2007-01-0808

Nonlinear Observer for Vehicle Velocity with Friction and Road Bank Angle Adaptation – Validation and Comparison with an Extended Kalman Filter

Lars Imsland, Håvard Fjær Grip*, Tor A. Johansen* and Thor I. Fossen*
SINTEF ICT, Applied Cybernetics, N-7465 Trondheim, Norway

Jens C. Kalkkuhl and Avshalom Suissa
DaimlerChrysler Research and Technology, 70456 Stuttgart, Germany

ABSTRACT

A nonlinear observer for automotive vehicle velocity estimation in presence of varying friction and road bank angles is presented, and validated and compared to an Extended Kalman Filter (EKF) implemented for the same purpose. The performance of the nonlinear observer is as good as the EKF, while having significantly lower computational complexity.

INTRODUCTION

Many active safety systems in automotive vehicles, for instance yaw stability systems such as ESC/ESP, depend on information about vehicle velocity, in particular lateral velocity or side-slip angle, to be able to function properly. However, the vehicle velocity is rarely measured directly due to issues of cost and reliability, and must therefore in general be inferred from other measurements, such as wheel speed, steering angle, yaw rate, and acceleration measurements.

Systems that use measurements and dynamic models to infer dynamic state variables are often called state estimators, or *observers*. Nonlinear dynamic models and measurement equations call for nonlinear techniques. The predominant nonlinear state estimation technique is the use of the Extended Kalman Filter (EKF). Although the EKF has proved its performance in countless applications, it has some potential drawbacks with respect to especially computational complexity, but also tuning and in some cases, lack of analytical stability guarantee.

This paper reports on the development of a nonlinear observer for vehicle velocity estimation, based on nonlinear analysis techniques. This is an overall velocity observer with adaptation of the maximum friction coefficient and the road bank angle. The detailed development of the different parts of the observer is described in [12, 11, 7, 6, 10], and will be briefly summarized herein. Furthermore, we want to assess, to a certain degree, its performance in terms of accuracy and robustness, tuning, and computational complexity, in comparison with the EKF [19] applied to the same problem. The sensor suite assumed provides us with the fairly standard measurements mentioned above, but in addition both the nonlinear observer and the EKF use longitudinal and vertical acceleration sensors.

Earlier work on observers for estimation of lateral velocity is mainly based on linear or quasi-linear techniques, for example [4, 22, 20, 3]. A nonlinear observer linearizing the observer error dynamics is proposed in [13, 14]. The same type of observer, in addition to an observer based on forcing the dynamics of the nonlinear estimation error to the dynamics of a linear reference system, are investigated in [9]. The problem formulation there assumes that the longitudinal wheel forces are known, similarly to the observer implemented in ESP [21]. In our work, we do not make this assumption, as such information is not always available.

An EKF is used for estimating vehicle velocity and tire forces in [17, 18], thus without the explicit use of friction models. A similar, but simpler, approach is suggested in [3]. An EKF based on a tire-road friction model which also includes estimation of the adhesion coefficient and road inclination angle is suggested in [19]. In [1], the use of an EKF based on a nonlinear tire-road friction model is considered, which also includes estimation of cornering stiffness. The strategy proposed in [15] combines dynamic and kinematic models of the vehicle with numerical band-limited integration of the equations to provide a side-slip estimate. In [8] the side-slip angle is estimated along with yaw rate in an approach which is similar to the one considered herein, but without yaw rate measurements. The approach is validated using experimental data, but no stability proofs are presented.

*Also affiliated with NTNU, Department of Engineering Cybernetics, N-7491 Trondheim, Norway

VEHICLE MODELING

For a more complete description of the vehicle model, we refer to [12, 11, 14]. Here, we briefly sum up the vehicle model in the horizontal plane as

$$\dot{v}_x = v_y r + a_x, \tag{1a}$$

$$\dot{v}_y = -v_x r + a_y - g\sin\phi_R, \tag{1b}$$

$$\dot{r} = \frac{1}{J_z}\sum_{i=1}^{4} \mathbf{g}_i^{\mathsf{T}} \mathbf{R}(\delta_i)\mathbf{F}_i(v_x, v_y, r, \delta_i, \omega_i; \mu_H), \tag{1c}$$

where v_x and v_y are longitudinal and lateral velocity of the vehicle given in a body-fixed coordinate system with the origin at the center of gravity (CG). The angular velocity about the vertical axis of this coordinate system (the *body-fixed coordinate system*) is r, often referred to as the yaw rate. Furthermore, a_x and a_y are (measured) longitudinal and lateral acceleration (in an inertial frame), respectively, g is the acceleration of gravity, and ϕ_R is the road bank angle. The vectors $\mathbf{g}_i$ give the placement of wheel i in the body-fixed coordinate system. The rotation matrix $\mathbf{R}(\delta_i)$, given by the wheel steering angle δ_i, rotates the forces $\mathbf{F}_i$ acting on wheel i, from the wheel coordinate system to the body-fixed one.

The friction forces at each wheel are typically modeled as functions of the individual wheel tire slips, $\mathbf{F}_i = \mathbf{F}_i(\lambda_{i,x}, \lambda_{i,y})$, where the tire slips $\lambda_{i,x}$ and $\lambda_{i,y}$ are measures of the relative difference in vehicle and tire longitudinal and lateral velocity for wheel i. See [12, 11, 16]. Since these tire slips are functions of the dynamic states and wheel steering angles and wheel speeds, we will often write $\mathbf{F}_i = \mathbf{F}_i(v_x, v_y, r, \delta_i, \omega_i; \mu_H)$, where we have added μ_H to denote dependence on the road conditions through the maximum tire-road friction coefficient. We will sometimes, for notational convenience, use $f_y = \sum_{i=1}^{4} (0\ 1)\,\mathbf{R}(\delta_i)\mathbf{F}_i$ for the overall tire forces acting on the vehicle in the lateral direction in the body-fixed coordinate system, and $f_r = \sum_{i=1}^{4} \mathbf{g}_i^{\mathsf{T}}\mathbf{R}(\delta_i)\mathbf{F}_i$ for the torque about the vertical axis generated by the friction forces.

LATERAL VELOCITY ESTIMATION: PROBLEM DESCRIPTION AND OUTLINE OF APPROACH

The main challenge is to construct an estimate of the lateral velocity v_y, which dynamics can be described by (1b). If we assume we have a good estimate of v_x, we can obtain an estimate of v_y by integrating this equation since r and a_y are measured. However, measurement and estimation errors will soon lead to divergence of this estimate, and we thus need some feedback (an injection term) to make the estimate converge. The only measurement that depends on v_y algebraically is a_y, through the tire friction model via Newton's law. We write

$$ma_y = f_y(t, v_y), \tag{2}$$

where m is the vehicle mass. We have the fortunate situation that typically (for many tires and/or driving situations) the partial derivative of f_y with respect to v_y is sign definite,

$$\frac{\partial f_y(t, v_y)}{\partial v_y} < 0, \tag{3}$$

which, assuming we know the friction model and that the road is flat, allows us to specify an observer by copying the system dynamics (1b) (with $\phi_R = 0$) and adding an injection term,

$$\dot{\hat{v}}_y = -v_x r + a_y - K_{v_y}(ma_y - f_y(t, \hat{v}_y)). \tag{4}$$

Convergence follows if $K_{v_y} > 0$ (see [12, 11] for details).

Having good estimates of v_y and v_x, the vehicle body side-slip angle β can be calculated as $\tan\beta = -\frac{v_y}{v_x}$.

However, the situation above is idealized since we have assumed that the friction model is known perfectly, and that the road is flat.

UNKNOWN FRICTION MODEL The road-tire friction will vary significantly on different types of road surfaces. Fortunately, most of the uncertainty in the friction model can be lumped into one parameter denoted μ_H (the maximum tire-road friction coefficient), which can be said to characterize the road surface; from $\mu_H \approx 0.1$ on wet ice to $\mu_H \approx 1.0$ on dry asphalt. Based on a friction model parameterized in μ_H, the above observer can be extended with adaptation of μ_H, using the techniques in [7, 6].

SLANTED ROADS On slanted roads, gravity affects the horizontal acceleration measurements. Disregarding road inclination, we will consider only road bank angle. A non-zero road bank angle influences the dynamics through the term $g\sin\phi_R$ in (1b). It should be noted that for the sake of notational simplicity, a_y denotes the reading of the acceleration sensor, which (theoretically) is equal to the "real" a_y when there is no road bank angle.

Since a road bank angle also introduces a gravity component in the wheel force balance, the effect of the road bank angle will cancel out in (2) and thus the road bank angle will not affect the injection term in (4). For this reason, one approach could be to use a large gain (a large K_{v_y}) in (4). However, this will amplify noise and friction model uncertainty, and is therefore not a satisfactory approach in practice.

One could imagine adapting ϕ_R similarly as for friction. However, as $g\sin\phi_R$ could be regarded an unknown input in (1b), we have chosen to use techniques from "unknown input observer"-theory, and combine this with an adaptation law. The theory behind this can be found in [10].

As we have relatively good measurements of longitudinal velocity through the wheel speeds, it turns out the effect of

non-zero road inclination angle is not as crucial for the v_x estimate (except, perhaps, for very small v_x) as non-zero road bank angle is for the v_y estimate. Hence we choose not to estimate the road inclination angle.

OVERVIEW OF THE NONLINEAR OBSERVER

The implemented observer can be described by a number of ordinary differential equations derived using nonlinear observer theory [12, 11, 7, 6] based on the dynamic model of the vehicle and the friction model. The sensors used are three-axis accelerometers (a_x, a_y and a_z) in addition to yaw rate r (angular rate about vertical axis). The steering wheel angle (used to calculate individual wheel angles, lumped into vector $\boldsymbol{\delta}$) and wheel speeds ($\boldsymbol{\omega}$) are also measured.

LONGITUDINAL VELOCITY OBSERVER The observer equation for longitudinal velocity is

$$\dot{\hat{v}}_x = \hat{v}_y r + a_x + \sum_{i=1}^{4} K_i(t)(v_{x,i} - \hat{v}_x). \tag{5}$$

The functions (the time-varying *gains*) $K_i(t)$ weights the difference between the wheel speed measurement translated to longitudinal velocity of the CG, $v_{x,i}$, and the estimate $\hat{v}_x$. The gains vary with time to let the influence of this difference vary, for example large longitudinal accelerations usually imply large slips, when $v_{x,i}$ is not a good measurement of v_x. See [12] for further details.

LATERAL VELOCITY OBSERVER The lateral velocity observer makes use of the friction model:

$$\dot{\hat{v}}_y = -r\hat{v}_x + a_y - \hat{w} + K_{v_y}\Lambda\xi\left(ma_y - \hat{f}_y\right) + \frac{\Gamma_2}{\Gamma_1}\zeta\left(r - \hat{r}\right),$$
$$\dot{\hat{r}} = \frac{1}{J_z}\hat{f}_r + K_r\left(r - \hat{r}\right),$$

where K_{v_y}, K_r, Γ_1 and Γ_2 are positive gains, Λ is a (possibly time-varying, state-dependent) positive scaling factor, and ξ and ζ can be interpreted as partial derivatives of f_y and f_r with respect to v_y (evaluated using the estimated states, see [6] for details). The friction forces on wheel i (in the wheel coordinate system) are calculated via a friction model, $\hat{\mathbf{F}}_i = \mathbf{F}_i(\hat{v}_x, \hat{v}_y, \hat{r}, \boldsymbol{\delta}, \boldsymbol{\omega}; \hat{\mu}_H)$, based on the estimated state variables and the measurements, and an estimate of the maximum friction coefficient, μ_H (see below and [7, 6]). Based on this, $\hat{f}_y = \sum_{i=1}^{4}(0\ 1)\,\mathbf{R}(\delta_i)\hat{\mathbf{F}}_i$ is the estimated force in lateral direction in the body-fixed coordinate system, and $\hat{f}_r = \sum_{i=1}^{4}\mathbf{g}_i^{\mathsf{T}}\mathbf{R}(\delta_i)\hat{\mathbf{F}}_i$ is the estimated torque about the vertical axis generated by the friction forces.

The estimate $\hat{w}$ in the equation for $\hat{v}_y$ is an estimate of $w = g\sin\phi_R$, the influence of the road bank angle ϕ_R (see below and [10]).

For low velocities, issues like decreasing signal-to-noise ratios and accuracy of the friction model deteriorate the lateral velocity observer performance. Thus, for longitudinal velocities below 5 m/s, we make a smooth transition of the observer estimates to a static estimate based on the average front wheel steering angle δ,

$$\tan\hat{\beta} = \kappa\delta, \tag{6}$$

where κ is a constant given by the placement of the CG relative to the front and rear axle.

FRICTION COEFFICIENT ADAPTATION LAW The friction coefficient adaptation law is

$$\dot{\hat{\mu}}_H = \Gamma_1 K_{v_y}\Lambda\xi_\mu\left(ma_y - \hat{f}_y\right) + \Gamma_2\zeta_\mu\left(r - \hat{r}\right), \tag{7}$$

where ξ_μ and ζ_μ are partial derivatives of f_y and f_r with respect to μ_H (evaluated using the estimated states). See [6]. The exact formulation of the adaptation law (that is, ξ_μ and ζ_μ) depends on the chosen parameterization of the friction model; either using a Taylor expansion [7] or a linear parameterization [6].

The friction coefficient (and thereby v_y) is only observable under driving maneuvers that are *persistently exciting* (PE). Supported by theory [6], it is possible to monitor an excitation condition, and only adapt when this condition is fulfilled. Loosely speaking, driving maneuvers with varying lateral velocity are PE. During periods of low excitation, the friction coefficient is attracted to 1 for robustness and safety reasons (a large friction coefficient generally gives lower lateral velocity estimates). Thus, friction adaptation only occurs in periods with significant changes in v_y, and one could argue that these are exactly the periods when it is needed, since the purpose of adapting the friction coefficient is to improve the v_y estimate. Moreover, it can be shown that when driving with $v_y = 0$ (straight ahead), the lateral velocity estimate will converge even though the friction coefficient estimate will not. Thus, theoretically, only indefinitely long maneuvers with constant non-zero v_y (that is, driving in circles) are problematic, but prolonged maneuvers of this kind are rare.

Another practical modification to the adaptation algorithm is the implementation of an upper and (possibly time-varying) lower bound on the friction coefficient. These aspects are covered together with stability and robustness analysis in [7, 6].

ESTIMATION OF ROAD BANK ANGLE The road bank angle can be viewed as an unknown input in the differential equation describing v_y, hence we borrow from Unknown Input Observer-theory [10]. The observer equa-

tions for $\hat{w}$ are

$$\dot{z} = -K_{v_y}\left(ma_y - \hat{f}_y\right), \tag{8a}$$

$$v_y^m = \left(\frac{\partial \hat{f}_y}{\partial v_y}\right)^{-1}\left(ma_y - \hat{f}_y + \frac{\partial \hat{f}_y}{\partial v_y}\hat{v}_y\right), \tag{8b}$$

$$\hat{w} = K_w(\hat{v}_y - z - v_y^m). \tag{8c}$$

Here, we note that (8b) can be interpreted as an approximate inversion of the friction model, (2). The notation $\frac{\partial \hat{f}_y}{\partial v_y}$ denotes the partial derivative of the friction model with the estimated states inserted, thus this is essentially the same as ξ above. When this partial derivative in (8b) becomes small (in absolute value), which happens when wheel side-slips are high, the friction model is not practically invertible. To avoid numerical problems at high side-slips, we saturate the partial derivative such that it does not become too small. This practical modification seem to provide reasonably good road bank angle estimates also in cases with high side-slips.

It has proven difficult to estimate the friction coefficient and road bank angle at the same time, since, at least during transients, a road bank angle is easily detected as low friction (see e.g. [21]). Our approach to this problem is to try to detect whether there is a non-zero road bank angle. For this, we also use a measurement of vertical acceleration, a_z. By default, the maximum road-tire friction coefficient is adapted, but if a (significant) non-zero road bank angle is detected, we estimate the bank angle instead, and the estimate of the maximum road-tire friction coefficient tends to 1 exponentially to ensure safer lateral velocity estimates (it is easy to argue that in most applications, too small estimates, in absolute value, are safer than too large). The drawback of this approach is that we cannot expect good estimates if we have both low friction and significant road bank angles. On low friction, inverting the friction model will in this case not produce correct results since the friction coefficient is wrong. It should also be pointed out that, in accordance with this approach, the developed theory only covers either friction or road bank angle adaptation, not changes in both simultaneously.

DISCRETIZATION For implementation in digital hardware, the ODEs are discretized. Due to the short sampling time, the use of moderate gains in the observer, and the scaling factor Λ which scales the partial derivatives of the friction model, we get good performance with simple forward Euler discretization.

EXTENDED KALMAN FILTER (EKF)

Since we will compare observer performance with an EKF (see e.g. [5] for a general introduction), we include a brief description of the discrete-time EKF used. This EKF is developed over many years, and is similar to the EKF described in [19]. The dynamic model used is a linear time-varying (LTV) one, obtained using the model (1) for v_x and v_y (using r as a known time-varying quantity), and first-order Markov models for the friction coefficient, road inclination angle, and road bank angle.

The measurement equations are given by force and torque balances via the nonlinear friction model for a_x, a_y and $\dot{r}$ (calculated from r using numerical differentiation). In addition, v_{ref}, a "reference velocity" calculated from wheel speeds (and using other information such as longitudinal acceleration and brake flags, ABS flags, etc), somewhat similar to (5), is used. The EKF also uses a_z, by letting non-zero a_z influence the process noise covariance matrix.

The fact that the model is LTV makes the implementation more efficient. Furthermore, the Bierman algorithm [2] is used for efficient, stable numerical implementation of the measurement update process.

VALIDATION AND COMPARISON WITH THE EKF

IMPLEMENTATION Both the observer and the EKF are implemented in Simulink as S-functions coded in C, ready to be downloaded to in-vehicle target hardware using the Real Time Workshop in Simulink[1]. The target hardware in this case is a dSPACE MicroAutoBox[2] running at 800 MHz. This hardware has floating point arithmetics. The sampling rate of 100 Hz is no problem for the nonlinear observer, nor the EKF, on this hardware.

COMPUTATIONAL COMPLEXITY The main properties of interest of an implementation of an observer, are the execution time and the memory usage. The execution times are tested on the hardware mentioned above, while some general comments are made about memory usage.

Execution times The execution times were obtained by marking the Simulink subsystem containing the nonlinear observer (and the one containing the EKF, respectively) as an atomic subsystem. Custom code is added to measure the execution time from the beginning to the end of the `update` part of the Simulink code for that subsystem. As the `output` part of the Simulink code both for the nonlinear observer and the EKF consist of delivering the estimated states to the parent Simulink block, these are not included in the measurement of the execution times.

Table 1 shows the average and maximum execution time for both approaches, for a 20 s test run. The ratio between the maximum execution times is close to three.

In interpretation of these numbers, it should be taken into

[1] Simulink, Matlab and Real Time Workshop are registered trademarks of The Mathworks, Inc., `http://www.mathworks.com/`.

[2] dSPACE GmbH, `http://www.dspace.com/`.

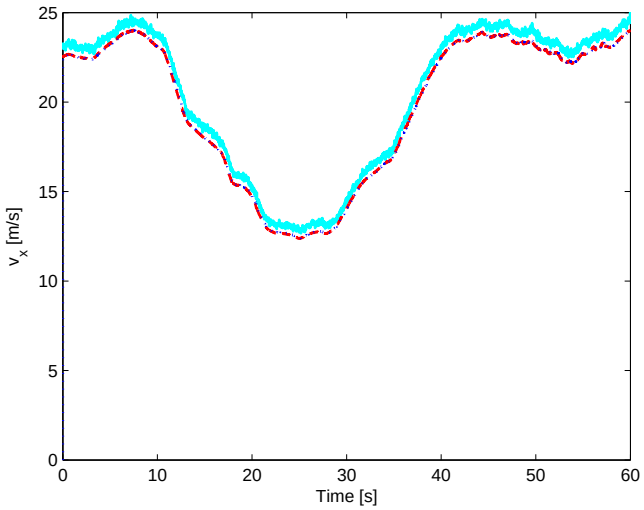

(a) Longitudinal velocity v_x

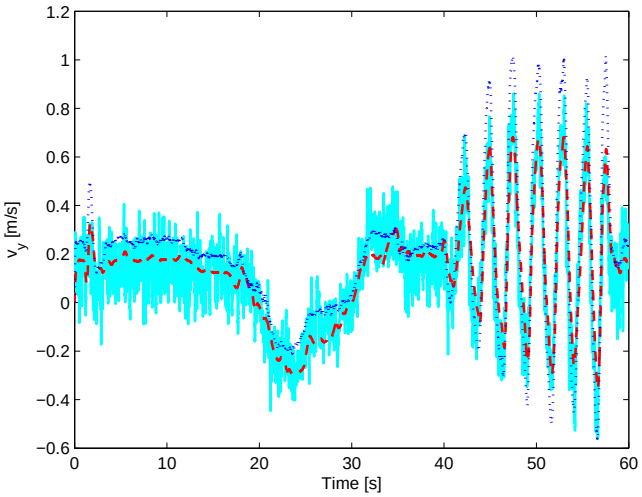

(b) Lateral velocity v_y

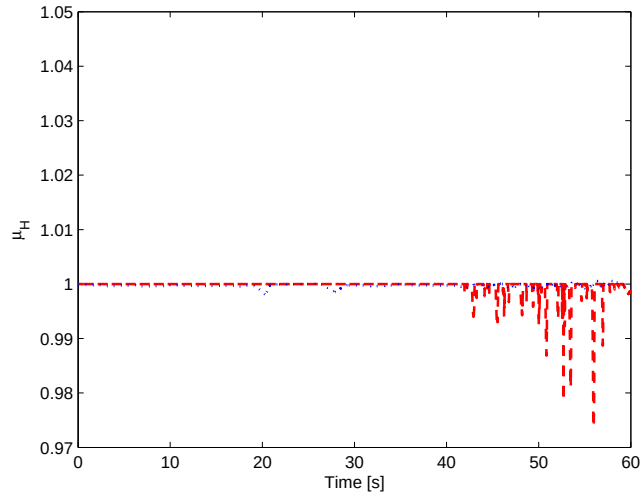

(c) Estimated friction parameter μ_H

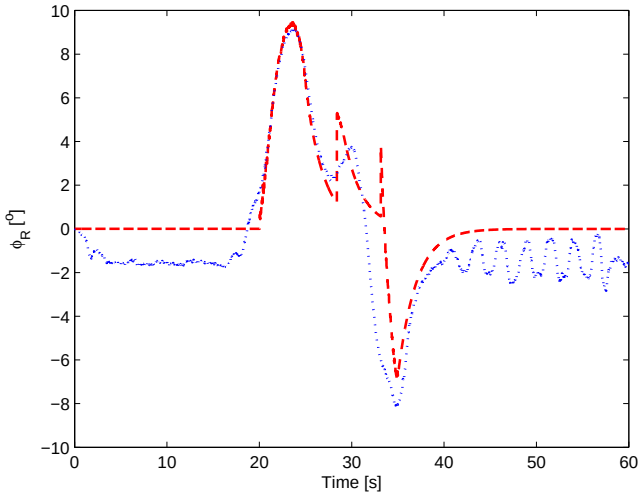

(d) Estimated road bank angle ϕ_R

Figure 1: Results for measurement set 1. The estimates from the nonlinear observer are plotted with a dashed line, while the estimates from the EKF are plotted with a dotted line. Velocity measurements are plotted with a whole line in (a) and (b).

	Nonlinear observer	EKF
Mean	84.8 μs	253.5 μs
Maximum	86.6 μs	258.1 μs

Table 1: Execution times for the nonlinear observer and the EKF

consideration that the EKF has been tested more extensively than the nonlinear observer, and contains more heuristics to tackle special cases. Nevertheless, the difference in execution times is significant, and considerably larger than can be explained by these heuristics alone.

It is interesting to note that the friction model, including its partial derivatives, dominates computations in the nonlinear observer. The same friction model is used for both the nonlinear observer and the EKF.

Memory usage We found no obvious way of measuring memory usage of the nonlinear observer and the EKF on the real time operating system.

However, some general comments about the differences can be made. The number of estimated states are the same in both applications, and both have the friction model implemented in the same way (the EKF needs one more partial derivative of the friction model than the nonlinear observer). The main difference is that the EKF must store the 15 dynamic states related to the solution of the Riccati equation. Therefore, we may expect that the memory requirements of the EKF are larger than for the nonlinear observer. The extra logics and heuristics of the EKF add to the memory requirements.

The code for the nonlinear observer is an order of magnitude smaller than the code for the EKF, in terms of lines

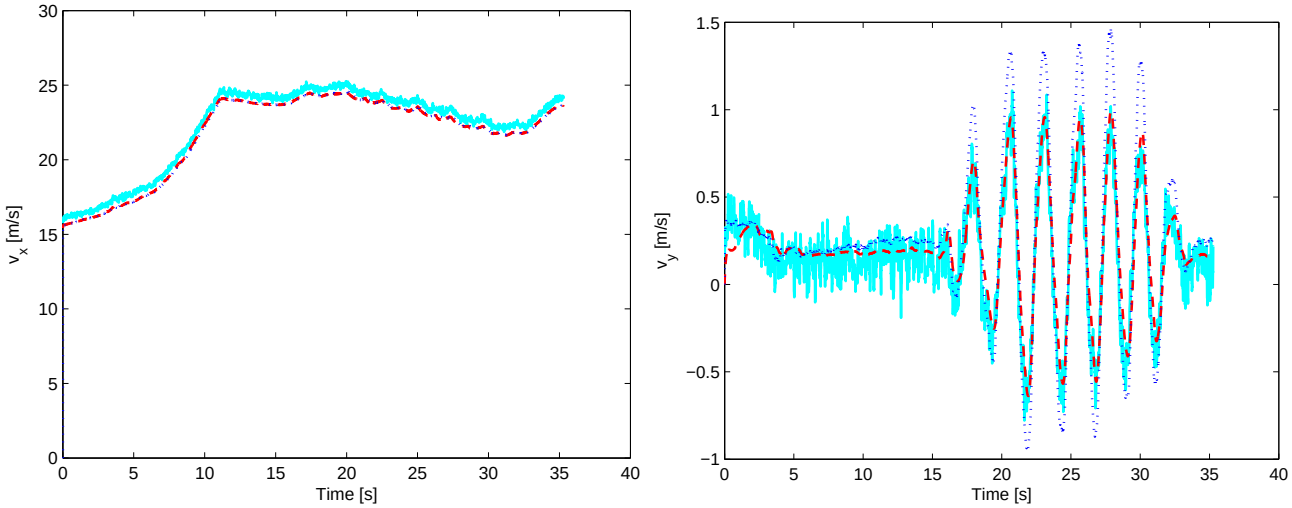

(a) Longitudinal velocity v_x

(b) Lateral velocity v_y

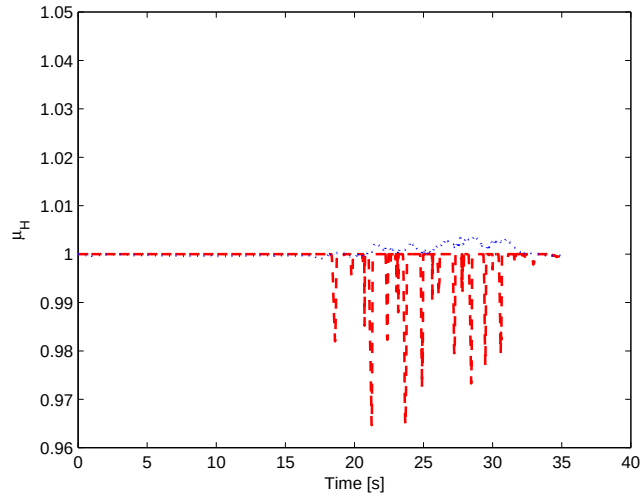

(c) Estimated friction parameter μ_H

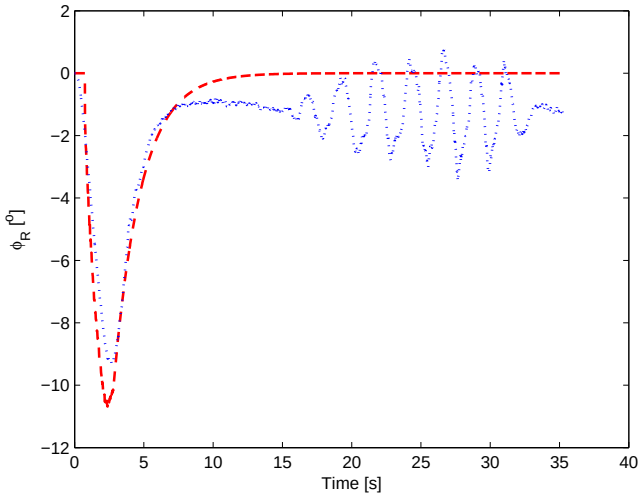

(d) Estimated road bank angle ϕ_R

Figure 2: Results for measurement set 2. The estimates from the nonlinear observer are plotted with a dashed line, while the estimates from the EKF are plotted with a dotted line. Velocity measurements are plotted with a whole line in (a) and (b).

of code.

As a last remark on computational complexity, we note that the differences found above are likely to be more significant for production hardware — typically fixed-point microprocessors — since most of the computation consist of floating-point operations.

OBSERVER TUNING The gains in the nonlinear observer are the following:

- K_{v_y}: A typical value is $K_{v_y} = 1/m$, where m is the vehicle mass. For this particular value, the v_y-part of the observer actually becomes independent of a_y, and therefore not directly influenced by errors in this measurement.
- K_r: A typical value is 40. This is a rather high value, which means that the estimate $\hat{r}$ will follow the measured r closely. To further reduce computational complexity, one may consider removing the estimation of r. Choices regarding the value of K_r and whether estimation of r should be included should be based on the quality of the torque estimate computed by the friction model, and issues related to graceful degradation.
- Γ_1: Used for friction adaptation. A typical value is 4.
- Γ_2: Used for friction adaptation. This is typically rather small, meaning that the yaw rate error has limited direct influence on the v_y and μ_H estimates.
- K_w: Used for road bank angle estimation. A typical value is 1.

The typical values mentioned above are the ones used in the implementation (that is, all the plots) shown in this

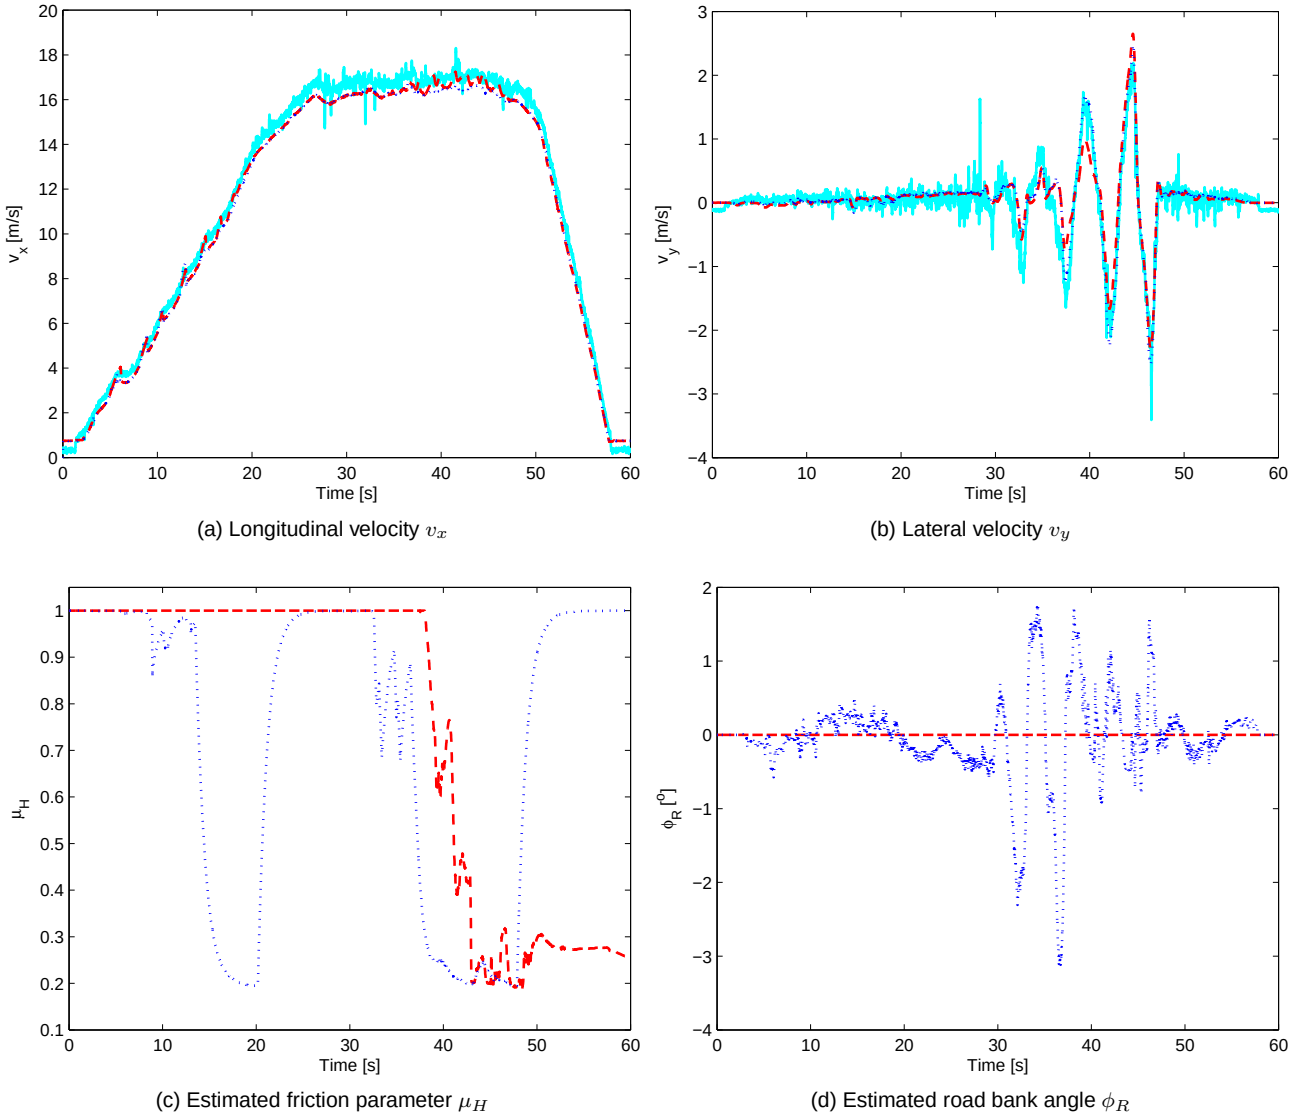

(a) Longitudinal velocity v_x

(b) Lateral velocity v_y

(c) Estimated friction parameter μ_H

(d) Estimated road bank angle ϕ_R

Figure 3: Results for measurement set 3. The estimates from the nonlinear observer are plotted with a dashed line, while the estimates from the EKF are plotted with a dotted line. Velocity measurements are plotted with a whole line in (a) and (b).

paper. Moreover, a scaling factor Λ is used, as explained in [6].

The tuning of the gains in the v_x-part of the observer is discussed in [12, 11]. We do not go into detail on this issue here, since the longitudinal velocity-part of the observer does not have a direct counterpart in the EKF. The EKF calculates a reference speed based on the wheel speed measurements externally (note that this calculation is included in the measurement of the EKF execution time). This can be compared to the logic behind the gains $K_i(t)$ in (5), and we thus keep this out of the tuning comparison.

The EKF estimates 5 states (including road inclination angle, but not r) and uses 4 measurements (a_y, a_x, r and the reference speed v_{ref}) in the update equations. The number of tuning-parameters for the observer is considerably smaller than for an EKF with the same number of estimated states, which needs a 5×5 process noise measurement matrix and a 4×4 measurement matrix. If diagonal measurement and process noise covariances are assumed (as is the case in the EKF implementation considered here), then the number of tuning parameters for the EKF is 9, while the observer has 5.

In addition, the EKF uses considerable logic to change these covariance matrices based on the driving condition and (other) measurements such as a_z, which results in more tuning parameters. The nonlinear observer (excluding the v_x-part) includes logical functions for choosing between adaptation of μ_H and ϕ_R, and monitoring of the PE condition for friction adaptation.

Both the EKF and the nonlinear observer include some linear filtering (mostly low-pass) for signal conditioning.

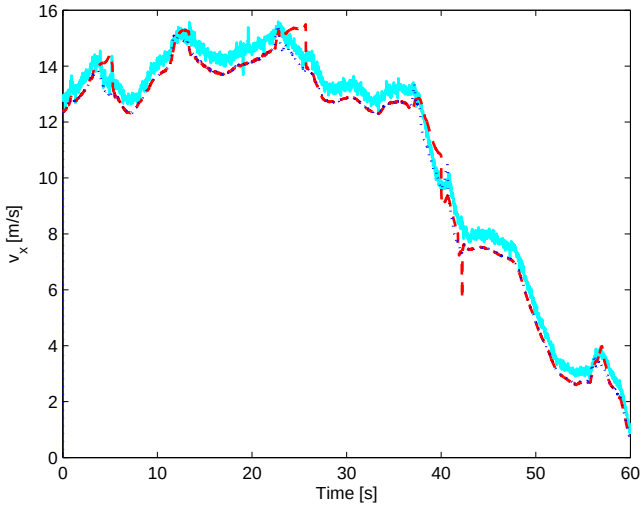

(a) Longitudinal velocity v_x

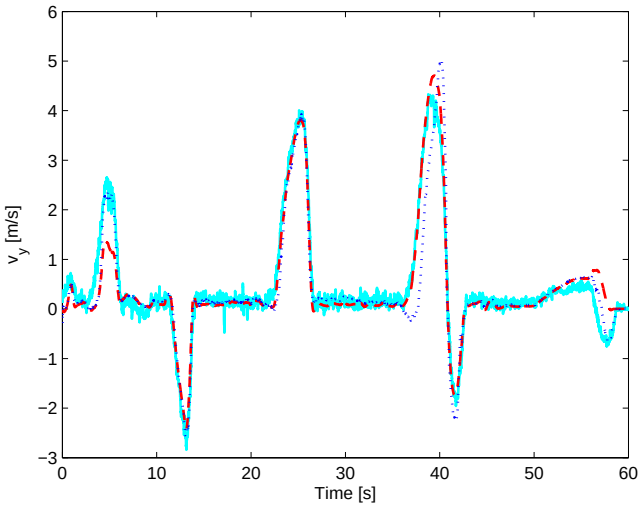

(b) Lateral velocity v_y

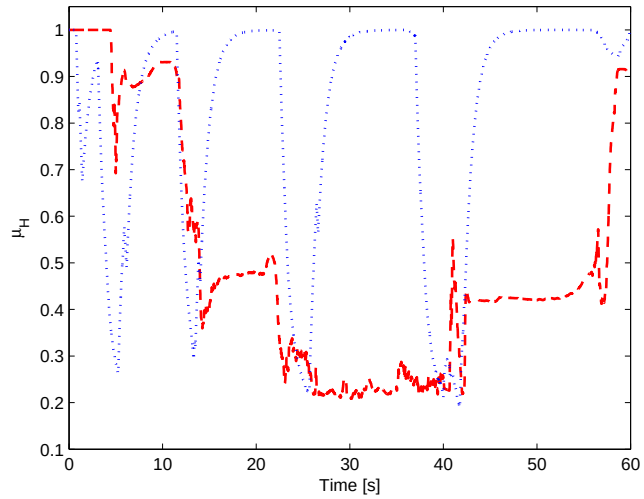

(c) Estimated friction parameter μ_H

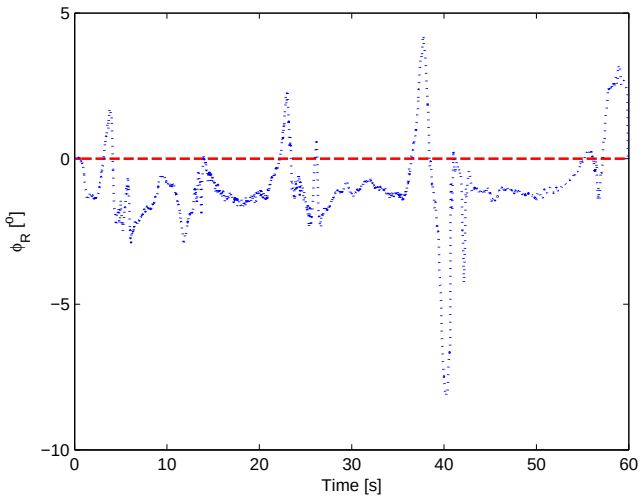

(d) Estimated road bank angle ϕ_R

Figure 4: Results for measurement set 4. The estimates from the nonlinear observer are plotted with a dashed line, while the estimates from the EKF are plotted with a dotted line. Velocity measurements are plotted with a whole line in (a) and (b).

EXPERIMENTAL PERFORMANCE Plots comparing the estimates from the nonlinear observer and the EKF are shown in Figures 1- 5. The test data used covers:

- straight driving and slalom on a high-friction surface with road bank angle (Figures 1 and 2);
- circle/slalom on a low-friction surface (Figure 3);
- sudden steering maneuvers on low-friction surface (Figure 4);
- slalom on a low-friction surface (Figure 5).

The data sets cover driving both on low friction surfaces and on banked roads (although, not at the same time). The data set in Figure 3 covers start-up from zero speed, and coming to a full stop. The general impression is that the performance of the nonlinear observer is as good as, or better, than the EKF.

Slalom maneuvers are generally easiest to estimate, and the observer handles this type of situation well. Driving in a circle is more of a challenge, because circle maneuvers with different radii and velocities can result in the same lateral acceleration. In terms of friction adaptation, PE becomes an issue in this type of situation, as explained earlier. However, the observer also handles most of these situations well, and actually considerably better on real data than in simulations. The main reason for this is that in reality it is hard to drive in perfect circles. Hence, in practical situations there will always be some excitation.

From Figure 4, it is clear that the strategy of letting the friction coefficient estimate go back to one in periods of low excitation differs for the two approaches. The EKF

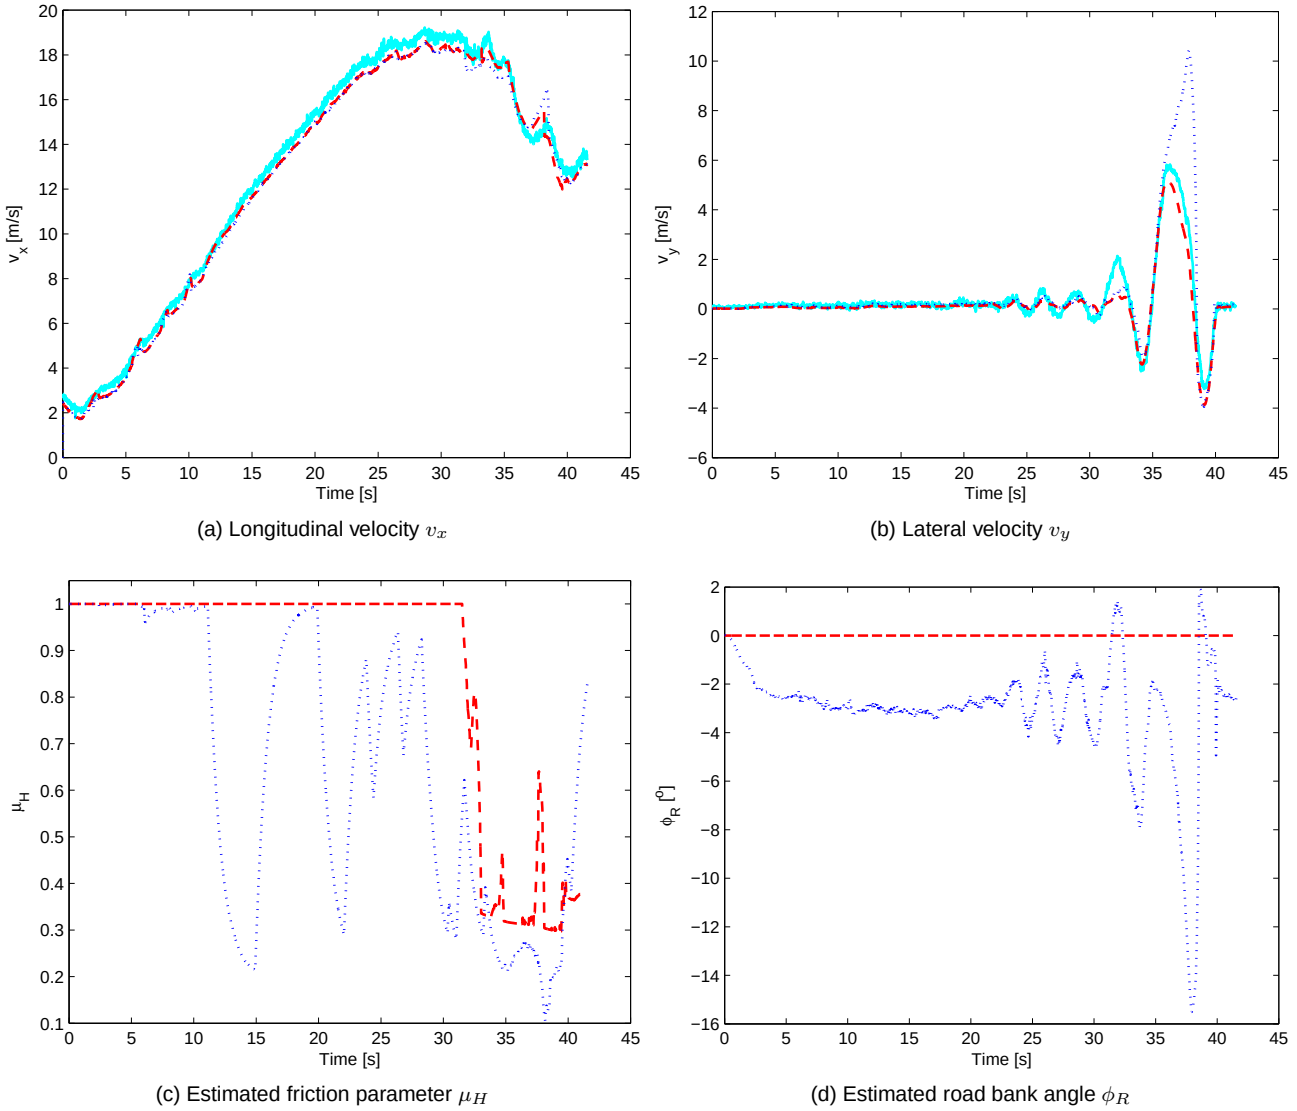

(a) Longitudinal velocity v_x

(b) Lateral velocity v_y

(c) Estimated friction parameter μ_H

(d) Estimated road bank angle ϕ_R

Figure 5: Results for measurement set 5. The estimates from the nonlinear observer are plotted with a dashed line, while the estimates from the EKF are plotted with a dotted line. Velocity measurements are plotted with a whole line in (a) and (b).

rapidly increases the friction coefficient estimate, while in the nonlinear observer it is increased more slowly, based on the excitation criteria. This could be looked upon as a performance vs. robustness issue. It should be mentioned that the EKF uses information from the ESP system for friction estimation, by increasing the gain in the friction adaptation when the ESP flag is set. The variation in the friction coefficient is largely explained from this. The nonlinear observer does not utilize ESP information.

ROBUSTNESS With adaptation of the friction coefficient and road bank angle, the observer becomes robust to changes in these, as can be seen in the plots in the previous section.

Although it seems the observer works fine for the data sets presented here, it should be noted that there is still room for improvement in certain regards. In particular, this is true when it comes to distinguishing between changes in friction and changes in road bank angle. If a road bank angle is not detected quickly enough, the adaptation of the friction coefficient may lead to significant errors in v_y. We can see an indication of this in Figure 1; before the road bank angle is detected, the observer estimates a low friction coefficient. If this continued for a longer period of time, the error in v_y would become large. Using a_z for detection of road bank angle seems to work fine in many cases, but further validation under a wider range of conditions is needed.

To test for robustness with respect to parameters such as mass and position of the CG, we the mass (in the observer) is increased by 200 kg, the yaw moment of inertia is increased by 20%, and the position of the CG is shifted backward by 20% of the length from the front axle. This

has little effect on performance. Figure 6 shows v_y and μ_H for a data set with increasing v_x and circle/slalom maneuver on low friction. The performance is good, only slightly worse than the results obtained using the correct parameters (not plotted). The only notable difference is that the friction adaptation becomes somewhat slower. The EKF is tested with the same parameter errors, and it is hard to see any influence (the error around 57 s is present also for nominal parameters).

Similar observations are made when parameters in the friction model are altered. The observer seems quite robust with respect to limited perturbations in the friction model, but the speed of the friction adaptation can be influenced.

FAULT TOLERANCE AND GRACEFUL DEGRADATION The tests of fault tolerance and grateful degradation are limited to setting the steering wheel angle measurement to zero (Figure 7), and setting the lateral acceleration measurement to zero (Figure 8) for the same data set as above. We consider this data set to be fairly representative.

Although the nonlinear observer misses the first lateral motion (at around 10 s) when the steering wheel angle is set to zero, its performance must be said to be good, and much better than the EKF. The somewhat surprisingly good performance can be attributed to the adaptation of the low friction coefficient. On higher friction surfaces we get larger errors, but other data sets confirm that the nonlinear observer handles this better than the EKF.

When the lateral acceleration measurement is missing, the EKF handles this in a safer manner, in the sense that the velocity estimates become smaller (in absolute value) than the real ones. The EKF also seems to extract more information about the friction coefficient than the nonlinear observer in this case. The reason for this may be higher feedback from the yaw rate measurement to the lateral velocity and friction coefficient estimates.

FREQUENCY DOMAIN COMPARISON It is interesting to compare the performance of the nonlinear observer and the EKF in the frequency domain. To do this, we identify linear models, using measurements of lateral acceleration and yaw rate as inputs, and estimates of the lateral velocity as output. The identification is done using the prediction error method implemented in Matlab's System Identification Toolbox (the pem function). All models identified are of order 12, based on the same data set (a slalom maneuver on a low friction surface).

The identified transfer functions are plotted in Figures 9 and 10. Note that all phases start from -180°, since there is a negative gain from a_y and r to v_y. The need for this negative gain can be deduced from (3) for a_y, and similarly for r.

From the transfer function from a_y to v_y (Figure 9), it appears that the EKF has a significantly higher steady-state/low-frequency gain. However, since the only low-frequency maneuver in practice (and in this data set) is driving along a straight path which means near-zero a_y and v_y, one should probably not attach much significance to this. It is more interesting to note that the gain of the nonlinear observer falls more quickly at high frequencies, indicating a more low-pass nature. Another indication of this, is that the phase of the nonlinear observer falls by 180°, while the phase of the EKF falls only by 90°. Looking closely at the earlier plots, it appears that the v_y estimates from the nonlinear observer are somewhat smoother than the corresponding ones from the EKF.

Since the v_y part of the nonlinear observer has a first-order nature from a_y to v_y, the second-order like transfer function demands some explanation. The likely cause is that for low slips, the term $K_{v_y}\Lambda\xi$ in the equation for $\dot{\hat{v}}_y$ with the present tuning is approximately equal to 1, meaning that a_y is canceled out of the expression. Thus, the lateral acceleration affects v_y mainly through the friction estimate, that is, through two differentiations.

Another interesting observation is that the bandwidths of the nonlinear observer and the EKF are very close, may be slightly higher for the nonlinear observer.

The transfer functions from r to v_y (Figure 10) seem to be rather similar for the EKF and the nonlinear observer. Also here the nonlinear observer and the EKF have similar bandwidths. It is interesting to note that both approaches utilize more high-frequency information from the lateral acceleration measurement than from the yaw rate measurement.

CONCLUSION

A nonlinear observer for vehicle velocity, robust with respect to changes in the road-tire friction and road bank angle, is presented, and validated and compared to an EKF designed for the same purpose.

The nonlinear observer performs well, as well as the EKF and sometimes better, when applied to a fairly wide range of test data sets, some of which are presented here. However, it should be mentioned that the EKF has been more extensively tested, and contains some logic and heuristics to handle special situations robustly, which the nonlinear observer may not have been tested for.

The main advantage of the nonlinear observer lies in its simplicity. This manifests itself in little code and easy-to-understand implementation, and in low computational complexity. In the implementation considered here, the

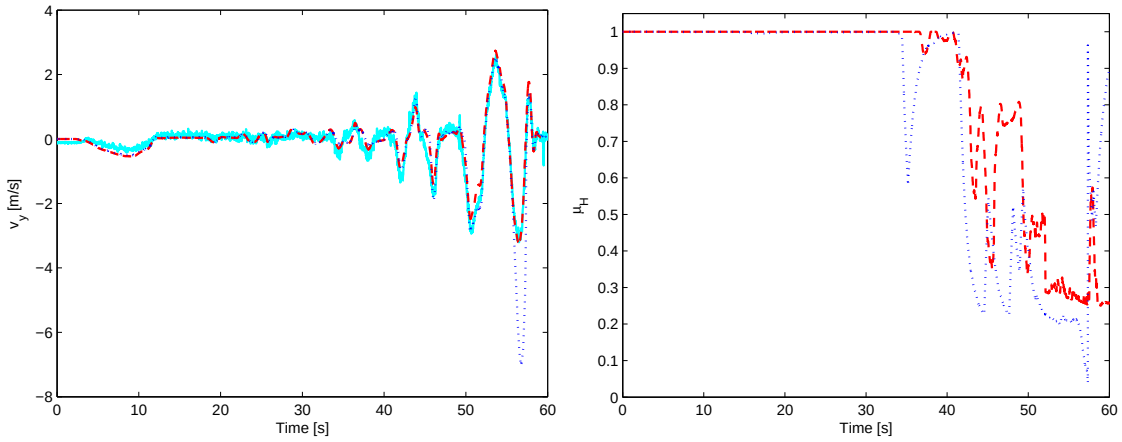

(a) Lateral velocity v_y (b) Estimated friction parameter μ_H

Figure 6: Results for measurement set 6 with parameter errors in the observer and the EKF. The measurements are plotted with a solid line, the nonlinear observer with a dashed line and the EKF with a dotted line.

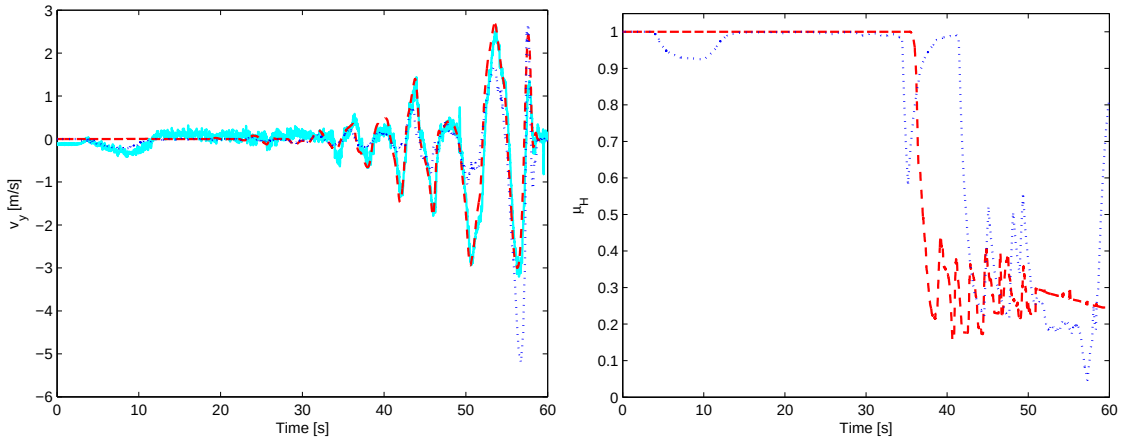

(a) Lateral velocity v_y (b) Estimated friction parameter μ_H

Figure 7: Results for measurement set 6 with steering angle measurement set to zero in the observer and the EKF. The measurements are plotted with a solid line, the nonlinear observer with a dashed line and the EKF with a dotted line.

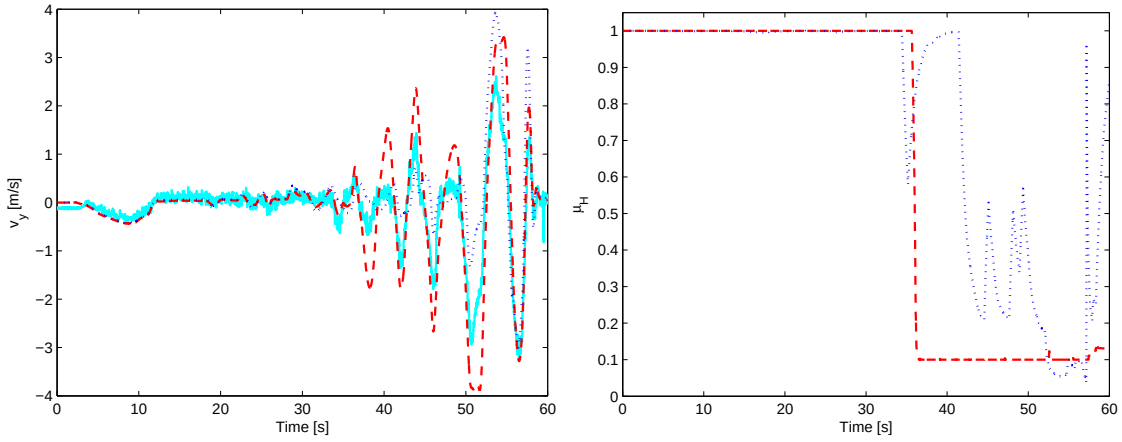

(a) Lateral velocity v_y (b) Estimated friction parameter μ_H

Figure 8: Results for measurement set 6 with lateral acceleration measurement set to zero in the observer and the EKF. The measurements are plotted with a solid line, the nonlinear observer with a dashed line and the EKF with a dotted line.

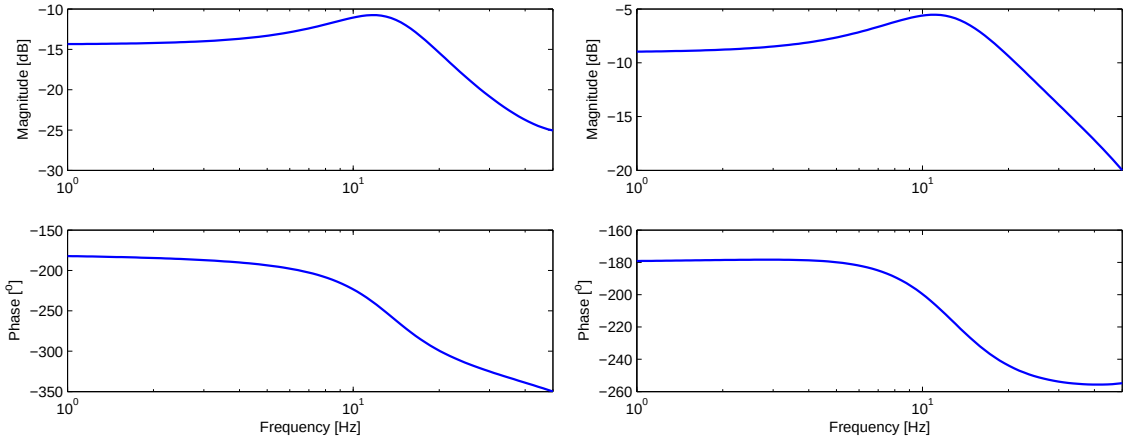

(a) The nonlinear observer (b) The EKF

Figure 9: Transfer function from a_y to v_y for the nonlinear observer and the EKF.

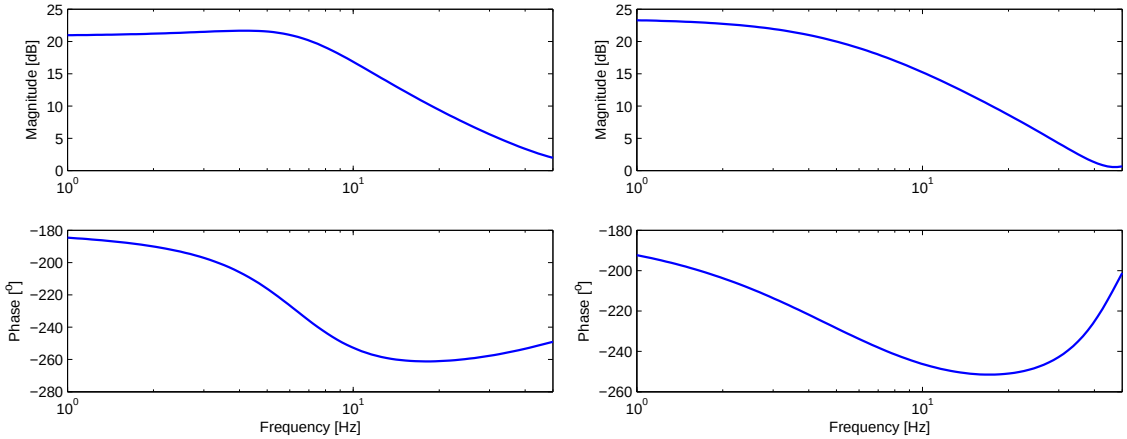

(a) The nonlinear observer (b) The EKF

Figure 10: Transfer function from r to v_y for the nonlinear observer and the EKF.

execution time of the nonlinear observer is about one-third of the execution time of the EKF. These advantages can be crucial when it comes to implementation on production hardware — typically cheap, fixed-point microprocessors.

A further advantage is simplicity of tuning, and in some cases, the analytical stability analysis will be considered an advantage.

ACKNOWLEDGMENTS

This research is supported by the European Commission STREP project CEMACS, contract 004175.

REFERENCES

[1] M. C. Best, T. J. Gordon, and P. J. Dixon. An extended adaptive Kalman filter for real-time state estimation of vehicle handling dynamics. *Vehicle System Dynamics*, 34:57–75, 2000.

[2] G. J. Bierman. *Factorization methods for discrete sequential estimation*. Academic Press [Harcourt Brace Jovanovich Publishers], New York, 1977. Mathematics in Science and Engineering, Vol. 128.

[3] J. Farrelly and P. Wellstead. Estimation of vehicle lateral velocity. In *Proceedings of the 1996 IEEE International Conference on Control Applications*, pages 552–557, 1996.

[4] Y. Fukada. Slip-angle estimation for stability control. *Vehicle Systems Dynamics*, 32:375–388, 1999.

[5] A. Gelb, editor. *Applied optimal estimation*. The MIT Press, Cambridge, Mass.-London, 1974. Written by the technical staff of The Analytic Sciences Corporation, Principal authors: Arthur Gelb, Joseph F. Kasper, Jr., Raymond A. Nash, Jr., Charles F. Price and Arthur A. Sutherland, Jr.

[6] H. F. Grip, L. Imsland, T. A. Johansen, T. I. Fossen, J. C. Kalkkuhl, and A. Suissa. Nonlinear vehicle sideslip estimation with friction adaptation. *Automatica*, 2006. Submitted.

[7] H. F. Grip, L. Imsland, T. A. Johansen, T. I. Fossen, J. C. Kalkkuhl, and A. Suissa. Nonlinear vehicle velocity observer with road-tire friction adaptation. In *Proc. 45th IEEE Conf. Decision Contr.*, San Diego, CA, 2006.

[8] A. Hac and M. D. Simpson. Estimation of vehicle side slip angle and yaw rate. In *SAE 2000 World Congress*, Detroit, MI, USA, 2000.

[9] M. Hiemer, A. von Vietinghoff, U. Kiencke, and T. Matsunaga. Determination of vehicle body slip angle with non-linear observer strategies. In *Proceedings of the SAE World Congress*, 2005. Paper no. 2005-01-0400.

[10] L. Imsland, H. F. Grip, T. A. Johansen, and T. I. Fossen. On nonlinear unknown input observers—applied to lateral vehicle velocity estimation on banked roads. *Int. J. Contr.*, 2006. Submitted.

[11] L. Imsland, T. A. Johansen, T. I. Fossen, H. F. Grip, J. C. Kalkkuhl, and A. Suissa. Vehicle velocity estimation using nonlinear observers. *Automatica*, 42(12), 2006.

[12] L. Imsland, T. A. Johansen, T. I. Fossen, J. Kalkkuhl, and A. Suissa. Nonlinear observer for vehicle velocity estimation. In *Proc. SAE World Congress*, 2006.

[13] U. Kiencke and A. Daiss. Observation of lateral vehicle dynamics. *Control Engineering Practice*, 5(8):1145–1150, 1997.

[14] U. Kiencke and L. Nielsen. *Automotive Control Systems*. Springer, 2000.

[15] J. Lu and T. A. Brown. Vehicle side slip angle estimation using dynamic blending and considering vehicle attitude information. Patent US 6671595, 2003.

[16] H. B. Pacejka. *Tyre and vehicle dynamics*. Butterworth-Heinemann, 2002.

[17] L. R. Ray. Nonlinear state and tire force estimation for advanced vehicle control. *IEEE Transactions on Control Systems Technology*, 3(1):117–124, 1995.

[18] L. R. Ray. Nonlinear tire force estimation and road friction identification: simulation and experiments. *Automatica*, 33(10):1819–1833, 1997.

[19] A. Suissa, Z. Zomotor, and F. Böttiger. Method for determining variables characterizing vehicle handling. Patent US 5557520, 1996.

[20] A. Y. Ungoren, H. Peng, and H. Tseng. A study on lateral speed estimation methods. *Int. J. Vehicle Autonomous Systems*, 2(1/2):126–144, 2004.

[21] A. T. van Zanten. Bosch ESP system: 5 years of experience. In *In Proceedings of the Automotive Dynamics & Stability Conference (P-354)*, 2000. Paper no. 2000-01-1633.

[22] P. J. T. Venhovens and K. Naab. Vehicle dynamics estimation using Kalman filters. *Vehicle System Dynamics*, 32:171–184, 1999.