THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING

Multi-Unit State Estimation

Real-Time Filtering of Articulated Heavy Vehicles

Axel Ceder

Department of Electrical Engineering
Chalmers University of Technology
Gothenburg, Sweden, 2026

Multi-Unit State Estimation

Real-Time Filtering of Articulated Heavy Vehicles

Axel Ceder

Acknowledgements, dedications, and similar personal statements in this thesis, reflect the author's own views.

© Axel Ceder 2026 except where otherwise stated.

Department of Electrical Engineering
Chalmers University of Technology
SE-412 96 Gothenburg, Sweden
Phone: +46 (0)31 772 1000

Cover:
Tractor–semitrailer with the quantities focused on in this thesis: fifth-wheel coupling forces (red), articulation angle and its uncertainty (blue), and tractor velocity components (green). See Section 1.2 and Figure 2.3. Illustration by the author, rendered in Blender.

Printed by Chalmers Digital Printing
Gothenburg, Sweden, October 2026

Multi-Unit State Estimation
Real-Time Filtering of Articulated Heavy Vehicles
Axel Ceder
Department of Electrical Engineering
Chalmers University of Technology

Abstract

Any choice needs information, and good choices need good information. That holds for picking a path on a misty mountain, for deciding whether to trust the weather forecast, and no less for a vehicle. For safety-critical decisions, the information must be not only accurate but consistent: its reported uncertainty must match its actual error. For an articulated heavy vehicle, part of this information is the motion state. Stability-critical states like the coupling force, the articulation angle, and the lateral velocities of both units are not measured. These are the quantities this thesis estimates. Estimators for articulated vehicles have approximated the coupling: they linearize, assume small angles, prescribe inputs, or relax it by introducing a state. Paper A takes this last approach, a damper formulation that carries the coupling force as a filter state. Paper B keeps the exact coupling constraint and recovers the coupling force with its uncertainty. Uncertain parameters are carried as consider parameters, fleet-level knowledge in place of per-vehicle tuning. In large-scale simulation, the estimator of Paper B reaches median errors below one centimeter per second in lateral velocity, hundredths of a degree in articulation angle, and a few hundred newtons in coupling force. It is consistent at the median there, but over-confident on the tractor lateral velocity in the single test-track case. It runs in real time for up to four units, though only two are evaluated for accuracy. The main limitation is the static vertical-load assumption, under which accuracy is lost when the load shifts almost entirely to one side. This work shows that the exact coupling constraint is a viable basis for real-time, consistent estimation. The vehicle then knows its footing in the mist, and how far to trust it.

Keywords: Articulated heavy vehicles, tractor-semitrailer, coupling force, state estimation, differential-algebraic equations, unscented Kalman filter, consider parameters, filter consistency, real-time estimation.

List of Publications

This thesis is based on the following publications:

[A] **A. Ceder**, L. Hammarstrand, M. Jonasson, M. Kumru, L. Laine, "Bayesian Motion Estimation for Articulated Heavy Vehicles; A Damper-Based Model for Coupling Force". Published in *Proc. 28th Int. Conf. Information Fusion (FUSION)*, Rio de Janeiro, Brazil, 2025, pp. 1–8.

[B] **A. Ceder**, L. Hammarstrand, M. Jonasson, M. Kumru, "Solving for the Fifth Wheel: Real-Time Exact-Constraint Coupling-Force Estimation for Articulated Heavy Vehicles". Submitted to *IEEE Trans. Veh. Technol.*, 2026.

Contents

Abstract i

List of Papers iii

Acknowledgements vii

Acronyms viii

I Overview 1

1 Introduction 3
- 1.1 Background and Motivation . . . 3
- 1.2 Problem Statement and Challenges . . . 4
- 1.3 Delimitations . . . 5
- 1.4 Thesis Outline . . . 6

2 Heavy Vehicle Dynamics Modeling 7
- Notation . . . 8
- 2.1 Model Fidelity . . . 8
- 2.2 Rigid Body . . . 10
- Tire Dynamics . . . 12

2.3 Articulation 14

3 State Estimation **19**
3.1 Problem Definition 19
3.2 Assumptions 20
3.3 Kalman Filtering with Consider Parameters 21
Kalman Filter 21
Extended and Unscented Kalman Filters 22
Consider Parameters 24
3.4 Consistency 24
3.5 The Constrained Model in the Filter 26

4 Summary of Appended Papers **29**
4.1 Paper A: Damper-Based Coupling Force Estimation 29
4.2 Paper B: Exact-Constraint Real-Time Estimation 30

5 Conclusions and Future Work **33**
5.1 Future Outlook 34
5.2 Alternative Paths Explored 35

References **37**

II Papers **41**

A Bayesian Motion Estimation for Articulated Heavy Vehicles; A Damper-Based Model for Coupling Force **A1**

B Solving for the Fifth Wheel: Real-Time Exact-Constraint Coupling-Force Estimation for Articulated Heavy Vehicles **B1**

Acknowledgments

To get through a PhD you need strategy, determination and focus. Much of this I lack, so I am very happy to have been under the supervision of Lars Hammarstrand for this first part of the work. It would not have been possible without him. The insights from Mats Jonasson have been very valuable during the project as I had very little vehicle dynamics background at the beginning. Thank you Nicolas Andersson for your support during my more trying times. And finally thank you to Murat Kumru for your insights and interest in this project, looking forward to working more with you going forward.

On the personal side I want to thank my soon-to-be wife Sophie for always being there, supporting me when I need it, pushing me over the finish line when I struggle getting over it myself. And finally, thank you to my family, who are always there for me even though we are now physically apart.

I also want to acknowledge my use of AI tools. They have improved rapidly over the years I have used them, and my way of working with them has changed with that. The ideas and the words in this thesis are still my own, but at this point it would be dishonest to call them uninfluenced.

Funding

This work was funded by Vinnova through the Strategic Vehicle Research and Innovation program (FFI), grant 2021-02570, and by Volvo Group Trucks Technology.

Acronyms

AHV:	Articulated heavy vehicle
CoM:	Center of mass
DAE:	Differential-algebraic equation
DoF:	Degree of freedom
EKF:	Extended Kalman filter
IFT:	Implicit function theorem
IMU:	Inertial measurement unit
KF:	Kalman filter
NEES:	Normalized estimation error squared
ODE:	Ordinary differential equation
UKF:	Unscented Kalman filter
VTM:	Volvo Transport Models

Part I

Overview

CHAPTER 1

Introduction

1.1 Background and Motivation

A large share of road freight in Europe is done by articulated heavy vehicles (AHVs), often a tractor unit towing a semitrailer through a single mechanical coupling, the fifth wheel. The articulated architecture makes the vehicle economically attractive, since one tractor serves many trailers, and keeps it relatively nimble. These benefits are why the trend is toward longer, heavier combinations thanks to their lower fuel consumption, emissions and cost per unit of cargo [1]. The articulation is also what makes the dynamics hazardous. The trailer is connected to the tractor at a single point, and the yaw dynamics of the two units interact through the force transmitted there. When this interaction becomes unstable, the result is one of the characteristic AHV accident modes [2], [3], [4]. These modes have been studied since at least the 1980s [3], and remain relevant today, as safety and automation raise the demands on onboard motion control [5]. Motion-control functions act on the vehicle state, so its accuracy and consistency are crucial.

For an articulated vehicle some of the governing states are the articulation angle, each unit's lateral velocity and the coupling force. However, these

quantities are hard to measure, and harder still to measure cheaply. Coupling force is one such case, where force transducers can be used, but they are costly and require calibration [6]. What vehicles can measure are their actuator inputs like steering angle and applied torque. They also have access to inexpensive sensors such as tachometers, accelerometers and gyroscopes. While these measurements are important, they do not directly say anything about the articulation angle, lateral velocity or coupling force. This is where *virtual sensing*, or *state estimation*, comes in.

A virtual sensor does not measure the quantity but infers it, and the error of the inference varies with the maneuver, the vehicle and the road. An estimator therefore reports both a best guess and how far from it the true value is likely to be. A control function that uses the estimate, for example a stability function, acts when the estimate crosses a threshold, and it needs a margin that reflects the uncertainty. If the reported uncertainty is smaller than the actual error, the estimator is over-confident, and the function trusts a wrong estimate and may act wrongly. If it is larger, the estimator is under-confident, and the function intervenes late or not at all, which costs availability. The reported uncertainty is therefore only useful if it matches the actual error, which is what it means for the estimator to be consistent (Section 3.4).

1.2 Problem Statement and Challenges

The problem is to estimate the motion state of the vehicle online, from the standard sensors of a production vehicle, and to report an uncertainty that is consistent with the actual error. The articulation angle, the lateral velocity of each unit and the coupling force are in focus, as they are the quantities that are hard to measure. This thesis examines how the coupling constraint can be kept in such a real-time estimator, and what accuracy and consistency this gives over a fleet of vehicles and over driving situations.

Three properties make the problem hard: the units are coupled, the vehicle is known only as a member of a fleet, and the driving situations excite the coupling differently and differ in what the sensors reveal. Each is a challenge the thesis addresses.

C1 *The coupling constraint.* Each unit is a body with its own mass and inertia, and the fifth wheel ties them together so they cannot move independently. The estimator predicts the motion step by step, but the

coupling does not give the motion of each unit directly, only the condition that the two stay joined. The challenge is to keep the coupling in the model at a cost a real-time estimator allows, without a cost to accuracy. Ideally in a form that addresses more units.

C2 *Fleet variation.* Each vehicle differs from others of the same model, and from itself over time, through its load, tires and wear, and these properties are not available to the estimator on board. The estimator is designed once for the fleet rather than tuned for each vehicle, so it has to stay accurate and report a consistent uncertainty without exactly knowing the properties of the vehicle it runs on.

C3 *Situations and sensors.* The driving situations vary from steady cornering to braking in a turn, slopes, low speed and reversing. They load the coupling in different ways, and in some of them the standard sensors say little, as the wheel speeds do near standstill. The challenge is to stay accurate and consistent over all of them with the sensor set of a production vehicle.

These challenges are revisited in Chapter 5. The contributions of this thesis are the two appended papers. Paper A (C1, C2) describes the coupling approximately and accounts for the unknown vehicle properties in the reported uncertainty. Paper B (C1–C3) keeps the exact coupling constraint in a real-time estimator and evaluates it in simulation and on real data. Chapter 4 summarizes each paper and how it addresses these challenges.

1.3 Delimitations

The main limitation of this thesis is the static vertical-load assumption, discussed in Chapter 5. The evaluation is further delimited to a range of scenarios, and scenarios outside it would require more complex estimators, simulations or data collection:

1. Low-friction surfaces, ice and snow, are not evaluated.

2. Only nominal driving is evaluated. Wheel lift, sliding and locked wheels are not considered.

3. Properties that vary within a run, such as tire temperature and load shifts, are not evaluated. Slowly drifting sensor offsets are not evaluated either, since the runs are too short for them to develop.
4. Accuracy and consistency are not evaluated for more than two units.
5. Additional sensors are not considered.

1.4 Thesis Outline

Chapter 2 outlines the vehicle dynamics field and how an articulated heavy vehicle can be described, and goes over the constraint imposed by the joint and why this is a problem. Chapter 3 details how the models used are combined with the sensors to produce the required estimation. Chapter 4 summarizes the two appended papers and states the author's contributions. Chapter 5 states the conclusions, outlines future work and describes other avenues that have been explored.

CHAPTER 2

Heavy Vehicle Dynamics Modeling

A tractor-semitrailer, the articulated heavy vehicle treated in this thesis, consists of two units joined at the fifth wheel. The fifth wheel is a pivot that transmits force between the units but leaves their relative yaw free. Each unit therefore has its own yaw motion, and the two interact only through the force at the coupling point. How strongly the two units influence each other depends largely on the ratio of their masses and on how the mass is distributed [7]. The articulation adds a trailer mode with low damping and low natural frequency [7].

This lightly damped mode makes the relative motion of the units easy to excite, which makes the vehicle susceptible to loss of control. The loss takes two yaw-instability forms, jackknife (tractor spins) and trailer swing (trailer spins) [8]. These forms are expressed in the articulation angle, but they start with a loss of lateral tire force on one unit, and the coupling force is what transmits the disturbance to the other. A laden trailer can also have a high center of mass, so that rollover is a real concern and can occur at lateral accelerations as low as 0.3 g [3].

The quantities that govern stability between units are the articulation angle, the coupling force and the lateral velocity of each unit. Though the

articulation angle is a clear indicator, it alone is not sufficient for identifying yaw instability [4]. Of these quantities, the articulation angle is the easiest to measure, but all are hard to measure, as stated in Chapter 1. This is the premise of the thesis: these quantities can be inferred from other measurements and modeling.

Section 2.1 places the model on a ladder of model fidelity, Section 2.2 describes a single unit as a rigid body, and Section 2.3 joins the units at the fifth wheel, where the coupling constraint turns the model into a differential-algebraic equation.

Notation

Following ISO 8855 [9], each unit carries a body-fixed frame with X forward, Y left, and Z up, forming a right-handed system. Two reference points per body are distinguished. The first-axle center A1 is a load-independent origin for body-fixed position vectors: unlike the center of mass, it does not move when the payload changes, which makes it a convenient anchor for geometry and for evaluation. The center of mass (CoM) O is where the equations of motion hold (Section 2.2).

Notation follows the appended papers: bold symbols are vectors, a parenthesized superscript gives the coordinate frame, component superscripts x, y, z select components, and subscripts give the unit and the position within it. Thus $v^{y}_{1,\mathrm{O}}$ is the lateral velocity at the tractor's center of mass and $\boldsymbol{v}^{(1)}_{2,\mathrm{C}}$ the trailer's coupling-point velocity expressed in the tractor frame. Figure 2.1 shows the frame and the notation on the tractor. The equations are written in vector form, one symbol per quantity with its frame and axis as superscripts, rather than component by component as is common in vehicle dynamics texts. The vector form is what the constraint formulation of Section 2.3 and its extension to more units need, and it is the form the model is implemented in.

2.1 Model Fidelity

Vehicle models can be ordered as a ladder of model fidelity, where each rung adds physical effects to the one below. Each added effect brings states or parameters, and computational cost. For an estimator the question is therefore not how much of the vehicle can be described, but which is the lowest rung

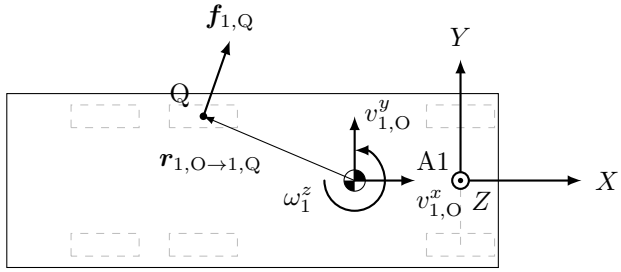


Figure 2.1: The body-fixed frame of the tractor at the first-axle center A1, with Z out of the page, the velocity and yaw rate at the center of mass, and the position $\boldsymbol{r}_{1,\mathrm{O}\to 1,\mathrm{Q}}$ of a wheel point Q where the tire force $\boldsymbol{f}_{1,\mathrm{Q}}$ acts.

that carries the quantities to be estimated and predicts the quantities that are measured.

The lowest rung is the linear single-track model, which lumps the wheels of each axle into one [10, p. 329]. It can be described purely kinematically, or dynamically to include forces, and then describes only yaw-plane motion, with roll, pitch and bounce neglected. The kinematic form assumes ideally tracking tires, so the lateral velocity is zero at each axle, and a vehicle with several rear axles is given an equivalent wheelbase, that of a two-axle vehicle which turns the same way [10, pp. 278, 285]. The lateral velocities that this thesis estimates are therefore absent from it. The next rung is simply discarding the lumping of the axles, describing each wheel individually as a dual-track model. From here the model can be made nonlinear, although such models are often linearized already in symbolic form, through linear tires and small steering and slip angles, which narrows the range where they are valid [11]. Suspension can be included, adding roll, pitch and heave [12]. Describing each movable part of the vehicle is another rung, where the multibody transport models used in industry for handling analyses of different trucks are found [13]. One step further, the bodies themselves can be made deformable, for example a frame that twists, which is rarely done outside specific evaluations.

Each rung also demands more knowledge of the vehicle and a more detailed parameterization. In an estimator every added parameter has to be known, or its distribution has to be carried by the filter (Section 3.3). Not every effect is worth that cost. Load transfer, the shift of vertical load between the wheels in

turns and under braking, has little influence on the vehicle in normal driving, and modeling it dynamically adds states that are expensive to compute [11]. The model in this thesis therefore keeps the vertical loads static.

In this thesis a relatively low rung is chosen. Each unit is a rigid body (Section 2.2), and the model is planar and nonlinear, with linear tires, but it is dual-track rather than single-track. The wheel speeds are measured at every wheel, so the model has to give the velocity of every wheel to compare with them, which a single-track model cannot. For an estimator, extending the model to both wheels is a minor cost. With static loads and linear tires it adds no states, only the track widths and a slip and force evaluation per wheel instead of per axle. In simulation a dual-track model usually comes with load transfer and nonlinear tires, and it is these that make it expensive there. Joining a second unit introduces a constraint between the two, the topic of Section 2.3.

2.2 Rigid Body

The frame does not twist, and the cab and suspension do not move relative to it, so every point of unit i shares the same angular velocity $\boldsymbol{\omega}_i$. Translational velocities $\boldsymbol{v}_i$ then transfer between points on the unit by the usual kinematic relation

$$\boldsymbol{v}_{i,\mathrm{Q}} = \boldsymbol{v}_{i,\mathrm{O}} + \boldsymbol{\omega}_i \times \boldsymbol{r}_{i,\mathrm{O}\to i,\mathrm{Q}}, \tag{2.1}$$

where Q is any point on the unit and $\boldsymbol{r}_{i,\mathrm{O}\to i,\mathrm{Q}}$ is the relative position between the CoM and the point. Figure 2.1 shows such a point and its position vector on the tractor, and Figure 2.2 the velocity, acceleration and yaw rate at the CoM of each unit.

The planar motion of each unit follows from Newton's second law for translation and Euler's equation for rotation about the CoM. The road is taken as a plane, possibly inclined, and the unit is assumed to stay aligned with it, so the inclination enters only through the component of gravity along the road. The inclination, the banking and grade of the road, is described by the roll and pitch angles ϕ_1^x and ϕ_1^y of the tractor, rotations about its X and Y axes, and the trailer is assumed to lie in the same plane. In the road plane this gives three equations per unit, two for the velocity and one for the yaw rate. The forces acting on the unit are the tire forces, including rolling resistance, gravity when the road is inclined, and, in an articulated vehicle, the coupling

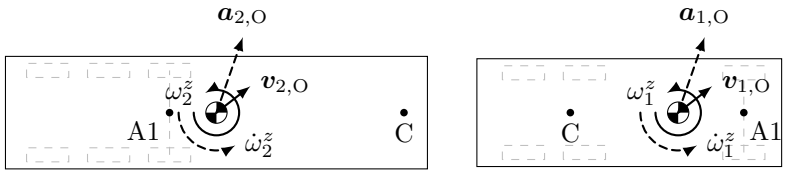


Figure 2.2: The tractor (right) and the semitrailer (left), each a rigid dual-track body, with the first-axle center A1, the coupling point C, and the center of mass O with its velocity $\boldsymbol{v}_{i,\mathrm{O}}$, acceleration $\boldsymbol{a}_{i,\mathrm{O}}$, yaw rate ω_i^z and yaw acceleration $\dot{\omega}_i^z$ (accelerations dashed).

force. Air resistance is not modeled. Its effect is left to the model uncertainty of the estimator (Chapter 3).

The articulation angle $\psi_{2\to1}^z$ (Section 2.3) is kept, while the absolute heading is not considered. Only the roll and pitch rows of the ZYX Euler-angle kinematics are needed to describe the dynamics of the vehicle. They map the angular velocity $\boldsymbol{\omega}_1 = [\omega_1^x,\ \omega_1^y,\ \omega_1^z]^\top$ of the tractor, resolved in its own frame, to the rates of the angles,

$$\begin{bmatrix} \dot{\phi}_1^x \\ \dot{\phi}_1^y \end{bmatrix} = \begin{bmatrix} 1 & \sin\phi_1^x \tan\phi_1^y & \cos\phi_1^x \tan\phi_1^y \\ 0 & \cos\phi_1^x & -\sin\phi_1^x \end{bmatrix} \boldsymbol{\omega}_1, \tag{2.2}$$

and the angles rotate gravity into the frames of the units,

$$\boldsymbol{g}_1 = \left(\mathbf{R}_y(\phi_1^y)\,\mathbf{R}_x(\phi_1^x)\right)^\top \boldsymbol{g}, \qquad \boldsymbol{g}_2 = \mathbf{R}_z(\psi_{2\to1}^z)\,\boldsymbol{g}_1, \tag{2.3}$$

with $\boldsymbol{g} = [0,\ 0,\ -g]^\top$, where $\mathbf{R}_x$, $\mathbf{R}_y$ and $\mathbf{R}_z$ are the elementary rotations about the respective axes, and $\psi_{2\to1}^z$ is the articulation angle (Section 2.3). Since the units are assumed to stay aligned with the plane, the angles describe the banking and grade of the road, not suspension roll. The roll and pitch rates ω_1^x and ω_1^y have no dynamics in the model, their derivatives are zero, and their drift is left to the process noise of the estimator (Chapter 3).

The velocities can be expressed either in ground-fixed or in vehicle-fixed directions. In ground-fixed directions the acceleration is simply the derivative of the velocity components [10, p. 335]. In vehicle-fixed directions the frame rotates with the unit, and the acceleration picks up centripetal terms,

$$\boldsymbol{a}_{i,\mathrm{O}} = \dot{\boldsymbol{v}}_{i,\mathrm{O}} + \boldsymbol{\omega}_i \times \boldsymbol{v}_{i,\mathrm{O}}, \tag{2.4}$$

where, in the plane, $\boldsymbol{\omega}_i \times \boldsymbol{v}_{i,\mathrm{O}} = [-\omega_i^z v_{i,\mathrm{O}}^y,\ \omega_i^z v_{i,\mathrm{O}}^x]^\top$ and the lateral term $\omega_i^z v_{i,\mathrm{O}}^x$ usually dominates, as $v^y \ll v^x$. The vehicle-fixed form is the common choice in vehicle dynamics, as the tire forces, the steering and the onboard sensors all act in vehicle directions. The planar equations of motion of unit i are then

$$\begin{aligned} m_i\, \boldsymbol{a}_{i,\mathrm{O}} &= \boldsymbol{f}_i, & \boldsymbol{f}_i &= \textstyle\sum_{\mathrm{Q}} \boldsymbol{f}_{i,\mathrm{Q}} + m_i \boldsymbol{g}_i, \\ I_i^z\, \dot{\omega}_i^z &= \tau_i^z, & \tau_i^z &= \textstyle\sum_{\mathrm{Q}} \big(\boldsymbol{r}_{i,\mathrm{O}\to i,\mathrm{Q}} \times \boldsymbol{f}_{i,\mathrm{Q}}\big)_z, \end{aligned} \tag{2.5}$$

with mass m_i and yaw inertia I_i^z. The resultant $\boldsymbol{f}_i$ sums, per direction, the forces $\boldsymbol{f}_{i,\mathrm{Q}}$ acting at the points Q and the in-plane components of gravity $m_i \boldsymbol{g}_i$ from (2.3), and τ_i^z is their yaw moment about the CoM. The equations hold at the CoM, which is why the velocity states are expressed there.

In the dual-track model every wheel is one of the points Q, at its axle position and on its side of the track. Its velocity from (2.1) gives the slip of its tire (Section 2.2), and the tire force acting at the wheel enters the force sums in (2.5), with its moment about the CoM in τ_i^z. The track positions set the slip of each tire, and the two wheels of an axle add up to one axle force whose yaw moment acts through the axle position. The vertical load of each wheel is static. The axle loads and the vertical load at the fifth wheel follow from the vertical force and moment balance of each unit at rest, given the masses and the CoM positions, with the axles of a tandem sharing the load equally and each axle load split equally between its two wheels. The inclination of the road does not change them.

Tire Dynamics

Apart from gravity and air resistance, every external force acting on the vehicle combination in the road plane passes through the contact patches of the tires. A tire develops force by deforming, and the deformation is described by slip, the motion of the wheel relative to the road. Longitudinal slip compares the rolling speed $R\omega_{\mathrm{w}}$ of the wheel, with radius R and rotational speed ω_{w}, to its forward speed, and lateral slip, or the slip angle, compares the lateral velocity of the wheel to its forward speed,

$$s^x = \frac{R\omega_{\mathrm{w}} - v^x}{|v^x|}, \qquad \alpha = \arctan \frac{v^y}{|v^x|}, \tag{2.6}$$

where v^x and v^y are the velocity components of the wheel center resolved in the directions of the wheel. Both are one of several definitions [10, p. 116]. On

a steered wheel these directions are turned by the steer angle δ about Z, so the body-frame velocity of the wheel center from (2.1) is resolved as

$$\boldsymbol{v}_{\mathrm{w}} = \mathbf{R}_z(\delta)^{\top} \boldsymbol{v}_{i,\mathrm{Q}}, \qquad \boldsymbol{f}_{i,\mathrm{Q}} = \mathbf{R}_z(\delta)\, \boldsymbol{f}_{\mathrm{w}}, \tag{2.7}$$

and the tire force $\boldsymbol{f}_{\mathrm{w}}$, computed in the wheel directions, is rotated back into the body frame before it enters (2.5). Only the front axle of the tractor is steered, and δ is an input, with a constant offset in it among the uncertain parameters (Section 3.3). The parameters of a tire model are only valid with the slip definition they were fitted with. Every definition divides by a speed, so the slip grows without bound as the vehicle comes to a stop [10, pp. 140–142]. Models that are to work at low speed therefore regularize the slip, or describe the tire by its deflection instead (see Section 5.2).

For small slip the lateral force is close to linear in the slip angle. Its slope at zero slip, the cornering stiffness, grows with the vertical load f^z, and the first approximation is that it is proportional to the load. The lateral force is then

$$f^y = -C_\alpha\, f^z\, \alpha, \tag{2.8}$$

with the cornering coefficient C_α, the cornering stiffness per unit load [10, pp. 157–158]. The sign follows ISO 8855, where a positive slip gives a negative force [9]. In reality the relation is not linear, meaning that the stiffness grows more slowly than the load. This has a consequence for an axle in a turn: the inner wheel loses load and the outer wheel gains it, but the outer wheel does not gain the stiffness the inner one loses, and so the cornering stiffness of the axle decreases as the load transfer grows [10, p. 307]. A model with static vertical loads, as the one in this thesis, does not capture this.

As the slip grows, the force saturates at the friction limit. Physical models, such as the brush model, derive the saturation from the elasticity of the tread and dry friction in the contact patch. Curve-fit models, most notably the Magic Formula [14], instead fit a function to measurements. They are accurate for the tire they are fitted to, but less reliable outside the conditions of that experiment.

For the estimation problem in this thesis the choice of tire model is a trade-off. Every parameter in the tire model is a parameter of the estimator, and the cornering stiffness is already uncertain between tires [15], with wear and with temperature. Both papers use a linear model with a load-proportional stiffness. The longitudinal force is set from the drive and brake torques, not from the

slip: the net torque divided by the wheel radius, without wheel rotational dynamics, less a rolling resistance proportional to the vertical load. When the tire both brakes and steers, the two directions share the same friction, and the resultant force is bounded by the friction circle

$$(f^x)^2 + (f^y)^2 \leq (\mu f^z)^2, \tag{2.9}$$

with the friction coefficient μ. A friction ellipse is more realistic, as the tire is not isotropic [10, p. 165]. In Paper B the slip angle is regularized at low speed, and the lateral force is capped by (2.9) with the longitudinal force given priority.

2.3 Articulation

For a single unit with static vertical loads, (2.5) with a tire model is an explicit ordinary differential equation (ODE): the forces follow from the state and the inputs, and dividing by the mass and inertia gives the state derivative. With load transfer the loads, and through them the tire forces, depend on the accelerations, and the equation becomes implicit in them. An articulated vehicle has one such set per unit. The frames of the units are related by the planar rotation $\mathbf{R}_z(\psi^z_{2\to1})$, which maps components from the tractor frame into the trailer frame, and whose transpose maps back. At the coupling point, where the two units meet,

$$\boldsymbol{v}^{(2)}_{1,\mathrm{C}} = \mathbf{R}_z(\psi^z_{2\to1})\, \boldsymbol{v}^{(1)}_{1,\mathrm{C}}. \tag{2.10}$$

The articulation angle $\psi^z_{2\to1} = \psi^z_1 - \psi^z_2$ is the difference of the units' headings and evolves with the difference of the yaw rates,

$$\dot{\psi}^z_{2\to1} = \omega^z_1 - \omega^z_2. \tag{2.11}$$

Figure 2.3 shows the two units and the coupling force between them.

At the fifth wheel the two units share a point, so the velocity of the coupling point must be the same whether it is computed from the tractor or from the trailer. Kinematically this is no harder than the single unit. Applying (2.1) on both units, the velocity of the trailer follows from the tractor velocity, the two yaw rates and the articulation angle. The dynamics are harder. The accelerations of both units depend on the coupling force, which acts on the

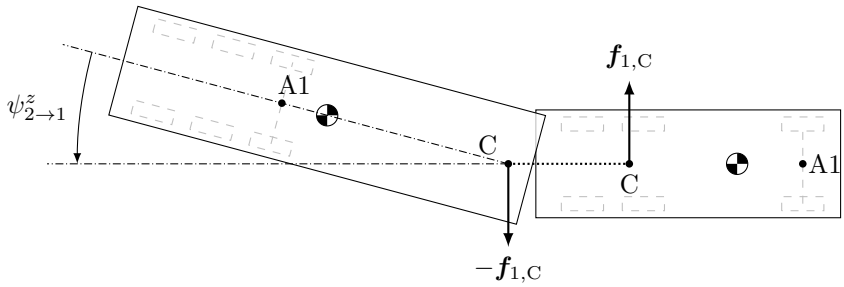


Figure 2.3: The tractor-semitrailer combination, drawn apart at the coupling point C, with its two drawings joined by the dotted line. The coupling force $\boldsymbol{f}_{1,\mathrm{C}}$ acts on the tractor and $-\boldsymbol{f}_{1,\mathrm{C}}$ on the trailer, and the articulation angle $\psi^z_{2\to1}$ lies between the longitudinal axes of the units (dash-dotted).

tractor and with opposite sign on the trailer. The force is unknown, and only the requirement that the units stay connected determines it. The equations of motion together with this constraint form a differential-algebraic equation (DAE), which does not give the state derivative directly.

The constraint can be expressed in two ways [16, pp. 120–122]. (i) In the augmented formulation each unit is a free rigid body with its own states, and the coupling is a constraint equation with the coupling force as the unknown that enforces it. (ii) In the embedding technique the constraint is used to express the dependent coordinates through a minimal set of independent ones, so it holds by construction and the coupling force drops out of the equations. The constraint is not removed but absorbed into the minimal model, where it enters the mass matrix and the forces through the articulation angle [16, pp. 121–122]. The coupling force is not part of the formulation, but it can be computed afterwards in a separate step [11]. Both formulations keep the constraint exact. The model satisfies it, as an equation in (i) and through the choice of coordinates in (ii).

In the augmented formulation the constraint is differentiated to the acceleration level and solved together with the equations of motion [16]. For a combination of N units joined by $N-1$ couplings,

$$\begin{bmatrix} \mathbf{M} & \mathbf{C}^\top \\ \mathbf{C} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \dot{\boldsymbol{w}} \\ \boldsymbol{f}_{\mathrm{C}} \end{bmatrix} = \begin{bmatrix} \boldsymbol{f} \\ \boldsymbol{\gamma} \end{bmatrix}, \tag{2.12}$$

where $\boldsymbol{w} = [\boldsymbol{v}_{1,\mathrm{O}}^\top\ \omega_1^z\ \cdots\ \boldsymbol{v}_{N,\mathrm{O}}^\top\ \omega_N^z]^\top$ stacks the velocities of all units, $\boldsymbol{f}_\mathrm{C}$ stacks the $N-1$ coupling forces, each expressed in the frame of the unit ahead of it, $\mathbf{M} = \mathrm{diag}(m_1, m_1, I_1^z, \ldots, m_N, m_N, I_N^z)$ holds the masses and inertias, $\boldsymbol{f}$ stacks the resultant forces $\boldsymbol{f}_i$ and yaw moments τ_i^z of (2.5) for each unit, without the coupling force, with the velocity-dependent terms of (2.4) moved to the force side, $\mathbf{C}$ is the $2(N-1) \times 3N$ Jacobian of the constraint that each coupling point has the same velocity on both units it joins, and $\boldsymbol{\gamma}$ is the velocity-dependent term that differentiating that constraint to the acceleration level produces. The structure is modular, with one block of $\mathbf{M}$ per unit and one block of two rows of $\mathbf{C}$ per coupling, so adding a unit adds a block rather than a new derivation.

For the tractor-semitrailer, $N = 2$, the unknowns are $\dot{\boldsymbol{w}} = [\dot{\boldsymbol{v}}_{1,\mathrm{O}}^\top\ \dot{\omega}_1^z\ \dot{\boldsymbol{v}}_{2,\mathrm{O}}^\top\ \dot{\omega}_2^z]^\top$ and $\boldsymbol{f}_\mathrm{C} = \boldsymbol{f}_{1,\mathrm{C}}$. The coupling point lies on the longitudinal axis of both units, so of $\boldsymbol{r}_1 = \boldsymbol{r}_{1,\mathrm{O}\to1,\mathrm{C}}$ and $\boldsymbol{r}_2 = \boldsymbol{r}_{2,\mathrm{O}\to2,\mathrm{C}}$ only the signed longitudinal components r_1^x and r_2^x are nonzero. Written in the tractor frame, the constraint is that the trailer's coupling-point velocity, rotated into the tractor frame, minus the tractor's is zero. With $\psi = \psi_{2\to1}^z$ for short, this gives

$$\mathbf{C} = \begin{bmatrix} -1 & 0 & 0 & \cos\psi & \sin\psi & r_2^x \sin\psi \\ 0 & -1 & -r_1^x & -\sin\psi & \cos\psi & r_2^x \cos\psi \end{bmatrix}, \tag{2.13}$$

and for two units the velocity term is

$$\boldsymbol{\gamma} = -\dot{\mathbf{R}}^\top \big(\boldsymbol{v}_{2,\mathrm{O}} + \boldsymbol{\omega}_2 \times \boldsymbol{r}_2\big), \tag{2.14}$$

where $\mathbf{R} = \mathbf{R}_z(\psi_{2\to1}^z)$ and $\dot{\mathbf{R}}$ follows from (2.11). The transpose $\mathbf{C}^\top$ then applies the coupling force $\boldsymbol{f}_{1,\mathrm{C}}$ to the tractor and its reaction $-\mathbf{R}\boldsymbol{f}_{1,\mathrm{C}}$ to the trailer, with their yaw moments, so the unknown is the coupling force itself. The accelerations and the coupling force come out of the same linear solve, which has a unique solution since the matrix is nonsingular for a physically meaningful system [16, p. 246]. With static vertical loads, $\boldsymbol{f}$ depends only on the state and the inputs, which keeps the system linear in the unknowns. With load transfer, the tire forces would depend on the accelerations through the vertical loads.

The minimal form (ii) follows from the same constraint. The trailer velocity follows from the tractor states through the shared coupling point, and differentiating it gives the trailer acceleration in terms of the tractor accelerations and the trailer yaw acceleration,

$$\dot{\boldsymbol{v}}_{2,\mathrm{O}} = \mathbf{R}\big(\dot{\boldsymbol{v}}_{1,\mathrm{O}} + \dot{\boldsymbol{\omega}}_1 \times \boldsymbol{r}_1\big) - \dot{\boldsymbol{\omega}}_2 \times \boldsymbol{r}_2 + \dot{\mathbf{R}}\big(\boldsymbol{v}_{1,\mathrm{O}} + \boldsymbol{\omega}_1 \times \boldsymbol{r}_1\big). \tag{2.15}$$

Solving the trailer's force balance for the coupling force, with (2.15) substituted, and inserting it in the tractor's equations and the trailer's yaw equation leaves a 4×4 system, in general $(N+2)\times(N+2)$ for N units, in the four independent accelerations $\dot{\boldsymbol{\nu}} = [\dot{v}^x_{1,\mathrm{O}},\ \dot{v}^y_{1,\mathrm{O}},\ \dot{\omega}^z_1,\ \dot{\omega}^z_2]^\top$, solved at every step,

$$\bar{\mathbf{M}}(\psi)\,\dot{\boldsymbol{\nu}} = \bar{\boldsymbol{b}}, \tag{2.16}$$

with

$$\bar{\mathbf{M}} = \begin{bmatrix} m_1+m_2 & 0 & 0 & -m_2 r_2^x \sin\psi \\ 0 & m_1+m_2 & m_2 r_1^x & -m_2 r_2^x \cos\psi \\ 0 & m_2 r_1^x & I_1^z + m_2 (r_1^x)^2 & -m_2 r_1^x r_2^x \cos\psi \\ -m_2 r_2^x \sin\psi & -m_2 r_2^x \cos\psi & -m_2 r_1^x r_2^x \cos\psi & I_2^z + m_2 (r_2^x)^2 \end{bmatrix},$$

$$\bar{\boldsymbol{b}} = \begin{bmatrix} \bar{f}_1^x + \tilde{f}_2^x \cos\psi + \tilde{f}_2^y \sin\psi \\ \bar{f}_1^y - \tilde{f}_2^x \sin\psi + \tilde{f}_2^y \cos\psi \\ \tau_1^z + r_1^x\big(-\tilde{f}_2^x \sin\psi + \tilde{f}_2^y \cos\psi\big) \\ \tau_2^z - r_2^x \tilde{f}_2^y \end{bmatrix}.$$

Here

$$\bar{\boldsymbol{f}}_i = \boldsymbol{f}_i - m_i \boldsymbol{\omega}_i \times \boldsymbol{v}_{i,\mathrm{O}}, \qquad \tilde{\boldsymbol{f}}_2 = \bar{\boldsymbol{f}}_2 - m_2 \dot{\mathbf{R}}\big(\boldsymbol{v}_{1,\mathrm{O}} + \boldsymbol{\omega}_1 \times \boldsymbol{r}_1\big), \tag{2.17}$$

where $\boldsymbol{f}_i$ is the resultant force of (2.5) on unit i, without the coupling force, $\bar{\boldsymbol{f}}_i$ moves the velocity-dependent term of the equations of motion to the force side, and $\tilde{\boldsymbol{f}}_2$ also removes the velocity term of (2.15). The yaw moment of the tire forces about the CoM is τ_i^z. This is the minimal-coordinate form (ii). The matrix has the structure

$$\bar{\mathbf{M}} = \mathrm{diag}(m_1, m_1, I_1^z, I_2^z) + m_2 \mathbf{J}_h^\top \mathbf{J}_h, \tag{2.18}$$

where $\mathbf{J}_h$ is the 2×4 Jacobian of (2.15) with respect to $\dot{\boldsymbol{\nu}}$. It is therefore symmetric and positive definite for every articulation angle, in any universe where masses and inertias are positive, and the solve is always well posed. The matrix depends on the state only through the articulation angle, and the constraint shows in it as the trailer mass carried at the coupling point, $m_2(r_1^x)^2$ in the tractor yaw, and as the trailer yaw inertia about the coupling point, $I_2^z + m_2(r_2^x)^2$. Paper B solves (2.16) numerically at every step and recovers the coupling force from the trailer's force balance, with the trailer acceleration $\dot{\boldsymbol{v}}_{2,\mathrm{O}}$ obtained from (2.15) at the solved $\dot{\boldsymbol{\nu}}$,

$$\boldsymbol{f}_{1,\mathrm{C}} = \mathbf{R}^\top\big(\bar{\boldsymbol{f}}_2 - m_2 \dot{\boldsymbol{v}}_{2,\mathrm{O}}\big). \tag{2.19}$$

Paper B runs the formulation for two to four units.

Together these parts give the process model. Its states are the tractor velocity and yaw rate, the trailer yaw rate, the articulation angle, and the roll and pitch angles of the road with their rates, and its inputs are the per-wheel torques and the steer angle. The measurement model uses the same parts: the wheel speeds follow from (2.1) and the wheel radius, the accelerometer from (2.4) with gravity from (2.3), and the gyroscope reads the angular rates directly. Chapter 3 takes the model from here.

The treatments of the constraint in the literature start from a model that keeps the constraint exact, as (i) and (ii) do, and simplify downward. In this thesis they are sorted into three groups. *Exact*: the constraint is satisfied exactly, as in both formulations above and in Paper B. Either can be made explicit by solving symbolically. For two units the matrix in (2.16) has a short closed-form inverse, but with the tire forces expanded in the states the explicit form becomes very large [10, pp. 354–355], and minimal-coordinate models are often left in implicit form because the explicit form is more expensive to evaluate [11]. The coupling force comes from the solve or is recovered afterwards. *Reduced*: the model is simplified so that the constraint becomes easier to handle. It can be linearized around small articulation angles, kept only kinematically so no forces appear, or replaced by prescribed inputs, where a coupled quantity such as the speed and its rate is imposed from outside. A linearized model approaches the exact one as the angles go to zero, while a kinematic model or prescribed inputs have no such limit and the coupling force is lost. *Relaxed*: the coupled quantity becomes a state of the filter, for example the coupling force as a random walk on a single unit, which drifts freely and is corrected only by the measurements, optionally driven by a penalty on the constraint mismatch. The damper model of Paper A is of this kind. The penalty approaches the exact constraint as its gain goes to infinity, but a high gain makes the system stiff, so that it needs very small integration steps. The coupling force becomes a state.

CHAPTER 3

State Estimation

This chapter sets up the estimation problem that Papers A and B solve, and the theory the papers use but do not derive. Chapter 2 derived the models that describe the deterministic approximation of an articulated heavy vehicle. This chapter is about how the inference is run with those models. The problem definition (Section 3.1) states the problem in the terms of state estimation, with the vehicle of Chapter 2 as the system. Section 3.2 states the assumptions made on the stochastics. Section 3.3 gives the filters and the consider parameters, Section 3.4 how to judge whether the covariance is right, and Section 3.5 how the constrained model of Chapter 2 enters the filter.

3.1 Problem Definition

The estimation problem follows the same form as stated in the papers. For a discrete system at time step k

$$\boldsymbol{x}_k = G(\boldsymbol{x}_{k-1}, \boldsymbol{u}_{k-1}, \boldsymbol{\theta}) + \boldsymbol{w}_{k-1}, \quad \boldsymbol{z}_k = H(\boldsymbol{x}_k, \boldsymbol{u}_k, \boldsymbol{\theta}) + \boldsymbol{v}_k, \tag{3.1}$$

the additive noise models are $\boldsymbol{w}_k \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{Q}_k)$ and $\boldsymbol{v}_k \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{R}_k)$. Five kinds of quantities are distinguished:

States $\boldsymbol{x}$ evolve over time and are estimated.

Inputs $\boldsymbol{u}$ are measured and treated as deterministic.

Measurements $\boldsymbol{z}$ are measured and treated as uncertain.

Parameters $\boldsymbol{\theta}$ are constant but uncertain, for example the cornering stiffness.

Outputs $\boldsymbol{y}$ are functions of the states, inputs and parameters, whose distributions follow from propagating the posterior.

The goal is the posterior $p(\boldsymbol{x}_k \mid \boldsymbol{z}_{1:k}, \boldsymbol{u}_{0:k})$, the distribution of the state given all measurements and inputs up to time k. It can be computed recursively in two steps [17]. The prediction propagates the previous posterior through the process model,

$$p(\boldsymbol{x}_k \mid \boldsymbol{z}_{1:k-1}) = \int p(\boldsymbol{x}_k \mid \boldsymbol{x}_{k-1})\, p(\boldsymbol{x}_{k-1} \mid \boldsymbol{z}_{1:k-1})\, \mathrm{d}\boldsymbol{x}_{k-1}, \tag{3.2}$$

and the update corrects it with the new measurement,

$$p(\boldsymbol{x}_k \mid \boldsymbol{z}_{1:k}) \propto p(\boldsymbol{z}_k \mid \boldsymbol{x}_k)\, p(\boldsymbol{x}_k \mid \boldsymbol{z}_{1:k-1}), \tag{3.3}$$

where the conditioning on the inputs and parameters is left out for brevity. In practice the posterior is summarized by its mean $\hat{\boldsymbol{x}}_{k|k}$ and covariance $\boldsymbol{P}_{k|k}$. The mean is the estimate, and the covariance is the reported uncertainty, which is only useful if it matches the actual error (Section 3.4). The integral has a closed form only in special cases, most notably the linear-Gaussian case, where the recursion is the Kalman filter [18]. The nonlinear case is the topic of Section 3.3.

3.2 Assumptions

The assumptions below are shared by the papers, unless stated otherwise, and are stated once here so the sections that follow need not repeat them.

- *Noise:* The process and measurement noise are additive, zero-mean, white and Gaussian, $\boldsymbol{w}_k \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{Q}_k)$ and $\boldsymbol{v}_k \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{R}_k)$, and independent of each other. The process noise, together with the consider parameters, absorbs what the model leaves out, for example load transfer and road disturbances. The sensors are independent of each other, so $\boldsymbol{R}_k$ is block diagonal.

- *Time:* The model is continuous and discretized with a first-order Euler step. The step need not be fixed. When the sensors arrive asynchronously, the filter predicts to each arrival and updates there, and the process noise is then given as a continuous-time density $\boldsymbol{Q}_c$ and scaled with the step T_s, $\boldsymbol{Q}_k = \boldsymbol{Q}_c T_s$, so that the variance accumulated per unit time does not depend on the step.
- *Parameters:* The parameters are constant over a run and known only within a prior distribution. They are not estimated, but their uncertainty reaches the state covariance (Section 3.3).
- *Initial state:* The prior is Gaussian with a diagonal covariance.

3.3 Kalman Filtering with Consider Parameters

This section gives the filters used in the papers in three steps: the Kalman filter for the linear case, its nonlinear extensions, and the consider parameters added on top.

Kalman Filter

If the process and measurement models are linear,

$$\boldsymbol{x}_k = \mathbf{G}\boldsymbol{x}_{k-1} + \boldsymbol{w}_{k-1}, \qquad \boldsymbol{z}_k = \mathbf{H}\boldsymbol{x}_k + \boldsymbol{v}_k, \tag{3.4}$$

with the Gaussian noise of Section 3.2, the posterior (3.3) is Gaussian, and the Kalman filter gives its mean and covariance exactly [18], by the recursion below [17, p. 57]. The inputs and parameters are left out for brevity. The prediction is

$$\hat{\boldsymbol{x}}_{k|k-1} = \mathbf{G}\hat{\boldsymbol{x}}_{k-1|k-1}, \qquad \boldsymbol{P}_{k|k-1} = \mathbf{G}\boldsymbol{P}_{k-1|k-1}\mathbf{G}^\top + \boldsymbol{Q}_{k-1}, \tag{3.5}$$

and the update is

$$\begin{aligned} \hat{\boldsymbol{z}}_k &= \mathbf{H}\hat{\boldsymbol{x}}_{k|k-1}, \\ \boldsymbol{S}_k &= \mathbf{H}\boldsymbol{P}_{k|k-1}\mathbf{H}^\top + \boldsymbol{R}_k, \\ \boldsymbol{K}_k &= \boldsymbol{P}_{k|k-1}\mathbf{H}^\top \boldsymbol{S}_k^{-1}, \\ \hat{\boldsymbol{x}}_{k|k} &= \hat{\boldsymbol{x}}_{k|k-1} + \boldsymbol{K}_k\big(\boldsymbol{z}_k - \hat{\boldsymbol{z}}_k\big), \\ \boldsymbol{P}_{k|k} &= \boldsymbol{P}_{k|k-1} - \boldsymbol{K}_k\boldsymbol{S}_k\boldsymbol{K}_k^\top, \end{aligned} \tag{3.6}$$

where $\hat{\boldsymbol{x}}_{k|k-1}$ is the mean at time k given the measurements up to $k-1$, $\hat{\boldsymbol{z}}_k$ the predicted measurement, $\boldsymbol{S}_k$ its covariance, and $\boldsymbol{K}_k$ the Kalman gain. Under these assumptions the Kalman filter is the minimum mean-square-error estimator. If the noise is not Gaussian, it is still the best estimator among those that are linear in the measurements [18, Theorem 2].

Extended and Unscented Kalman Filters

If the process or measurement model is nonlinear, as is the case for the articulated vehicle of Chapter 2, the posterior is no longer Gaussian. Its mean and covariance no longer describe it fully, and they can no longer be propagated exactly. Every practical filter is an approximation of that propagation, and the choice of filter is a choice of approximation.

The *extended Kalman filter (EKF)* linearizes the models around the current estimate and applies the Kalman filter to the result [17, p. 69]. With the Jacobians

$$\mathbf{J}_G = \left.\frac{\partial G}{\partial \boldsymbol{x}}\right|_{\hat{\boldsymbol{x}}_{k-1|k-1}}, \qquad \mathbf{J}_H = \left.\frac{\partial H}{\partial \boldsymbol{x}}\right|_{\hat{\boldsymbol{x}}_{k|k-1}}, \tag{3.7}$$

the mean and the predicted measurement pass through the models themselves,

$$\hat{\boldsymbol{x}}_{k|k-1} = G\big(\hat{\boldsymbol{x}}_{k-1|k-1}, \boldsymbol{u}_{k-1}, \boldsymbol{\theta}\big), \qquad \hat{\boldsymbol{z}}_k = H\big(\hat{\boldsymbol{x}}_{k|k-1}, \boldsymbol{u}_k, \boldsymbol{\theta}\big), \tag{3.8}$$

while $\mathbf{J}_G$ takes the place of $\mathbf{G}$ in the covariance prediction of (3.5) as

$$\boldsymbol{P}_{k|k-1} = \mathbf{J}_G \boldsymbol{P}_{k-1|k-1} \mathbf{J}_G^\top + \boldsymbol{Q}_{k-1}, \tag{3.9}$$

and $\mathbf{J}_H$ that of $\mathbf{H}$ in the innovation covariance and the gain of (3.6).

The *unscented Kalman filter (UKF)* takes another route. Instead of linearizing the model, it approximates the moments of the transformed variable with a set of weighted points, called sigma points. The points are chosen so that their sample mean and covariance equal those of the state, and each point is passed through the model [19]. For a random variable of dimension n with mean $\hat{\boldsymbol{x}}$ and covariance $\boldsymbol{P}$, the $2n+1$ sigma points and their weights are [19]

$$\begin{aligned} \boldsymbol{\chi}_0 &= \hat{\boldsymbol{x}}, & W_0 &= \frac{\kappa}{n+\kappa}, \\ \boldsymbol{\chi}_i &= \hat{\boldsymbol{x}} + \left(\sqrt{(n+\kappa)\boldsymbol{P}}\right)_i, & W_i &= \frac{1}{2(n+\kappa)}, \\ \boldsymbol{\chi}_{n+i} &= \hat{\boldsymbol{x}} - \left(\sqrt{(n+\kappa)\boldsymbol{P}}\right)_i, & W_{n+i} &= \frac{1}{2(n+\kappa)} \end{aligned} \tag{3.10}$$

where $i = 1, \dots, n$ and $(\cdot)_i$ is the ith column of a matrix square root of $(n+\kappa)\boldsymbol{P}$. The weighted mean and covariance of the transformed points agree with a Taylor expansion of the true moments to second order for any distribution of $\boldsymbol{x}$, and to third order when $\boldsymbol{x}$ is Gaussian, since the odd moments then vanish [19], [20]. The parameter κ is a design choice that tunes the higher moments, and for a Gaussian state the heuristic is $n + \kappa = 3$ [19]. For $n > 3$ this makes κ negative, which can give a covariance that is not positive semidefinite [19]. The scaled form of the transform [20] adds two parameters and reduces to the form above when they are one and zero.

The UKF applies the transform in both the prediction and the update. With the sigma points $\boldsymbol{\chi}_{i,k-1}$ drawn from $\hat{\boldsymbol{x}}_{k-1|k-1}$ and $\boldsymbol{P}_{k-1|k-1}$, the prediction is

$$\begin{aligned} \boldsymbol{\chi}_{i,k|k-1} &= G\big(\boldsymbol{\chi}_{i,k-1}, \boldsymbol{u}_{k-1}, \boldsymbol{\theta}\big), \\ \hat{\boldsymbol{x}}_{k|k-1} &= \sum_{i=0}^{2n} W_i \boldsymbol{\chi}_{i,k|k-1}, \\ \boldsymbol{P}_{k|k-1} &= \sum_{i=0}^{2n} W_i \big(\boldsymbol{\chi}_{i,k|k-1} - \hat{\boldsymbol{x}}_{k|k-1}\big)\big(\boldsymbol{\chi}_{i,k|k-1} - \hat{\boldsymbol{x}}_{k|k-1}\big)^\top + \boldsymbol{Q}_{k-1}, \end{aligned} \tag{3.11}$$

and the update draws new sigma points $\boldsymbol{\chi}^-_{i,k}$ from $\hat{\boldsymbol{x}}_{k|k-1}$ and $\boldsymbol{P}_{k|k-1}$ and passes them through the measurement model,

$$\begin{aligned} \boldsymbol{\zeta}_{i,k} &= H\big(\boldsymbol{\chi}^-_{i,k}, \boldsymbol{u}_k, \boldsymbol{\theta}\big), \\ \hat{\boldsymbol{z}}_k &= \sum_{i=0}^{2n} W_i \boldsymbol{\zeta}_{i,k}, \\ \boldsymbol{S}_k &= \sum_{i=0}^{2n} W_i \big(\boldsymbol{\zeta}_{i,k} - \hat{\boldsymbol{z}}_k\big)\big(\boldsymbol{\zeta}_{i,k} - \hat{\boldsymbol{z}}_k\big)^\top + \boldsymbol{R}_k, \\ \boldsymbol{C}_k &= \sum_{i=0}^{2n} W_i \big(\boldsymbol{\chi}^-_{i,k} - \hat{\boldsymbol{x}}_{k|k-1}\big)\big(\boldsymbol{\zeta}_{i,k} - \hat{\boldsymbol{z}}_k\big)^\top, \end{aligned} \tag{3.12}$$

with the gain $\boldsymbol{K}_k = \boldsymbol{C}_k \boldsymbol{S}_k^{-1}$ and the mean and covariance updated as in (3.6). This is the additive form of the UKF [17, pp. 87–88], which applies since the noise is additive.

For this thesis the most important property of the UKF is not its accuracy, but that "the process and observation models can be treated as 'black boxes'" [19]. The filter only needs to evaluate the process model at the sigma

points, so the model can contain a solve at every evaluation. Section 3.5 shows that the EKF can reach the same model.

Consider Parameters

The filter so far estimates the state with the parameters fixed at their nominal values. The parameters $\boldsymbol{\theta}$ are, however, uncertain, and an error in them biases the state without the covariance knowing about it. One option is to estimate them jointly with the state, or in a dual filter [20], [21], but parameters that are weakly observable from the sensors are then hard to separate from the model error. Augmenting the state with the parameters also makes their dynamics singular, which can make nonlinear Kalman filters diverge, and adding artificial noise to the parameters to avoid this turns them into time-varying quantities [17, p. 186]. The consider, or Schmidt-Kalman, filter [22] instead appends the parameters to the state, but never updates them. In its unscented form, the unscented Schmidt-Kalman filter (USKF) [23], the aim is to carry the uncertainty of the model and measurement parameters into the state covariance without estimating the parameters [23]. The parameter prior seeds additional sigma points and the cross-covariance between state and parameters is kept, but the parameter mean and covariance are left unchanged by both the prediction and the update.

What the consider parameters buy is a covariance that carries the parameter uncertainty the measurements cannot remove. They do not remove the parameter error, so the bias it causes remains. The cross-covariance does enter the gain, so the estimate differs from that of a filter that ignores the parameter uncertainty, but the main benefit is that the reported uncertainty is honest across a fleet of vehicles whose cornering stiffness, mass and inertia are known only as distributions.

3.4 Consistency

A physical sensor comes with a datasheet. Its accuracy is known before it is mounted, and a function that uses it can set its margins once. A virtual sensor has no datasheet. Its accuracy depends on the maneuver, on the vehicle it runs on and on how observable the quantity is at the moment, and the only statement it makes about its own accuracy is the covariance it reports. For

the coupling force there is not even a production sensor to compare against, so the covariance is the only quality signal a downstream function will ever get. A stability function that acts on an estimated articulation angle or coupling force sets its thresholds with a margin, and that margin is only as good as the reported uncertainty. A filter that is over-confident reports a force within limits when it is not, and one that is under-confident makes the function intervene late or not at all. Accuracy, measured as RMSE over a fleet, says nothing about either. This is why challenge C2 asks for consistency and not only for accuracy, and why the covariance is validated with the same care as the mean.

An estimator is consistent when its errors are unbiased and their second moment equals the covariance it reports [24]. The test statistic is the normalized estimation error squared (NEES),

$$\epsilon_{x,k} = \boldsymbol{e}_{x,k}^{\top} \boldsymbol{P}_{k|k}^{-1} \boldsymbol{e}_{x,k}, \qquad \boldsymbol{e}_{x,k} = \boldsymbol{x}_k - \hat{\boldsymbol{x}}_{k|k}. \tag{3.13}$$

If the filter is consistent and the errors are Gaussian, $\epsilon_{x,k}$ is χ^2 distributed with n_x degrees of freedom, n_x being the number of states. Averaged over an ensemble of independent runs at a fixed time k, the statistic gets two-sided tail bounds for a chosen error rate. A value below the lower bound means that the filter is pessimistic, or under-confident, since the reported covariance is larger than the actual error, and a value above the upper bound that it is optimistic, or over-confident [24]. The thesis uses under-confident and over-confident.

The ensemble test needs the true state and many realizations at the same time instant, which in practice means a truth-model simulation [24]. When the runs differ by design, because consistency is asked across vehicles and maneuvers, there is no common time to test across runs. The statistic is then averaged over time within a run instead, which assumes that the error process is ergodic [24]. Using the diagonal of $\boldsymbol{P}$ gives a per-state NEES,

$$\mathrm{NEES}_j = \frac{1}{N} \sum_{k=1}^{N} \frac{e_{j,k}^2}{P_{jj,k}}, \tag{3.14}$$

where $e_{j,k}$ is the jth component of $\boldsymbol{e}_{x,k}$ and $P_{jj,k}$ the corresponding diagonal element of $\boldsymbol{P}_{k|k}$. For a consistent filter each term is χ_1^2 distributed with expectation one, so the target is $\mathrm{NEES}_j = 1$. Within a run the errors are correlated in time, so the χ^2 bounds of the ensemble test do not carry over to the time average. The thesis instead judges the per-state NEES against

the band $[0.5, 2]$, a factor of two in either direction, which flags a filter whose reported standard deviation is off by more than $\sqrt{2}$. It tests the marginals only, not the cross-covariance between states. On real data without a true state, the normalized innovation squared plays the same role on the measurement residuals [24], and where only a reference exists, the per-state NEES is limited to the referenced states.

3.5 The Constrained Model in the Filter

Chapter 2 gives a model that is continuous, deterministic and constrained, while the filters of Section 3.3 need the discrete, stochastic model (3.1). The noise is additive, so the model only has to supply G, the next state from the current one, and H, the measurement at the current state. The filters use these only pointwise. The UKF evaluates them at the sigma points, and the EKF also needs their Jacobians at the estimate. Neither needs them in closed form.

In multibody estimation the DAE is converted to an ODE before filtering, by projection onto independent coordinates or by a penalty formulation, both of which remove the constraint force from the equations [25]. In terms of the differential index, the multibody equations with the position-level constraint are index 3, the acceleration-level system (2.12) is index 1, with the coupling force as an algebraic variable fixed by one linear solve, and the minimal form (2.16) is an ODE in the state, index 0, whose right-hand side is evaluated by that solve rather than written out, since the force is eliminated and the trailer velocity is computed from the tractor states rather than integrated. The damper model of Paper A reaches an ODE by the other route, a penalty that turns the constraint into stiff dynamics with the force as a state. Paper B instead keeps the constraint and solves (2.16) numerically wherever the filter evaluates G.

Assembled into the process model, the solution of (2.16) gives an explicit ODE whose right-hand side contains the linear solve,

$$\dot{\boldsymbol{x}} = F(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{\theta}) = \begin{bmatrix} \bar{\mathbf{M}}(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{\theta})^{-1}\, \bar{\boldsymbol{b}}(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{\theta}) \\ F_{\text{expl}}(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{\theta}) \end{bmatrix}, \tag{3.15}$$

where the first block is the solution of a linear system whose coefficients depend on the state, the inputs and the parameters, and F_{expl} holds the derivatives

that are available in closed form. The form applies to any model with such a subsystem. For the exact-constraint model the first block is the accelerations $\dot{\boldsymbol{\nu}}$, with $\bar{\mathbf{M}}$ depending on the state only through the articulation angle and on the parameters through the masses, inertias and coupling positions, and not on the inputs, $\bar{\boldsymbol{b}}$ carrying the tire forces and velocity terms, and F_{expl} the kinematic rates, such as the articulation rate (2.11). The Euler step of Section 3.2 then gives the process model

$$G(\boldsymbol{x}_{k-1}, \boldsymbol{u}_{k-1}, \boldsymbol{\theta}) = \boldsymbol{x}_{k-1} + T_s\, F(\boldsymbol{x}_{k-1}, \boldsymbol{u}_{k-1}, \boldsymbol{\theta}). \tag{3.16}$$

The measurement model H uses the same solve when it includes accelerations, as an accelerometer does.

The EKF also needs the Jacobians of G and H, and in (3.15) the first block is the solution of a linear system, not a closed-form expression. Differentiating the linear system behind that block and applying the implicit function theorem, valid since $\bar{\mathbf{M}}$ is nonsingular (Section 2.3), gives

$$\frac{\partial F}{\partial \boldsymbol{x}} = \begin{bmatrix} \bar{\mathbf{M}}^{-1}\Big(\dfrac{\partial \bar{\boldsymbol{b}}}{\partial \boldsymbol{x}} - \dfrac{\partial \bar{\mathbf{M}}}{\partial \boldsymbol{x}}\,\bar{\mathbf{M}}^{-1}\bar{\boldsymbol{b}}\Big) \\ \dfrac{\partial F_{\text{expl}}}{\partial \boldsymbol{x}} \end{bmatrix}, \tag{3.17}$$

where every term carries the arguments $(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{\theta})$ of (3.15), dropped for brevity, the derivative of $\bar{\mathbf{M}}$ is taken one component of $\boldsymbol{x}$ at a time, and the same expression holds with $\boldsymbol{u}$ or $\boldsymbol{\theta}$ in place of $\boldsymbol{x}$. The factorization that solves the system is reused for its Jacobian. From (3.16) the process Jacobian of Section 3.3 is

$$\mathbf{J}_G = \mathbf{I} + T_s\,\frac{\partial F}{\partial \boldsymbol{x}}. \tag{3.18}$$

The measurement Jacobian $\mathbf{J}_H$ is formed the same way. Where H includes accelerations, as an accelerometer does, the chain rule brings in the first block of (3.17), and the remaining rows are in closed form. An EKF whose Jacobians are obtained this way is called the IFT-EKF in the rest of the thesis.

The constraint also has to survive the noise. The state of (2.16) holds only the independent coordinates, and the trailer velocity is computed from them rather than carried as a state, so every value of $\boldsymbol{x}$, with or without $\boldsymbol{w}_k$, satisfies the coupling constraint. Had the trailer velocity been a state, the additive noise would push the estimate off the constraint, and the filter would need a projection or a constrained update [26]. The elimination therefore matters for

the stochastic model, not only for the solve: the constraint is imposed through the choice of coordinates, and no constrained filter is needed. The coupling force is recovered from the trailer's force balance (2.19), so it is an output $\boldsymbol{y}$ (Section 3.1) rather than a state. It has no process noise of its own, and its mean and covariance follow from the propagated sigma points or from the linearization, consider parameters included. In the damper model of Paper A the force is a state with its own entry in $\boldsymbol{Q}$ and drifts between updates. The process noise of the exact model then stands for unmodeled forces on the independent coordinates, such as load transfer and road disturbances, and not for the joint itself. Play or compliance in the fifth wheel is not representable in this model.

CHAPTER 4

Summary of Appended Papers

4.1 Paper A: Damper-Based Coupling Force Estimation

Axel Ceder, Lars Hammarstrand, Mats Jonasson, Murat Kumru, Leo Laine

Bayesian Motion Estimation for Articulated Heavy Vehicles; A Damper-Based Model for Coupling Force

Published in Proc. 28th Int. Conf. Information Fusion (FUSION), Rio de Janeiro, Brazil, 2025, pp. 1–8.

©2025 IEEE. DOI: 10.23919/FUSION65864.2025.11123934

Summary

Paper A estimates the articulation angle, the coupling force and the lateral velocity of each unit of a tractor-semitrailer with a UKF. The coupling is modeled by a penalty approximation, a directional damper between the units with the coupling force as a filter state. The paper names this damper-based coupling after the mechanism it was derived from. In effect, the damper

enforces a balanced force on each unit. Cornering stiffness and mass are known only as distributions and enter as consider parameters. The paper addresses a gap in the literature, where the coupling force has received limited attention and its uncertainty even less. The evaluation is in simulation on a few hundred cases of steady cornering, sinusoidal steering and a slope, with limited variation of vehicle properties. Of the challenges in Section 1.2, the paper addresses C1 with the penalty approximation of the constraint and C2 with the consider parameters. It considers two units only.

Main Findings

Lateral velocity and the longitudinal coupling force are captured well. The lateral coupling force is estimated less well, and the articulation angle estimate is over-confident in the circle scenario. The missing lateral load transfer of the planar model is the suspected cause. Low speed, reversing and heavy braking were left out after preliminary tests showed that the formulation did not handle them.

Author Contributions

Axel Ceder developed the model and the estimator, implemented and ran the simulations and wrote the paper. Lars Hammarstrand, Mats Jonasson, Murat Kumru and Leo Laine supervised the work and reviewed the manuscript. Volvo provided the base VTM simulator.

4.2 Paper B: Exact-Constraint Real-Time Estimation

Axel Ceder, Lars Hammarstrand, Mats Jonasson, Murat Kumru
Solving for the Fifth Wheel: Real-Time Exact-Constraint Coupling-Force Estimation for Articulated Heavy Vehicles
Submitted journal manuscript, 2026.

Summary

Paper B keeps the coupling constraint exact. The process model is the minimal-coordinate system of Section 2.3, solved numerically at every evaluation with

the coupling force recovered from the trailer's force balance, and the paper shows that this is tractable for online estimation. The damper formulation of Paper A is carried along as the baseline. The exact model is used in two filters, the UKF, where the solve is a point evaluation at each sigma point, and the IFT-EKF, where the Jacobians follow from the implicit function theorem. For two units the minimal-coordinate system also has a closed form. Cornering stiffnesses, trailer mass, yaw inertias and CoM positions, among others, are consider parameters. The evaluation is a held-out Monte Carlo study over eight maneuvers, including braking turns, grade and banking, low speed and reversing, with varied load, cornering stiffness and noise, and a comparison on real tractor-semitrailer data, where only the kinematic states have a reference and the force has none. Run time is measured for two to four units. The paper addresses all three challenges in Section 1.2: C1 with the exact constraint and its run time at two to four units, C2 with the consider parameters, and C3 with the maneuver set and the real data.

Main Findings

The exact form has lower error than the damper form and is cheaper per step at two units. At the median over the held-out cases the errors are small, below a centimeter per second in lateral velocity, hundredths of a degree in articulation angle and a few hundred newtons in coupling force. Divergence is rare and occurs in braking turns. Under lateral excitation the reported uncertainty is consistent or somewhat under-confident at the median, but it is not uniform across cases. The coupling force is weakly observable near standstill, where the wheel speeds drop out. On the test track the estimator is over-confident on the tractor lateral velocity (NEES 2.95), just outside the band for the articulation angle (2.04) and within it for the trailer lateral velocity (1.45), also when the start of the sequence, where the inner wheels are near lift-off, is excluded. The formulation is modular per unit and runs in real time on a desktop CPU at three and four units, with the IFT-EKF the cheapest at every number of units.

What the estimator needs falls into three groups: the specification (geometry and sensor mounting), quantities calibrated once on the vehicle (axle loads at standstill, and on the test vehicle also the sensor noise and the wheel radius from its own recordings, as the supplied values were not reliable), and fleet-wide priors carried as consider parameters (cornering stiffness, yaw inertias, trailer mass, CoM positions, the wheel radius's deviation from its calibrated mean

and its left-right difference, rolling resistance, and the steering offset). The consider parameters keep the estimator consistent over the fleet, but they do not reduce the parameter error itself, so the bias in an individual vehicle remains.

Author Contributions

Axel Ceder developed the model and the estimators, built the simulation and evaluation pipeline, ran the experiments and wrote the paper. Lars Hammarstrand, Mats Jonasson and Murat Kumru supervised the work and reviewed the manuscript. Volvo provided the base VTM simulator. The real data is from tests run by Volvo and Revere.

CHAPTER 5

Conclusions and Future Work

This thesis set out to estimate the articulation angle, the coupling force and the lateral velocities of a tractor-semitrailer from the standard sensor set, with the exact coupling constraint kept in a real-time estimator. Paper A skirts the constraint with a damper formulation, and Paper B shows that the constraint can be used as it is. The exact formulation keeps the constraint without simplifying the physics, runs faster than the damper formulation for two units, and reports an uncertainty that holds at the median in simulation. The formulation is modular per unit and runs in real time at three and four units. Fleet-level knowledge, carried as consider parameters, is what makes the reported uncertainty consistent over varying vehicle properties.

The first challenge was the coupling constraint. For two units the coupling can be described exactly while allowing online estimation. The exact form is more accurate and cheaper per step than the damper form, and the formulation is modular per unit and runs in real time at three and four units, though its accuracy beyond two units was not evaluated. The second challenge was fleet variation. The estimator runs from the specification, from fleet-wide priors carried as consider parameters, and from a few quantities calibrated once on the vehicle. Its reported uncertainty is consistent at the median in simulation,

erring on the under-confident side, but it is not uniform across cases, and it does not hold on the test track. The consider parameters do not reduce the parameter error itself, so the bias in an individual vehicle remains, and more accuracy for the individual vehicle requires more information. The third challenge was situations and sensors. The estimator is accurate across the maneuvers of Paper B in nominal, high-friction driving. It loses accuracy when the vertical load moves almost entirely to one side, and the coupling force is weakly observable near standstill, where the wheel speeds drop out.

5.1 Future Outlook

This section lists the directions that follow from the estimators and findings in this work.

Load transfer and roll: The physical model is in many respects a simple model, which can be improved upon. Some improvements follow from the delimitations, where the model has only been evaluated at high friction and in nominal driving conditions. The main improvement, however, lies in load transfer and roll. As mentioned, the estimator loses accuracy when the vertical load moves almost entirely to one side. Attempts to make a simplified version were made but did not give the desired results. In a more complete model, where the vertical forces depend on the unknowns, the constraint is more complex. On initial inspection the closed form is expected to be lost under drive and brake torque, requiring an iterative solve per step, which the UKF and IFT-EKF can still use.

Longer combinations and sensor selection: The described model is already modular and runs in real time at both three and four units. A full evaluation has not been done, but early results suggest similar performance with some tuning. The A-double, a tractor with a semitrailer, dolly and second semitrailer, is of particular interest, and a sensor sweep could show which sensors are valuable.

Low speed and friction: The coupling force is weakly observable near standstill, where the wheel speeds drop out. The tire model has only been evaluated at high friction and low load transfer, and low-friction surfaces remain outside the evaluation.

Real data and parameters: Consistency at the median holds in simulation but not on the test track. The test-track vehicle is assumed to belong to the same fleet as the simulated population, which may not fully hold. More

data has to be collected, both to estimate the parameter distributions and for validation. The consider parameters have been used to capture the fleet-wide uncertainty. If a parameter is observable, it is a candidate for online estimation instead.

The simulator and the real data are released in association with Paper B. The simulated dataset can be regenerated from the simulator with a single Python call, no license required. Simulator and data are meant to give other researchers in the field a benchmark against which solutions, ours included, can be directly compared.

5.2 Alternative Paths Explored

This section reports approaches that were tried and set aside, as they may save others the effort.

FD-EKF on the exact model: Before the IFT-EKF was implemented, a finite-difference EKF was in its place. It used the same exact-constraint model, but instead calculated the Jacobians by numerical differentiation. It turned out to be about as fast as the UKF and two to three times slower than the IFT-EKF for every number of units. It also lost accuracy faster than the alternatives in highly nonlinear situations. For a proper EKF solution, the IFT-EKF beats it in everything except compilation time.

Slip as displacement: The idea is simple. Instead of a slip angle, the tire is described by its deflection. Per wheel, the deflection $\boldsymbol{s} = [s^x,\, s^y]^\top$ of the tire at the contact is the state. It follows the first-order relaxation-length model of the single-contact-point transient tire model [14, Sec. 7.2.1], applied in both directions,

$$\dot{\boldsymbol{s}} = \boldsymbol{v}_s - \frac{|v_r|}{\sigma}\,\boldsymbol{s}, \qquad \boldsymbol{v}_s = \begin{bmatrix} v_w^x - v_r \\ v_w^y \end{bmatrix}, \qquad \boldsymbol{f} = -\boldsymbol{C}\boldsymbol{s}, \tag{5.1}$$

with $\boldsymbol{v}_s$ the slip velocity of the wheel, $v_r = \omega R$ the rolling speed, $\boldsymbol{C} = \mathrm{diag}(C_x, C_y)$ the deflection stiffness and σ the relaxation length. Pacejka [14, Sec. 7.2.1] uses $|V_x|$, whereas here $|v_r|$ is used. With that replacement at steady state $\boldsymbol{s} = \sigma\,\boldsymbol{v}_s/|v_r|$, so $\boldsymbol{f} = -\boldsymbol{C}\sigma\,\boldsymbol{v}_s/|v_r|$: $\boldsymbol{C}\sigma$ is the slip stiffness, and the usual slip quantities are approximately recovered, exactly at zero longitudinal slip, without dividing by a velocity in the dynamics. As $v_r \to 0$, (5.1) reduces to $\dot{\boldsymbol{s}} = \boldsymbol{v}_s$: the deflection integrates the displacement of the wheel relative to

the ground, a spring to ground instead of a slip angle. This then avoids the singularity of the slip angle without a regularized slip $\max(|v_w^x|, v_\varepsilon)$, and does not misrepresent transients as steady-state. At high speed the lag becomes small and the models converge. At low speeds each approach has its own weaknesses, but the results are similar. The caveat is one state per direction and wheel, so the estimator carries a considerably larger state space. The approach may still be useful. For example, a cornering stiffness estimator might be served well with a transient description of force instead of the steady-state assumption.

Pseudo lateral load transfer: To consider load transfer without committing to a full three-dimensional vehicle combination, the load transfer was described as a combination of $v^x\omega^z$ and a consider parameter. This did not improve results, and added complexity. A description that includes the roll dynamics is likely required.

References

[1] International Transport Forum, "High capacity transport: Towards efficient, safe and sustainable road freight," OECD Publishing, Paris, International Transport Forum Policy Papers 69, 2019.

[2] A. L. Dunn, "Jackknife stability of articulated tractor semitrailer vehicles with high-output brakes and jackknife detection on low coefficient surfaces," Ph.D. dissertation, The Ohio State University, 2003.

[3] C. B. Winkler, "Improving the dynamic performance of multitrailer vehicles: A study of innovative dollies," University of Michigan Transportation Research Institute, Ann Arbor, MI, Tech. Rep., 1986.

[4] U. Erdinc, M. Jonasson, M. Sadeghi Kati, B. Jacobson, J. Fredriksson, and L. Laine, "Safe operating envelope based on a single-track model for yaw instability avoidance of articulated heavy vehicles," *Veh. Syst. Dyn.*, vol. 64, no. 8, pp. 2138–2161, 2023.

[5] K. B. Singh, M. A. Arat, and S. Taheri, "Literature review and fundamental approaches for vehicle and tire state estimation," *Veh. Syst. Dyn.*, vol. 57, no. 11, pp. 1643–1665, 2019.

[6] A. Habibnejad Korayem, E. Hashemi, A. Khajepour, and B. Fidan, "Estimation of Vehicle–Trailer Hitch-Forces and Lateral Tire Forces Independent of Trailer Type and Geometry," *J. Dyn. Syst. Meas. Control*, vol. 144, no. 5, p. 051 005, 2022.

[7] W. (Deng and X. Kang, "Parametric Study on Vehicle-Trailer Dynamics for Stability Control," *SAE Trans.*, vol. 112, pp. 1411–1419, 2003.

[8] M. Bouteldja, A. Koita, V. Dolcemascolo, and J. Cadiou, "Prediction and Detection of Jackknifing Problems for Tractor Semi-Trailer," in *Proc. IEEE Veh. Power Propuls. Conf. (VPPC)*, 2006, pp. 1–6.

[9] "ISO 8855:2011: Road vehicles — Vehicle dynamics and road-holding ability — Vocabulary," International Organization for Standardization, Geneva, Switzerland, Standard, 2011.

[10] B. J. H. Jacobson, I. Johansson, M. Jonasson, E. Drenth, and P. Johannessen, "Compendium in Vehicle Motion Engineering," Chalmers University of Technology, Gothenburg, Sweden, Tech. Rep., 2021.

[11] T. Ghandriz, B. Jacobson, P. Nilsson, L. Laine, and N. Fröjd, "Computationally Efficient Nonlinear One- and Two-Track Models for Multitrailer Road Vehicles," *IEEE Access*, vol. 8, pp. 203 854–203 875, 2020.

[12] O. Lindgärde and M. Gäfvert, "A 9-DOF tractor-semitrailer dynamic handling model for advanced chassis control studies," *Veh. Syst. Dyn.*, vol. 41, no. 1, pp. 51–82, 2004.

[13] N. Fröjd, *Handling Analysis and Control Development of Commercial Trucks with Volvo Transport Models, Matlab EXPO*, 2021.

[14] H. Pacejka, *Tyre and Vehicle Dynamics* (Automotive engineering). Elsevier Science & Technology, 2006, ISBN: 978-0-7506-6918-4.

[15] M. Hjort, S. Kharrazi, N. Fröjd, and T. Siltanen, "Tyre modelling for high capacity vehicle simulations," in *HVTT16*, 2021.

[16] A. A. Shabana, *Dynamics of Multibody Systems*, 3rd ed. Cambridge, UK: Cambridge University Press, 2005, ISBN: 978-0-521-85011-7.

[17] S. Särkkä, *Bayesian Filtering and Smoothing*. Cambridge, UK: Cambridge University Press, 2013, ISBN: 978-1-107-03065-7.

[18] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[19] S. J. Julier and J. K. Uhlmann, "New extension of the Kalman filter to nonlinear systems," in *Proc. SPIE*, 1997, p. 182.

[20] E. Wan and R. Van Der Merwe, "The unscented Kalman filter for nonlinear estimation," in *Proc. IEEE Adapt. Syst. Signal Process., Commun., Control Symp. (AS-SPCC)*, 2000, pp. 153–158.

[21] C. Cheng and D. Cebon, "Parameter and state estimation for articulated heavy vehicles," *Veh. Syst. Dyn.*, vol. 49, no. 1-2, pp. 399–418, 2011.

[22] S. F. Schmidt, "Application of state-space methods to navigation problems," in *Advances in Control Systems*, C. T. Leondes, Ed., vol. 3, New York: Academic Press, 1966, pp. 293–340.

[23] J. Stauch and M. Jah, "Unscented Schmidt–Kalman Filter Algorithm," *J. Guid. Control Dyn.*, vol. 38, no. 1, pp. 117–123, 2015.

[24] Z. Chen, C. Heckman, S. Julier, and N. Ahmed, "Weak in the NEES?: Auto-tuning kalman filters with bayesian optimization," in *ICIF*, 2018.

[25] R. Pastorino, D. Richiedei, J. Cuadrado, and A. Trevisani, "State estimation using multibody models and nonlinear Kalman filters," in *The 2nd Joint International Conference on Multibody System Dynamics*, Stuttgart, Germany, May 2012.

[26] D. Simon, "Kalman Filtering With State Constraints: A Survey of Linear and Nonlinear Algorithms," *IET Control Theory Appl.*, vol. 4, no. 8, pp. 1303–1318, 2010.