

Vehicle Sideslip Estimation

Design, implementation, and experimental validation

Håvard Fjær Grip, Lars Imsland, Tor A. Johansen,
Jens C. Kalkkuhl, and Avshalom Suissa

Control systems that help the driver avoid accidents, or limit the damage in case of an accident, have become ubiquitous in modern passenger cars. For example, new cars typically have an anti-lock braking system (ABS), which prevents the wheels from locking during hard braking, and they often have an electronic stability control system (ESC), which stabilizes the lateral motion of the vehicle to prevent skidding. Collision warning and avoidance, rollover prevention, crosswind stabilization, and preparation for an impending accident by adjusting seat positions and seat belts are additional examples of control systems for automotive safety.

These systems rely on information about the state of the vehicle and its surroundings. To obtain this information, modern cars are equipped with various sensors. For a typical car with an ESC system, necessary measurements include the steering wheel angle, wheel angular velocities, lateral acceleration, and the rate of rotation around the vertical body-fixed axis, known as the *yaw rate*. These measurements alone contain a great deal of information about the state of the vehicle. The speed of the car can be estimated using the wheel angular velocities, and a linear reference model taking the speed, steering wheel angle, and additional measurements as inputs can be used to predict the behavior of the car under normal driving conditions. The predicted behavior can be compared to the actual behavior of the car; ESC systems, for example, use the brakes to correct the deviation from a yaw reference model when the vehicle starts to skid [1].

Although some quantities are easily measured, others are difficult to measure because of high cost or impracticality. When some quantity cannot be measured directly, it is often necessary to estimate it using the measurements that are available. Observers combine the available measurements with dynamic models to estimate unknown dynamic states. Often, dynamic models of sufficient accuracy are not available, and must be carefully constructed as part of the observer design. The observer estimates can be used to implement control algorithms, as Figure 1 illustrates.

A. Vehicle Sideslip Angle

When a car is driven straight on a flat surface, the direction of travel at the center of gravity (CG) remains the same as the orientation of the vehicle (that is, the direction of the longitudinal axis). When the car turns, however, it exhibits a yaw rate, causing the orientation to change, and a lateral acceleration directed toward the center of the turn. The car also exhibits a velocity component perpendicular to the orientation, known as the lateral velocity. Nonzero lateral velocity means

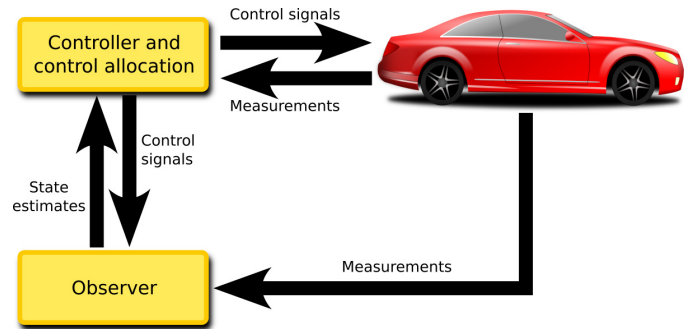


Fig. 1. The relationship between the car, an observer, and a controller with control allocation. The observer combines measurements from the car with a dynamic model, to provide estimates of unmeasured states to the controller.

that the orientation of the vehicle and the direction of travel are no longer the same. The lateral velocity differs from one point on the vehicle body to another; as a point of reference, we use the vehicle's CG.

The angle between the orientation of the vehicle and the direction of travel at the CG is called the *vehicle sideslip angle*. In production cars, the vehicle sideslip angle is not measured, because this measurement requires expensive equipment such as optical correlation sensors. Under normal circumstances, when the car is driven safely without danger of losing road grip, the vehicle sideslip angle is small, not exceeding $\pm 2^\circ$ for the average driver [1]. Moreover, for a given speed in normal driving situations, the steering characteristics specify a tight connection between the steering wheel angle, yaw rate, lateral acceleration, and vehicle sideslip angle. The vehicle sideslip angle can therefore be estimated using a static model or a simple linear, dynamic model. In extreme situations, however, when the vehicle is pushed to the physical limits of adhesion between the tires and the road surface, the behavior of the car is highly nonlinear, and the tight coupling of the vehicle sideslip angle to various measured quantities is lost. This behavior is due to the nonlinearity of the friction forces between the tires and the road surface. In these situations, the vehicle sideslip angle can become large, and knowledge about it is essential for a proper description of vehicle behavior. Accurate online estimation of the vehicle sideslip angle has the potential to enable development of new automotive safety systems, and to improve existing algorithms that use information about the vehicle sideslip angle, such as ESC.

B. Previous Research

Many designs for estimating velocity and sideslip angle are found in the literature. These designs are typically based on linear or quasi-linear techniques [2]–[5]. A nonlinear observer linearizing the observer error dynamics is presented in [6], [7]. An observer based on forcing the dynamics of the nonlinear estimation error to follow the dynamics of a linear reference system is investigated in [8], [9]. In [8], [9] availability of the longitudinal road-tire friction forces is assumed, whereas in [10] the longitudinal forces used in the estimation algorithm are calculated from the brake pressure, clutch position, and throttle angle. An extended Kalman filter (EKF) based on a road-tire friction model, which includes estimation of a road-tire friction coefficient, the inclination angle, and the bank angle, is developed in [11]. The term inclination angle refers to the road sloping upward or downward along the orientation of the vehicle; bank angle refers to the road sloping toward the left or the right. Alternative EKFs are used in [12], which estimates velocity and tire forces without the explicit use of a road-tire friction model; and in [13], which is based on a linear model of the road-tire friction forces with online estimation of road-tire friction parameters. In [14] a linear observer for the vehicle velocity is used as an input to a Kalman filter based on the kinematic equations of motion. In [15] the sideslip angle is estimated along with both the yaw rate and a road-tire friction coefficient without the use of a yaw rate measurement.

A majority of designs, including [1]–[15], are based on a vehicle model, usually including a model of the road-tire friction forces. The main argument against using such a model is its inherent uncertainty. Changes in the loading of the vehicle and the tire characteristics, for example, introduce unknown variations in the model. A different direction is taken in [16], where a six-degree-of-freedom inertial sensor cluster is used in an EKF that relies mainly on open-loop integration of the kinematic equations, without a vehicle or friction model. The main measurement equation comes from the longitudinal velocity, which is calculated separately based primarily on the wheel speeds. This idea is similar to the kinematic observer approach in [2]. Designs of this type are sensitive to unknown *sensor bias* and *drift*, as well as misalignment of the sensor cluster. Sensor bias refers to a constant error in the measurement signal; drift refers to a slowly varying error. Drift in inertial sensors is primarily caused by variations in temperature.

To estimate the vehicle sideslip angle in real-world situations, some information about the surroundings of the vehicle is usually needed. In particular, the road bank angle has a significant effect on the lateral velocity. Furthermore, when a design includes modeling of the road-tire friction forces, as in [1]–[11], [13]–[15], information about the road surface conditions is needed. When estimating the longitudinal velocity, knowledge of the inclination angle is useful when the wheel speeds fail to provide high-quality information. Some model-based designs, such as [3], [5], [11], [12], [13], [15], take unknown road surface conditions into account, but it is more commonly assumed that the road surface conditions are known, or that the vehicle is driven in a way that minimizes

the impact of the road surface conditions. Similarly, although a horizontal road surface is typically assumed, [3], [5], [11], [16] consider inclination and bank angles. Estimation of the road bank angle is considered in [17], which is based on transfer functions from the steering angle and road bank angle to the yaw rate and lateral acceleration, and in [18], which uses an EKF to estimate the sideslip angle, which is in turn used in a linear unknown-input observer to estimate the road bank angle.

Inertial measurements are combined with GPS measurements in [19]–[23] to provide estimates of the sideslip angle. Designs of this type often include estimates of the bank angle and inclination angle.

C. Goal of this Article

The goal of this article is to develop a vehicle sideslip observer that takes the nonlinearities of the system into account, both in the design and theoretical analysis. Design goals include a reduction of computational complexity compared to the EKF in order to make the observer suitable for implementation in embedded hardware, and a reduction in the number of tuning parameters compared to the EKF. The design is based on a standard sensor configuration, and is subjected to extensive testing in realistic conditions.

Parts of the theoretical foundation for the observer design, as well as some preliminary experimental results, are found in [24], which assumes known road-tire friction properties and a horizontal road surface, and in [25], where the approach from [24] is extended to take unknown road surface conditions into account. In [26] a comparison is made between an EKF and an observer modified from [24], [25], with an added algorithm for bank-angle estimation developed in [27]. In the present article, we extend the approach of [24], [25] by estimating the inclination and bank angles. We refer to the resulting observer as the *nonlinear vehicle sideslip observer* (NVSO). The NVSO has undergone extensive and systematic testing in a variety of situations. Tests have been performed on both test tracks and normal roads, on high- and low-friction surfaces, and with significant inclination and bank angles. Based on the results of this testing, we discuss strengths and weaknesses of the NVSO design, as well as practical implementation and tuning. We compare the experimental results with those of an EKF. Since much of the discussion concerns modeling accuracy, observability, and constraints imposed by the sensor configuration, the results and observations are applicable in a wider sense to alternative model-based designs with similar sensor configurations.

D. Sensor Configuration

A crucial design consideration in estimation problems of this type is the choice of sensor configuration. Typical automotive-grade sensor configurations are characterized by the need to minimize costs, which often translates into sensors with low resolution, narrow range, and significant noise, bias, and drift. The availability and quality of sensors place fundamental constraints on the accuracy that can be expected from any estimation scheme. Bias and drift in inertial sensors is

particularly limiting, because it prohibits accurate integration of kinematic equations, except over short time spans.

We focus on a standard sensor configuration found in modern cars with an ESC system, consisting of measurements of the longitudinal and lateral accelerations, the yaw rate, the steering wheel angle, and the wheel speeds. Although GPS measurements would be a valuable addition to the measurements mentioned above, GPS navigation systems are not yet standard equipment, even in high-end passenger cars. Moreover, GPS signals are sometimes unavailable or degraded, for example, when driving through tunnels, under bridges, or near large structures such as steep mountains or tall buildings. We therefore do not consider GPS measurements to be available.

The ESC-type sensor configuration lacks measurements of the vertical acceleration and the angular rates around the vehicle's longitudinal and lateral axes, known as the *roll* and *pitch rates*, respectively. In addition, ESC-type sensors may have significant bias and drift, which prohibits a design based exclusively on kinematic equations, such as [16]. We therefore supplement the kinematic equations with a vehicle model that includes a model of the road-tire friction forces.

I. VEHICLE MODEL

For a car driving on a horizontal surface, the longitudinal and lateral velocities at the CG are governed by the equations of motion [7]

$$\dot{v}_x = a_x + \dot{\psi}v_y, \quad (1)$$

$$\dot{v}_y = a_y - \dot{\psi}v_x, \quad (2)$$

where v_x and v_y are the longitudinal and lateral velocities, a_x and a_y are the longitudinal and lateral accelerations, and $\dot{\psi}$ is the yaw rate. Each of the equations (1), (2) includes an acceleration term and a term involving the yaw rate $\dot{\psi}$. The accelerations are directly related to the forces acting on the vehicle, attributable mainly to road-tire friction. The terms involving the yaw rate appear because the coordinate system in which we resolve the velocities and accelerations is fixed to the car, which rotates with respect to an inertial coordinate system. The vehicle is illustrated in Figure 2, where the velocities and yaw rate are shown together with the vehicle sideslip angle β , which is given as $\beta = \arctan(v_y/v_x)$.

A. Road-Tire Friction

When the driver turns the steering wheel to make a regular turn, the tires on the front axle of the car become misaligned with the direction of travel, and we obtain a *tire-slip angle*. The tire-slip angle is conceptually similar to the vehicle sideslip angle, except that the relevant frame of reference is associated with a single tire rather than the vehicle body. In particular, the tire-slip angle is defined as the angle between the velocity vector at the center of the wheel and the orientation of the tire. This definition is illustrated for the front left wheel in Figure 2, where α_1 denotes the tire-slip angle. A nonzero tire-slip angle implies a nonzero velocity in the lateral direction of the tire, relative to the road surface. Because the tire is elastic, however, it does not simply slide laterally over the road surface. As a point on the tire tread rolls into contact




Fig. 2. Illustration of the vehicle velocity, yaw rate, vehicle sideslip angle, tire-slip angle, and road-tire friction forces. The vehicle velocity vector v at the CG is decomposed into the longitudinal velocity v_x , which is positive in the forward direction, and the lateral velocity v_y , which is positive in the leftward direction. The angle β between the longitudinal x -axis and the velocity vector v is the vehicle sideslip angle, which is defined as positive in the counterclockwise direction. For the front left wheel, the velocity vector v_1 at the wheel center is shown. The angle α_1 between v_1 and the longitudinal axis of the tire is the tire-slip angle, which is defined as positive in the clockwise direction. The positive tire-slip angle shown in the figure generates a positive lateral tire force F_{y1} . The longitudinal tire force F_{x1} is generated by a nonzero longitudinal tire slip. Together, F_{x1} and F_{y1} make up the tire force vector F_1 . Similar road-tire friction forces are also generated for the other three tires.

with the surface, its path is deflected and it briefly grips the surface as it passes through the contact patch. This process results in a deformation of the tire, and the tire's resistance to this deformation generates lateral forces that lead the car to start turning [28, Ch. 6]. As the car starts to turn, tire-slip angles are also built up for the rear tires. After an initial transient, a steady state is reached where the road-tire friction forces balance to give zero net moment on the vehicle body, as well as constant lateral acceleration and lateral velocity.

We also define the longitudinal tire slip, as the normalized difference between the circumferential speed of the tire and the speed of the wheel center along the orientation of the tire. Longitudinal tire slip gives rise to longitudinal friction forces, as illustrated in Figure 2 by F_{x1} . Collectively, we refer to the longitudinal tire slip and lateral tire-slip angle as the tire slips.

During normal driving the road-tire friction forces are

approximately linear with respect to the tire slips. The lateral friction forces are then modeled as $F_{y1} = C_y \alpha_1$, where the constant C_y is the *cornering stiffness*. In extreme situations, however, the tire slips may become so large that this linearity is lost. Beyond a certain point, the tire loses road grip and the road-tire friction forces begin to saturate, meaning that an increase in the tire slips does not result in a corresponding increase in the friction forces. This effect can be seen in Figure 3, where the lateral road-tire friction force is plotted against the tire-slip angle after being normalized by dividing it by the vertical contact force. The curves in Figure 3 correspond to varying degrees of road grip due to various road surface conditions, ranging roughly from ice for the lower curve to dry asphalt for the upper one. The road surface conditions are represented by a friction coefficient μ_H , which is lower for more slippery surfaces. The region where the friction forces are linear with respect to the tire-slip angle is called the *linear region*. Beyond the linear region, where the tire-slip angles are larger, is the *nonlinear region*.

When the friction forces saturate for the front tires before they do so at the rear, the car stops responding properly to steering inputs, a situation known as understeer or plowing. Conversely, when the friction forces saturate for the rear tires first, the car becomes severely oversteered and may become unstable, a situation known as fishtailing. The point at which severe under- or oversteer occurs depends on the driving maneuver, the tires, the design of the vehicle, and the properties of the road surface. As Figure 3 makes clear, road-tire friction forces saturate sooner on low-friction surfaces, such as ice, than on high-friction surfaces, such as asphalt. For a more comprehensive description of road-tire friction, and more detailed definitions, see [7], [29].

1) *Modeling the Road-Tire Friction Forces*: As with many alternative designs for vehicle sideslip estimation [1]–[11], [13]–[15], a key component of the NVSO is a road-tire friction model. The road-tire friction model takes measurements and observer estimates as inputs, and returns estimates of the road-tire friction forces. The expression $F_{y1} = C_y \alpha_1$ is an example of a linear friction model, but to account for the nonlinearity of road-tire friction forces for large tire slips, a nonlinear model is needed. A widely used nonlinear road-tire friction model is the *magic formula* [29]. Like linear road-tire friction models, nonlinear models such as the magic formula are based on tire slips. But instead of increasing linearly with the tire slips, the friction forces in nonlinear models are made to follow curves similar to those seen in Figure 3.

The NVSO is not designed with a particular road-tire friction model in mind. Instead, we use a nonlinear friction model satisfying specific physical properties, and we base the design and analysis on these properties. The curves in Figure 3 are created with a friction model for which the physical road-tire friction curves are approximated by fractional polynomial expressions, similar to [7, Ch. 7.2].

B. The Lateral Acceleration

According to Newton's second law, the vehicle acceleration along each direction is equal to the total force acting on the

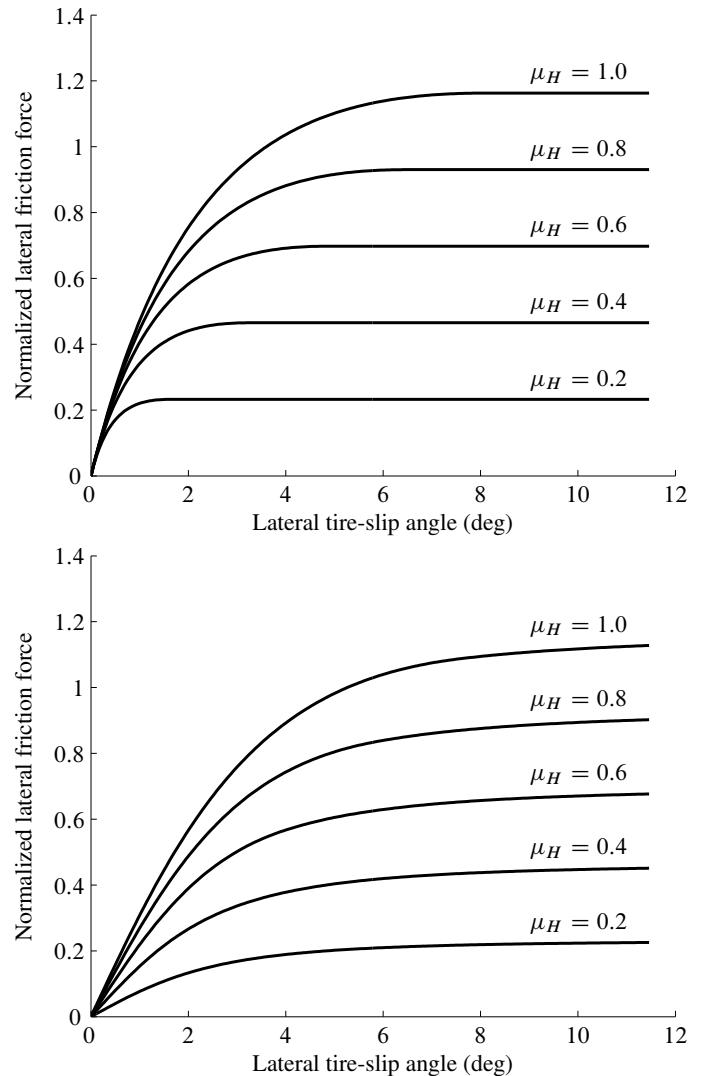


Fig. 3. Nonlinear road-tire friction curves. The lateral tire force, which is normalized by dividing it by the vertical contact force, is plotted against the lateral tire-slip angle for various road surface conditions. Where the tire-slip angles are small, the forces increase linearly, and we call this the linear region. For larger tire-slip angles, in the nonlinear region, the forces saturate. The curves to the left are created with the longitudinal tire slip kept at zero. The curves to the right are created with a large longitudinal tire slip.

vehicle in that direction, divided by the mass. When the road surface is slanted, rather than horizontal, gravity acts on the vehicle in the tangent plane of the road surface, which affects the vehicle velocity. Assume for now that the road surface is horizontal. The dominant forces acting in the plane are the road-tire friction forces; we ignore smaller influences such as wind and air resistance. The road-tire friction forces are algebraic functions of the tire slips, which in turn are algebraic functions of the vehicle velocity. Consequently, measurements of the vehicle accelerations depend algebraically on the vehicle velocities, and the accelerations can therefore be used as indirect measurements of the velocities. In particular, the lateral acceleration a_y contains valuable information about the lateral velocity.

To see how the relationship between the lateral acceleration a_y and the lateral velocity v_y can be used, we consider what

happens to the lateral road-tire friction forces when we perturb v_y while keeping everything else constant. From Figure 2, we see that if v_y is increased, the lateral component of the velocity vector v_1 at the center of the front left wheel is also increased, and the tire-slip angle α_1 for the front left tire is therefore decreased (note that the sign convention for the tire-slip angle is opposite from the sign convention for the vehicle sideslip angle). This decrease in the tire-slip angle leads to a decrease in the lateral road-tire friction force F_{y1} , as indicated by the shape of the friction curves in Figure 3. Put differently, F_{y1} is perturbed in a negative direction in response to a positive perturbation in v_y . The lateral road-tire friction forces for the remaining three tires are affected in the same way by a perturbation in v_y .

The above discussion suggests that the total lateral force on the vehicle, and therefore the lateral acceleration a_y , is perturbed in a negative direction in response to a positive perturbation in v_y . Indeed, it is demonstrated in [24] that, except in some particular cases that we discuss below, the partial derivative $\partial a_y / \partial v_y$ is less than some negative number when a_y is considered as a function of v_y . We therefore consider a_y to be strictly decreasing with respect to v_y . This monotonicity is a key property in the NVSO design.

1) *Roll Correction:* During a turn, the vehicle body is typically at a slight roll angle ϕ around the x -axis, relative to the road surface. The roll angle is roughly proportional to the lateral acceleration and thus can be approximated as $\phi \approx p_\phi a_y$, where the constant p_ϕ is the roll-angle gradient and a_y is the lateral acceleration in the tangent plane of the road surface, as above. The roll angle ϕ causes the lateral acceleration measurement to be influenced by an additive gravity component $\sin(\phi)g \approx \phi g$, where g is the acceleration of gravity. The measured lateral acceleration is therefore approximately $(1 + p_\phi g)a_y$. The gravity influence due to roll angle is undesirable, and to remove it we divide the acceleration measurement by the constant factor $1 + p_\phi g$. Throughout the article, we therefore assume that a_y is a measurement that is not influenced by the roll angle.

II. OBSERVER DESIGN

To estimate the vehicle sideslip angle, we estimate the longitudinal and lateral velocities v_x and v_y at the CG, and then calculate the vehicle sideslip angle from the velocities using the expression $\beta = \arctan(v_y/v_x)$. We assume for now that there is no uncertainty in the friction model regarding the road surface conditions. We discuss the estimation of longitudinal and lateral velocity separately, ignoring at first the coupling between the two estimates. Throughout the rest of this article, a hat indicates an estimated quantity, and a tilde indicates an estimation error; thus \hat{v}_x is an estimate of v_x , and $\tilde{v}_x := v_x - \hat{v}_x$. The goal is to stabilize the estimation error and make it vanish asymptotically, for example, by making the origin of the observer error dynamics asymptotically or exponentially stable.

A. Longitudinal-Velocity Estimation

A conventional speedometer approximates the longitudinal vehicle velocity by using the wheel speeds. The wheel

speeds usually provide good measurements of the longitudinal velocity, but the accuracy is sometimes severely reduced, for example, during strong acceleration or braking, and in situations where one or more of the wheels are spinning or locked.

To estimate the longitudinal velocity we use the observer

$$\dot{\hat{v}}_x = a_x + \dot{\psi} \hat{v}_y + K_{v_x}(t)(v_{x,\text{ref}} - \hat{v}_x), \quad (3)$$

where \hat{v}_y is an estimate of the lateral velocity as specified below, $K_{v_x}(t)$ is a time-varying gain, and $v_{x,\text{ref}}$ is a reference velocity calculated from the wheel speeds, which plays the role of measurement. The observer (3) is a copy of the system equation (1) with an added injection term $K_{v_x}(t)(v_{x,\text{ref}} - \hat{v}_x)$. When $v_{x,\text{ref}}$ represents the true longitudinal velocity, the error dynamics obtained by subtracting (3) from (1) are

$$\dot{\tilde{v}}_x = \dot{\psi} \tilde{v}_y - K_{v_x}(t) \tilde{v}_x. \quad (4)$$

Ignoring for the moment the error \tilde{v}_y , global exponential stability of the origin of (4) can be verified if $K_{v_x}(t)$ is larger than some positive number, that is, $K_{v_x}(t) \geq K_{v_x,\text{min}} > 0$. The primary challenge therefore lies in producing a reference velocity $v_{x,\text{ref}}$ and selecting a sensible gain $K_{v_x}(t)$.

1) *Reference Velocity and Gain:* The four wheel speeds, together with the steering wheel angle and the yaw rate, can be used to create four separate measurements of v_x . It is possible to create a reference velocity by taking a weighted average of these four measurements, with weightings determined by various factors, such as the spread in the four measurements and whether the car is braking or accelerating [24], [25]. To make speed estimation as reliable as possible in every situation, however, car manufacturers have developed sophisticated reference velocity algorithms that use information from a variety of sources, such as the accelerations, brake pressures, engine torque, and status information from ABS, ESC, and drive-slip control systems. The goal is to determine which wheel speeds provide accurate information about the longitudinal velocity at any given time, for example, during heavy braking or acceleration. We therefore assume that a reference velocity is available to the observer.

No matter how much effort is put into the reference velocity algorithm, the accuracy of $v_{x,\text{ref}}$ varies. We take this variation into account through the time-varying gain $K_{v_x}(t)$, which is intended to reflect the accuracy of the reference velocity. When the accuracy is deemed to be low, the gain $K_{v_x}(t)$ is reduced to make the observer less reliant on the reference velocity and more reliant on integration of the system equations. In creating $K_{v_x}(t)$, all of the sources of information that are used to create $v_{x,\text{ref}}$ can be used. A measure of the accuracy of $v_{x,\text{ref}}$ may indeed be naturally available from the inner workings of the reference velocity algorithm. For the experimental results presented in this article, we use a simple algorithm that is based on reducing the nominal gain depending on the variance of the four measurements, on the theory that agreement between the measurements indicates high accuracy. Design of $K_{v_x}(t)$ is heuristic and must therefore be empirically based.

B. Lateral-Velocity Estimation

The wheel speeds can be used to create measurements of the longitudinal velocity. For the lateral velocity, however, we do not have a similar source of information, and estimating lateral velocity is therefore more difficult. We start by introducing $\hat{a}_y(t, \hat{x})$, which denotes an estimate of the lateral acceleration a_y . We use \hat{x} to denote the vector of estimated velocities \hat{v}_x and \hat{v}_y , and x to denote the vector of actual velocities v_x and v_y . The estimate $\hat{a}_y(t, \hat{x})$ is formed by using the nonlinear friction model for each wheel, where measurements of the steering wheel angle, yaw rate, and wheel speeds, as well as the estimated velocities \hat{v}_x and \hat{v}_y , are used as inputs. The friction forces modeled for each wheel are added up in the lateral direction of the vehicle and divided by the mass, resulting in the lateral acceleration estimate $\hat{a}_y(t, \hat{x})$. The time argument in $\hat{a}_y(t, \hat{x})$ denotes the dependence of \hat{a}_y on time-varying signals such as the steering wheel angle and yaw rate. We also define $\tilde{a}_y(t, \tilde{x}) := a_y - \hat{a}_y(t, \hat{x})$, where $\tilde{x} := x - \hat{x}$. To write $\tilde{a}_y(t, \tilde{x})$ as a function of t and \tilde{x} , we replace \hat{x} with $x - \tilde{x}$ and absorb x in the time argument.

To estimate the lateral velocity we use the observer

$$\dot{\hat{v}}_y = a_y - \dot{\psi}\hat{v}_x - K_{v_y}(a_y - \hat{a}_y(t, \hat{x})), \quad (5)$$

where K_{v_y} is a positive gain. The observer consists of a copy of the original system (2) and an injection term $-K_{v_y}(a_y - \hat{a}_y(t, \hat{x}))$. The error dynamics obtained by subtracting (5) from (2) are

$$\dot{\tilde{v}}_y = -\dot{\psi}\tilde{v}_x + K_{v_y}\tilde{a}_y(t, \tilde{x}). \quad (6)$$

To see why the injection term asymptotically stabilizes the observer error, assume for the moment that the longitudinal-velocity estimate \hat{v}_x is exact, meaning that the lateral-velocity estimate \hat{v}_y is the only uncertain input to $\hat{a}_y(t, \hat{x})$. Assuming furthermore that the friction model is continuously differentiable, the assumption that a_y is strictly decreasing with respect to v_y means that we can use the mean value theorem to write

$$\tilde{a}_y(t, \tilde{x}) = -\eta(t, \tilde{x})\tilde{v}_y, \quad (7)$$

where $\eta(t, \tilde{x}) \geq \eta_{\min} > 0$ for some constant η_{\min} . The error dynamics can therefore be written as $\dot{\tilde{v}}_y = -K_{v_y}\eta(t, \tilde{x})\tilde{v}_y$. It is then easy to verify that the origin of the error dynamics is globally exponentially stable.

We have so far ignored the coupling between the longitudinal and lateral observer components. For each observer component, a quadratic Lyapunov function can be used to verify the global exponential stability property. Taking the coupling between the observer components into account, we use the sum of the two Lyapunov functions as a Lyapunov-function candidate for the full error dynamics (4), (6). With an additional assumption that the partial derivative of the friction model with respect to v_x is bounded, we can verify that the origin of (4), (6) is globally exponentially stable, provided $K_{v_x}(t)$ is chosen large enough to dominate the cross terms that occur because of the coupling. The observer is also input-to-state stable with respect to errors in the reference velocity $v_{x,\text{ref}}$ [24].

III. UNKNOWN ROAD SURFACE CONDITIONS

The observer presented in the previous section depends on the construction of a lateral acceleration estimate $\hat{a}_y(t, \hat{x})$ as a function of measured signals and velocity estimates. The approach works well when the road-tire friction model from which $\hat{a}_y(t, \hat{x})$ is computed is accurately parameterized. However, the friction model is sensitive to changes in the road surface conditions, as seen in Figure 3. Since information about the road surface conditions is not available, we modify the observer to take this uncertainty into account by introducing a friction parameter to be estimated together with the velocities.

A. Parameterization

The friction parameter can be defined in several ways. In [26], [30] the road-tire friction coefficient μ_H is chosen as the parameter to be estimated. We instead scale the friction forces by redefining \hat{a}_y as $\hat{a}_y(t, \hat{x}, \theta) = \theta\hat{a}_y^*(t, \hat{x})$, where θ is the friction parameter. The value $\hat{a}_y^*(t, \hat{x})$ is given by

$$\hat{a}_y^*(t, \hat{x}) = \frac{\hat{a}_y(t, \hat{x}; \mu_H^*)}{\mu_H^*},$$

where $\hat{a}_y(t, \hat{x}; \mu_H^*)$ represents the lateral acceleration estimate calculated using a nominal value μ_H^* of the friction coefficient. In words, $\hat{a}_y^*(t, \hat{x})$ is a normalized lateral acceleration estimate based on a fixed road-tire friction coefficient, and θ is a scaling factor that changes depending on the road surface. As is evident from Figure 3, variation in μ_H does not correspond precisely to a scaling of the friction forces. The friction parameter is nevertheless closely related to μ_H , and, for large tire-slip angles, where the curves in Figure 3 flatten out, $\theta \approx \mu_H$. Because the friction parameter appears linearly, the design and analysis is simplified compared to estimating μ_H directly.

The friction parameter is defined with respect to the lateral acceleration of the vehicle, rather than the friction forces at each wheel. This definition means that, on a nonuniform road surface, the parameter represents an average of the road-tire friction properties over all the wheels. In the stability analysis we assume that the friction parameter is positive and constant.

B. Modified Observer Design

Since the longitudinal-velocity estimation does not depend on the friction model, we consider only the lateral-velocity estimation. We modify the lateral-velocity estimate to obtain

$$\dot{\hat{v}}_y = a_y - \dot{\psi}\hat{v}_x + K_{v_y}\Lambda(t, \hat{x})\xi(t, \hat{x})(a_y - \hat{a}_y(t, \hat{x}, \hat{\theta})), \quad (8)$$

and introduce an estimate of θ , given by

$$\dot{\hat{\theta}} = K_{\theta}\Lambda(t, \hat{x})\hat{a}_y^*(t, \hat{x})(a_y - \hat{a}_y(t, \hat{x}, \hat{\theta})). \quad (9)$$

The value $\xi(t, \hat{x})$ in (8) is an approximate slope of the line between $(\hat{v}_y, \hat{a}_y^*(t, \chi_1))$ and $(v_y, \hat{a}_y^*(t, \chi_2))$, with $\chi_1 = [\hat{v}_x, \hat{v}_y]^T$ and $\chi_2 = [\hat{v}_x, v_y]^T$. The slope is negative, according to (7), and the approximation can be made in several ways [25]. We choose $\xi(t, \hat{x}) = [\partial\hat{a}_y^*/\partial\hat{v}_y](t, \hat{x})$, which is sufficiently accurate to give good performance in practical experiments. In

the analysis below, we assume that $\xi(t, \hat{x})$ represents the true slope. The value $\Lambda(t, \hat{x})$ is a strictly positive scaling that can be chosen freely. We choose $\Lambda(t, \hat{x}) = (\xi^2(t, \hat{x}) + \hat{a}_y^{*2}(t, \hat{x}))^{-1/2}$ as a normalization factor to prevent large variations in the magnitude of the gain on the right-hand side of (8), (9), because such variations can cause numerical problems when implementing the observer. In the following discussion of the lateral-velocity and friction estimation, we ignore the longitudinal-velocity error \tilde{v}_x .

Analysis of the observer (8), (9) starts with the Lyapunov-function candidate

$$V = \theta \tilde{v}_y^2 + \frac{K_{v_y}}{K_\theta} \tilde{\theta}^2.$$

The time derivative of V is negative semidefinite, satisfying

$$\dot{V} \leq -k \tilde{a}_y^2(t, \tilde{x}, \tilde{\theta}),$$

for some positive k , where $\tilde{a}_y(t, \tilde{x}, \tilde{\theta}) = a_y - \hat{a}_y(t, \hat{x}, \hat{\theta})$. From the negative semidefiniteness of \dot{V} we conclude that the origin of the error dynamics is uniformly globally stable. Stability is not enough, however; we also need to show that the error vanishes with time. The bound on \dot{V} suggests that this behavior occurs if the observer error is observable from $\tilde{a}_y(t, \tilde{x}, \tilde{\theta})$. We therefore use the results from [31] to prove that if $\tilde{a}_y(t, \tilde{x}, \tilde{\theta})$ satisfies a property known as *uniform δ -persistency of excitation* with respect to \tilde{v}_y and $\tilde{\theta}$, then the origin is uniformly globally asymptotically stable.

The excitation condition takes the form of an inequality. Specifically, for each $(\tilde{x}, \tilde{\theta})$ we assume that there exist positive constants T and ε such that

$$\int_t^{t+T} \tilde{a}_y^2(\tau, \tilde{x}, \tilde{\theta}) d\tau \geq \varepsilon(\tilde{v}_y^2 + \tilde{\theta}^2) \quad (10)$$

holds for all $t \geq 0$. The inequality (10) requires the error $\tilde{a}_y(t, \tilde{x}, \tilde{\theta})$ to contain information about both the lateral-velocity error and the friction-parameter error, when taken over sufficiently long time windows. An implication of (10) is that $\tilde{a}_y(t, \tilde{x}, \tilde{\theta})$ cannot remain zero for an extended period of time unless both the lateral-velocity error and the friction-parameter error are also zero. This implication hints at the invariance-like origins of the stability proof; the results applied from [31] are based on Matrosov's theorem, a counterpart to the Krasovskii-LaSalle invariance principle that is applicable to nonautonomous systems. Note that, unlike some nonlinear persistency of excitation conditions, the condition in (10) does not depend on knowledge about the trajectories of the observer error. This property allows the condition to be evaluated based on the physical behavior of the car in different situations.

Under conditions given in [25], the excitation condition can also be used directly in the Lyapunov function, by letting

$$V = \theta \tilde{v}_y^2 + \frac{K_{v_y}}{K_\theta} \tilde{\theta}^2 - \mu \int_t^\infty e^{t-\tau} \tilde{a}_y^2(\tau, \tilde{x}, \tilde{\theta}) d\tau,$$

where μ is a small positive number. In this case, \dot{V} is negative definite in a region around the origin, yielding local exponential stability.

C. Physical Interpretation

The excitation condition (10) has a clear physical interpretation. This condition is satisfied when there is sufficient variation in the lateral movement so that the road surface conditions influence the behavior of the car; it is not satisfied when there is no variation in the lateral movement. We can therefore estimate the lateral velocity and the friction parameter simultaneously during dynamic maneuvers, but we cannot expect to do so during steady-state maneuvers. This conclusion has intuitive appeal. It is clear, for example, that we cannot determine the road surface conditions from the lateral acceleration when the car is driven straight for an indefinitely long time. An analogy to this situation is a person trying to determine how slippery the pavement is without moving around, and without the aid of visual information. A secondary result regarding the observer is that when the car is driven straight for an indefinitely long time, the velocity errors \tilde{v}_x and \tilde{v}_y still converge to zero, even though the friction-parameter error $\tilde{\theta}$ does not [25].

D. When to Estimate Friction

In normal driving situations, the tire slips are well within the linear region, and the road surface conditions have little impact on the behavior of the car [15]. Moreover, the excitation condition shows that variation in the lateral movement of the car is needed to estimate the friction parameter. These observations suggest that it is often undesirable to estimate the friction parameter, and we therefore estimate it only when necessary and possible.

When the friction parameter is not estimated, we choose to let it be exponentially attracted to a default value θ^* corresponding to a high-friction surface. There are two reasons for choosing θ^* high, rather than low. The first reason is that driving on high-friction surfaces such as asphalt and concrete is more common than driving on low-friction surfaces such as ice and snow. The second reason is that using a friction parameter that is too low results in vehicle sideslip estimates that are too high, often by a large amount.

To determine when to estimate friction, we use a linear reference model for the yaw rate to determine when the car becomes over- or understeered. We also estimate \dot{v}_y , according to (2), by $a_y - \dot{\psi} \hat{v}_x$, highpass-filtered with a 10-s time constant. A high estimate of \dot{v}_y indicates a fast-changing sideslip angle, which in turn indicates a high level of excitation and that some of the tires might be in the nonlinear region. When the \dot{v}_y estimate is high, and the reference yaw rate is above a threshold value, we turn the friction estimation on. We also turn the friction estimation on when the car is oversteered; and when the vehicle's ESC system is active, but not due to steady-state understeer. A small delay in turning the friction estimation off reduces chattering in the friction-estimation condition.

It is worth examining whether friction estimation is sometimes necessary, but not possible due to a lack of excitation. This situation can occur during long, steady-state maneuvers with one or more of the tires in the nonlinear region. On high-friction surfaces, maneuvers of this type are easily carried out.

A standard identification maneuver consists of driving along a circle while slowly increasing the speed until the circle can no longer be maintained. Toward the end of this circle maneuver, it is common for the front tires to be well within the nonlinear region, leading to severe steady-state understeer. Recall, however, that the friction parameter is attracted to a high default value when friction estimation is turned off. For steady-state maneuvers on high-friction surfaces, the default friction parameter is therefore approximately correct. On slippery surfaces such as snow and ice, maintaining a steady-state maneuver for a long time with some of the tires in the nonlinear region is more difficult. If such a situation occurs, however, it can lead to estimation errors, typically in the direction of underestimated vehicle sideslip angle.

The difficulty in handling low-excitation situations is not specific to the NVSO. Because the problem is fundamentally a lack of information, all estimation strategies suffer in low-excitation situations. It is always possible to perform open-loop integration of the kinematic equations, but the accuracy is then entirely dependent on the sensor specifications.

Experimental results indicate that the friction-estimation condition has a large impact on the performance of the observer, in particular on low-friction surfaces. If friction estimation is turned on too late, the observer might fail to capture sudden changes in the vehicle sideslip angle. If friction estimation is turned on without sufficient excitation, modeling errors and sensor bias might cause the friction parameter to be estimated as too low, leading to large errors in the estimated vehicle sideslip angle. The friction-estimation condition therefore requires careful tuning, involving some tradeoffs in performance for different situations.

IV. INCLINATION AND BANK ANGLES

The above discussion assumes that the road surface is horizontal. This assumption is usually not accurate, however, and we now consider nonzero inclination and bank angles. To define these angles, we say that the orientation of the road surface is obtained from the horizontal position by an inclination-angle rotation Θ around the vehicle's y -axis and a subsequent bank-angle rotation Φ around the x -axis. The inclination and bank angles cause gravity components to appear in (1), (2), which become [16]

$$\dot{v}_x = a_x + \dot{\psi}v_y + g \sin(\Theta), \quad (11)$$

$$\dot{v}_y = a_y - \dot{\psi}v_x - g \cos(\Theta) \sin(\Phi). \quad (12)$$

In (11), (12), a_x and a_y denote the accelerations measured by the accelerometers, which equal the total road-tire friction forces acting on the vehicle divided by the mass, as before. We assume that the inclination and bank angles vary slowly enough compared to the dynamics of the system to be modeled as constants.

We return to the observer without friction estimation as the basis for adding inclination- and bank-angle estimation. The approximate effect of a nonzero inclination or bank angle is to create an estimation bias, which is more pronounced for the lateral-velocity estimate. Inspired by the theory of nonlinear unknown-input observers, one method for estimating

the bank angle is developed in [27]; however, the nature of the disturbance suggests that something similar to standard integral action is appropriate. We present such a solution here, at first dealing with the inclination and bank angles separately.

A. Inclination Angle

We define $\theta_i = \sin(\Theta)$, and introduce an estimate $\hat{\theta}_i$ of θ_i . Typical inclination angles are small enough that $\Theta \approx \theta_i$, and the estimate $\hat{\theta}_i$ is therefore considered an estimate of the inclination angle. We use $\hat{\theta}_i$ to compensate for the disturbance, by letting

$$\dot{\hat{v}}_x = a_x + \dot{\psi}\hat{v}_y + g\hat{\theta}_i + K_{v_x}(t)(v_{x,\text{ref}} - \hat{v}_x), \quad (13)$$

where $\hat{\theta}_i$ is given by

$$\dot{\hat{\theta}}_i = K_{\theta_i}K_{v_x}(t)(v_{x,\text{ref}} - \hat{v}_x), \quad (14)$$

with K_{θ_i} a positive gain. The total gain in (14) is $K_{\theta_i}K_{v_x}(t)$. The reason for including the time-varying gain $K_{v_x}(t)$ in (14) is the same as in the longitudinal-velocity estimation (13), namely, that we wish to rely less on the reference velocity $v_{x,\text{ref}}$ whenever it is of poor quality.

To justify this approach, we use the theory of absolute stability [32, Ch. 7.1]. We split the gain $K_{v_x}(t)$ into a constant part and a time-varying part by writing $K_{v_x}(t) = a + (K_{v_x}(t) - a)$, where $0 < a < K_{v_x,\text{min}}$, and $K_{v_x,\text{min}}$ is a lower bound on $K_{v_x}(t)$. Ignoring for the moment the lateral-velocity error \tilde{v}_y , and assuming that $v_{x,\text{ref}}$ represents the true longitudinal velocity, we write the error dynamics of the longitudinal velocity and inclination-angle estimates as the interconnection of the linear time-invariant system

$$\dot{\tilde{v}}_x = -a\tilde{v}_x + g\tilde{\theta}_i + u,$$

$$\dot{\tilde{\theta}}_i = -K_{\theta_i}a\tilde{v}_x + K_{\theta_i}u$$

with the time-varying sector function $u = -(K_{v_x}(t) - a)\tilde{v}_x$. With K_{θ_i} chosen sufficiently small compared to a , the transfer function from u to \tilde{v}_x is strictly positive real. Using the circle criterion [32, Th. 7.1], we can conclude that the origin of the error dynamics is globally exponentially stable.

B. Bank Angle

We handle the bank angle in roughly the same way as the inclination angle. We define $\theta_b = \cos(\Theta) \sin(\Phi)$, and introduce an estimate $\hat{\theta}_b$ of θ_b . Typical bank and inclination angles are small enough that $\Phi \approx \theta_b$, and $\hat{\theta}_b$ is therefore considered an estimate of the bank angle. Using $\hat{\theta}_b$ for compensation in the observer without friction estimation, we obtain

$$\dot{\hat{v}}_y = a_y - \dot{\psi}\hat{v}_x - g\hat{\theta}_b - K_{v_y}(a_y - \hat{a}_y(t, \hat{x})). \quad (15)$$

Since no counterpart to $v_{x,\text{ref}}$ is available for the lateral velocity, we again use the lateral acceleration as an indirect measurement of the lateral velocity, letting

$$\dot{\hat{\theta}}_b = K_{\theta_b}(a_y - \hat{a}_y(t, \hat{x})). \quad (16)$$

Ignoring the effect of the longitudinal-velocity error \tilde{v}_x , we write the error dynamics of the lateral-velocity and bank-angle

estimates as the interconnection of the linear time-invariant system

$$\begin{aligned}\dot{\tilde{v}}_y &= -K_{v_y}\eta_{\min}\tilde{v}_y - g\tilde{\theta}_b + K_{v_y}u, \\ \dot{\tilde{\theta}}_b &= K_{\theta_b}\eta_{\min}\tilde{v}_y - K_{\theta_b}u\end{aligned}$$

with the time-varying sector nonlinearity $u = -(\eta(t, \hat{x}) - \eta_{\min})\tilde{v}_y$. With K_{θ_b} chosen sufficiently small compared to K_{v_y} and η_{\min} , we conclude, as above, that the origin of the error dynamics is globally exponentially stable.

Absolute stability provides two separate Lyapunov functions, V_1 for the $(\tilde{v}_x, \tilde{\theta}_i)$ subsystem and V_2 for the $(\tilde{v}_y, \tilde{\theta}_b)$ subsystem. Following the proof of the circle criterion [32, Th. 7.1], an excess term $-(K_{v_x, \min} - a)\tilde{v}_x^2$ appears in the time derivative \dot{V}_1 . This term can be made arbitrarily negative by increasing the difference between $K_{v_x, \min}$ and a . We thus form a new Lyapunov-function candidate as the weighted sum $V_1 + cV_2$, where c is a positive constant. By choosing c sufficiently large and using the excess term to dominate the cross terms, global exponential stability of the overall error dynamics, including the subsystem coupling that we have so far ignored, is proven.

Exponential stability is a useful property because it implies a certain level of robustness to perturbations [32, Ch. 9]. For example, an accelerometer bias added to (15), (16) can perturb the solutions, typically producing a bias in the bank-angle estimate, but it does not cause the estimates to drift off or become unstable.

V. COMBINED APPROACH

The above discussion extends the initial observer design in two separate directions, namely, by adding estimation of a friction parameter and by adding estimation of the inclination and bank angles. Including these extensions simultaneously causes no problem in the longitudinal direction, because the estimation of the longitudinal velocity and inclination angle does not depend on the friction model. In the lateral direction, however, more careful consideration is needed.

The lateral-velocity estimate changes depending on whether friction estimation is turned on or off, as can be seen from (5) and (8). Because of this change, we also modify the bank-angle estimate when friction estimation is turned on, by replacing (16) with

$$\dot{\hat{\theta}}_b = -K_{\theta_b}\Lambda(t, \hat{x})\xi(t, \hat{x})(a_y - \hat{a}_y(t, \hat{x}, \hat{\theta})). \quad (17)$$

With this modification the error dynamics for the lateral-velocity and bank-angle estimates can still be written in the form required by the absolute-stability analysis.

In previous sections tunable gains are used to dominate cross terms that appear in the Lyapunov analysis. When estimating friction, however, the stability margin depends on the excitation condition placed on the error $\tilde{a}_y(t, \hat{x}, \hat{\theta})$, specifically, T and ε in (10), which cannot be modified using tunable gains. We therefore cannot find a particular set of gains to guarantee stability in combination with bank-angle estimation. From a practical point of view, it can be difficult to distinguish the effect of low friction and a bank angle, as noted in [1]. This observation suggests that, with the chosen sensor

configuration, the estimation problem is poorly conditioned in some situations, and the difficulty in analyzing the combined approach reflects this problem.

The simplest strategy for avoiding interference between the friction estimation and the bank-angle estimation would be to disable the bank-angle estimation whenever friction estimation is turned on. This strategy, however, may cause problems when driving on slippery surfaces in hilly terrains, where combinations of large vehicle sideslip angles and large bank angles may occur. For this reason, we need to perform bank-angle estimation when estimating friction, but we set the gain lower when one of three conditions holds. The first condition is $|a_y - \hat{a}_y(t, \hat{x}, \hat{\theta})| > c_1$ for some $c_1 > 0$. The second condition is $|\ddot{\psi} - \hat{\psi}(t, \hat{x}, \hat{\theta})| > c_2$ for some $c_2 > 0$, where $\ddot{\psi}$ is the yaw acceleration found by numerical differentiation of the yaw rate, and $\hat{\psi}(t, \hat{x}, \hat{\theta})$ is calculated using the friction model in the same way as $\hat{a}_y(t, \hat{x}, \hat{\theta})$ with the forces scaled by $\hat{\theta}$. The third condition is $\text{sign}(\dot{\psi}_{\text{ref}}) = \text{sign}(\dot{\psi})$ and $(|\dot{\psi}_{\text{ref}}| - |\dot{\psi}|)\hat{v}_x > c_3$ for some $c_3 > 0$, where $\dot{\psi}_{\text{ref}}$ is the reference yaw rate used in the friction-estimation condition. From experimental data, these conditions are found to reduce the negative interference of the bank-angle estimation on low-friction surfaces. We emphasize, however, that the conditions are heuristic and that alternative conditions may be equally good or better, in particular, when different types of vehicles are used.

The full observer that merges the friction-parameter estimation and the inclination- and bank-angle estimation has two modes. In the first mode, with friction estimation turned off, the observer equations are

$$\dot{\hat{v}}_x = a_x + \dot{\psi}\hat{v}_y + g\hat{\theta}_i + K_{v_x}(t)(v_{x, \text{ref}} - \hat{v}_x), \quad (18)$$

$$\dot{\hat{v}}_y = a_y - \dot{\psi}\hat{v}_x - g\hat{\theta}_b - K_{v_y}(a_y - \hat{a}_y(t, \hat{x}, \hat{\theta})), \quad (19)$$

$$\dot{\hat{\theta}} = K_e(\theta^* - \hat{\theta}), \quad (20)$$

$$\dot{\hat{\theta}}_i = K_{\theta_i}K_{v_x}(t)(v_{x, \text{ref}} - \hat{v}_x), \quad (21)$$

$$\dot{\hat{\theta}}_b = K_{\theta_b}(t)(a_y - \hat{a}_y(t, \hat{x}, \hat{\theta})). \quad (22)$$

We see from (20) that the friction-parameter estimate is exponentially attracted to the default value θ^* . In the second mode, with friction estimation turned on, the observer equations are

$$\dot{\hat{v}}_x = a_x + \dot{\psi}\hat{v}_y + g\hat{\theta}_i + K_{v_x}(t)(v_{x, \text{ref}} - \hat{v}_x), \quad (23)$$

$$\begin{aligned}\dot{\hat{v}}_y &= a_y - \dot{\psi}\hat{v}_x - g\hat{\theta}_b \\ &\quad + K_{v_y}\Lambda(t, \hat{x})\xi(t, \hat{x})(a_y - \hat{a}_y(t, \hat{x}, \hat{\theta})),\end{aligned} \quad (24)$$

$$\dot{\hat{\theta}} = K_{\theta}\Lambda(t, \hat{x})\hat{a}_y^*(t, \hat{x})(a_y - \hat{a}_y(t, \hat{x}, \hat{\theta})), \quad (25)$$

$$\dot{\hat{\theta}}_i = K_{\theta_i}K_{v_x}(t)(v_{x, \text{ref}} - \hat{v}_x), \quad (26)$$

$$\dot{\hat{\theta}}_b = -K_{\theta_b}(t)\Lambda(t, \hat{x})\xi(t, \hat{x})(a_y - \hat{a}_y(t, \hat{x}, \hat{\theta})). \quad (27)$$

In both modes, $K_{\theta_b}(t)$ is reduced according to the practical conditions described above. Based on experimental results, we set the gain K_{θ_b} , which is used for inclination-angle estimation, to be lower for the second mode than for the first. The best choice of observer gains and other tunable parameters, such as threshold values for logical conditions, are likely to vary depending on the vehicle type and model.

VI. REMARKS ABOUT ROAD-TIRE FRICTION

The above development assumes a strictly decreasing relationship between the lateral acceleration and the lateral velocity. The curves in Figure 3 nevertheless flatten out for large tire-slip angles, suggesting that, when the friction forces are simultaneously saturated for all four tires, the strictly decreasing relationship might not hold for limited periods of time. Reflecting this possibility, the stability results in [24] are stated as regional. We can prove stability of the observer with friction estimation (8), (9) using the weaker condition of a nonstrictly decreasing relationship, because $\eta(t, \tilde{x}) = 0$ is acceptable in (7) for limited periods of time [25]. However, the observer tends to act as an open-loop integrator when the friction forces become saturated for all of the tires. If the situation persists for a long time, the achievable performance is therefore dictated by the sensor specifications, and the estimates may start to drift off due to sensor bias. It is difficult to overcome this deficiency because the vehicle model provides no useful information for stabilizing the estimates when all of the friction forces are saturated. In some applications it is desirable to detect large vehicle sideslip angles, in excess of 90° , which means that the tire forces may remain saturated for a considerable amount of time.

It is also common for the actual friction curves to decrease slightly, rather than to flatten out, for large tire-slip angles. Even the weakened condition of a nonstrictly decreasing relationship may therefore fail to hold in some extreme situations. Nonlinear friction models, like the magic formula, typically model this decrease. Not modeling this decrease may lead to limited modeling errors in some situations, but has a stabilizing effect on the observer dynamics itself. More to the point, experimental results indicate that this discrepancy does not cause a deterioration in the quality of the estimates.

VII. IMPLEMENTATION AND PRACTICAL MODIFICATIONS

For implementation in real-time hardware, the observer (18)–(27) is discretized using the forward Euler method with a sample time of 10 ms. At each time step the reference velocity $v_{x,\text{ref}}$ and the gain $K_{v_x}(t)$ are updated. The friction model used in Figure 3 is also used in the implementation, with nominal friction coefficient $\mu_H^* = 1$. The default friction parameter is set to $\theta^* = 1$.

The friction model needs information about the vertical contact forces between the tires and the road, known as the wheel loads. These loads change during driving, depending mainly on the acceleration of the vehicle. We therefore use the measured longitudinal and lateral accelerations to calculate the wheel loads, similar to [7, Ch. 7.4].

The friction model specifies an algebraic relationship between the tire slips and the road-tire friction forces. In reality, however, dynamic effects known as tire relaxation dynamics cause a small phase lag in the buildup of forces. Although the tire relaxation dynamics are negligible at high speeds, taking them into account can yield improvements at low speeds. We approximate the effect of the tire relaxation dynamics using an approach similar to that in [29, Ch. 1] by filtering the output of the friction model using the first-order transfer function

$1/(Ts + 1)$, where $T = T_e + l_e/\hat{v}_x$ is a speed-dependent time constant, which is small for large velocities.

To calculate the steering angles for the front wheels, we use a lookup table based on a steering transmission curve, with the measured steering wheel angle as the input. We do not use the output of the lookup table directly, however, because the steering angles are also affected by elasto-kinematic effects in the wheel suspension system [33]. A simplified way to account for these effects is the introduction of *caster*, which is the distance between the wheel contact point and the point where the steering axis intersects the road [34]. The caster acts as a moment arm causing the lateral road-tire friction forces to exert a torque on the steering axis. Positive caster causes the steering wheel to naturally return to the neutral position when it is released, thereby contributing to the stability of the steering system. To account for the effect of caster, the output from the lookup table must be reduced in proportion to the lateral road-tire friction forces. Following [29, Ch. 1], we calculate the steering angle for each of the front wheels using the expression $\delta_y = \delta_u - F_{y_f} e/c_l$, where δ_y is the steering angle, δ_u is the output of the lookup table, F_{y_f} is the sum of the lateral road-tire friction forces for the front wheels, e is the caster length, and c_l is a steering stiffness constant. We calculate F_{y_f} directly from the output of the friction model, which in turn depends on the steering angles, meaning that the expression for δ_y specifies a nonlinear algebraic equation. For practical purposes, this problem is solved by using a delay of one time step for the value F_{y_f} used in the calculation of δ_y , to avoid an algebraic loop in the solution.

We restrict the estimated friction parameter $\hat{\theta}$ to the interval $[0.05, 1.1]$, which spans the range of road surfaces that one can expect to encounter. We thereby obtain a discrete-time equivalent of a continuous-time parameter projection, which is theoretically justified in the continuous-time analysis [30]. We also ensure that $1.2g\hat{\theta} \geq (a_x^2 + a_y^2)^{1/2}$, since $1.2\hat{\theta}g$ is the approximate maximum acceleration achievable from the friction model when multiplied by the friction parameter. This maximum must never be smaller than the actual acceleration of the vehicle.

VIII. EXPERIMENTAL TESTING

Using data from a rear-wheel-drive passenger car, we test the observer in various situations, including dynamic and steady-state maneuvers on high- and low-friction surfaces, with different tires and at different speeds. The measurements are taken from automotive-grade sensors with no additional bias correction, and are filtered with a 15-Hz discrete lowpass filter before entering the observer. The inertial sensors are different for the high-friction and low-friction tests. For the high-friction tests, the yaw-rate sensor specifications indicate that bias and drift together may add up to a maximum of $3.5^\circ/\text{s}$, and the acceleration sensor specifications indicate that bias and drift may add up to a maximum of 1.0 m/s^2 . For the low-friction tests, the corresponding numbers are $3.5^\circ/\text{s}$ for the yaw-rate sensor and 0.6 m/s^2 for the acceleration sensors.

As a reference, we use an optical correlation sensor to obtain independent vehicle velocities. The tuning of the observer and

the parameterization of the friction model are the same for all tests, even though the tires are different.

The required accuracy of the vehicle sideslip estimate depends on the application. Typically, the error must be less than about 1° if the estimate is to be used for feedback control. When the main goal is to detect severe skidding, for example, to tighten seat belts before an impending accident, the requirements are less stringent. It is then sufficient to estimate large vehicle sideslip angles, approximately in the range $10\text{--}130^\circ$. An error of $2\text{--}3^\circ$ degrees, or approximately 10% for vehicle sideslip angles larger than 30° , is then acceptable.

A. Comparison with EKF

We compare the observer estimates to those of an EKF similar to the one presented in [11]. The states in the EKF are the longitudinal and lateral velocities, the inclination and bank angles, and the friction coefficient μ_H . The dynamic model used in the EKF is based on (11), (12), where the yaw rate is considered a known, time-varying quantity. The derivatives of the friction coefficient, inclination angle, and bank angle are modeled as filtered white noise. The resulting dynamic model is linear time varying, which improves the accuracy of the EKF prediction compared to using a nonlinear dynamic model. The measurement equations are given by the force and moment balances through the nonlinear friction model for the longitudinal acceleration a_x , the lateral acceleration a_y , and the yaw acceleration $\ddot{\psi}$, which is calculated by numerically differentiating $\dot{\psi}$. In addition, an externally calculated reference velocity is used as a measurement. The EKF switches between slow and fast estimation of the friction coefficient, depending on the situation. The switching rule is the same as for the nonlinear observer, and the same friction model is used in both designs. As with the NVSO, the continuous-time model is discretized using the forward Euler method with a sample time of 10 ms. For efficient and stable numerical implementation, the Bierman algorithm [35] is used in the measurement update process. The execution time of the NVSO is approximately one-third of that of the EKF. In the NVSO, the friction model dominates other computations.

B. Dynamic Maneuvers on High-Friction Surfaces

We begin with two examples of dynamic, high-excitation maneuvers on dry asphalt. The first maneuver consists of a series of steps in the steering wheel angle, carried out at a constant, high speed of 200 km/h. The lateral acceleration alternates between approximately $\pm 0.4g$. Figure 4 shows the results for the vehicle sideslip angle, friction parameter, and bank angle. The bank angle is estimated at approximately 2.2° , whereas the actual bank angle is approximately 1.5° , the difference being attributable to sensor bias. The resulting estimate of the vehicle sideslip angle is accurate to within approximately 0.3° . The friction estimation is activated briefly at each step, but is otherwise inactive.

The second example is an ISO-standard double lane-change maneuver, carried out at approximately 120 km/h, on a surface with a slight bank angle of approximately 0.5° . The car reaches




Fig. 4. Experimental results for periodic steps in the steering angle at 200 km/h. The uppermost plot shows the actual vehicle sideslip angle (blue) and the vehicle sideslip angles estimated by the EKF (green) and the NVSO (red). The middle plot shows the friction coefficient estimated by the EKF (green) and the friction parameter estimated by the NVSO (red). The lowest plot shows the bank angles estimated by the EKF (green) and the NVSO (red).




Fig. 5. Experimental results for an ISO-standard double lane-change maneuver at 120 km/h. The upper plot shows the actual vehicle sideslip angle (blue) and the vehicle sideslip angles estimated by the EKF (green) and the NVSO (red). The lower plot shows the friction coefficient estimated by the EKF (green) and the friction parameter estimated by the NVSO (red).

a lateral acceleration of $-0.90g$ during the maneuver. Figure 5 shows the vehicle sideslip angle and the friction parameter. During this maneuver, the tires are brought well into the nonlinear region. Even though the lane change is carried out on a high-friction surface, the friction parameter is estimated far below its default value $\theta^* = 1$, suggesting that there are significant inaccuracies in the friction model during this maneuver. A sideslip estimation error of approximately 1.4° is briefly reached.

In general, for dynamic maneuvers on high-friction surfaces,




Fig. 6. Experimental results for circle maneuvers with a 40-m radius and different tires. Each plot shows the actual vehicle sideslip angle (blue) and the vehicle sideslip angles estimated by the EKF (green) and the NVSO (red). In the upper plot, 18-inch summer tires are used. In the lower plot, 17-inch winter tires are used.

the high degree of excitation means that inaccuracies in the friction model can to some extent be compensated by the friction parameter, as happens in the lane-change example.

C. Steady-State Maneuvers on High-Friction Surfaces

We now consider two examples of a steady-state circle maneuver on dry asphalt. In the first example, 18-inch summer tires are used; in the second example, 17-inch winter tires are used. The tests are otherwise the same. The car is driven counterclockwise along a circle with a 40-m radius, while the longitudinal velocity is slowly increased until the circle can no longer be maintained because of severe understeer. The maximum lateral accelerations reached toward the end of the maneuvers are 0.84g and 0.93g, respectively. Figure 6 shows the vehicle sideslip angle for both maneuvers. The estimated vehicle sideslip angle is more accurate in the upper plot than in the lower one, where the sideslip estimation error reaches approximately 1.4° toward the end. This result indicates that the friction model more accurately matches the tires used for the first example than for the second. Because of the low level of excitation, the friction parameter cannot be used to compensate for errors in the friction model.

The EKF performs poorly for the circle maneuvers, mainly because it brings the estimated friction coefficient slightly below the initial value early in the maneuver, whereas the NVSO estimates friction only at the very end. The difference in results therefore has more to do with implementation details and tuning than with the fundamental properties of each method. Nevertheless, the examples illustrate the sensitivity of friction-based estimation designs to errors in the friction model for low-excitation maneuvers of this kind.

D. Low-Friction Surfaces

We now look at two examples from low-friction surfaces. In the first example, the car is driven around a circle on a lake covered with ice and snow. Figure 7 shows the longitudinal velocity, vehicle sideslip angle, and friction parameter. The car




Fig. 7. Experimental results for driving on an ice- and snow-covered lake. The uppermost plot shows the actual longitudinal velocity (blue) and the longitudinal velocities estimated by the EKF (green) and the NVSO (red). The middle plot shows the actual vehicle sideslip angle (blue) and the vehicle sideslip angles estimated by the EKF (green) and the NVSO (red). The lowest plot shows the friction coefficient estimated by the EKF (green) and the friction parameter estimated by the NVSO (red).

exhibits very large vehicle sideslip angles at several points, and the sideslip estimation error remains smaller than 3° for most of the test. A notable exception is the period between 140 s and 160 s, when a large vehicle sideslip angle is sustained for a long time. Reduced excitation results in inaccurate estimation of the vehicle sideslip angle in this situation.

The final example is carried out on snow-covered mountain roads. Figure 8 shows the vehicle sideslip angle, inclination angle, and bank angle. In contrast to the lake example, large inclination and bank angles occur in this example. We do not know the true values of the inclination and bank angles, but the magnitudes of the estimates are plausible. The vehicle sideslip angle reaches approximately -14° , and the estimation error remains smaller than 1.5° . In this example, estimating the bank angle at the same time as the friction parameter is necessary to avoid deterioration of the sideslip estimate.

IX. CONCLUDING REMARKS

Based on extensive experimental tests we conclude that, overall, the NVSO performs as well as the EKF, while achieving a significant reduction in execution time. The reduction is likely to be even more significant on production hardware with a fixed-point microprocessor, since the computations largely consist of floating-point operations. The NVSO is based on nonlinear analysis, which contributes toward understanding the strengths and limitations of the observer, the EKF, and alternative designs with similar sensor configurations. The number of observer gains in the NVSO is smaller than the number of tunable elements in the EKF. Nevertheless, because




Fig. 8. Experimental results for driving on snow-covered mountain roads. The uppermost plot shows the actual vehicle sideslip angle (blue) and the vehicle sideslip angles estimated by the EKF (green) and the NVSO (red). The middle plot shows the inclination angles estimated by the EKF (green) and the NVSO (red). The lowest plot shows the bank angles estimated by the EKF (green) and the NVSO (red).

of the switching between different operating regimes and the various thresholds involved in the switching logic, tuning the NVSO is nontrivial.

The observer estimates are typically of high quality. As illustrated by the examples, however, the design is sensitive to errors in the friction model, in particular during steady-state maneuvers, and the accuracy of the friction model depends to some extent on the tires. A full, systematic analysis of the response of the observer to various model uncertainties and sensor inaccuracies is a formidable task that is yet to be carried out.

Furthermore, problems remain in some situations involving low-friction surfaces. These problems can largely be attributed to the difficulty of distinguishing a nonzero bank angle from low friction. The NVSO's method of combining friction estimation and bank-angle estimation is a weakness in the present design, and the ad hoc approach of selectively reducing the bank-angle gain based on various criteria is likely to need revision and improvement before the NVSO can reach production quality. The EKF is in general better at automatically distinguishing the influences of low friction and a nonzero bank angle, giving it an advantage in some situations, and suggesting that there is still room for improving the NVSO.

A central strategy in the NVSO design is to estimate the total force acting in the lateral direction of the car by using a friction model, and to subtract this estimate from the total force measured through the lateral acceleration. The resulting difference produces the quantity $\tilde{a}_y(t, \tilde{x}, \hat{\theta})$, when divided by the mass. By using the yaw acceleration $\dot{\psi}$ it is possible to separate between the lateral friction forces on the front and

rear axles. We can then produce two quantities similar to $\tilde{a}_y(t, \tilde{x}, \hat{\theta})$, one for the front axle and one for the rear axle. The NVSO design can be carried out in the same way using these two quantities, resulting in two injection terms in the derivatives of the lateral-velocity and bank-angle estimates. We can furthermore use separate friction parameters for the front and rear axles. This strategy shows some promise with respect to improving the separation between nonzero bank angles and low-friction surfaces. Since we have not obtained a consistent improvement in quality, however, we have not used this type of design for the experimental results presented in this article.

A limitation imposed by the sensor configuration is that derivative information about the inclination and bank angles is not available in the form of roll- and pitch-rate measurements. The estimates of the inclination and bank angles are therefore produced by a form of integral action, which, in the case of the bank-angle estimation, depends on the friction model. With this solution, the inclination- and bank-angle estimates tend to capture low-frequency disturbances, which also includes sensor bias. The estimates have a limited ability to respond to rapid changes, which is underscored by upper limits on the allowable gains K_{θ_i} and K_{θ_b} . Preliminary results indicate that significant improvement can be obtained in this respect if derivative information is made available by using a six-degree-of-freedom sensor cluster.

Fundamentally, the level of accuracy that can consistently be achieved with any estimation strategy depends on the sensor configuration. As discussed at several points, the achievable performance is dictated by the sensor specifications in situations where the vehicle model provides no useful information. It is our opinion that, to achieve a consistent vehicle sideslip error of less than 1° , it is necessary to improve on the standard ESC-type sensor configuration considered in this article.

ACKNOWLEDGMENTS

The research presented in this article is supported by the European Commission STREP project *Complex Embedded Automotive Control Systems*, contract 004175, and the Research Council of Norway, and was partly conducted while Håvard Fjær Grip, Lars Imsland, and Tor A. Johansen were associated with SINTEF ICT, Applied Cybernetics, NO-7465 Trondheim, Norway.

REFERENCES

- [1] A. T. van Zanten, "Bosch ESP systems: 5 years of experience," in *Proc. Automot. Dyn. Stabil. Conf.*, Troy, MI, 2000, paper no. 2000-01-1633.
- [2] J. Farrelly and P. Wellstead, "Estimation of vehicle lateral velocity," in *Proc. IEEE Int. Conf. Contr. Appl.*, Dearborn, MI, 1996, pp. 552–557.
- [3] Y. Fukada, "Slip-angle estimation for vehicle stability control," *Vehicle Syst. Dyn.*, vol. 32, no. 4, pp. 375–388, 1999.
- [4] P. J. TH. Venhovens and K. Naab, "Vehicle dynamics estimation using Kalman filters," *Vehicle Syst. Dyn.*, vol. 32, no. 2, pp. 171–184, 1999.
- [5] A. Y. Ungoren, H. Peng, and H. E. Tseng, "A study on lateral speed estimation methods," *Int. J. Veh. Auton. Syst.*, vol. 2, no. 1/2, pp. 126–144, 2004.
- [6] U. Kiencke and A. Daiß, "Observation of lateral vehicle dynamics," *Contr. Eng. Pract.*, vol. 5, no. 8, pp. 1145–1150, 1997.
- [7] U. Kiencke and L. Nielsen, *Automotive Control Systems: For Engine, Driveline, and Vehicle*. Springer, 2000.

- [8] M. Hiemer, A. von Vietinghoff, U. Kiencke, and T. Matsunaga, "Determination of the vehicle body slip angle with non-linear observer strategies," in *Proc. SAE World Congress*, Detroit, MI, 2005, paper no. 2005-01-0400.
- [9] A. von Vietinghoff, M. Hiemer, and U. Kiencke, "Nonlinear observer design for lateral vehicle dynamics," in *Proc. IFAC World Congress*, Prague, Czech Republic, 2005, pp. 988–993.
- [10] A. von Vietinghoff, S. Olbrich, and U. Kiencke, "Extended Kalman filter for vehicle dynamics determination based on a nonlinear model combining longitudinal and lateral dynamics," in *Proc. SAE World Congress*, Detroit, MI, 2007, paper no. 2007-01-0834.
- [11] A. Suissa, Z. Zomotor, and F. Böttiger, "Method for determining variables characterizing vehicle handling," US Patent 5,557,520, 1994, filed Jul. 29, 1994; issued Sep. 17, 1996.
- [12] L. R. Ray, "Nonlinear tire force estimation and road friction identification: Simulation and experiments," *Automatica*, vol. 33, no. 10, pp. 1819–1833, 1997.
- [13] 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 Syst. Dyn.*, vol. 34, no. 1, pp. 57–75, 2000.
- [14] H. Lee, "Reliability indexed sensor fusion and its application to vehicle velocity estimation," *J. Dyn. Syst. Meas. Contr.*, vol. 128, no. 2, pp. 236–243, 2006.
- [15] A. Hac and M. D. Simpson, "Estimation of vehicle side slip angle and yaw rate," in *Proc. SAE World Congress*, Detroit, MI, 2000, paper no. 2000-01-0696.
- [16] W. Klier, A. Reim, and D. Stapel, "Robust estimation of vehicle sideslip angle – an approach w/o vehicle and tire models," in *Proc. SAE World Congress*, Detroit, MI, 2008, paper no. 2008-01-0582.
- [17] H. E. Tseng, "Dynamic estimation of road bank angle," *Vehicle Syst. Dyn.*, vol. 36, no. 4, pp. 307–328, 2001.
- [18] C. Sentouh, Y. Sebsadji, S. Mammari, and S. Glaser, "Road bank angle and faults estimation using unknown input proportional-integral observer," in *Proc. Eur. Contr. Conf.*, Kos, Greece, 2007, pp. 5131–5138.
- [19] J. Ryu and J. C. Gerdes, "Integrating inertial sensors with Global Positioning System (GPS) for vehicle dynamics control," *J. Dyn. Syst. Meas. Contr.*, vol. 126, no. 2, pp. 243–254, 2004.
- [20] D. M. Bevly, J. C. Gerdes, and C. Wilson, "The use of GPS based velocity measurements for measurement of sideslip and wheel slip," *Vehicle Syst. Dyn.*, vol. 38, no. 2, pp. 127–147, 2002.
- [21] D. M. Bevly, "Global Positioning System (GPS): A low-cost velocity sensor for correcting inertial sensor errors on ground vehicles," *J. Dyn. Syst. Meas. Contr.*, vol. 126, no. 2, pp. 255–264, 2004.
- [22] D. M. Bevly, J. Ryu, and J. C. Gerdes, "Integrating INS sensors with GPS measurements for continuous estimation of vehicle sideslip, roll, and tire cornering stiffness," *IEEE Trans. Intell. Transport. Syst.*, vol. 7, no. 4, pp. 483–493, 2006.
- [23] J. A. Farrell, *Aided Navigation: GPS with High Rate Sensors*. McGraw-Hill, 2008.
- [24] L. Imsland, T. A. Johansen, T. I. Fossen, H. F. Grip, J. C. Kalkkuhl, and A. Suissa, "Vehicle velocity estimation using nonlinear observers," *Automatica*, vol. 42, no. 12, pp. 2091–2103, 2006.
- [25] H. F. Grip, L. Imsland, T. A. Johansen, T. I. Fossen, J. C. Kalkkuhl, and A. Suissa, "Nonlinear vehicle side-slip estimation with friction adaptation," *Automatica*, vol. 44, no. 3, pp. 611–622, 2008.
- [26] L. Imsland, H. F. Grip, T. A. Johansen, T. I. Fossen, J. C. Kalkkuhl, and A. Suissa, "Nonlinear observer for vehicle velocity with friction and road bank angle adaptation – validation and comparison with an extended Kalman filter," in *Proc. SAE World Congress*, Detroit, MI, 2007, paper no. 2007-01-0808.
- [27] L. Imsland, T. A. Johansen, H. F. Grip, and T. I. Fossen, "On nonlinear unknown input observers – applied to lateral vehicle velocity estimation on banked roads," *Int. J. Contr.*, vol. 80, no. 11, pp. 1741–1750, 2007.
- [28] P. Haney, *The Racing & High-Performance Tire: Using Tires to Tune for Grip & Balance*. SAE International, 2003.
- [29] H. B. Pacejka, *Tire and Vehicle Dynamics*, 2nd ed. SAE International, 2005.
- [30] 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. IEEE Conf. Dec. Contr.*, San Diego, CA, 2006, pp. 3603–3608.
- [31] A. Loría, E. Panteley, D. Popović, and A. R. Teel, "A nested Matrosov theorem and persistency of excitation for uniform convergence in stable nonautonomous systems," *IEEE Trans. Automat. Contr.*, vol. 50, no. 2, pp. 183–198, 2005.
- [32] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Prentice Hall, 2001.
- [33] W. Matschinsky, *Radführungen der Straßenfahrzeuge: Kinematik, Elasto-Kinematik und Konstruktion*, 3rd ed. Springer, 2007.
- [34] Bosch, *Automotive Handbook*, 6th ed. SAE International, 2005.
- [35] G. J. Bierman, *Factorization Methods for Discrete Sequential Estimation*. Dover Publications, 2006.

AUTHOR BIOGRAPHIES

Håvard Fjær Grip received the M.Sc. degree in control engineering from the Norwegian University of Science and Technology (NTNU) in 2006. He has worked as a research scientist for SINTEF ICT, Applied Cybernetics, on the EU project Complex Embedded Automotive Control Systems (CEmACS), and on a vehicle state observation project at Daimler Group Research and Advanced Engineering in 2007–2008, where he participated in the development of active crosswind stabilization for passenger cars. He is pursuing Ph.D. studies at NTNU and is currently a visiting scholar at the School of Electrical Engineering and Computer Science at Washington State University. His research interests include nonlinear and adaptive estimation and control with automotive applications.

Lars Imsland received the Ph.D. degree in electrical engineering from the Department of Engineering Cybernetics at NTNU in 2002. As part of his Ph.D. studies, he was a visiting scholar at the Institute for Systems Theory in Engineering at the University of Stuttgart, Germany. After his Ph.D. studies, he worked for two years as a postdoctoral researcher at the Gas Technology Center NTNU–SINTEF, before joining SINTEF ICT, Applied Cybernetics, in Trondheim, Norway, as a research scientist. He currently works as a specialist for Cybernetica AS in Trondheim. His main research interests are in the theory and application of nonlinear and optimizing control and estimation.

Tor A. Johansen received the Dr.Ing. (Ph.D.) degree in electrical and computer engineering from NTNU in 1994. From 1995 to 1997 he was a research engineer with SINTEF Electronics and Cybernetics. Since 1997 he has been a professor in engineering cybernetics at NTNU. He has been a research visitor at the University of Southern California, Technical University in Delft, and University of California in San Diego. He is a former associate editor of *Automatica* and *IEEE Transactions on Fuzzy Systems*. His research interests include optimization-based control and industrial application of advanced control. He has published more than 60 journal papers.

Jens C. Kalkkuhl is a research scientist and project leader with the Autonomous Systems Laboratory of Daimler Group Research and Advanced Engineering. He obtained his Dipl.Ing. and his Dr.Ing. degrees in electrical engineering from the Technical University of Dresden in 1988 and 1992, respectively. From 1988 to 1992 he worked as a research assistant with the Institute of Automation at the Technical University of Dresden. From 1992 to 1994 he was a postdoctoral research fellow with the Industrial Control Centre at the University of Strathclyde, Glasgow. He was the coordinator of the EU project Complex Embedded Automotive Control Systems (CEmACS). Since 2002 he has held an adjunct professorship at the Hamilton Institute, National University of Ireland, Maynooth. His research interests are in vehicle dynamics control, nonlinear, hybrid, and multivariable control.

Avshalom Suissa graduated in aerospace engineering at the Technion in Haifa, Israel, in 1981. Since 1987 he has been a research scientist with Daimler Group Research and Advanced Engineering, where he is currently leading the vehicle dynam-

ics control team. He did research on automotive drive-by-wire systems for which he received the DaimlerChrysler Research Award in 1998. He holds more than 100 patents in automotive engineering. His research interests are in vehicle dynamics and nonlinear control.