

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

A Distributed Framework for Coordinated Heavy-duty Vehicle Platooning

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

Mathematics and Computer Science Division

Preprint ANL/MCS-P5100-0214

February 2014

¹Support for this work was provided through the SciDAC program funded by U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research.

A Distributed Framework for Coordinated Heavy-duty Vehicle Platooning

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

(Invited Paper)

Abstract—Heavy-duty vehicles traveling in a single file with small intervehicle distances experience a reduced aerodynamic drag and therefore have an 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 savings surpassing 5% when only few thousand vehicles participate in the system. These results are corroborated by an analysis of real-world heavy-duty vehicle data that shows significant platooning opportunities currently exist, suggesting a slightly invasive network of distributed controllers, such as the one proposed in this paper, can yield considerable savings.

I. INTRODUCTION

HEAVY-DUTY vehicles (HDVs) driving in a platoon use significantly less fuel than when driving separately. By traveling in a single file with small intervehicle distances, trailing vehicles experience reduced aerodynamic drag. This reduces the total fuel consumed, and therefore reduces a significant cost to vehicle owners. For example, transport in Europe uses approximately 180 million tons of diesel fuel [1] every year, and therefore even modest decreases in fuel use can yield drastic savings. Also, since HDVs are a significant contributor to greenhouse gas emissions (road transport generates 16% of the CO₂ pollution in Europe [2]), platooning vehicles can also have a notable environmental impact. A photograph of platooning vehicles can be seen in Fig. 1a.

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 has largely been neglected up to 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. Secondly, there is no global coordinator with the authority to suggest routes which provide platooning opportunities. Lastly, even if complete knowledge of every HDV was available and a global coordinator could direct vehicles along any route, the general problem of finding the 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 which coordinate platoon formation by slightly adjusting HDVs' speeds 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 to be 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 the vehicles coordinate their arrival at the intersection). The controller only advises a velocity adjustment if the savings from platooning are more than the cost of adjusting speeds to meet at the intersection. The controller naturally can be constrained so as to ensure 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. Though the amount of fuel saved clearly

K.-Y. Liang, and K. H. Johansson are with the ACCESS Linnaeus Centre and the Department of Automatic Control at KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden. {kyliang, kallej}@kth.se

J. Larson is with the Mathematics and Computer Science Division at Argonne National Laboratory, Argonne, IL 60439. jmlarson@anl.gov This work was performed while he was with KTH.

K.-Y. Liang is also with Scania CV AB, SE-151 87 Södertälje, Sweden. kuo-yun.liang@scania.com

This work was supported by VINNOVA within the FFI program, under contract 2011-01144, EU FP7 Project COMPANION, the Swedish Research Council, the Swedish Foundation for Strategic Research, and the Knut and Alice Wallenberg Foundation.

Manuscript received Date; revised Date.

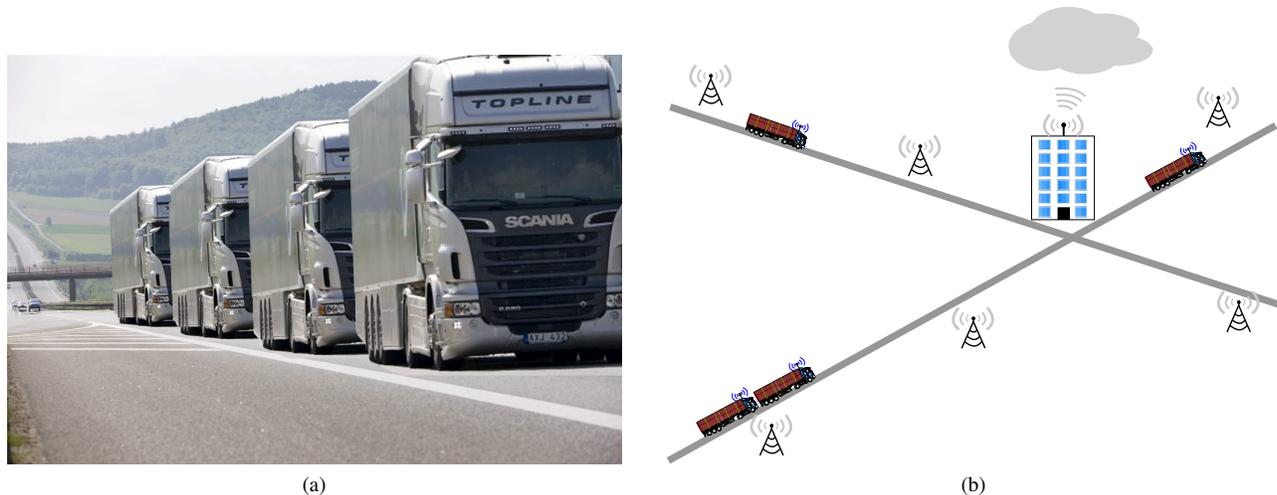


Fig. 1. (a) Four platooning heavy-duty vehicles. (b) A single HDV and a platoon approaching a local virtual controller at a highway road junction.

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 two thousand participating vehicles can reduce their fuel consumption by over 5% when using our approach. Given that Germany has over 400,000 registered HDVs [16], this is a relatively small number of vehicles. We believe 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 due to legality concerns surrounding driving HDVs at small intervehicle distances. Though 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 7,500 vehicles in a section of Europe during a 24 hour period. By observing each vehicle's longitude and latitude at 5-10 minute intervals, we can infer the number of platooning opportunities available from a small group of HDVs. Our data shows 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 only 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] where minimum-fuel-use routes are desired over minimum-distance routes. The routes suggested by the controller may differ slightly from the individual HDV's shortest path if considerable platooning savings can be achieved. Lastly, we mention the work of [20] which attempts to increase platooning throughout a network

by using data-mining techniques to identify common routes where 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 the paper is as follows. Section II formally describes the model and introduces the local controller algorithm at the heart of our distributed control architecture. We develop an efficient algorithm for solving this local control problem for two HDVs in Section II-C and present one possible method for the multiple HDV case in Section II-D. 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. Lastly, Section V concludes the paper.

II. COORDINATED HDV PLATOONING

In the following, we present a framework for modeling vehicles traveling on a road network and 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

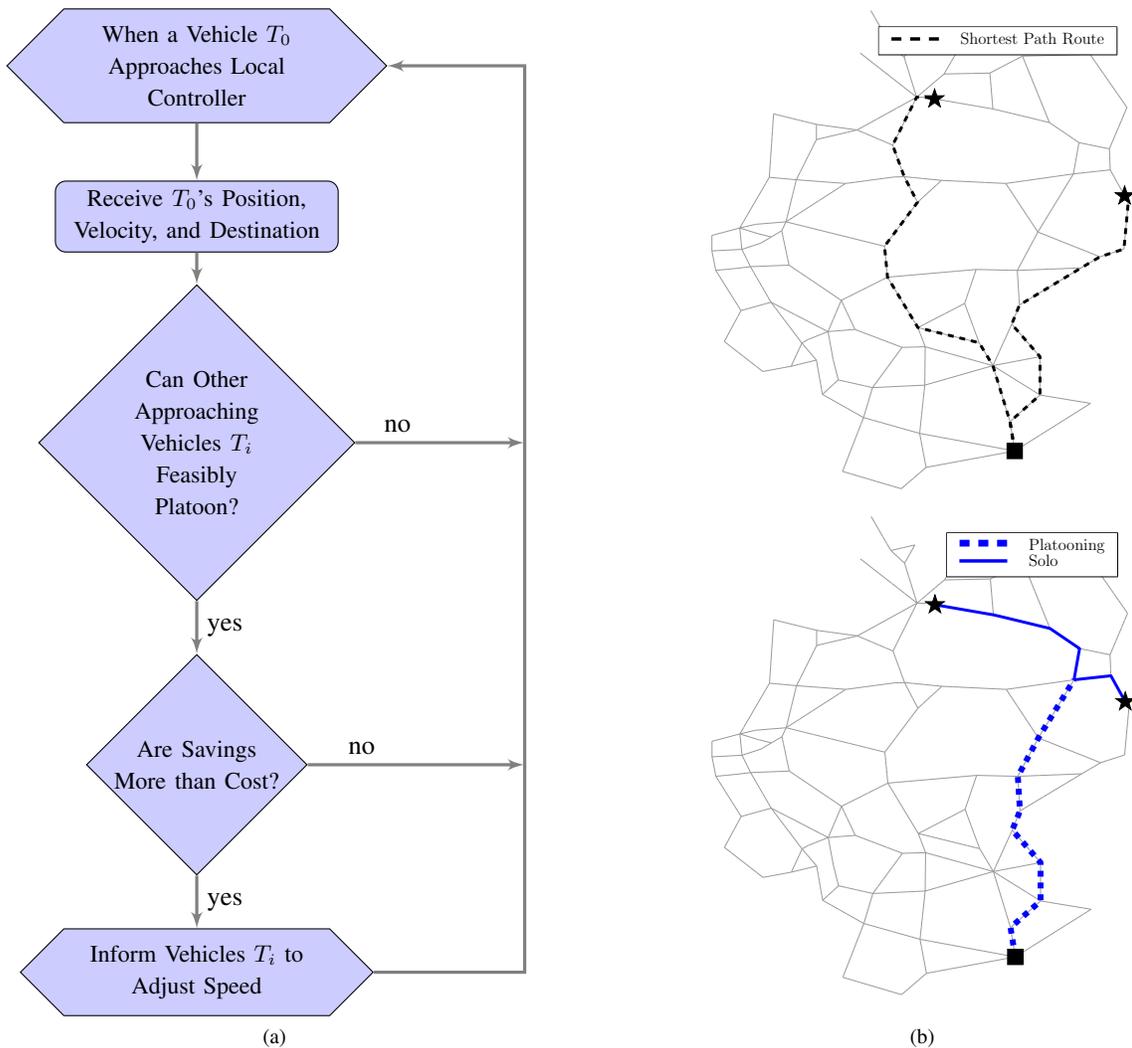


Fig. 2. (a) The controller's logical flowchart. (b) Two HDVs taking their shortest paths (top) versus the fuel-optimal route using platooning (bottom).

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, 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 the HDVs are 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. It should be noted 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 the trailing vehicles, so we assume 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 all HDVs reach their destinations on time. To accurately describe reality, there must be additional constraints on the solution space. Real-world HDV drivers will not go significantly out of their way in order 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 T does not travel more time than $\mathcal{T}(s_T, d_T) + m_T$. For the majority of this work, we consider $m_T = 0$ as 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 take 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

As argued previously, 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. 1b. 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 speeds 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, meaning the HDVs are not traveling at their maximum speed. Allowing vehicles to slow down provides more opportunities for platoon formation. Also, 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 the paper (aside from Section III-C) we only consider vehicles taking routes that are the same time as the time required to traverse their shortest path. We also do not consider the possibility of splitting existing platoons so a given vehicle can increase its speed to form a platoon in the future. Though 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 can be found in Fig. 2a. 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 their speeds to meet (e.g., without violating posted speed limits), and 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(s) to modify their speed. If the approaching vehicle can not 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.

C. Shortest Paths vs. Fuel-optimal Routes

We now present a single instance of a local controller facilitating HDV platooning. Consider the map of Germany in Fig. 2b 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. 2b (top). However, if one HDV slightly adjusts its speed so both arrive at the junction simultaneously, they could form a platoon. The local controller, located at the square, must decide if the additional fuel required to form the platoon will be offset by the savings from platooning.

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

$SV = \mathcal{F}(s, d_1) + \mathcal{F}(s, d_2) - Best$;

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 in order to follow the recommendations from the local controller. We consider increasing m_i slightly in Section III-C. We see in Fig. 2b (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 in order for the platoon to be formed (i.e., $\frac{D_1}{v_1} > \frac{D_2}{v_2}$). 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 velocity¹. For the platoon to be formed

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

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 if the savings SV from Algorithm 1 is more than the cost ΔF_1

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

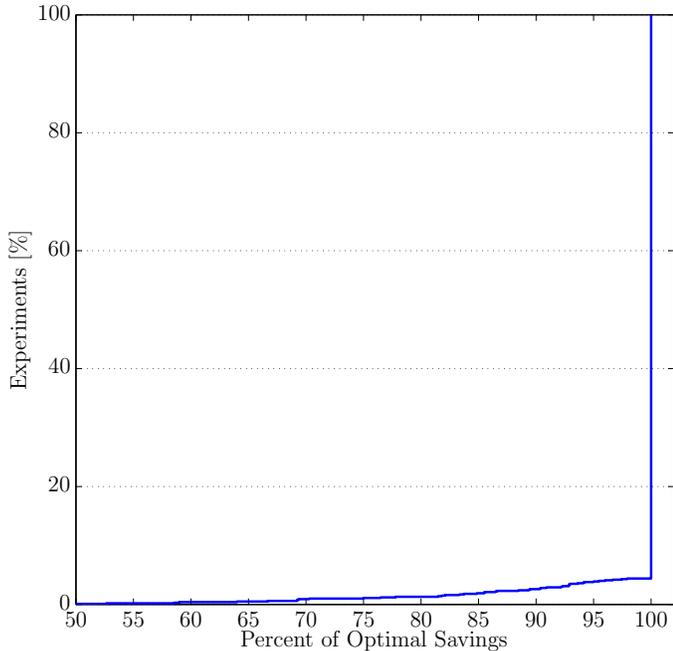


Fig. 3. Cumulative distribution function of the percentage of the amount of fuel saved by the local controller compared to the optimal solution for 4 HDVs, 1,000 random experiments.

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 units, as we have assumed that edges have unit length, travelling 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 $g \geq 4$, but 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 quickly be 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.

Though 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 4 HDVs is attainable within reasonable time, and a situation with 5 or more HDVs occurred rarely in our simulation, we chose to evaluate the heuristic for this number. We place 4 HDVs at a random node s in the Germany network (seen in Fig. 4c) 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 1,000 times for 4 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 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 over 90% of the experiments, our proposed heuristic computes the path with optimal savings.

III. SIMULATION OF THE GERMAN ROAD NETWORK

The following is a variety of results showing 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 a simplified graph of the German autobahn network with 647 nodes, 695 edges, and 12 destinations in Fig. 4c. 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 only speed up to catch another vehicle 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 towards 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 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

²MATLAB code to recreate all simulation data, and duplicate all figures can be found here: people.kth.se/~jeffrey1/Platooning/

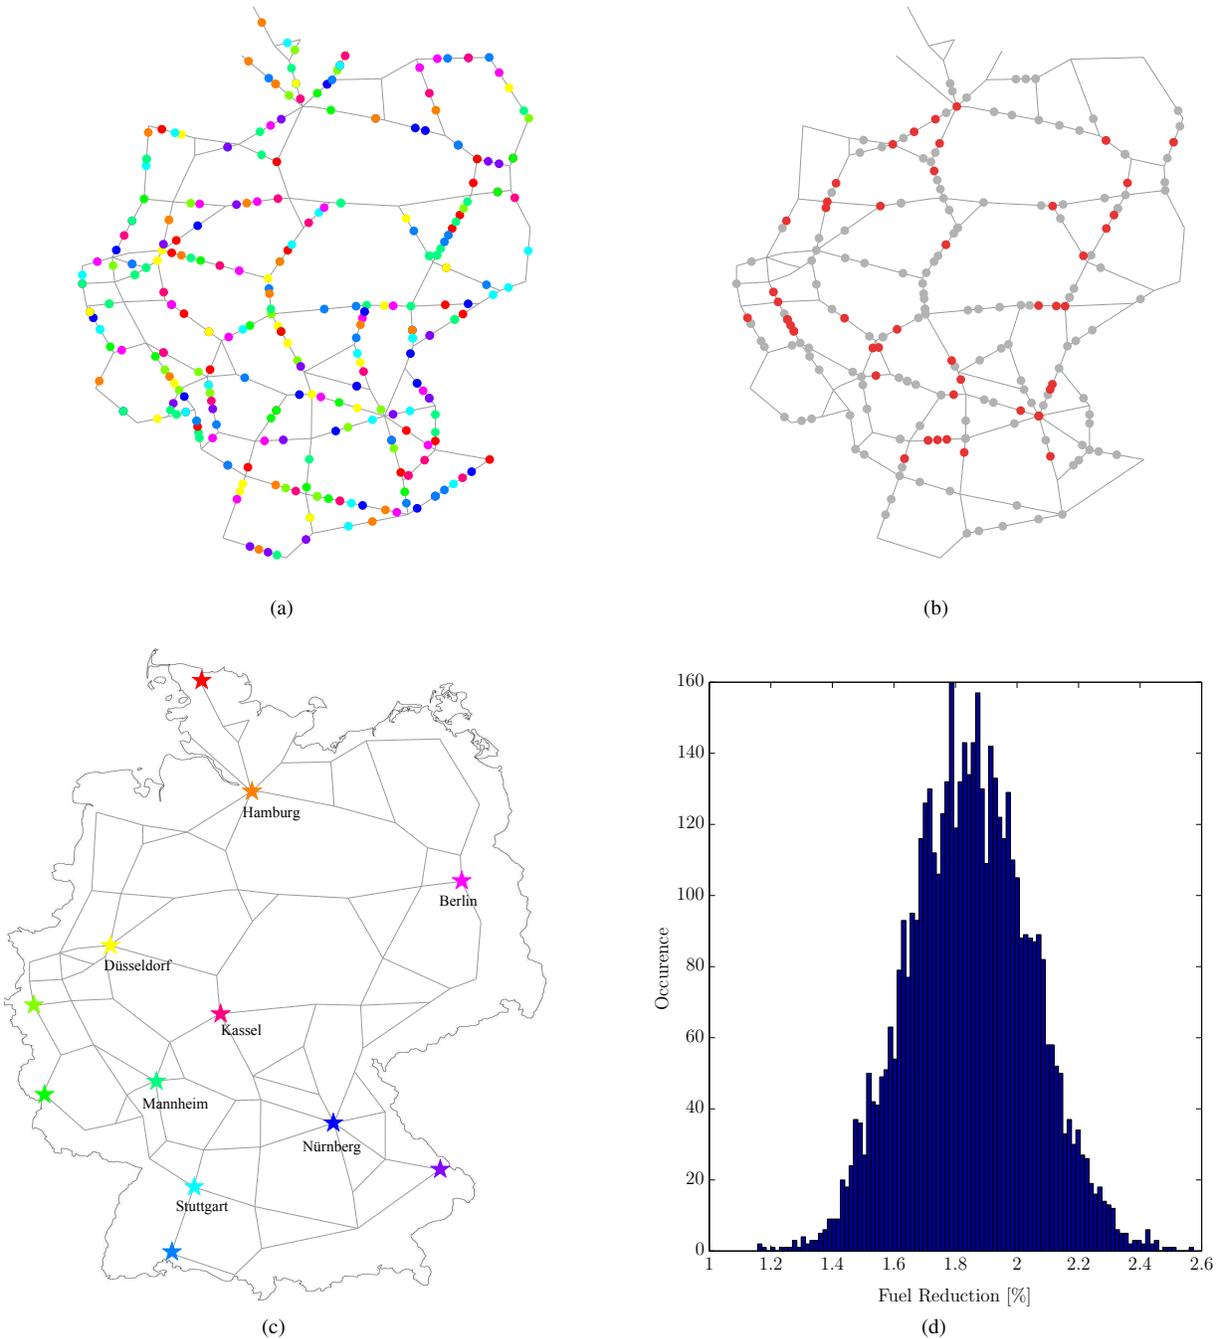


Fig. 4. Large-scale simulation results: (a) A snapshot of an initial state of our system, with dots representing HDVs. The color of an HDV designates its destination according to Fig. 4c. (b) A snapshot of the same system as in Fig. 4a after 10 time steps. Platoons are depicted as red dots, single vehicles 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 to every HDV taking its shortest path, for 300 HDVs, repeated 5,000 times with random starting points and destination.

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. 4a, with HDVs represented by dots, where the color matches the destination's color in Fig. 4c. After starting the simulation using the initial state shown in Fig. 4a, we pause after 10 time steps and observe the network in Fig. 4b. We can see the number of platooning vehicles in red and solo HDVs in gray.

Approximately 30% of the vehicles have formed platoon at this stage.

We repeat the simulation 5,000 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. 4d we see a histogram of the percentage of total fuel saved by our approach in each simulation (compared to every HDV taking its shortest path).

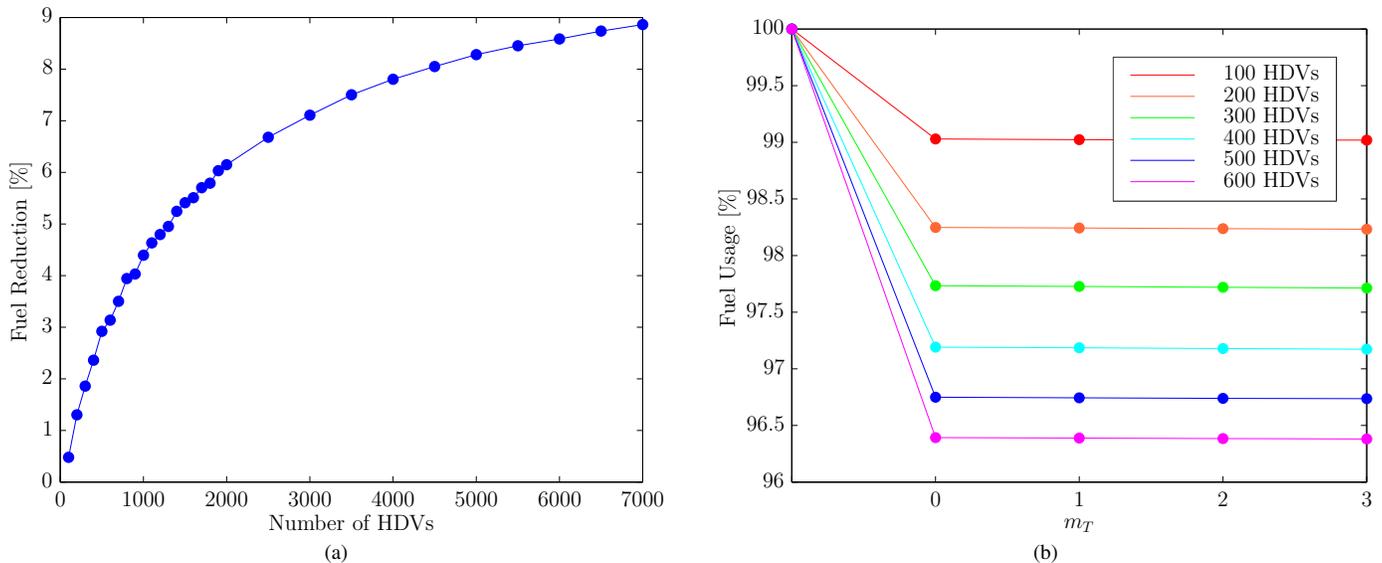


Fig. 5. (a) The percentage of fuel saved by our method compared to every HDV taking its shortest path and not platooning. Each data point is an average of the savings from 5 random instances. (The point (300, 1.82) mirrors the results in Fig. 4d.) (b) The number of additional edges an HDV can travel 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.)

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

In this section 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 so 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 this intuition is true (at least for the autobahn) in Fig. 5a. Average fuel savings increase rapidly between 0 and 2,000 vehicles. As the network becomes “saturated” with vehicles, nearly every edge can be traveled in a platoon, so nearly every HDV uses 10% less fuel (compared to the fuel use when driving its shortest path alone), and adding more HDVs will only result in marginal savings.

C. Increasing Allowable Detours

All previous results assume vehicle routes are 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 its start to destination.

Intuitively, adding extra edges a vehicle could travel would have quickly diminishing returns. Allowing an HDV to travel 10 or 20 km extra will help to improve the average fuel

use because more platooning options will be available. But 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 in order to offset the costs of platoon formation.

We find this intuition holds in Fig. 5b, where we partially resimulate our German road network with a slight modification to the 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 which 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 is seen in Fig. 5b. 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. 5b is then the average of 5 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 which 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 which does not significantly modify routes they are already traveling.

In another light, it might seem surprising that allowing every

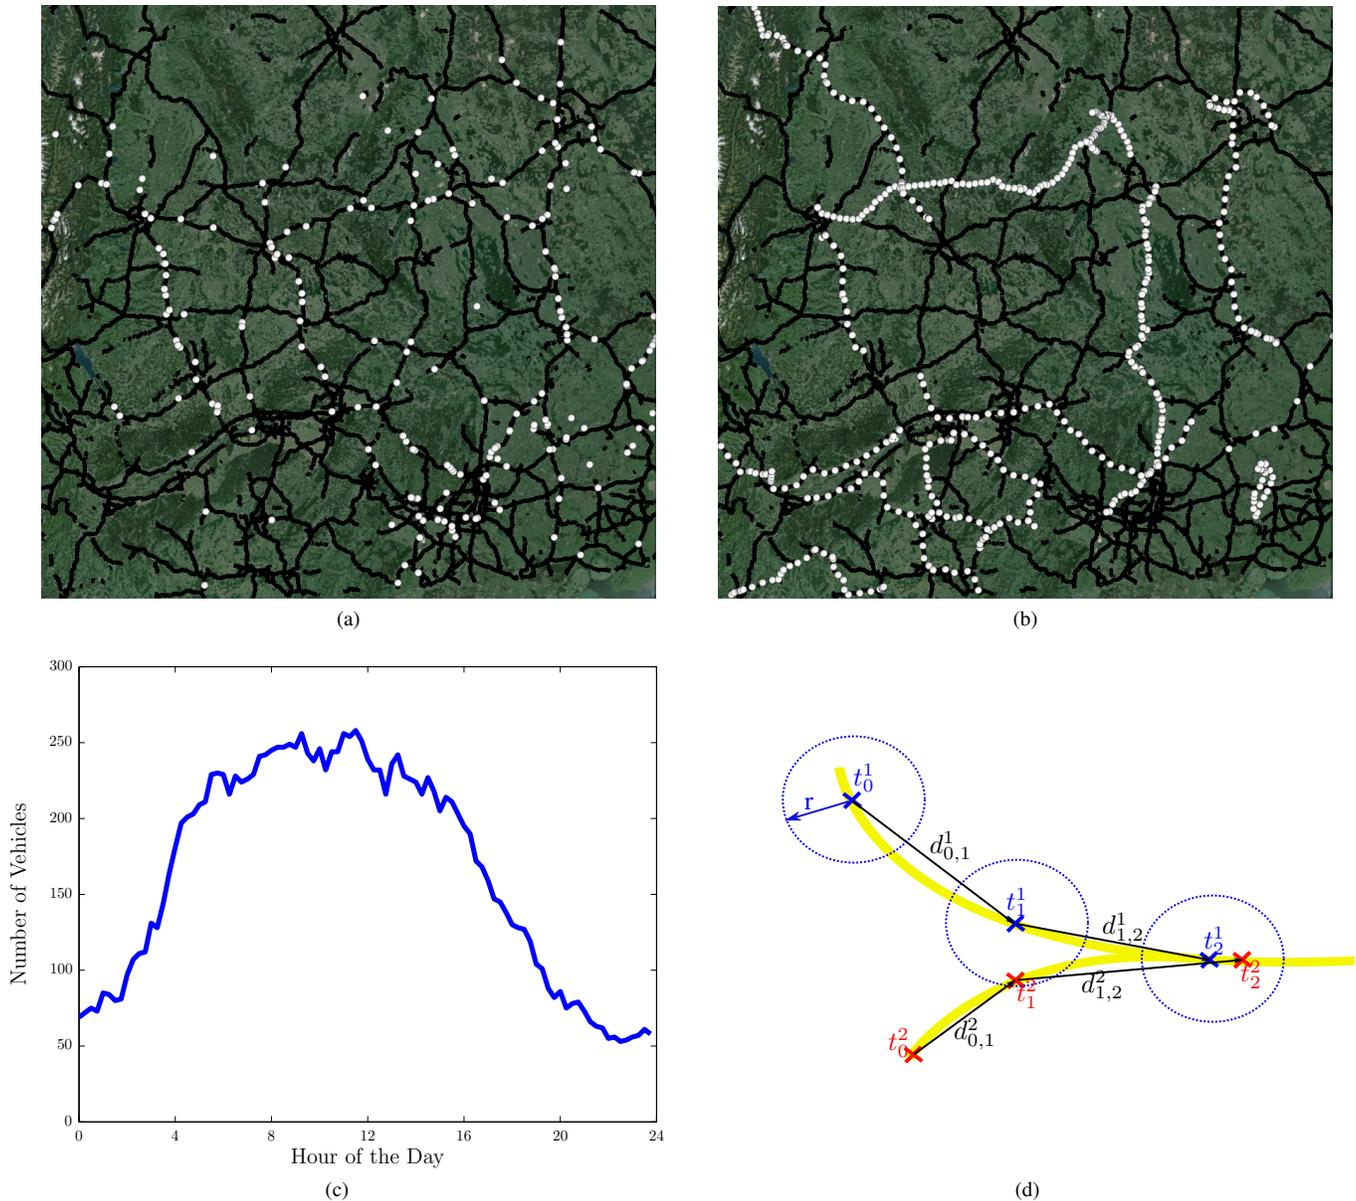


Fig. 6. Real-time HDV data study: (a) A snapshot of the probe data. Black lines are the roads and white dots show each vehicle's location. (b) A few selected vehicles' probe data trajectories. (c) Total number of vehicles moving at a given point throughout the 24-hour period. (d) An example of the probe data with a radius r used to determine if another vehicle is nearby.

HDV to possibly detour 10 km does not allow for a relatively larger savings than when no detour is allowed. Though 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.

The impact of increasing m_T naturally 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 where m_T is an increasing function of $\mathcal{T}(s_T, d_T)$ (HDVs with longer travel times can tolerate more detours), but find the fuel savings to be nearly identical to the constant m_T case.

IV. SUPPORTING ANALYSIS OF HDV DATA

In this section, we show preliminary results that we believe substantiate the results obtained from our simulations. In order to understand if our approach is feasible, we collected position data from HDVs throughout a single day. Our data shows 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 comes from vehicle probe data (commonly called floating car data) from a collection of HDVs. The probe data consists of the location and time of a vehicle and was collected by the in-cab GPS. Since the data is somewhat sparse, the conclusions that can be drawn are limited, but indicate some

potential. Nevertheless, we believe the data supports the result of our simulations.

A. Data Description

The probe data was obtained from HDVs from a single day in a region in Europe. The probe data comes from Scania HDVs within a 500,000 km² area over a 24-hour period in the spring of 2013. For the 7,634 HDVs in the region, including both long-haulage and local-distribution vehicles, the following information was collected at 5-10 minute intervals using the on-board GPS: vehicle ID, time stamp, longitude, and latitude. A snapshot of the data at one point during the day can be seen in Fig. 6a and some vehicles' trajectories can be seen in Fig. 6b.

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). 875 of the 7,634 vehicles satisfy the 150 km threshold. We can see how many of these 875 vehicles are moving at any point in time in Fig. 6c.

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. This was done by first interpolating a line between each vehicle's latitude and longitude at each data point, and then adding data points every 5 minutes along this line. To determine if a platooning opportunity exists for a given vehicle, we look to see if there is a data point for any other vehicle within a radius r . If there is another vehicle with a 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. 6d. As with the rest of the paper, we assume this platooning allows one vehicle to reduce its fuel cost by 10% on the stretch of road in question. We assume that the fuel cost is proportional to the distance driven for simplicity, hence the fuel cost will be 1 unit/km when driving alone and 0.9 units/km when platooning. If n vehicles are found to be within a radius r of each other between two points in time, then $n - 1$ vehicles will obtain the reduced fuel cost.

Admittedly, this approach can introduce errors for a variety of reasons. It is possible that the two vehicles are driving on different roads for the period of time in question. Also, it is possible to miss platooning opportunities if the road curvature and asynchronous timing prevents 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 km 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 platooning naturally today. HDVs that are within a radius of 1 km are likely not platooning

and the driver is unlikely to know that another HDV is close by, but a slight amount of coordination would allow for these vehicles to platoon. Lastly, 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 $r = 1 \text{ km}$ and $r = 5 \text{ km}$, 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 the radius r . The results of the study are as follows:

- $r = 0.2 \text{ km}$
 - 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.
- $r = 1 \text{ km}$
 - 241 out of 875 vehicles platooned at least once during their daily route.
 - 0.38% of total fuel saved by the platooned vehicles.
 - 4,369 km platooning.
- $r = 5 \text{ km}$
 - 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. 7a, 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. 7b shows the percentage of fuel that would have been saved by vehicles which had at least one opportunity to platoon at some point during the day. Note that since Fig. 6c 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 from Fig. 5a for a similar number of vehicles.

C. Interpretation of Results

This case study is brief, but provides an indication of how many platooning opportunities exist in reality, and how much effort might be required to coordinate their formation. As we can see from the $r = 0.2 \text{ km}$ 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 days travel when $r = 5 \text{ km}$.

V. CONCLUSION

In this paper we 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 together. This can be accomplished by

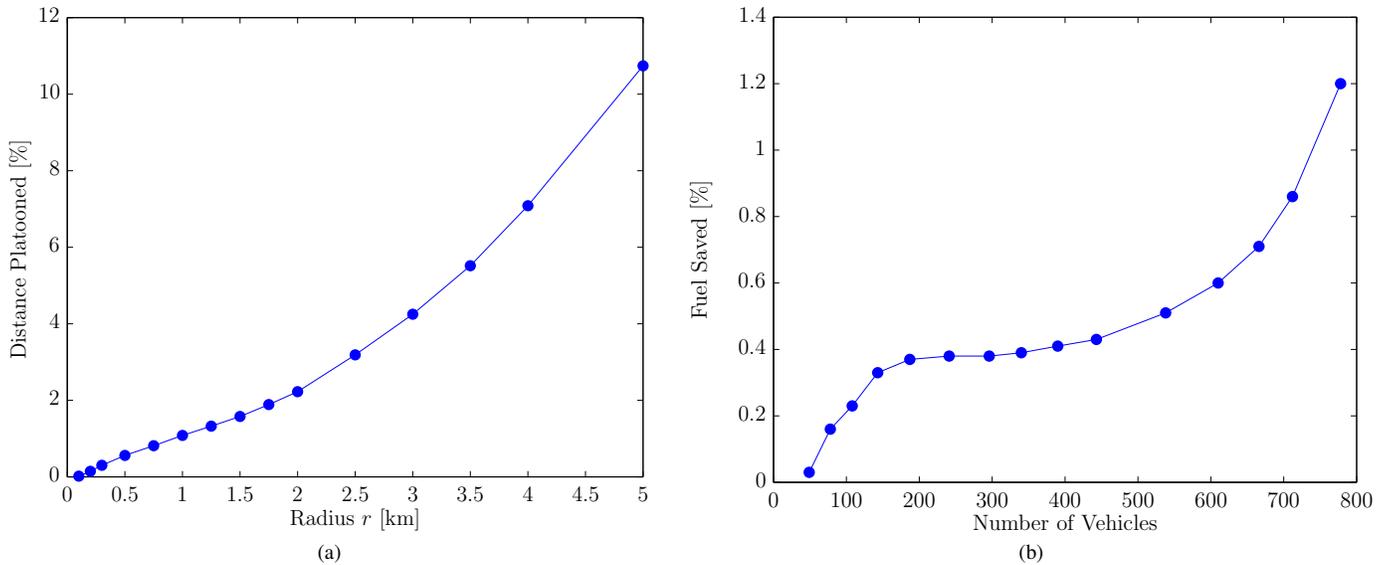


Fig. 7. (a) Percentage of total distance (403,413 km) that the 875 vehicles have platooned for a given radius. (b) Percentage of fuel saved by the vehicles that platooned at least once.

keeping vehicles on their shortest path from start to destination, but slightly adjusting their speeds in order to synchronize travel with other HDVs. These results assume a static network where 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 due to capacity overloaded roads. Lastly, we 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 experimental evaluation of the approach proposed in this paper. For example, HDV platooning experiments are currently being performed with 25 HDV's traveling regularly between Södertälje in Sweden and Zwolle in the Netherlands.

ACKNOWLEDGMENT

The authors would like to thank Christoph Kammer for the German Autobahn road network simulation code. This work was supported by the U.S. Department of Energy, Office of Science, under Contract DE-AC02-06CH11357.

REFERENCES

- [1] Eurostat, "Energy balance sheets," Tech. Rep., 2013.
- [2] European Commission, "Road Freight Transport Vademecum," Tech. Rep. March, 2009.
- [3] W. Levine and M. Athans, "On the Optimal Error Regulation of a String of Moving Vehicles," *IEEE Transactions on Automatic Control*, vol. 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 Transactions on Vehicular Technology*, vol. 46, no. 2, pp. 528–536, May 1997.
- [6] A. Broggi, M. Bertozzi, A. Fascioli, C. Guarino Lo Bianco, and A. Piazzini, "Visual Perception of Obstacles and Vehicles for Platooning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 1, no. 3, pp. 164–176, 2000.
- [7] F. Michaud, P. Lepage, P. Frenette, D. Letourneau, and N. Gaubert, "Coordinated Maneuvering of Automated Vehicles in Platoons," *IEEE Transactions on Intelligent Transportation Systems*, 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, 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 *Proceedings of the 17th ITS World Congress*, Busan, Korea, 2010.
- [10] J. Mårtensson, A. Alam, S. Behere, M. A. A. Khan, J. Kjellberg, K.-Y. Liang, H. Pettersson, and D. Sundman, "The Development of a Cooperative Heavy-Duty Vehicle for the GCDC 2011: Team Scoop," *IEEE Transactions on Intelligent Transportation Systems*, 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," *Intelligent Vehicle Technology*, no. 724, 2000.
- [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*, ser. Lecture Notes in Computer Science, M. Xie, Y. Xiong, C. Xiong, H. Liu, and Z. Hu, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, vol. 5928, pp. 135–146.
- [13] S. Tsugawa, "An Overview on an Automated Truck Platoon within the Energy ITS Project," in *7th IFAC Symposium on Advances in Automotive Control*, Tokyo, Japan, Sep. 2013, pp. 41–46.
- [14] R. Horowitz and P. Varaiya, "Control design of an automated highway system," *Proceedings of the IEEE*, vol. 88, no. 7, pp. 913–925, Jul. 2000.
- [15] E. Larsson, G. Sennton, and J. Larson, "Computational Complexity and Heuristic Solvers for Maximizing Platooning Fuel Savings," *Preprint*, 2013. [Online]. Available: http://www.optimization-online.org/DB_HTML/2013/09/4030.html
- [16] European Automobile Manufacturers Association, "European Motor Vehicle PARC 2010," Tech. Rep. April, 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 *7th IFAC Symposium on Advances in Automotive Control*, Tokyo, Japan, Sep. 2013, pp. 738–743.
- [18] E. Ericsson, H. Larsson, and K. Brundell-Frej, "Optimizing route choice for lowest fuel consumption – Potential effects of a new driver support tool," *Transportation Research Part C: Emerging Technologies*, 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 *15th World Congress on Intelligent Transport Systems*, 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 *Proceedings of the International Conference on Heavy Vehicles*, Paris, France, 2008.
- [21] L. D. Baskar, B. De Schutter, and H. Hellendoorn, "Hierarchical Traffic Control and Management with Intelligent Vehicles," in *IEEE Intelligent Vehicles Symposium*. Istanbul, Turkey: IEEE, 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 *IEEE Conference on Intelligent Transportation Systems*. St. Louis, MO, USA: IEEE, Oct. 2009, pp. 1–6.

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.

Jeffrey Larson received a Ph.D. in applied mathematics from the University of Colorado Denver. He then served as a Postdoctoral Researcher with the Automatic Control Laboratory and the ACCESS Linnaeus Centre in the School of Electrical Engineering at KTH Royal Institute of Technology in Stockholm, Sweden. He is currently a Postdoctoral Researcher with the Mathematics and Computer Science Division at Argonne National Laboratory in Argonne, IL.

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

Karl H. Johansson Karl H. Johansson is Director of the KTH ACCESS Linnaeus Centre and Professor at the School of Electrical Engineering, Royal Institute of Technology, Sweden. He is a Wallenberg Scholar and Director of the Stockholm Strategic Research Area ICT The Next Generation. He received MSc and PhD degrees in Electrical Engineering from Lund University. He has held visiting positions at UC Berkeley (1998-2000) and California Institute of Technology (2006-2007). His research interests are in networked control systems, hybrid and embedded system, and applications in smart transportation, energy, and automation systems. He has been a member of the IEEE Control Systems Society Board of Governors and the Chair of the IFAC Technical Committee on Networked Systems. He has been on the Editorial Boards 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 the *European Journal of Control*. He has received many awards, including the triennial Young Author Prize from IFAC. He is a Fellow of the IEEE.