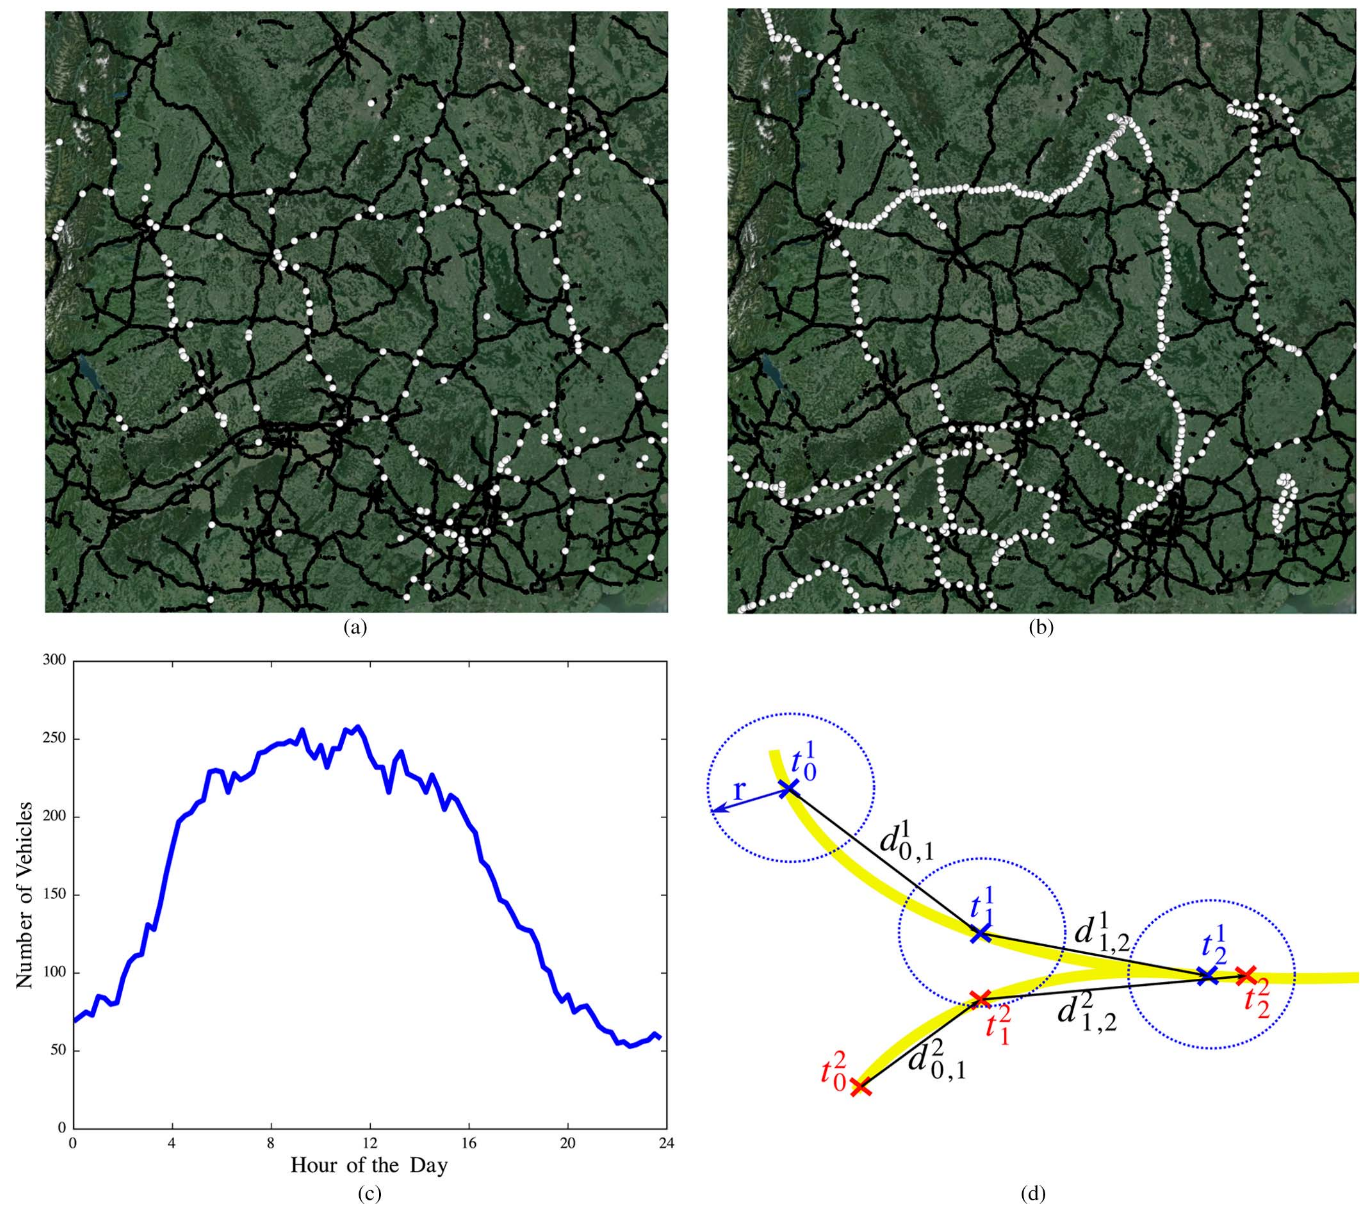

Fig. 6. Real-time HDV data study. (a) Snapshot of the probe data. Black lines are the roads and white dots show each vehicle's location. (b) Selected vehicles' probe data trajectories. (c) Total number of vehicles moving at a given point throughout the 24-h period. (d) Example of the probe data with a radius r used to determine whether another vehicle is nearby.

The impact of increasing m_T depends on the structure of the graph; if many similar length routes exist, allowing slight detours will likely produce greater savings. We have also simulated more realistic scenarios in which m_T is an increasing function of $\mathcal{T}(s_T, d_T)$ (HDVs with longer travel times can tolerate more detours); however, we found the fuel savings to be nearly identical to the constant m_T case.

IV. Supporting Analysis of HDV Data

Here, we show preliminary results that we believe substantiate the results obtained from our simulations. To understand whether our approach is feasible, we collected position data from HDVs throughout a single day. Our data show that many platooning possibilities exist today; HDVs in Europe are often very close to other HDVs. Therefore, a minimally invasive method, such as the proposed network of distributed controllers, can facilitate significant platooning opportunities.

Our data come from vehicle probe data (commonly called floating car data) from a collection of HDVs. The probe data consist of the location and time of a vehicle and were collected by the in-cab GPS. Since the data are somewhat sparse, the conclusions that can be drawn are limited but indicate some potential. We believe that the data support the result of our simulations.

A. Data Description

The probe data were obtained from HDVs from a single day in a region in Europe. The probe data come from Scania HDVs within a 500 000 km^2 area over a 24-h period in the spring of 2013. For the 7634 HDVs in the region, including both long-haulage and local-distribution vehicles, the following information was collected at 5–10-min intervals using the on-board

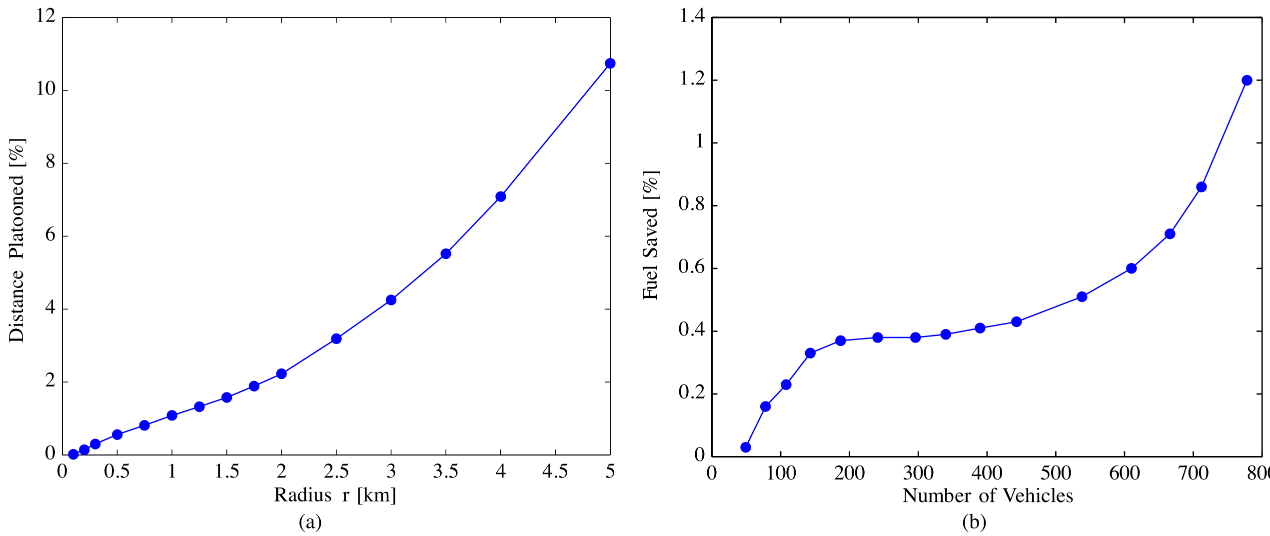

Fig. 7. (a) Percentage of total distance (403 413 km) that the 875 vehicles have platooned for a given radius. (b) Percentage of fuel saved by the vehicles that platooned at least once.

GPS: vehicle ID, timestamp, longitude, and latitude. A snapshot of the data at one point during the day can be seen in Fig. 6(a), and some vehicles' trajectories can be seen in Fig. 6(b).

We consider vehicles that travel more than 150 km from their start to the end to be long-haulage vehicles and, therefore, relevant to our study. This cutoff likely categorizes some long-haulage vehicles falsely as local distributors (if they drove far away and then returned to a home base) and likely incorrectly labels some local distributors as long-haulage vehicles (if they deliver along a main thoroughfare with many stops). Of the 7634 vehicles, 875 satisfy the 150-km threshold. In Fig. 6(c), we can see how many of these 875 vehicles are moving at any point in time in Fig. 6(c).

B. Measuring Current Platooning Potential

To measure the potential for platoon formation, we synchronized the probe data to deduce how many vehicles are traveling in close proximity to other HDVs. First, we interpolated a line between each vehicle's latitude and longitude at each data point. Then, we added data points every 5 min along this line. To determine whether a platooning opportunity exists for a given vehicle, we look to see if there is a data point for any other vehicle within radius r. If there is another vehicle with radius r at time t_1 and t_2, then we say that the two vehicles could have platooned between time t_1 and t_2 [for an example, see Fig. 6(d)]. As with the rest of the paper, we assume that this platooning allows one vehicle to reduce its fuel cost by 10% on the stretch of road in question. For simplicity, we assume that the fuel cost is proportional to the distance driven; hence, the fuel cost will be 1 unit/km when driving alone and 0.9 unit/km when platooning. If n vehicles are found to be within radius r of each other between two points in time, then $n-1$ vehicles will obtain reduced fuel cost.

Admittedly, this approach can introduce errors for a variety of reasons. Two vehicles may be driving on different roads for the period of time in question. Moreover, platooning opportunities can be missed if the road curvature and asynchronous timing prevent us from recognizing that the vehicles are driving nearby (errors can be both positive and negative).

We analyzed our data for radii between 0.1 and 5 km. We highlight the data for $r = \{0.2\ \text{km}, 1\ \text{km}, 5\ \text{km}\}$ below. We consider the radius of 0.2 km to be sufficiently small so as to count the number of vehicles naturally platooning today. HDVs that are within a radius of 1 km are likely to not be platooning, and the driver is unlikely to know that another HDV is close by; however, a slight amount of coordination would allow for these vehicles to platoon. A radius of 5 km would measure the platooning opportunities that could be generated by a moderate amount of help, such as the proposed local controller network. Notice that for the case of $r = 1$ km and $r = 5$ km, we did not consider the additional cost for speed changes that would be required to coordinate platoon formation; we are merely counting the vehicles within radius r. The results of the study are as follows:

- $r = 0.2$ km
  - 78 out of 875 vehicles platooned at least once during their daily route.
  - 0.16% of total fuel saved by the platooned vehicles.
  - 585-km platooning out of total 403 413 km driven.
- $r = 1$ km
  - 241 out of 875 vehicles platooned at least once during their daily route.
  - 0.38% of total fuel saved by the platooned vehicles.
  - 4369-km platooning.
- $r = 5$ km
  - 778 out of 875 vehicles platooned at least once during their daily route.
  - 1.2% of total fuel saved by the platooned vehicles.
  - 43 325-km platooning.

The percentage of distance platooned as well as the percentage of fuel that could have been saved for a range of radii is presented in Fig. 7. In Fig. 7(a), as the radius approaches 5 km, nearly all vehicles have at least one platooning opportunity at some point during the day. Total fuel saved is directly proportional to total distance traveled, but Fig. 7(b) shows the percentage of fuel that would have been saved by vehicles that had at least one opportunity to platoon at some point during the day. Note that since Fig. 6(c) shows that, at most, 250 vehicles are on the road at any period in time, the fuel savings of 1.2% for a radius of 5 km match the simulation results in Fig. 5(a) for a similar number of vehicles.

C. Interpretation of Results

This case study is brief but provides an indication as to how many platooning opportunities exist in reality and how much effort might be required to coordinate their formation. As we can see from the $r = 0.2$ km case, few vehicles are platooning or even close enough to easily recognize platooning opportunities today. However, we can clearly see that if r is increased, the potential savings and platooning rate increases appreciably. For example, nearly 90% of vehicles in the study had an opportunity to platoon at some point in their day's travel when $r = 5$ km.

V. Conclusion

In this paper, we have developed a distributed method for platoon formation. Using virtual local controllers throughout a road network, we showed how significant fuel savings can be achieved by platooning vehicles. This can be accomplished by keeping vehicles on their shortest path from start to destination but slightly adjusting their speeds to synchronize travel with other HDVs. These results assume a static network in which the number of participating vehicles does not affect the time required to traverse an edge. It is an open problem to route vehicle platoons in a dynamic network where the speeds and catch-up possibilities might decrease because of capacity-overloaded roads. We also analyzed a real-world data set to show that while few vehicles are currently platooning, many opportunities to platoon exist. Therefore, if vehicles are willing to adjust their speeds a small amount, considerable savings can be achieved. Future work includes an experimental evaluation of the approach proposed in this paper. For example, HDV platooning experiments are currently being performed with 25 HDVs regularly traveling between Södertälje, Sweden, and Zwolle, The Netherlands.

Acknowledgment

The authors would like to thank C. Kammer for the German Autobahn road network simulation code and four anonymous referees for their comments that led to an improved version of this paper.

References

[1] "Energy balance sheets," Brussels, Belgium, Tech. Rep., 2013.

[2] "Road freight transport vademecum," Geneva, Switzerland, Tech. Rep., Mar. 2009.

[3] W. Levine and M. Athans, "On the optimal error regulation of a string of moving vehicles," *IEEE Trans. Automat. Control*, vol. AC-11, no. 3, pp. 355–361, Jul. 1966.

[4] S. Melzer and B. Kuo, "Optimal regulation of systems described by a countably infinite number of objects," *Automatica*, vol. 7, no. 3, pp. 359–366, May 1971.

[5] T. Tank and J.-P. Linnartz, "Vehicle-to-vehicle communications for AVCS platooning," *IEEE Trans. Veh. Technol.*, vol. 46, no. 2, pp. 528–536, May 1997.

[6] A. Broggi, M. Bertozzi, A. Fascioli, C. Guarino Lo Bianco, and A. Piazzi, "Visual perception of obstacles and vehicles for platooning," *IEEE Trans. Intell. Transp. Syst.*, vol. 1, no. 3, pp. 164–176, Sep. 2000.

[7] F. Michaud, P. Lepage, P. Frenette, D. Letourneau, and N. Gaubert, "Coordinated maneuvering of automated vehicles in platoons," *IEEE Trans. Intell. Transp. Syst.*, vol. 7, no. 4, pp. 437–447, Dec. 2006.

[8] F. Browand, J. McArthur, and C. Radovich, "Fuel saving achieved in the field test of two tandem trucks," California PATH, Richmond, CA, Canada, Final Report UCB-ITS-PRR-2004-20, 2004.

[9] T. Robinson, E. Chan, and E. Coelingh, "Operating Platoons on public motorways: An introduction to the SARTRE platooning programme," in *Proc. 17th ITS World Congr.*, Busan, Korea, 2010, pp. 1–11.

[10] J. Mårtensson *et al.*, "The development of a cooperative heavy-duty vehicle for the GCDC 2011: Team scoop," *IEEE Trans. Intell. Transp. Syst.*, vol. 13, no. 3, pp. 1033–1049, Sep. 2012.

[11] C. Bonnet and H. Fritz, "Fuel consumption reduction in a platoon: Experimental results with two electronically coupled trucks at close spacing," in *Proc. Future Transp. Technol. Conf. Expo.*, Costa Mesa, CA, USA, 2000, 724.

[12] R. Kunze, R. Ramakers, K. Henning, and S. Jeschke, "Organization and operation of electronically coupled truck platoons on German Motorways," in *Intelligent Robotics and Applications*, vol. 5928, M. Xie, Y. Xiong, C. Xiong, H. Liu, and Z. Hu, Eds. Berlin: Springer-Verlag, 2009, ser. Lecture Notes in Computer Science, pp. 135–146.

[13] S. Tsugawa, "An overview on an automated truck platoon within the Energy ITS Project," in *Proc. 7th IFAC Symp. Adv. Autom. Control*, Tokyo, Japan, Sep. 2013, pp. 41–46.

[14] R. Horowitz and P. Varaiya, "Control design of an automated highway system," *Proc. IEEE*, vol. 88, no. 7, pp. 913–925, Jul. 2000.

[15] E. Larsson, G. Sennton, and J. Larson, The Vehicle Platooning Problem: Computational Complexity and Heuristics, *Preprint,* 2014. [Online]. Available: http://www.optimization-online.org/DB_HTML/2013/09/4030.html

[16] "European motor vehicle PARC 2010," Brussels, Belgium, Tech. Rep., Apr. 2012.

[17] K.-Y. Liang, J. Mårtensson, and K. H. Johansson, "When is it fuel efficient for a heavy duty vehicle to catch up with a platoon?" in *Proc. 7th IFAC Symp. Adv. Autom. Control*, Tokyo, Japan, Sep. 2013, pp. 738–743.

[18] E. Ericsson, H. Larsson, and K. Brundell-Freij, "Optimizing route choice for lowest fuel consumption—Potential effects of a new driver support tool," *Transp. Res. Part C: Emerging Technol.*, vol. 14, no. 6, pp. 369–383, Dec. 2006.

[19] T. Kono, T. Fushiki, K. Asada, and K. Nakano, "Fuel consumption analysis and prediction model for 'Eco' route search," in *Proc. 15th World Congr. Intell. Transp. Syst.*, Washington, DC, USA, 2008, pp. 1–9.

[20] P. Meisen, K. Henning, and T. Seidl, "A data-mining technique for the planning and organization of truck platoons," in *Proc. Int. Conf. Heavy Veh.*, Paris, France, May 2008.

[21] L. D. Baskar, B. De Schutter, and H. Hellendoorn, "Hierarchical traffic control and management with intelligent vehicles," in *IEEE Intell. Veh. Symp.*, Istanbul, Turkey, Jun. 2007, pp. 834–839.

[22] L. D. Baskar, B. De Schutter, and J. Hellendoorn, "Optimal routing for intelligent vehicle highway systems using a macroscopic traffic flow model," in *Proc. IEEE Conf. Intell. Transp. Syst.*, St. Louis, MO, USA, Oct. 2009, pp. 1–6.

Jeffrey Larson received the Ph.D. degree in applied mathematics from University of Colorado, Denver, CO, USA.

He was a Postdoctoral Researcher with the Automatic Control Laboratory and the ACCESS Linnaeus Centre, School of Electrical Engineering, KTH Royal Institute of Technology, Stockholm, Sweden. He is currently a postdoctoral appointee with the Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL, USA.

Kuo-Yun Liang received the M.Sc. degree in electrical engineering from KTH Royal Institute of Technology, Stockholm, Sweden, in 2011. He is currently working toward an Industrial Ph.D. degree with the Automatic Control Laboratory and the ACCESS Linnaeus Centre, School of Electrical Engineering, KTH, and with the Research and Development Department, Scania CV AB, Södertälje, Sweden.

Karl H. Johansson (F'13) received the M.Sc. and Ph.D. degrees in electrical engineering from Lund University, Lund, Sweden.

He has held visiting positions with University of California, Berkeley, CA, USA, from 1998 to 2000 and with the California Institute of Technology, Pasadena, CA, USA, from 2006 to 2007. He is Director of the KTH ACCESS Linnaeus Centre and a Professor with the School of Electrical Engineering, Royal Institute of Technology, Stockholm, Sweden. He is a Wallenberg Scholar and the Director of the Stockholm Strategic Research Area ICT The Next Generation. His research interests include networked control systems, hybrid and embedded systems, and applications in smart transportation, energy, and automation systems.

Dr. Johansson has been a member of the *IEEE Control Systems Society* Board of Governors and the Chair of the International Federation of Automatic Control (IFAC) Technical Committee on Networked Systems. He has been on the editorial board of several journals, including *Automatica*, IEEE TRANSACTIONS ON AUTOMATIC CONTROL, and *IET Control Theory and Applications*. He is currently on the editorial board of IEEE TRANSACTIONS ON CONTROL OF NETWORK SYSTEMS and *European Journal of Control*. He has received many awards, including the triennial Young Author Prize from IFAC.