Title
Longitudinal Control of Commercial Heavy Vehicles: Experimental Implementation

Authors
Tan, Yaolong
Kanellakopoulos, Ioannis

Publication Date
2002-08-01

CALIFORNIA PATH PROGRAM
INSTITUTE OF TRANSPORTATION STUDIES
UNIVERSITY OF CALIFORNIA, BERKELEY

Longitudinal Control of Commercial Heavy Vehicles: Experimental Implementation

Yaolong Tan
Ioannis Kanellakopoulos

California PATH Research Report
UCB-ITS-PRR-2002-25

This work was performed as part of the California PATH Program of the University of California, in cooperation with the State of California Business, Transportation, and Housing Agency, Department of Transportation; and the United States Department of Transportation, Federal Highway Administration.

Final Report for MOU 314

August 2002

ISSN 1055-1425

Longitudinal Control of Commercial Heavy Vehicles: Experimental Implementation

Final Report for MOU 314

Yaolong Tan
Ioannis Kanellakopoulos

Department of Electrical Engineering
University of California, Los Angeles
Los Angeles, CA 90095-1594

California PATH Research Report

This work was performed as part of the California PATH Program of the University of California, in cooperation with the State of California Business, Transportation, and Housing Agency, Department of Transportation; and the United States Department of Transportation, Federal Highway Administration.

April 2000

Longitudinal Control of Commercial Heavy Vehicles: Experimental Implementation

Yaolong Tan and Ioannis Kanellakopoulos
April 2000

Summary

The report describes the results of the project funded under MOU 314.

The main result is the experimental implementation of various longitudinal control algorithms that were developed under MOU124 and MOU 240. In order to conduct these experiments, in collaboration with Professor Tomizuka's group and PATH Personnel, we first outfitted a Class-8 tractor-trailer commercial heavy-duty vehicle on loan from Freightliner Corporation with the necessary sensors and actuators for fully automated operation. These sensors and actuators included the vehicle speed sensors, air-brake pressure sensors, air-brake actuators, and the throttle actuator. Then, a series of open-loop and closed-loop experiments were carried out at Crow's Landing test field. In the open-loop experiments, we identified and collected important parameters for vehicle dynamics such as the working ranges for the brake and fuel actuators, the vehicle speed signals, and the air-brake pressure signals. These parameters enabled us to evaluate the vehicle dynamics and adjust the control algorithms accordingly, tuning their parameters off-line in preparation for the closed-loop experiments. In addition, because of the noise levels present in the sensor data, we designed low-pass filters to smooth out the speed signal and the brake/fuel command signal.

In the first stage of the closed-loop experiments, we focused on speed control. We implemented two speed control algorithms: PID and PIQ. These two algorithms were discretized based on the sampling rate and coded as longitudinal control modules on an on-board computer to regulate the throttle and brake actuators in order to track the desired speed profile. Several speed tracking closed-loop experiments were performed with different profiles, some of which were very demanding. Then we moved to the virtual vehicle following stage, where an imaginary vehicle is assumed to lead the experimental vehicle. This enabled us to compute the distance between the leading vehicle and the experimental vehicle without actually putting a vehicle in front of the experimental vehicle. The availability of the distance data enabled us to test the vehicle following controllers, which use both the relative velocity and the relative distance. Through the virtual vehicle following closed-loop experiments, we successfully validated the concepts of the nonlinear spacing policies: variable time headway and variable separation error gain.

We proved experimentally that these nonlinear spacing policies can significantly improve the system response and performance without adding significant complexity to the controller design; this had been only validated through simulation results in previous projects. Our closed-loop experiments for vehicle following also proved that simple PID and PIQ controllers are able to achieve fairly good performance when combined with nonlinear spacing policies.

1. Introduction

Very often, the benefits of vehicle and highway automation are almost exclusively considered from the point of view of the average motorist commuting to and from work on a daily basis. In the last few years, however, the ITS community has been paying increasing attention to the automation of heavy-duty vehicles. The primary driving factor in this trend is that transportation authorities and equipment manufacturers are becoming more aware of the potential economic benefits of partial of full automation of commercial heavy-duty vehicles at all geographic levels (regional, state, interstate, national, and international), at all user levels (individual owner/operator, truck/bus fleets, national associations).

If one might wonder why it is important to separately study a class of vehicles like commercial heavy-duty trucks, which accounts for only 1% of the total number of registered vehicles in the United States, let us note that the average commercial heavy-duty truck travels **six (6) times more miles** and consumes **twenty-seven (27) times more fuel** than the average passenger car. From the point of view of the operator, this makes automation, with its potential fuel savings, a very attractive option, since the noticeable increase in operating profits would significantly shorten the time to recovery of the initial cost. Furthermore, the truck industry places a huge premium on anything that has to do with making the job of the driver safer or more comfortable; truck fleets advertise the latest technological features of their trucks to attract better drivers, and depend on them to reduce their accident rates and hence their insurance premiums; this in turn makes the truck manufacturers more than willing to include in their vehicles any new technology that can offer a competitive advantage to their customers. And in Europe, where traffic regulations on trucks are so strict that in many countries they are banned from using the highways on weekends during the summer, any technology that could convince the public and consequently the policymakers to view truck traffic as less threatening is eagerly sought after.

As a result, in the last few years there has been significant research and development on the problem of automating commercial truck traffic. For example, the Combi-Road project in the Netherlands is aimed at automating the transport of cargo containers to and from the Rotterdam harbor, by loading them onto driverless vehicles that travel under automatic control on a separate corridor running along the highway. The Chauffeur project, carried out by a consortium led by DaimlerChrysler and Fiat and funded by the European Union, successfully demonstrated the *electronic towbar* concept, in which two or more commercial heavy-duty vehicles are connected into an "electronic train", with the lead vehicle driven manually and the following ones follow without human intervention. The Chauffeur final demonstration in June 1999 included two such "electronic trains", each comprising two vehicles.

The concept of using completely driverless vehicles is irrelevant for passenger cars, but it is very important in commercial vehicle operations, where driver-associated costs account for 50% of the total operational cost. The potential utility and cost-effectiveness of this concept is considerably enhanced by the fact that commercial traffic follows much more regular patterns than passenger car traffic, and that most of that traffic is conducted on highways; in other words, trucks usually travel on well-established commercial routes, mostly between major cities, and they use highways to get there. In cases where it is possible to build special-purpose roads for the exclusive use of automated truck traffic, as in the Combi-Road concept, the control task is greatly simplified. But even in cases where this is not possible, the electronic towbar concept

can still be used effectively, since fleet operators can easily compose platoons of trucks with the same origination and destination points which will travel together for the entire trip. One version of this scenario assumes the existence of departure/arrival stations in major cities, which will be operated by transit authorities but will be utilized by freight transport companies. Individual trucks will be driven manually from all over the city to the station, where they will join the platoon departing for their destination city. Drivers not needed for the departing platoons that will manually drive the trucks from arriving platoons to their individual points destination points within the city. The absence of the requirement for split/merge maneuvers and the fact that these routes will involve only highway travel are factors that significantly simplify the control problem for such platoons and make this concept feasible even in the absence of automation infrastructure on the highway.

These potential economic benefits explain why government and industry are becoming increasingly interested in heavy vehicle automation, but they do not justify why this area is quickly gaining a prominent position in the ITS efforts of many research organizations. This should be attributed to the significant difficulties associated with the implementation of heavy vehicle automation, which are due to factors that are not present in passenger cars: the much larger size and mass of these vehicles, as well as the much more significant variations in their configuration and operating conditions present formidable challenges to the control designer. In particular, the control scheme has to be able to cope with significant fuel, brake, and steering actuator delays, very low power-to-mass ratio, high sensitivity to wind gusts and changes in the road grade, strong coupling between lateral and longitudinal dynamics, drastic changes in the vehicle mass (as much as 400% depending on the cargo load), and tires that are optimized for load-carrying capacity and long tread life rather than traction.

Recognizing the importance of commercial vehicle automation and the difficulties associated with its successful implementation, for the last several years the California PATH Program has been providing continuous support research on heavy-duty vehicles to the groups of Professor Tomizuka at UC Berkeley and Professor Kanellakopoulos at UCLA. The lateral modeling and control effort is carried out by Professor Tomizuka's group at UC Berkeley, while the longitudinal modeling and control is carried out by our group at UCLA. Our group has developed suitable models and used them to design several novel nonlinear control schemes for longitudinal control of CHVs, which significantly outperform schemes developed for passenger cars and applied to heavy-duty vehicles. We have also combined our models with the lateral models developed by the Berkeley team, to produce complete models of single-unit and articulated heavy vehicles, which are currently used in the design of integrated lateral/longitudinal control schemes. Furthermore, we have developed a software package that allows the visualization of automated trucks and is used to evaluate both lateral and longitudinal controllers.

The next step in this process, and the goal of the proposed research, was the experimental evaluation of our control schemes, which is necessary for the eventual implementation and deployment of AVCS in commercial traffic. Longitudinal control of CHVs is generally viewed as a very challenging problem due to difficulties that arises in its implementation. Hence, our approach to this project was structured in a fashion that allowed us to gradually identify, understand, and overcome the obstacles that arose in the process.

2. Hardware and Software Configuration

Throughout the experimental phase, our efforts were closely coordinated with those of Professor Tomizuka's group at UC Berkeley, who were working on a parallel track on lateral control issues. This coordination produced integrated lateral/longitudinal control models for Commercial Heavy Vehicles, a joint chapter in an edited volume on Automated Highway Systems, and software for dynamic simulation and visualization of CHVs. But the first one of the most important results of this coordination was achieved in June 1997, when the combined multi-year efforts of Professors Tomizuka and Kanellakopoulos resulted in Freightliner Corporation agreeing to loan a Class-8 trailer to PATH free of charge for use in experimental validation of our research in commercial vehicle automation. The tractor physically arrived at PATH in September 1997; this important development was followed by a flurry of very closely coordinated activity by both research groups.

The main concern was to find suitable vendors and partners who could provide all the necessary ingredients for a fully operational automated tractor-semitrailer experimental vehicle. All the necessary components were either donated or sold to us at deeply discounted prices by vendors who were interested in our research and willing to help us. In particular, NSK in Japan donated the steering actuator, Oregon Western Trailer Sales provided a used Great Dane 45-ft semitrailer at cost, ISE Research in San Diego built and installed the electric-over-hydraulic brake actuation system also at cost, and Valley Detroit Diesel donated the manuals for the engine and the transmission and provided instruction on how to use the built-in electronic throttle circuits.

During the next year, the PATH staff supporting this experimental effort and the research groups of Professors Tomizuka and Kanellakopoulos worked closely together, providing specifications for components, negotiating with vendors, installing and testing hardware, writing and debugging software, and fixing problems that arose with the tractor and the trailer. Finally, in September 1998, we had the experimental vehicle up and running. The resulting experimental setup has several components that are used exclusively for lateral control; they are described in detail in the companion report of MOU 313 submitted by Professor Tomizuka. As far as longitudinal control is concerned, the major parts of the hardware configuration are the fuel and brake actuators and the on-board computer that controls all the signals to and from the actuators:

- The on-board computer uses an Intel Pentium 100MHz processor running the QNX real-time operating system. It connects with the sensors through a PC-TIO-10 interface to read measurements of the speed of each individual wheel and engine speed, and an AT-MIO-64E-3 card to read measurements for the acceleration and the yaw rate.
- Electronic fuel actuator: The experimental truck is powered by a Detroit Diesel DDEC III engine, which has an electronic fuel actuator as original equipment; the gas pedal is connected to a potentiometer, which sends a voltage signal directly to the injection system. In contrast to passenger cars, where these actuators are rare and appear only in a few high-end luxury sedans, Class-8 trucks have had electronic throttle as standard equipment for over 10 years, because it helps reduce fuel consumption and engine wear. In order to utilize this actuator, we connected the circuit that reads the voltage output of the potentiometer directly to the voltage output of our D/A card that corresponds to the fuel command.

(a) Brake actuator

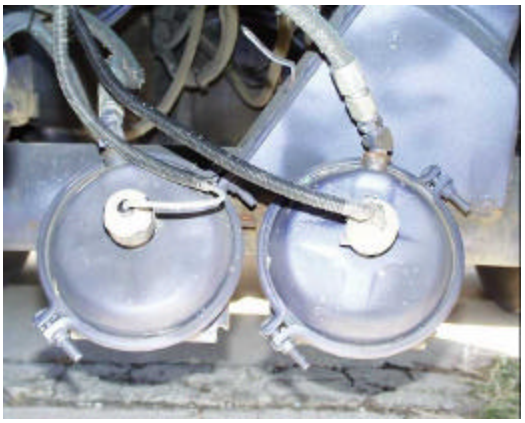

(b) Air pressure transducer

(c) Brake control circuit box

(d) Air reservoir

Figure 1: Electric-over-hydraulic brake actuation system components.

- Brake actuator: The air brake system of the vehicle was modified for electronic actuation by ISE Research Corporation. The resulting system uses four (4) proportional actuators (one for the front tractor brakes, one for the rear tractor brakes, one for the left trailer brakes, and one for the right trailer brakes) and one (1) on/off actuator (between the front tractor brakes and the trailer brakes, which are activated by the same signal), as well as ten (10) pressure transducers, one for each of the ten brakes. There are also two master-off switches for the whole system: one on the dash and one connected to the brake pedal so that all the brake automation shuts off if the driver steps on the brake pedal. The brake signal is transmitted by wires from the on-board computer's I/O ports to each proportional actuator. When the brake actuators are activated, they open the relay valves to suck the air from the air tank to the spring brake until the air pressure in the spring brake balances the commanded air pressure from the brake actuator. Figure 1 shows illustrative pictures of the brake actuator, air pressure transducer, brake control circuit box, and air reservoir.

- Transmission: Implementing automated longitudinal control on a vehicle with a manual transmission is a very challenging problem, especially if the vehicle is a Class-8 truck with a 13- or 20-speed gearbox. Even though automatic transmissions are not nearly as commonplace in heavy trucks as they are in passenger cars, we were very fortunate because Freightliner had already retrofitted the experimental truck they loaned to us with a Detroit Allison 6-speed automatic Heavy-Duty World Transmission (HD-WT) with a transmission retarder. The function of this retarder is to slow down the vehicle by converting kinetic energy into heat: when the dash-mounted retarder switch is turned on, the transmission shaft drives an auxiliary pump which circulates the transmission fluid in the reverse direction, thus

pumping it against the vanes of the hydraulic torque converter; this causes the transmission fluid to heat up, and therefore it is then circulated through a cooling unit, whose compressor is also driven by the engine. Thus, kinetic energy is used to drive by the pump, the fluid, and the compressor, resulting in the equivalent of nearly 700hp of braking power (according to the transmission mechanic of Valley Detroit Diesel). All of the kinetic energy absorbed by the transmission retarder is dissipated as heat to the environment through the engine cooling air exchange. Our longitudinal control algorithms did not utilize any information on the shifting schedule of this transmission, nor did they make use of the transmission retarder.

- Engine brake: The engine brake is turned on and off by a switch on the dash of the tractor cabin, but it can be switched on only when there is no fuel command to the engine, that is, when the voltage input to the fuel circuit is below the 0.5V threshold. When the engine brake is on, no fuel is injected into the cylinders, and the timing of the intake and exhaust valves of the engine is altered: during the air intake phase the intake valves are opened, to allow the cylinders to fill up with fresh air; then they are closed, to allow the piston to compress the air; at the top of the compression phase, the exhaust valves are opened to let all the compressed air escape, and then the cycle is repeated again. As a result, the engine operates essentially like a 2-stroke air compressor, extracting kinetic energy for the compression work performed by the pistons, and transforming it into heat that is released through the compressed air escaping from the exhaust valves. In our experimental vehicle, this provides about 300hp of braking power (also according to Valley Detroit Diesel), but does so in a very abrupt and non-progressive (either full-on or full-off) manner.

Finally, the real-time C software that runs the whole system was developed by software engineers in the California PATH Program; our research group coded the longitudinal control algorithms as subroutines of the larger program, which also includes lateral control.

3. Controller Design for Vehicle Following

3.1 Vehicle Longitudinal Control Model

The vehicle longitudinal control model in general is highly nonlinear due to the effect of the diesel engine with the turbocharger and intercooler, the automatic transmission with the torque converter, and the air brake dynamics including the delays. These complex nonlinear dynamics need to be captured in the simulation model that is used to design and pre-test the control algorithms. However, the resulting vehicle longitudinal dynamic model is far too complex to be used as the basis for control design. Therefore, for control design we used a simplified version of the full longitudinal dynamic model, which retains only the most important dynamic characteristics that significantly affect the vehicle speed; this allowed us to design controllers with moderate complexity. In particular, we used several linear models, each linearized around a different operating point determined by the fuel command and vehicle mass combinations. Each of these models is described by the equation:

$$\dot{v}_f = a(v_r + kd) + bu + d$$

where v_f is the velocity of the following vehicle, $v_r = v_l - v_f$ is the relative velocity between the leading vehicle and the following vehicle, $d = x_r - s_d$ is the separation error between the two vehicles (x_r is the distance between the leading and following vehicle, s_d is the desired vehicle separation), k is the separation error gain, u is the fuel/brake command, d incorporates external disturbances and modeling errors as well as the unknown nominal value of the control, and a and b are vehicle-specific parameters, whose values are dependent on the fuel command and the vehicle mass.

This reduced-order linearized model is the starting point for the design of our control schemes and it implies that our controllers do not explicitly rely on the particular details of the vehicle model and are thus inherently robust to modeling uncertainties. The controllers are also easier to experimentally implement and debug, because they do not contain highly complicated algebraic or dynamic equations. The disadvantage of simplifying the longitudinal control model is that we can only achieve good performance with designed controllers if the required maneuvers are slower than the slowest neglected modes of the full nonlinear model. However, since in these experiments we were only interested in maneuvers with time constants of several seconds, this constraint did not represent a significant limitation, due to the fact that the dynamics of the automatic transmission with torque converter, air brakes, and drivetrain, all have sub-second time constants.

The control objective is to regulate both the relative velocity and the separation error to zero. It can be shown that when $v_r + kd \equiv 0$, both the relative velocity and the separation error are regulated: $v_r \to 0$ and $d \to 0$.

3.2 Controller Design and Implementation

3.2.1 Fixed-Gain PID Controller

Since our simplified longitudinal model is a first-order linear system with a large time constant in the vehicle dynamics with respect to the brake and fuel inputs, we used a fixed-gain PID controller as the starting point for our closed-loop experiments, because it is the simplest candidate scheme that we would expect to work in an experimental setting. The PID can achieve regulation of the relative velocity v_r to zero, while providing some robustness with respect to unmeasured disturbances. However, the "D" term of the PID controller requires measurement of the vehicle acceleration, for which we do not have a sensor in our current hardware configuration. Therefore, we approximate this measurement by filtering the (fairly noisy) error signal with the first-order filter $\frac{s}{st_d + 1}$, essentially implementing a "dirty derivative". Thus, the PID controller we used for the closed-loop experiments is given by the following expression:

$$u = k_p(v_r + kd) + k_i \frac{1}{s}(v_r + kd) + k_d \frac{s}{st_d + 1}(v_r + kd).$$

Since the fuel and brake commands are sent out through electronic output ports approximately every 20ms, the PID controller needs to be implemented as a discrete-time controller. Denoting the sampling time by T_s and the output of the derivative term by u_d, we have:

$$u_d = \frac{s}{st_d + 1}(v_r + kd) = \frac{1}{t_d}(v_r + kd) - \frac{1/t_d}{st_d + 1}(v_r + kd_d).$$

A discrete-time realization for u_d is thus given by

$$u_d[(n+1)T_s] = \frac{1}{t_d}(v_r[(n+1)T_s] + kd[(n+1)T]) - \frac{1}{t_d}u_{dd}[(n+1)T_s]$$

where u_{dd} is the discrete-time output of the realization of the first-order system $\frac{1}{st_d + 1}$:

$$u_{dd}[(n+1)T_s] = e^{-T_s/t_d}u_{dd}[nT_s] + t_d\left(1 - e^{-T_s/t_d}\right)(v_r[nT_s] + kd[nT_s])$$

Thus, we can implement the PID controller with dirty derivative through the following discrete-time formula:

$$u[nT_s] = k_p(v_r[nT_s] + kd[nT_s]) + k_i u_i[nT_s] + u_d[nT_s],$$

where u_i is the dynamics for the integral term and is given by

$$u_i[(n+1)T_s] = u_i[nT_s] + T_s(v_r[(n+1)T_s] + kd[(n+1)T_s]).$$

In the first stage of our experiments, we focused on speed control. Thus, the leading vehicle speed was replaced by a desired velocity signal, and the longitudinal model above was modified by setting $k = 0$ to eliminate the dependence on the (non-existent) separation error and using v_l as the commanded speed; v_r then becomes the velocity error that the controller needs to regulate to zero. In this case, the discrete-time PID controller can be further simplified as follows:

$$u[nT_s] = k_p v_r[nT_s] + k_i u_i[nT_s] + u_d[nT_s]$$

$$u_i[(n+1)T_s] = u_i[nT_s] + T_s v_r[(n+1)T_s]$$

$$u_d[(n+1)T_s] = \frac{1}{t_d}v_r[(n+1)T_s] - \frac{1}{t_d}u_{dd}[(n+1)T_s]$$

$$u_{dd}[(n+1)T_s] = e^{-T_s/t_d}u_{dd}[nT_s] + t_d\left(1 - e^{-T_s/t_d}\right)v_r[nT_s]$$

3.2.2 FIXED-GAIN PIQ CONTROLLER

The lower actuation-to-weight ratio of the commercial heavy-duty vehicles requires a controller that is more aggressive at larger tracking errors but does not have the undesirable side-effect of overshoot. This can be achieved by adding a signed quadratic term of the form

$(v_r + kd)|v_r + kd|$ to the PI controller, which thus becomes the so-called nonlinear PIQ controller:

$$u = k_p(v_r + kd) + k_i \frac{1}{s}(v_r + kd) + k_q(v_r + kd)|v_r + kd|.$$

The quadratic nonlinear term becomes very small when the tracking error $(v_r + kd)$ is small, but grows fast as the error grows. Therefore, it has the effect of generating aggressive control action to reduce large errors quickly, yet it turns itself off for small errors. A discrete-time realization for the PIQ controller is

$$u[nT_s] = k_p(v_r[nT_s] + kd[nT_s]) + k_i u_i[nT_s] + k_q(v_r[nT_s] + kd[nT_s])|(v_r[nT_s] + kd[nT_s])|,$$

where the integral term u_i is given by

$$u_i[(n+1)T_s] = u_i[nT_s] + T_s(v_r[(n+1)T_s] + kdv_r[(n+1)T_s]).$$

In the case of speed control experiment, we set the separation gain k to zero, which produces the following discrete-time speed control scheme:

$$u[nT_s] = k_p v_r[nT_s] + k_i u_i[nT_s] + k_q v_r[nT_s]|v_r[nT_s]|$$

$$u_i[(n+1)T_s] = u_i[nT_s] + T_s v_r[(n+1)T_s]$$

3.2.3 NONLINEAR SPACING POLICIES

Nonlinear spacing policies can be used to further improve performance. In our closed-loop experiments, we tested two nonlinear spacing policies: variable time headway h, i.e., a headway that varies with the relative speed v_r between adjacent vehicles, and variable separation error gain k.

The intuition behind the concept of the variable time headway is as follows: If the relative speed between the two vehicles is positive, that is, if the preceding vehicle is moving faster than the follower, then it is safe to reduce the time headway. On the other hand, if the preceding vehicle is moving slower then it would be advisable to increase the time headway to increase the distance. This leads to the following choice of time headway as a function of the relative velocity v_r:

$$h = sat(h_0 - c_h v_r) = \begin{cases} 1 & \text{if } h_0 - c_h v_r \geq 1 \\ h_0 - c_h v_r & \text{if } 0 < h_0 - c_h v_r \leq 1 \\ 0 & \text{otherwise.} \end{cases}$$

Another nonlinear spacing policy is to introduce a variable separation error gain k. Our control objective is to regulate the error signal $v_r + kd$, where d is the separation error. The intuition behind choosing such a control objective is as follows: If two adjacent vehicles are closer than desired $(d < 0)$ but the preceding vehicle's speed is larger than that of the follower $(v_r > 0)$, then the controller does not need to take drastic action. And the same is true of the vehicles are farther apart than desired $(d > 0)$ but the preceding vehicle's speed is lower than the follower's

($v_r < 0$). However, when the separation error gain k is constant, the controller would try to reduce a very large spacing error d through a very large relative velocity v_r of opposite sign. Therefore, if the follower falls far behind the preceding vehicle, the controller would react aggressively by accelerating to a very high speed. This behavior is not desirable, since it increases fuel consumption and can even lead to collisions in some cases. It is also counter-intuitive; it would be much more natural for the controller to accelerate to a speed somewhat higher than that of the preceding vehicle's, and therefore reduce the spacing error smoothly and progressively. To achieve such an intuitive response, the separation error gain needs to be reduced when d becomes large and positive, making sure that it remains above some reasonable positive lower bound. The following choice of k satisfies the above requirement:

$$k = c_k + (k_0 - c_k)e^{-sd^2},$$

where $0 < c_k < k_0$ and $s \geq 0$ are design constants.

4. Experiments

4.1 Open-loop Experiments

After the completion of the instrumentation phase, in September 1998 we performed open-loop longitudinal experiments at Richmond Field Station to determine the dynamic characteristics of the experimental vehicle, specifically those needed for the first set of closed-loop experiments, which involved vehicle speed control (no distance regulation from a preceding vehicle). The two most important characteristics for vehicle speed control are the dynamics of the fuel actuator and those of the brake actuators. With these experiments we were able to approximate the dynamics (transfer function) from a fuel or brake command to the vehicle speed. Giving a voltage signal as input and measuring the vehicle velocity as output, we measured the step response of the system. About 12 sets of experiments were conducted with the fuel signal ranging from 1.3V to 5.0V; Figure 2 shows a typical response curve.

At the beginning of this run, we release the brake pedal and wait for the tractor to reach its steady-state speed with the engine running at idle. Then we issue a fuel command of 2.0V through the computer output port to the engine control module. As can be seen in Figure 2, the dynamics from the fuel command to the vehicle speed can be locally approximated as a first-order system with appropriate choices of parameters. This validates our use of a first-order model as the basis for longitudinal control design in our theoretical work. Also, it is worth noting that nonlinear behavior occurs during gearshifts.

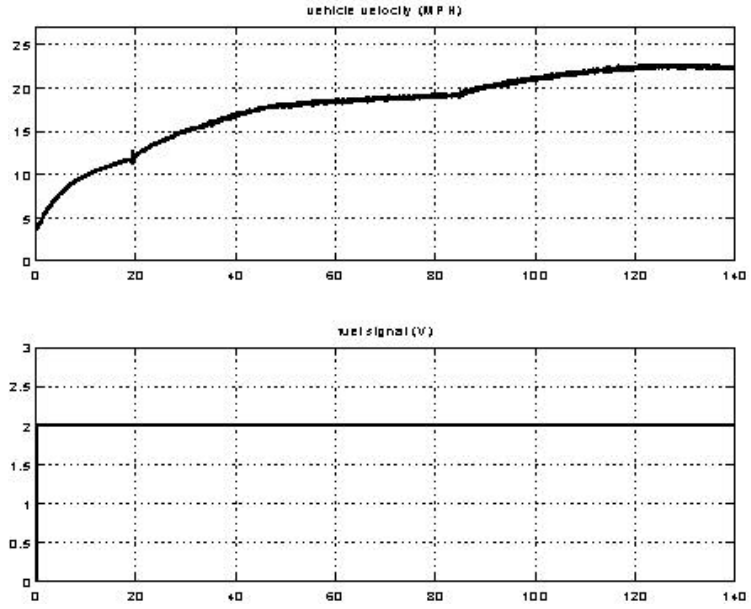

Figure 2: Vehicle velocity response to a step fuel command of 2.0V.

Similarly to the testing of the fuel actuator, we also used the step-input brake command to test the brake dynamics. A typical step input brake response is shown in Figure 3.

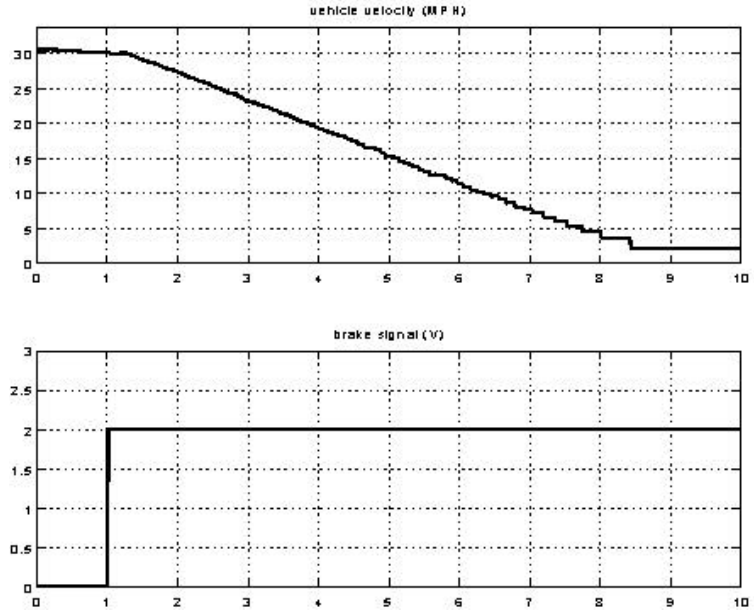

Figure 3: Vehicle velocity response to a step brake command of 2.0V.

At time t=0, we release both the brake and accelerator pedal and put the truck in automatic mode, which explains the small velocity drop at the beginning. At time t=1, a brake command of 2.0V is issued to all 4 brake actuators. Thus, the vehicle starts to decelerate until it stops. Because of the vehicle velocity sensor limitation, we cannot measure speeds lower than 1.1888mph or

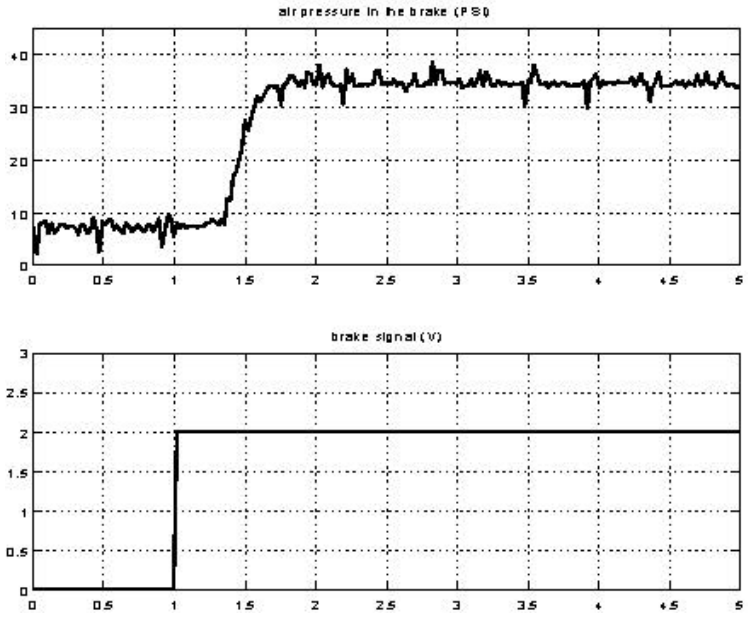

Figure 4: Air brake pressure in response to a step brake command of 2.0V.

0.5283m/s. That explains why the vehicle velocity appears not to reach zero in steady state, while in reality it does. The deceleration is constant, which means that the braking torque does not fluctuate during this braking maneuver. We also notice that the vehicle speed starts to drop abruptly about 0.3s after the brake command was issued. This means that despite the fact that in our modified brake system the brake signals are transmitted electronically and not by air traveling through the brake pipes, there is still a delay of 0.3s in the brake response. This delay is due to the fact that we are still using an air brake system with slow dynamics, where air pressure needs to build up before the brakes are activated. This is clear from Figure 4, which shows the behavior of the air pressure in the tractor brakes, as measured by the pressure transducers attached to the brakes. In Figure 4 we see that there is a delay of approximately 0.3s, followed by 0.5s during which the actual pressure builds up to its commanded value.

We should note that at this stage in our experiments we were still operating the longitudinal controller separately from the lateral controller. For that reason, we were not concerned with the effects of different brake torques on different wheels (differential braking), only with the total brake torque acting on the vehicle, and that is why we sent the same brake command to all brake actuators.

4.2 Closed-loop Experiments

4.2.1 SPEED CONTROL

In the first stage of our experiments, we focused on speed control with the two simplest control schemes we have designed, namely the fixed-gain PID and PIQ controllers described in Sections 3.2.1 and 3.2.2. As explained there, for the purposes of speed control the leading vehicle is replaced by a desired velocity signal, and the model above is modified by setting $k = 0$ to eliminate the dependence on the (non-existent) separation error, and using v_l as the commanded speed; v_r now becomes the velocity error that the controller needs to regulate to zero. Since our

simplified longitudinal model is a first-order linear system with a large time constant in the vehicle dynamics with respect to the brake and fuel inputs, we used a fixed-gain PID controller as the starting point for our closed-loop experiments, because it is the simplest candidate scheme that we would expect to work in an experimental setting. The PID can achieve regulation of the relative velocity v_r to zero, while providing some robustness with respect to unmeasured disturbances. However, the "D" term of the PID controller requires measurement of the vehicle acceleration, for which we do not have a sensor in our current hardware configuration. Therefore, we approximate this measurement by filtering the (fairly noisy) velocity signal with the first-order filter $\frac{s}{st_d + 1}$, essentially implementing a "dirty derivative". Thus, the PID controller we used for the closed-loop experiments is given by the following expression:

$$u = k_p v_r + k_i \frac{1}{s} v_r + k_d \frac{s}{st_d + 1} v_r .$$

Since the fuel and brake commands are sent out through electronic output ports approximately every 20ms, the PID controller needs to be implemented as a discrete-time controller. Denoting the sampling time by T_s and the output of the derivative term by u_d, we have:

$$u_d = \frac{s}{st + 1} v_r = \frac{1}{t_d} v_r - \frac{1/t_d}{st_d + 1} v_r .$$

A discrete-time realization for u_d is thus given by

$$u_d[(n+1)T_s] = \frac{1}{t_d} v_r[(n+1)T_s] - \frac{1}{t_d} u_{dd}[(n+1)T_s]$$

where u_{dd} is the discrete-time output of the realization of the first-order system $\frac{1}{st_d + 1}$:

$$u_{dd}[(n+1)T_s = e^{-T_s/t_d} u_{dd}[nT_s] + t_d \left(1 - e^{-T_s/t_d}\right) v_r[nT_s]$$

Thus, we can implement the PID controller with dirty derivative with the following discrete-time formula:

$$u[nT_s] = k_p v_r[nT_s] + k_i u_i[nT_s] + u_d[nT_s],$$

where u_i is the dynamics for the integral term and is given by

$$u_i[(n+1)T_s] = u_i[nT_s] + T_s v_r[(n+1)T_s] .$$

The PIQ controller replaces the derivative term of the PID controller with a signed quadratic (Q) term of the form $v_r |v_r|$. This nonlinear term becomes very small when the speed error v_r is small, but grows fast as the error grows. Therefore, it has the effect of generating aggressive control action to reduce large errors quickly, yet it turns itself off for small errors, thus avoiding the occurrence of undesirable overshoot. This feature helps achieve tight control despite the low actuation-to-weight ratio of heavy-duty vehicles. The corresponding expression is:

$$u = k_p v_r + k_i \frac{1}{s} v_r + k_q v_r \mid v_r \mid .$$

A discrete-time realization for the PIQ controller is

$$u[nT_s] = k_p v_r[nT_s] + k_i u_i[nT_s] + k_q v_r[nT_s] v_r[nT_s] ,$$

where u_i is given by

$$u_i[(n+1)T_s] = u_i[nT_s] + T_s v_r[(n+1)T_s] .$$

These two controllers were implemented in our closed-loop experiments, with hard limits that incorporated our subjective preferences for comfort and smoothness. To obtain the actual fuel and brake signals to the actuators, we pass the controller output through a linear mapping which converts the controller output calculated by the control scheme to the physical voltage values corresponding to the actual fuel and brake commands. After this conversion, the signals are sent out to the computer I/O ports and finally reach the fuel and brake actuators. The determination of the effective working ranges for the fuel and brake actuators should be based on ride quality, response speed, and safety. For example, a higher upper limit on the fuel would give us faster response during acceleration, but it can also cause undesirable jerk and possibly overshoot; a higher upper limit on the brake would give us quicker deceleration and shorter brake distance, but it would also deteriorate the ride quality and increase brake wear. In our closed-loop experiments, we chose the working range of the fuel command to be 1.25–4.5V and that of the brake command to be 1.2–3.0V. The sign of the controller output is used to decide whether a fuel or brake command should be issued, and then the value of the controller output is linearly mapped to the physical fuel or brake signal according to our pre-set working ranges.

Closed-loop speed control with PID controller. Even the simple PID controller can achieve fairly good tracking performance if the parameters are chosen appropriately, since the longitudinal vehicle dynamics is essentially a slow linear process. One closed-loop experimental result with PID controller is shown in Figure 5.

The parameters of the PID controller we used in this set of closed-loop experiments are $k_p = 30, k_i = 1, k_d = 5, t_d = 5$. The speed profile is the same as the speed profile used in our simulations; this makes it possible to compare our experimental results with our previous simulation result, which is shown in Figure 6. First, the vehicle starts out at an initial speed of 27mph and stays at this speed for 10s. At time t=10s the vehicle is given a command to accelerate at 0.3m/s^2 for 10s so that the vehicle velocity is 33.75mph after the acceleration. Then at time t=35s, a deceleration command of 3m/s^2 is issued for 3s. In our figures, the speed profile is shown with a dash line, while for the actual velocity a solid line is used.

Closed-loop speed control with PIQ controller. The PID and PIQ controllers have similar behavior and performance as we can see from the PIQ closed-loop results in Figure 7 and Figure 8. The parameters of the PIQ controller we used in this set of closed-loop experiment are $k_p = 10, k_i = 1, k_q = 10$. The PIQ controller was tested with the same speed profile as the PID controller.

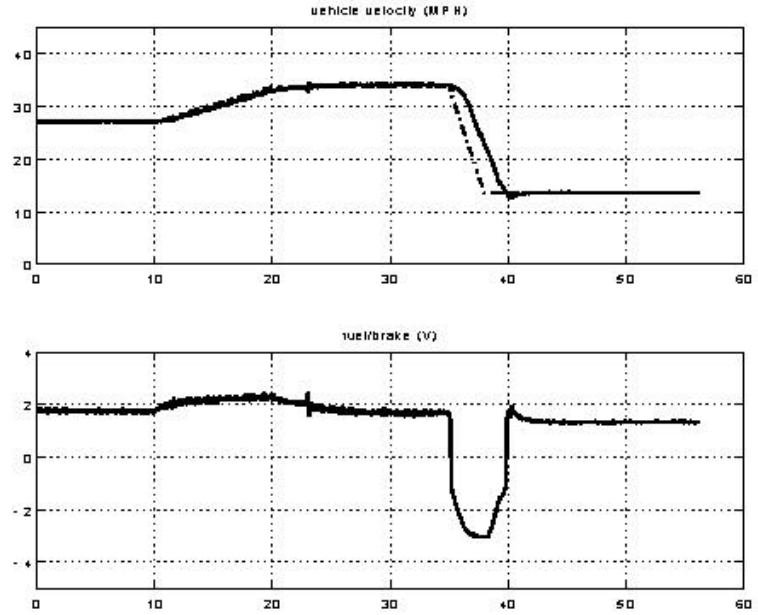

Figure 5: Closed-loop experiment with PID controller.

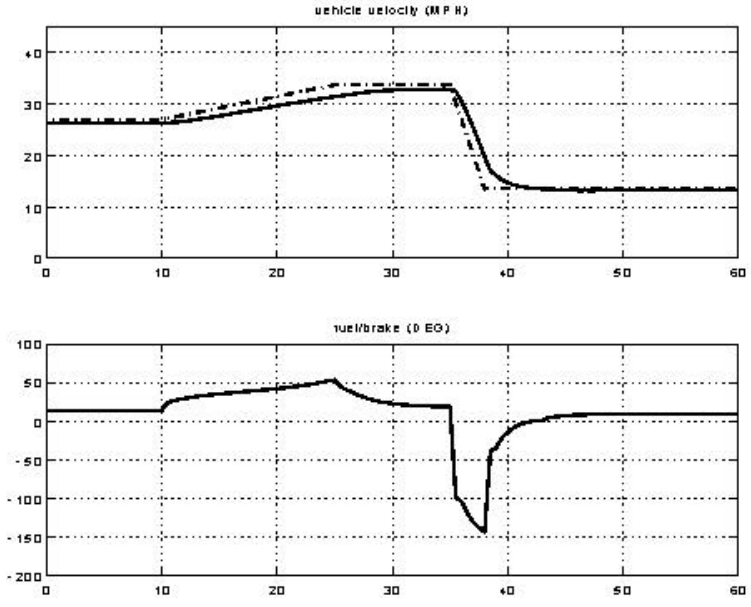

Figure 6: Closed-loop simulation with PID controller.

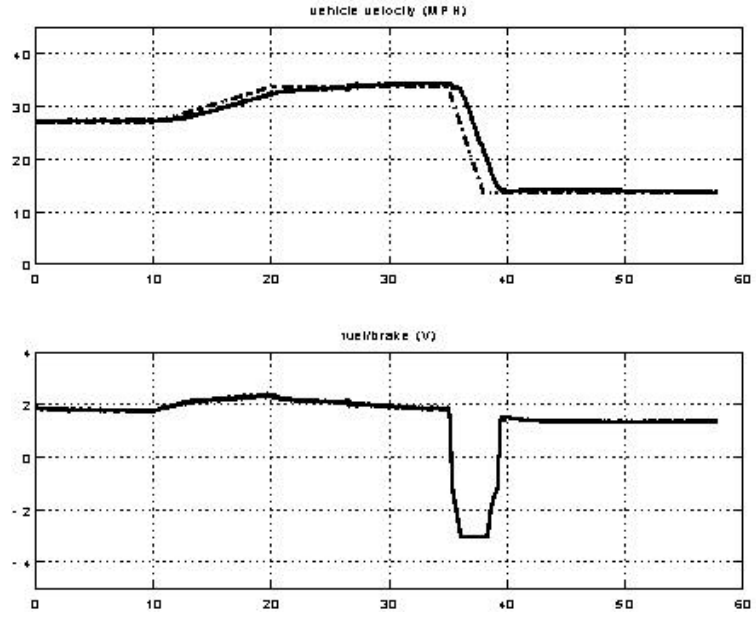

Figure 7: Closed-loop experiment with PIQ controller.

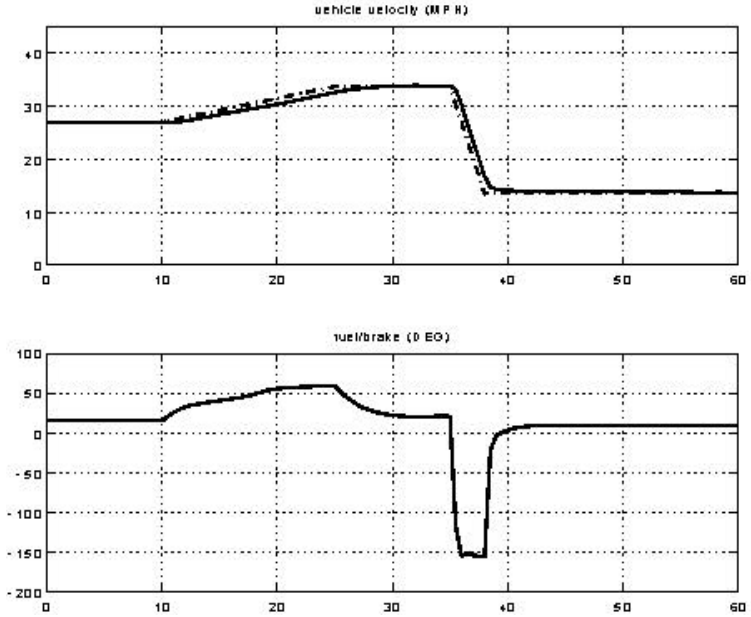

Figure 8: Closed-loop simulation with PIQ controller.

As can be seen from Figures 5–8, the agreement between our predicted performance through simulation and the actual experimental performance is good, though not perfect. The presented curves are quite similar in their low-frequency content, but the high frequency components exhibit non-negligible differences between simulation and experiment.

Robustness of fixed-gain PID and PIQ controllers. Our first experimental results were encouraging, especially since the speed profile that was used included a very abrupt deceleration command of $3 m/s^2$, yet the tracking performance was good. Another desirable feature of the implemented controllers, which may not be so obvious from the presented data, is that throughout these experiments the ride in the vehicle was very smooth and comfortable for both the driver and the passenger. However, these experiments also revealed several problems that we need to overcome in the future in order to be confident that our automated vehicle can be operated safely under all possible conditions. In particular, the experiments revealed drastic differences in the closed-loop performance as a result in changes in the vehicle operating parameters, such as the vehicle load. As an example, to investigate the effect of vehicle load variations on closed-loop performance, we tuned the parameters of the PID controller for good performance with the tractor only, and then we attached the trailer and ran the controller with the same parameters. As seen in Figure 9, this significant change in the truck mass resulted in unacceptable oscillations in the speed tracking response. Of course, this oscillation can be eliminated by returning the PID controller gains to the values tuned for tractor/trailer operation, but this would cause a deterioration of performance when the trailer is detached. Similar problems existed also with the PIQ controller.

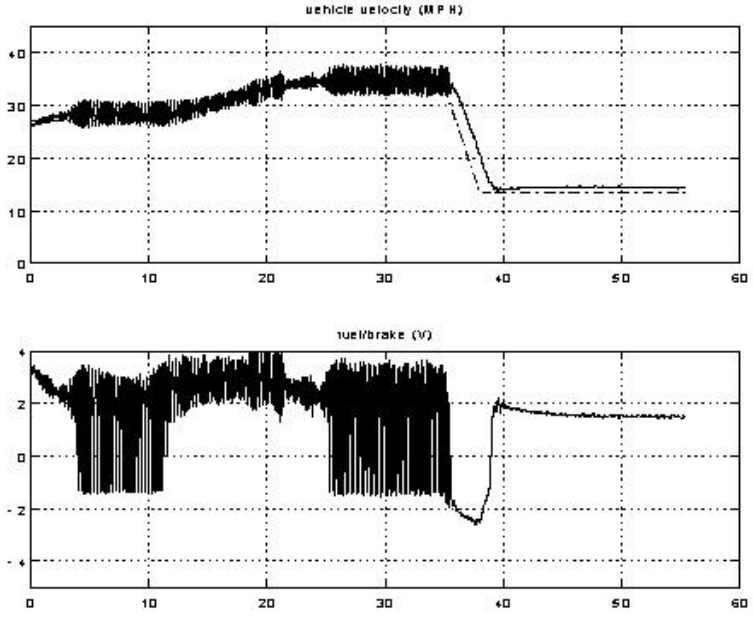

Figure 9: Oscillation of PID controller with trailer attached.

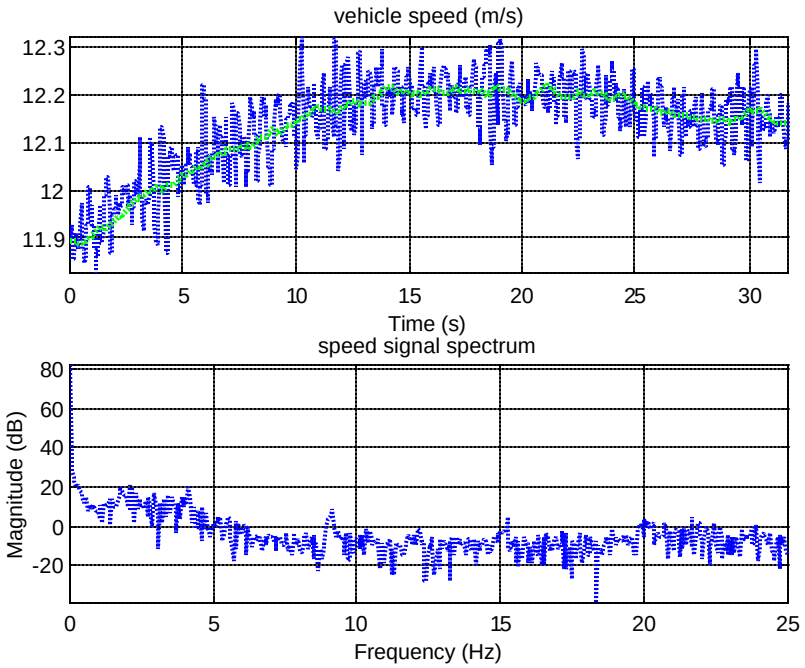

Figure 10: Original and filtered vehicle speed signal and its spectrum.

4.2.2 VIRTUAL VEHICLE FOLLOWING CONTROL

Speed sensor filter design. In the first stage of speed control experiments, we observed that the speed sensor signal contained a significant amount of noise. Therefore, we decided to design a low-pass filter to filter the noisy speed signal. An example of the vehicle speed signal is shown in Figure 10. We can see that most the signal energy stays in the 5Hz frequency range. On the other hand, 5Hz frequency is fairly high compared to the vehicle dynamics. To avoid exciting the unmodeled system dynamics, we chose 5Hz as the cutoff frequency of our speed signal filter. In order to keep the gain equal to one at DC, we used a third-order Chebyshev filter, whose transfer function is designed as:

$$H(z) = \frac{0.9 - 0.5z^{-1} - 0.5z^{-2} + 0.9z^{-3}}{1.0 - 2.8z^{-1} + 2.6z^{-2} - 0.8z^{-3}}.$$

Figure 11 shows the frequency response of the designed notch filter.

Fuel/brake output filter design. To further reduce possible oscillations due to high-frequency components in the fuel/brake command, we added a fuel/brake output filter to smoothen the outputs to the throttle and brake actuators. The filter we used is a second-order Chebyshev low-pass filter with cutoff frequency at 3.5 Hz

$$H(z) = \frac{0.1 - 0.2z^{-1} + 0.1z^{-2}}{11 - 1.8z^{-1} + 0.8z^{-2}}.$$

Figures 12 and 13 show the frequency response of the output filter and the original and filtered fuel/brake commands. In Figure 13, the black (green) line is the throttle control signal while the

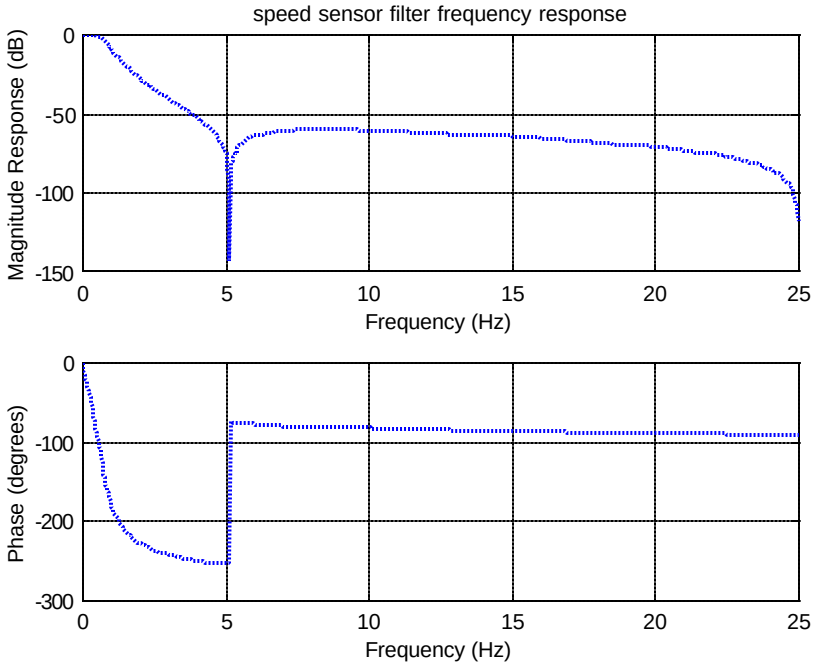

Figure 11. Speed sensor filter frequency response.

gray (red) line is the brake actuator signal. We can see that the filtered fuel/brake command is much cleaner than the original one.

Inner loop brake control. From the open-loop experiment, we observed that there is a significant delay due to the fact that the air brake system on the experimental truck has slow dynamics. The air pressure needs to build up before the brakes are activated. There is approximately 0.3s of delay, followed by 0.5s during which the actual pressure builds up to its commanded value. Therefore, we designed an inner loop PID controller to improve the dynamics of the air brake system and reduce the delay. This inner loop PID controller is able to increase the amount of the air flowing into the air brake system when the air pressure measured by the pressure transducers is much less than the commanded air pressure (corresponding to the brake command). As a result, the air pressure is built up faster. On the other hand, when the actual air pressure in the air brake system is significantly higher than the commanded air pressure, the PID controller can reduce the air pressure faster by setting the release valves to the wide-open position. In order to eliminate the measurement noise in the air pressure signals, we used a notch filter with a 7.5Hz cutoff frequency, whose magnitude and phase plots are shown in Figure 14. The effect of the inner loop brake control is shown in Figure 15. The light gray (green) lines in the upper and lower plots are the commanded air pressure converted from the brake command signal. The solid dark (blue) line in the upper plot is the original brake system air pressure response, from which we can see that it has a very slow dynamics response and steady-state error. The filtered air pressure signal is the dark gray (red) line going through the middle of the original air pressure. The solid dark (blue) line in the lower plot is the air pressure response with the inner loop brake controller, which shows a significant improvement in terms of the response time and the steady-state error.

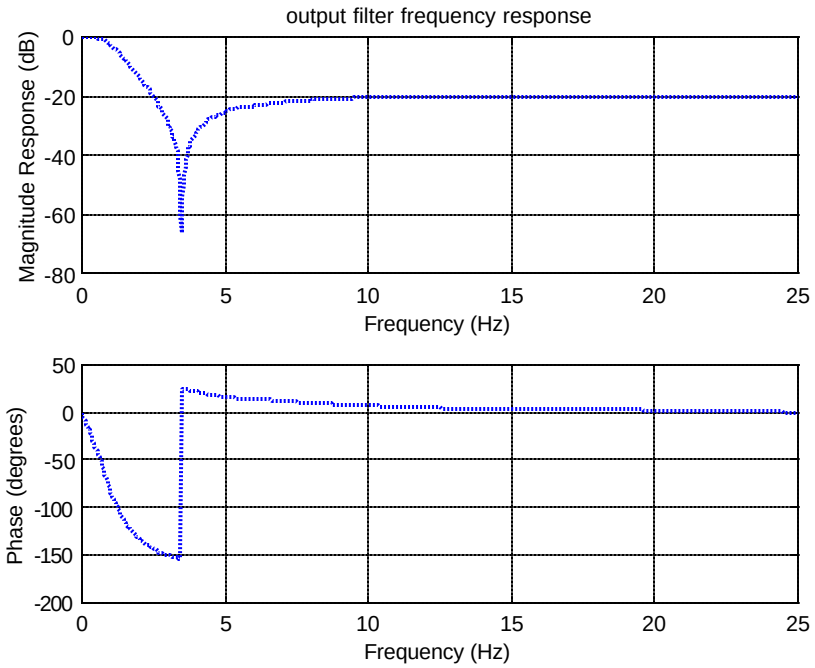

Figure 12: Fuel/brake output filter frequency response.

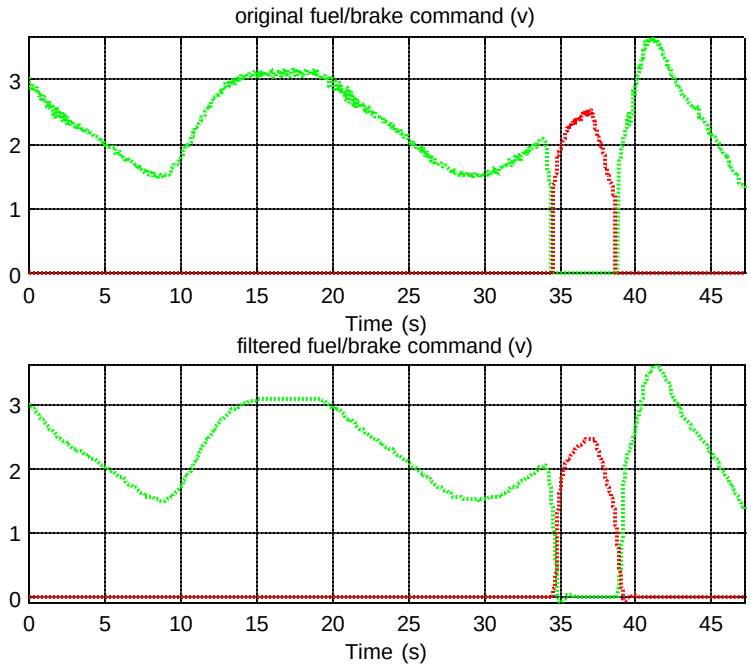

Figure 13: Original and filtered fuel/brake commands.

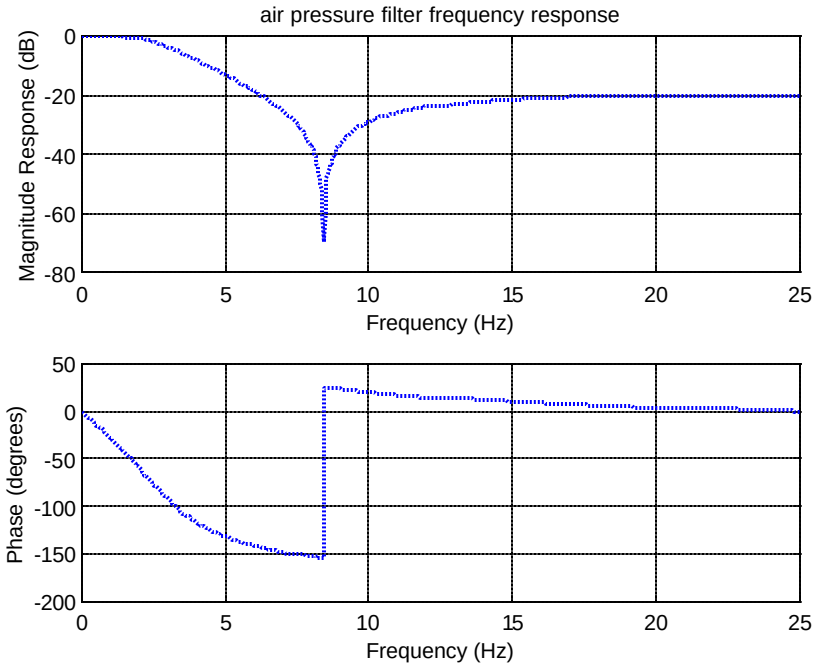

Figure 14: Air pressure filter frequency response.

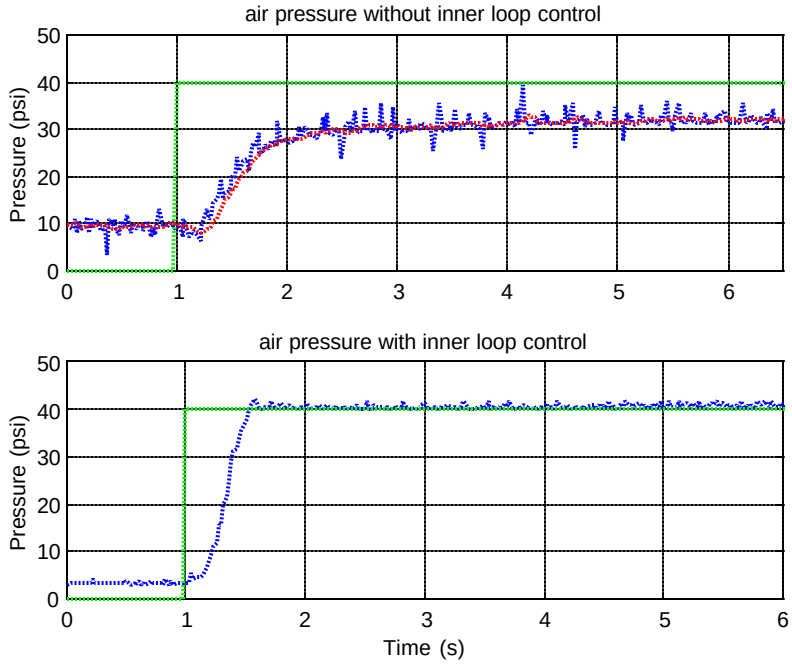

Figure 15: The effect of the inner loop brake control.

Virtual vehicle following. The design of the speed sensor filter, the fuel/brake command filter, the air pressure filter, and the inner loop brake control provided the necessary foundation for implementing the vehicle following control schemes described in detail in Sections 3.2.1 and 3.2.2. In order to carry out the closed-loop experiments for testing the performance of the vehicle following controllers, we needed to measure the separation between the leading vehicle and the follower. Since we only had one experimental truck, we utilized the concept of "virtual vehicle following" to circumvent this problem. Hence, we assumed that the leading vehicle is running at a certain ideal speed profile $v_l(t)$. By integrating the speed profile in time, we are able to compute the traveling distance of the leading vehicle

$$x_l(t) = \int_{m=0}^{t} v_l(m)dm .$$

There are two ways to obtain the traveling distance of the experimental vehicle (the follower). One way is to integrate the speed signal measured by the speed sensor of the truck, which gives us

$$x_f(t) = \int_{m=0}^{t} v_f(m)dm .$$

Therefore, the virtual distance between two vehicles can be computed with the following formula:

$$\begin{aligned} x_r(t) &= x_l(t) - x_f(t), \\ &= \int_{m=0}^{t} \left(v_l(m) - v_f(m)\right)dm. \end{aligned}$$

Similarly, the separation error d can be computed as

$$d = x_r(t) - s_d .$$

Another way is to use the magnetic markers that are embedded in the ground on the test track with about one meter separation. The magnetic field signals from the markers can be detected by the magnetometer array installed on the back of the truck trailer, which means that we are able to compute the traveling distance by counting how many markers have been passed. One disadvantage of the second method is that the traveling distance precision is limited by one meter due to the separation between magnetic markers.

In the closed-loop experiments for virtual vehicle following, we used the first method.

Effectiveness of the nonlinear spacing policies. The two nonlinear spacing policies described in Section 3.2.3, namely variable time headway and variable separation error gain, were implemented and tested. Figures 16 and 17 show the performance of the controller with and without variable time headway, respectively. The dark (blue) and gray (green) lines in the upper plot are the speeds of the leading and following vehicles. The lower plot shows the separation error. We can see that the transient separation error is significantly decreased with the use of variable time headway. When the following vehicle detects that the leading vehicle is moving

faster, it reduces its time headway according to the variable time headway spacing policy, which therefore decreases the separation error. On the other hand, when the following vehicle detects that its speed is higher that that of the leading vehicle, it increases its time headway for improved safety. These results provide experimental verification of the theoretical result that the variable time headway is effective in improving vehicle following performance.

As previously discussed, the introduction of variable separation gain relaxes the requirement on reducing the separation error and therefore prevents abrupt controller reactions. This is shown in Figures 18 and 19. In Figure 18, s is set to 0, resulting in a fixed separation error gain $k = k_0$, while in Figure 19 s is set to 0.05, resulting in the variable separation error gain:

$$k = c_k + (k_0 + c_k)e^{-0.05d^2}$$

We can see that accordingly the separation error becomes larger, in accordance with the relaxation on the spacing error requirement.

Closed-loop experiment for vehicle following. The PID and PIQ controllers were extensively tested in the closed-loop vehicle-following experiments with several different combinations of nonlinear spacing policies. The tuning of the parameters was based on the open-loop testing results aimed at characterizing the throttle and brake dynamics as well as the on-site experiments. The experiments were all carried at Crow's Landing test field. The total length of the test track is 2200 meters with three curvatures of 800 meters radius. Figure 20 shows a typical response, obtained with the PID controller with the following parameter setting for the gains and the nonlinear spacing policies:
$k_p = 3,\ k_i = 0.5,\ k_d = 5,\ h_0 = 0.5,\ c_h = 0.2,\ c_k = 0.5,\ k_0 = 1,\ s = 0.1.$

5. Conclusions

Our results are the first experimental results that prove the effectiveness of nonlinear spacing policies in improving the vehicle following performance of relatively simple PID and PIQ controllers. These controllers were originally designed to provide string stability and improved performance for autonomous vehicle operation (without intervehicle communication) and without excessive complexity in their practical implementation. The experiments carried out under MOU 314 verified that our theoretical predictions were correct, despite all the unavoidable differences between our simulated vehicle models and the actual experimental vehicle.

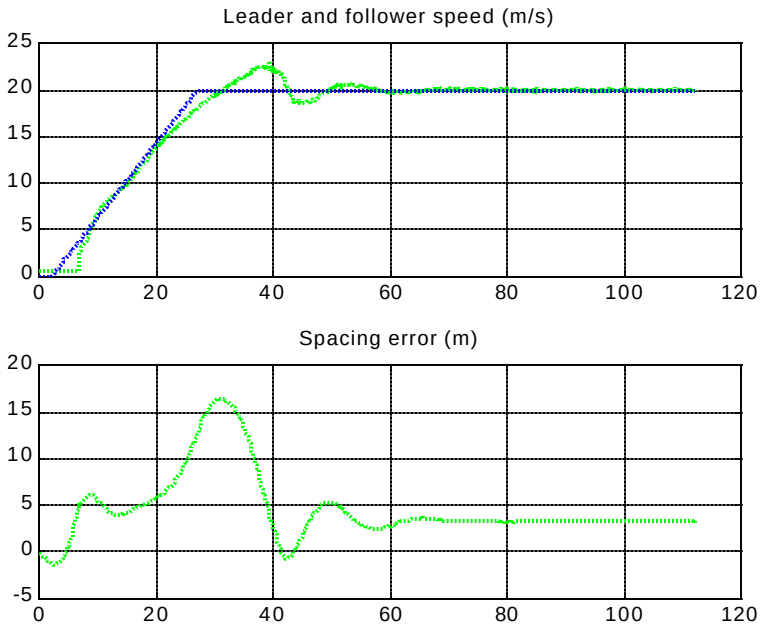

Figure 16: Performance with fixed time headway.

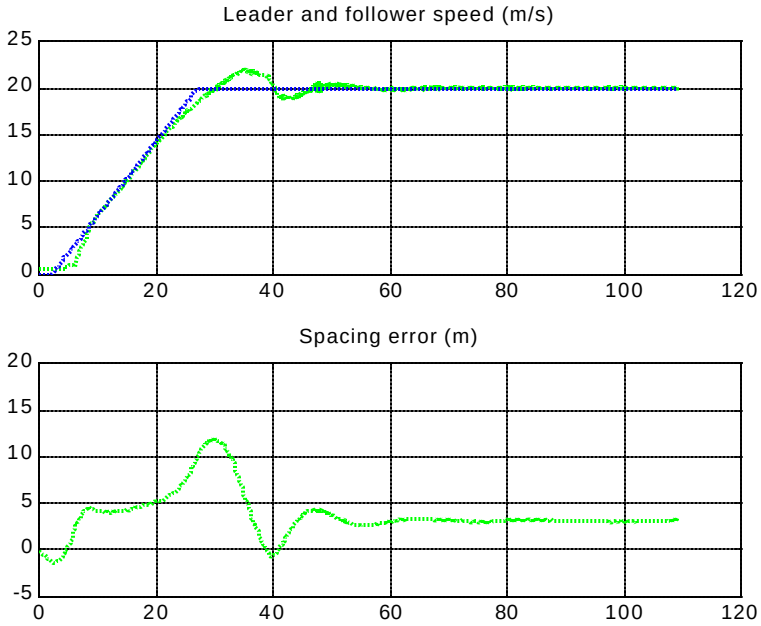

Figure 17: Performance with variable time headway.

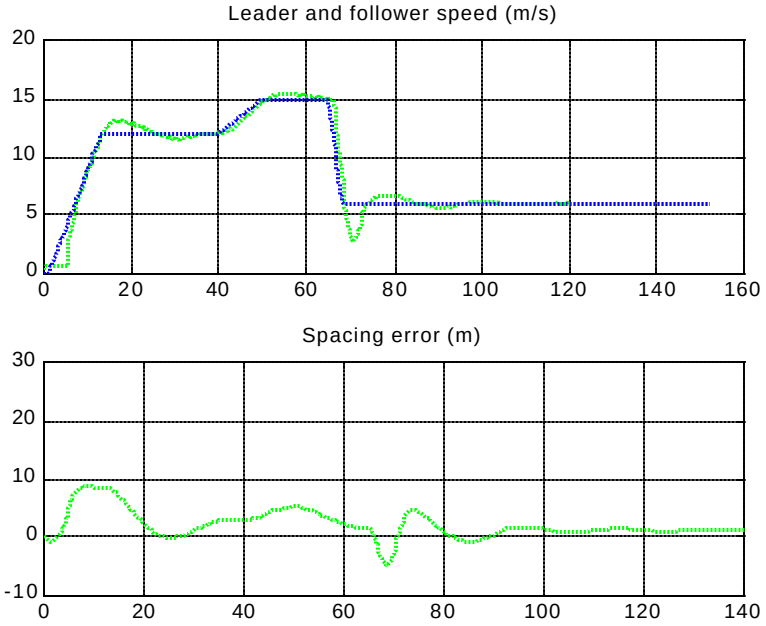

Figure 18: Performance with fixed separation error gain.

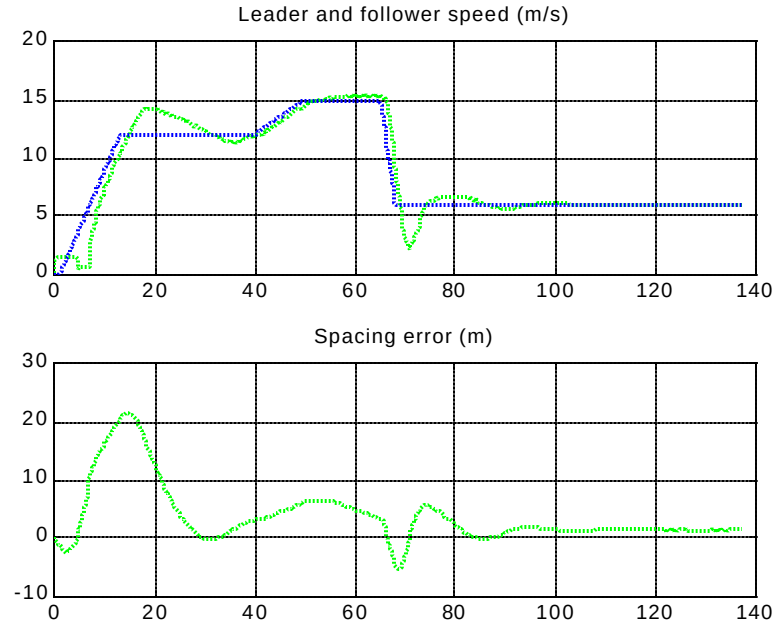

Figure 19: Performance with variable separation error gain.

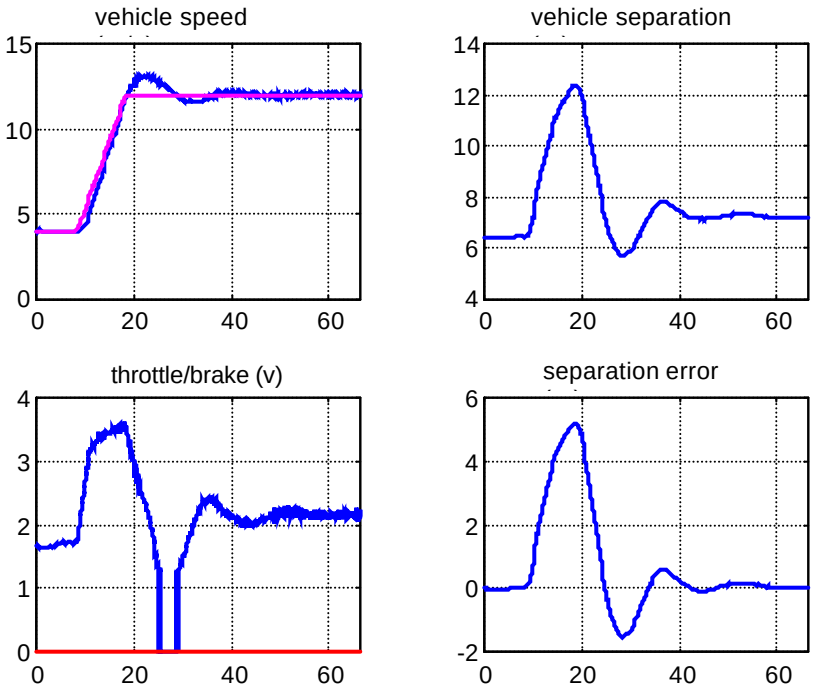

Figure 20: Typical closed-loop vehicle following experiment.