

Chapter 8

Identification of Tyre–road Friction Conditions

8.1 Introduction

Tyre–road friction characteristics are deeply interlaced with all vehicle safety oriented control systems as road conditions strongly affect the controlled system behaviour. Thus, the capability of estimating in real-time the friction conditions may provide a valuable source of information for any active vehicle control system. In particular, friction information can be used to enhance the performance of wheel slip control systems.

In this chapter we address three different problems related with friction estimation. Specifically, Section 8.2 illustrates an approach that is capable of estimating the sign of the slope of the friction curve, thereby allowing one to detect if the system is operating in the stable or in the unstable region of the friction curve. In fact, as largely discussed in this book, the equilibrium points associated with the wheel braking dynamics are stable for values of the wheel slip before the peak and unstable for those beyond the peak.

Hence, an online detection of the slope of the friction curve can be exploited to adapt and to optimise the closed-loop performance of wheel slip control systems. The advantage of this identification method is that it can be implemented also with a very limited set of sensors.

Secondly, in Section 8.3 an approach to the problem of estimating both the slip value corresponding to the peak of the friction curve and the parameters of the Burckhardt friction model (see Section 2.1) is presented. This is done by setting up a curve fitting problem which is then solved by two different identification approaches, namely a least squares and a maximum likelihood approach, arising from different parametrisations of the friction curve. A detailed analysis of the merits and drawbacks of the two approaches is also provided, which considers both the obtained accuracy in the estimated parameters and the convergence issues which have to do with the length of the available data set.

Finally, Section 8.4 presents an approach for estimating the instantaneous vertical and longitudinal forces from in-tyre acceleration measurements. Specifically, an appropriate set of sensors and regressors is illustrated, based on the measurements provided both by standard vehicle sensors (wheel encoders) and an accelerometer mounted directly in the tyre. Such estimates are based on the idea of extracting information from the phase shift between the wheel hub and the tyre, which is due to the transmission of traction and braking forces exerted on the tyre itself.

8.2 Detection of the Friction-curve Peak by Wheel-deceleration Measurements

In this section we focus on estimating the sign change of the slope of the curve $\mu(\lambda)$, which is responsible for the stability properties of the open-loop equilibrium points of the wheel dynamics. In particular, recall that for constant values of the braking torque, *i.e.*, for $T_b = \bar{T}_b$, the open-loop equilibrium points associated with slip values beyond the peak of the tyre–road friction curve are unstable (see Section 2.5.1).

The identification algorithm is formulated based on two different sensors configurations, *i.e.*, with and without wheel slip measurement (or estimation) available. This allows us to cover all the possible equipments available on commercial vehicles. Moreover, it is illustrated how such an identification algorithm can be employed within a supervisory control logic to enhance safety properties and performance of active braking systems.

To develop the identification approach, the considered dynamical model is the single-corner model discussed in Section 2.3. Once again, we treat the vehicle speed as a slowly-varying parameter and concentrate on the wheel dynamics only.

As discussed in Section 2.5.1, if we linearise the wheel dynamics (see the first equation of system (2.18)), the transfer functions $G_\lambda(s)$ from δT_b to $\delta\lambda$ and $G_\eta(s)$ from δT_b to $\delta\eta$ can be obtained. Specifically, they have the form

$$G_\lambda(s) = \frac{\frac{r}{J\bar{v}}}{s + \frac{\mu_1(\bar{\lambda})F_z}{m\bar{v}} \left((1 - \bar{\lambda}) + \frac{mr^2}{J} \right)}, \quad (8.1)$$

and

$$G_\eta(s) = \frac{\frac{r}{Jg} \left(s + \frac{\mu_1(\bar{\lambda})F_z}{m\bar{v}} (1 - \bar{\lambda}) \right)}{s + \frac{\mu_1(\bar{\lambda})F_z}{m\bar{v}} \left((1 - \bar{\lambda}) + \frac{mr^2}{J} \right)}. \quad (8.2)$$

For what follows, it is worth recalling that $G_\lambda(s)$ and $G_\eta(s)$ are both first-order transfer functions, characterised by the same pole. Moreover, recall also that the vehicle speed value \bar{v} considered for the linearisation acts only as a scaling parameter on such pole, but it does not affect its sign.

In fact, as was shown in Section 2.5.1, the (real) pole of the transfer functions (8.1) and (8.2) is negative if and only if the following inequality holds

$$\frac{\mu_1(\bar{\lambda})F_z}{m\bar{v}} \left((1 - \bar{\lambda}) + \frac{mr^2}{J} \right) > 0,$$

which can be reduced to

$$\mu_1(\bar{\lambda}) > 0,$$

where $\mu_1(\bar{\lambda})$ is as in (2.37) and it represents the slope of the tyre-road friction curve. Thus, the linear systems with transfer functions $G_\lambda(s)$ and $G_\eta(s)$ are unstable if the equilibrium point $\bar{\lambda}$ is *beyond* the peak of the friction curve. Accordingly, to detect the stability properties of wheel dynamics equilibrium points it is necessary to monitor the sign-change of $\mu_1(\bar{\lambda})$, *i.e.*, the slope of the tyre-road friction curve. It is worth pointing out that detecting a change of stability also gives a means for detecting the pair $(\mu_{\text{Max}}, \lambda_{\text{Max}})$ *i.e.*, the maximum tyre-road friction coefficient available on the current road condition and the corresponding value of the wheel slip.

8.2.1 Online Detection of the Sign of the Friction Curve Slope

We will now propose a strategy to monitor online the sign-change of $\mu_1(\bar{\lambda})$, *i.e.*, the slope of the tyre-road friction curve. To this end, we deal directly with the wheel dynamics, trying to find approximate static algebraic relationships linking measurable variables to the slope sign itself.

Differentiating the wheel dynamics given in the first equation of the single-corner model (2.15), we obtain

$$\begin{aligned} \ddot{\omega} &= \frac{d}{dt} \left(\frac{rF_x}{J} - \frac{T_b}{J} \right) = \frac{d}{dt} \left(\frac{rF_z\mu(\lambda)}{J} - \frac{T_b}{J} \right) \\ &= \frac{rF_z}{J} \frac{d\mu(\lambda)}{d\lambda} \Big|_{\lambda=\bar{\lambda}} \dot{\lambda} - \frac{\dot{T}_b}{J} = \frac{rF_z}{J} \mu_1(\bar{\lambda}) \dot{\lambda} - \frac{\dot{T}_b}{J}. \end{aligned}$$

We then define

$$H := \ddot{\omega} + \frac{1}{J} \dot{T}_b = \frac{rF_z}{J} \mu_1(\bar{\lambda}) \dot{\lambda} = \gamma \mu_1(\bar{\lambda}) \dot{\lambda}, \quad \gamma \in \mathbb{R}^+. \quad (8.3)$$

Equation (8.3) shows that we can relate the sign of $\mu_1(\bar{\lambda})$ to the wheel deceleration derivative (*i.e.*, the wheel *Jerk*) and to the braking torque derivative.

The above relation, though, has to be further analysed before being directly employed. In fact, the sign of H depends both on the sign of $\mu_1(\bar{\lambda})$ and on the sign of the wheel slip derivative $\dot{\lambda}$. In fact, it holds that

$$H > 0 \Leftrightarrow (\mu_1(\bar{\lambda}) > 0 \text{ AND } \dot{\lambda} > 0) \text{ OR } (\mu_1(\bar{\lambda}) < 0 \text{ AND } \dot{\lambda} < 0),$$

while

$$H < 0 \Leftrightarrow (\mu_1(\bar{\lambda}) < 0 \text{ AND } \dot{\lambda} > 0) \text{ OR } (\mu_1(\bar{\lambda}) > 0 \text{ AND } \dot{\lambda} < 0).$$

Based on this, two alternative strategies for solving such ambiguity and for deriving an expression for H suitable for online detection of the sign of the friction curve slope can be derived. The two strategies differ one from the other in that they are devised assuming that wheel slip measurement (or estimation) is either available or not available.

Let us start by assuming that the wheel slip can be either measured or estimated (for wheel slip estimation see Chapter 5). This assumption implies that the proposed strategy can be implemented on all passenger cars equipped with ABS and ESC sensors.

In this case, the slope sign can be derived as

$$\text{sign}(H) = \text{sign}\left(\ddot{\omega} + \frac{1}{J}\dot{T}_b\right) \text{sign}(\dot{\lambda}) = \text{sign}(\mu_1(\bar{\lambda})). \quad (8.4)$$

Hence, slope sign detection becomes a pure signal processing problem.

To cope with the disturbances affecting the measured signals, in the implementation of the friction curve slope sign detection strategy based on Equation 8.4, one should not use a pure zero-crossing detection algorithm, but it is necessary to define a small negative threshold on H . Such a threshold must be experimentally tuned according to the measurement noise in the available data and should fit all road conditions.

We now discuss how to estimate the slope sign when no slip measurement (or estimation) is available. The motivation for this analysis is to devise an identification strategy that can be implemented in any passenger car equipped only with standard ABS sensors (*i.e.*, wheel encoders and pressure sensors, but no longitudinal accelerometer).

Consider the expression of the wheel slip derivative

$$\dot{\lambda} = \frac{r}{v} \left(\frac{\omega}{v} \dot{v} - \dot{\omega} \right) = \frac{r}{v} \left(\frac{\omega}{v} \dot{v} + \frac{g}{r} \eta \right).$$

Assuming that the rate of change of the vehicle acceleration \dot{v} is negligible with respect to that of the wheel deceleration $\dot{\omega}$, one has

$$\dot{\lambda} = -\dot{\omega} \frac{r}{v} = \eta \frac{g}{v}.$$

Accordingly, the slope sign estimator becomes

$$\text{sign}(H) = \text{sign}\left(\ddot{\omega} + \frac{1}{J}\dot{T}_b\right) \text{sign}(\dot{\lambda}) \cong \text{sign}\left(\ddot{\omega} + \frac{1}{J}\dot{T}_b\right) \text{sign}(\eta). \quad (8.5)$$

Thus, the estimation of the sign of $\mu_1(\bar{\lambda})$ can be recast as the detection of the sign of two components, both of which are measurable on every passenger car.

It is worth analysing the effects of the approximation used in deriving (8.5). As we are concerned with braking manoeuvres, by approximating the slip velocity $\dot{\lambda}$ with the normalised wheel acceleration η , we neglect a negative term dependent on \dot{v} . Hence the following inequality holds:

$$\frac{\omega}{v} \dot{v} + \frac{g}{r} \eta < \frac{g}{r} \eta. \tag{8.6}$$

Inequality (8.6) shows that η changes sign before $\dot{\lambda}$ when the wheel slip is increasing, while the opposite is true when the wheel slip is decreasing (*i.e.*, during the pedal release phase). In both cases, there is a time interval during which η and $\dot{\lambda}$ lose their concordance and the algorithm based on Equation 8.5 fails in detecting the slope sign. This means that expression (8.5) can be error prone at the beginning of very hard braking manoeuvres.

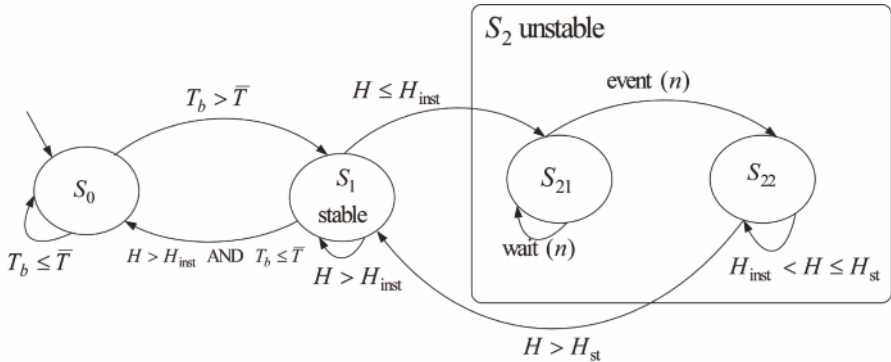


Figure 8.1 FSM representation of the proposed ABS supervisory control logic

8.2.2 ABS Supervisory Control Logic

To enhance passenger safety, it is possible to employ the slope sign detection strategy as a supervisory control logic, that commands, for example, the ABS system control actions. To this end, we are interested in detecting the slope sign change in both directions, *i.e.*, from positive to negative and *vice versa*.

In order for the proposed logic to be integrated with standard ABS systems, we assume that no slip measurement (or estimation) is available. Hence, the proposed results are based on Equation 8.5. An FSM representation of a supervisory control logic is depicted in Figure 8.1.

As can be seen, it is composed of three states, each with an associated control action. Specifically, one has:

- S_0 : the initial and final state. No control action is sent to the core ABS system.
- S_1 : the system is in the stable region of the tyre–road friction curve, hence the ABS system is allowed to increase the braking torque T_b .
- S_2 : the system is in the unstable region of the tyre–road friction curve, hence the ABS system is commanded to decrease the braking torque T_b . S_2 is further decomposed into S_{21} and S_{22} , whose meaning will be clarified below.

The overall supervisory logic works as follows: the system remains in S_0 while no braking occurs, *i.e.*, while $T_b < \bar{T}$, where \bar{T} is a properly tuned threshold. When $T_b \geq \bar{T}$, a transition to the stable region S_1 occurs. The system remains in this state until $H > H_{\text{inst}}$ where H_{inst} is the (negative) threshold identifying a negative slope sign.

Once $H \leq H_{\text{inst}}$, the system signals the crossing of the stability boundary and enables the transition to the state S_2 . Actually, the system enters the substate S_{21} , where it waits for n ms before entering S_{22} . The rationale under the *wait* action is that even if the release action is triggered by the transition from S_1 to S_2 , current ECU network topologies do not allow for an immediate system response.


Accordingly, the substate S_{21} ensures that the detection of a positive slope sign is not missed by evaluating H at the same time instant at which the release action takes place. While the instability condition persists (*i.e.*, as long as the estimated slope sign is negative), the system remains in S_{22} and the release phase is continued by the ABS system. When $H \geq H_{\text{st}}$ (H_{st} being the positive threshold identifying positive slope sign), stability is re-gained and a transition to S_1 takes place. Finally, when the braking manoeuvre is over, the system goes back to the final state S_0 .

8.2.3 Experimental Results


To analyse the performance of the identification algorithm, test drives have been carried out by performing strong braking manoeuvres on two different road conditions, *i.e.*, a high-grip asphalt road and a low-grip off-road.

All the measured signals (wheel encoders, longitudinal accelerometer and pressure sensors) are assumed to come from properly calibrated sensors. For a discussion on the processing issues related to wheel encoders the reader is referred to Appendix B.

Figures 8.2(a) and 8.2(b) show the time histories of the wheel slip and of the detection of the stability boundary crossings obtained *via* Equation 8.4, *i.e.*, assuming that wheel slip measurements are available, on high-grip and



(a)



(b)

Figure 8.2 Time histories of the wheel slip (top) and of the detection of the stability boundary crossings (bottom) obtained with (8.4). (a) high-grip road, (b) low-grip off-road

low-grip roads, respectively. Notice that the slope sign detection is very accurate on both road conditions. In fact, analysing the results obtained with the friction curve slope sign change detection algorithm based on Equation 8.4 (averaging the results obtained in all tests on both road conditions), we have $E[\lambda_{Max}] = 0.0854$ on the high-grip road and $E[\lambda_{Max}] = 0.1357$ on the low-grip off-road, which is compatible with the tyre–road friction conditions of the test grounds.

To visually inspect the performance of the estimation algorithms when no wheel slip measurement (or estimation) is available, refer to Figures 8.3(a) and 8.3(b), which depict the time histories of the wheel slip and of the detection of the stability boundary crossings obtained via Equation 8.5 on high-grip and low-grip roads, respectively. Analysing the results obtained with the slope sign detection algorithm based on (8.5), (averaging the results obtained in all tests on both road conditions) one finds $E[\lambda_{Max}] = 0.0648$ on the high-grip road and $E[\lambda_{Max}] = 0.1283$ on the low-grip off-road. Thus, the slope sign detection can also be reliably performed with no direct information on the wheel slip.

Finally, it is worth investigating the adaptation of the identification algorithm to the supervisory control logic presented in Section 8.2.2. We recall that it asks us to detect the slope sign change in both directions.

Figures 8.4(a) and 8.4(b) show the results obtained on dry and off-roads, respectively. Apparently, the proposed supervisory logic can be profitably employed both to enhance the safety properties of ABS systems and to optimise their performance. In fact, the parameters of the ABS controller might be adaptively tuned according to the current system stability properties.

8.3 Real-time Identification of Tyre–road Friction Conditions

In this section, we consider the problem of estimating the wheel slip value corresponding to the abscissa of the peak of the tyre–road friction curve combined with the capability of identifying also the full parametrisation of the tyre–road friction curve itself.

The characteristics of the friction estimation strategies are analysed and tested both in simulation and on experimental data. In the latter case, the performance of the identification techniques is tested also in combination with the wheel speed estimation algorithm discussed in Section 5.3, in order to discuss its effects on the accuracy of the final results.

Specifically, the identification approach is based on the solution of a curve fitting problem formulated using a properly parameterised friction model. The slip value corresponding to the curve peak is subsequently estimated from the fitted curve. To this aim, the Burckhardt friction model (see Section 2.1) is considered.





Figure 8.3 Time histories of the wheel slip (top) and of the detection of the stability boundary crossings (bottom) obtained with (8.5). (a) high-grip road, (b) low-grip off-road




Figure 8.4 Time histories of the wheel slip (*solid line*) and detection of the stability boundary crossings (*dashed line*). (a) high-grip road, (b) low-grip off-road

In presenting this topic, we assume that the reader has a basic familiarity with linear system identification methods, specifically with the least squares (LS) approach and with the maximum likelihood (ML) approach. For more details on the identification algorithms, the interested reader is referred to, *e.g.*, [61, 62, 99].

Since the Burckhardt model is nonlinear in the parameters, an iterative ML estimation technique is applied first. Further, to evaluate the practical applicability of the method for online use, a recursive version of the algorithm is also investigated. Secondly, a (recursive) LS method is introduced, based on an *ad hoc* parametrisation of the friction model which is linear in the parameters.

To set up the identification problem, let us assume that the signals ω , v , λ and T_b are available. As such, one can work on the single-corner dynamics (see also Equation 2.18)

$$\dot{\lambda} = -\frac{1}{v} \left(\frac{(1-\lambda)}{m} + \frac{r^2}{J} \right) F_z \mu(\lambda) + \frac{r}{vJ} T_b, \quad (8.7)$$

and invert it to obtain data points of the type $\mu(\omega, v, T_b, \lambda)$, which constitute the basis for collecting data suitable for the estimation of the tyre–road friction curve $\mu(\lambda)$. Actually, note that the angular wheel velocity ω and the braking torque T_b can be assumed to be measurable (in particular, T_b is assumed to be proportional to the pressure in the hydraulic braking circuit by a known coefficient, see Equation 1.1), while v is generally not available and needs to be estimated. Finally, λ can be determined from ω and v .

Namely, expressing v as $\omega r / (1 - \lambda)$, $\mu(\lambda)$ can be easily obtained from (8.7) as

$$\mu(\lambda; t) = \frac{\frac{1-\lambda(t)}{J\omega(t)} T_b(t) - \dot{\lambda}(t)}{\frac{1-\lambda(t)}{\omega(t)r} F_z(t) \left(\frac{1-\lambda(t)}{m} + \frac{r^2}{J} \right)}, \quad (8.8)$$

where the dependence on time is made explicit in order to point out that (8.8) holds at each time instant. Note that the right-hand side of (8.8) depends on $\omega(t)$, $\lambda(t)$, $\dot{\lambda}(t)$, r , $T_b(t)$ and $F_z(t)$. Thus, in principle, if noise-free measurements of such variables were available, one could estimate the corresponding sample of $\mu(\lambda; t)$ exactly.

However, $\omega(t)$ and $T_b(t)$ are measured from noisy sensors, $\lambda(t)$ is derived *via* speed estimation and hence it is affected by the estimation error and the vertical load $F_z(t)$, together with the wheel radius r , differs from its static value depending on the braking (or traction) induced pitch angle and tyre characteristics (see also Appendix B). Note, however, that according to the longitudinal tyre force model given in (2.12), the vertical load $F_z(t)$ acts only as a scaling factor on the friction curve, thus not altering the abscissa of the maximum.

In summary, (8.8) is in fact a highly nonlinear function of data and noise, which can only provide approximate and noisy samples of $\mu(\lambda; t)$. Accordingly, to set up a curve fitting problem with the aim of identifying tyre–road friction conditions, (8.8) is used to compute the observations from data, while suitable parametrisations based on Burckhardt model in (2.13) are employed to formulate the tyre–road friction relation (see the next section).

8.3.1 Identification Strategies

Based on the problem setting given in the previous section, a curve fitting problem can be set up to estimate $\mu(\lambda)$, and thus to obtain λ_{Max} from the estimated friction curve. The first of these two tasks is addressed by minimisation of a standard quadratic error function of the form

$$J(\vartheta) = \frac{1}{N} \sum_{i=1}^N \epsilon(\lambda_i, \vartheta)^2, \quad (8.9)$$

where N is the number of samples and $\epsilon(\lambda_i, \vartheta) = \mu(\lambda_i) - \widehat{\mu}(\lambda_i, \vartheta)$ is the estimation error. The next subsections focus on the derivation of suitable methods for fitting the tyre-road friction curve.

8.3.1.1 The Maximum Likelihood Approach

The most direct approach for the estimation of the tyre-road friction curve involves the use of the Burckhardt model (2.13), where ϑ_r is the vector of unknown parameters to be determined. Unfortunately, the structure of that model is nonlinear in the parameters and simple algorithms like the LS method cannot be applied to this case. The computationally more intensive ML approach must be used instead, [99].

The method performs an iterative minimisation of the fitting error criterion (8.9) by means of a quasi-Newton method [99], which ultimately amounts to adjusting the parameter vector at each iteration on the basis of the following expression:

$$\vartheta^{(r+1)} = \vartheta^{(r)} + \left(\sum_{i=1}^N \xi(\lambda_i) \xi(\lambda_i)^T \right)^{-1} \left(\sum_{i=1}^N \xi(\lambda_i) \epsilon(\lambda_i) \right), \quad (8.10)$$

where

$$\xi(\lambda_i) = -\frac{d\epsilon(\lambda_i)}{d\vartheta} = [1 - e^{-\lambda_i \vartheta_{r2}} \quad \vartheta_{r1} \lambda_i e^{-\lambda_i \vartheta_{r2}} \quad -\lambda_i]^T.$$

A crucial step in the ML approach is the initialisation of the parameter vector. An *ad hoc* empirical approach has been adopted for this purpose. Actually, for low values of the wheel slip, the Burckhardt model can be approximated as

$$\mu(\lambda, \vartheta_r) \cong (\vartheta_{r1} \vartheta_{r2} - \vartheta_{r3}) \lambda - \vartheta_{r1} \vartheta_{r2}^2 \lambda^2. \quad (8.11)$$

Now, since the parameter ϑ_{r3} varies in an extremely small range, from 0.06 for snow to 0.67 for cobblestone (see Table 8.1), an average value of 0.4 can be reasonably assumed as initial value $\widehat{\vartheta}_{r3}(0)$. Then, using the first available data, which appropriately correspond to low slip values, one can interpolate the quadratic approximation above to derive the initial values $\widehat{\vartheta}_{r1}(0)$ and $\widehat{\vartheta}_{r2}(0)$.

Several recursive versions of the general iterative batch algorithm described above have been developed in the literature, see *e.g.*, [62]. These exploit the structural similarity between the iteration equation of the ML method and the well-known LS equation and recover the same computational schemes of the

recursive LS (RLS) methods. Notice, however that in this case the recursive scheme mixes the data recursion with the iterative mechanism inherent in the ML approach, so that convergence may be harder to obtain with recursive ML (RML), as opposed to RLS.

8.3.1.2 The Least Squares Approach

In view of the potential convergence problems of the RML approach discussed in the previous section (which will be thoroughly analysed on the data of interest in Section 8.3.2), an alternative estimation scheme is presented based on a linear regression reformulation of the friction model and LS estimation.

The Burckhardt model (2.13) is nonlinear by way of the exponential term in ϑ_{r2} , which varies in a relatively small range of values depending on the road conditions. As an alternative, such a term can be approximated using a linear combination of fixed exponentials, suitably dispersed to cover the whole range of interest [33]. This results in the model

$$\hat{\mu}(\lambda) = a_1 e^{-b_1 \lambda} + a_2 e^{-b_2 \lambda} + \dots + a_{n-2} e^{-b_{n-2} \lambda} + a_{n-1} \lambda + a_n, \quad (8.12)$$

where a_{n-1} and a_n equal $-\vartheta_{r3}$ and ϑ_{r1} in model (2.13). If the $b_i, i = 1, \dots, n-2$ exponents are fixed *a priori*, model (8.12) can be recognised as a linear regression

$$\hat{\mu}(\lambda) = \varphi(\lambda)^T \theta,$$

where

$$\begin{aligned} \varphi(\lambda) &= [e^{-b_1 \lambda}, \dots, e^{-b_{n-2} \lambda}, \lambda, 1]^T \\ \theta &= [a_1, \dots, a_n]^T, \end{aligned}$$

to which LS estimation methods can be directly applied. The number of exponentials in the approximate model has been set to four *via* a trial and error process, aimed at establishing a satisfactory compromise between model flexibility and size (to avoid overfitting).

Parameters $b_i, i = 1, \dots, 4$ have been chosen as uniformly distributed in the range [4, 100], which includes all the values of ϑ_{r2} in the considered road conditions. More precisely, $b_1 = 4$, $b_2 = 36$, $b_3 = 68$ and $b_4 = 100$.

8.3.2 Numerical Analysis

In order to evaluate the behaviour of the estimation strategies described in the previous section, it is important to preliminarily discuss which are the performance levels to be sought in the considered application.

Specifically, accuracy requirements are set by the characteristics of typical vehicle electronic and control devices. More precisely, as far as the expected performances are concerned, a reasonable objective is to obtain an overall accuracy in the estimation of the abscissa of the peak of the friction curve of $\pm 5\%$, which is considered appropriate for ABS control [90]. Furthermore, the estimator should provide reliable information within the time instant at which the braking controller would be activated. Since this time interval can be approximately quantified to be at most of 500 ms, and data in these applications are typically sampled at 200 Hz, the estimation algorithms have been tested with 100 samples.

To investigate the performance of the two estimation strategies introduced in the previous section, simulation tests, which were carried out based on the single-corner model dynamics, are now presented. In order to consider a more realistic setting, though, all identification methods have been tested on noisy simulation data.

Specifically, zero mean white noises have been added to the wheel speed ω , to the vehicle speed v and to the braking torque T_b , with $\sigma_\omega^2 = 0.005 \text{ rad}^2/\text{s}^2$, $\sigma_v^2 = 0.25 \text{ m}^2/\text{s}^2$ and $\sigma_{T_b}^2 = 10 \text{ N}^2\text{m}^2$, respectively. In the case of v and ω , which represent the quantities based on which the wheel slip is computed, the noise variance values should correctly model the fact that the noise acting on the wheel speed ω is only due to measurement noise (ω is directly measured by means of a wheel encoder), while the noise acting on the vehicle speed is also due to estimation errors, see also the discussion in Chapter 6.

Let us first analyse the results obtained with the ML approach. To do this, consider Table 8.1, which reports the results on the parameter estimation obtained with ML and RML algorithms. Actually, while an almost perfect curve fitting (and a quite accurate parameter estimation) is obtained with the batch ML approach in all tested road conditions, the same does not apply to the computationally more affordable RML method.

Table 8.1 Summary of the $\mu(\lambda)$ estimation obtained with ML and RML on noisy simulation data

		Dry asphalt	Wet asphalt	Cobblestone	Snow
True values	ϑ_1	1.28	0.86	1.37	0.19
	ϑ_2	23.99	33.82	6.46	94.13
	ϑ_3	0.52	0.35	0.67	0.06
ML	$\hat{\vartheta}_1$	1.29	0.86	1.48	0.18
	$\hat{\vartheta}_2$	22.69	31.94	5.79	90.95
	$\hat{\vartheta}_3$	0.57	0.36	0.85	0.04
	# iter.	4	4	5	4
RML	$\hat{\vartheta}_1$	1.24	0.86	1.12	0.22
	$\hat{\vartheta}_2$	27.34	35.36	9.73	60.07
	$\hat{\vartheta}_3$	0.38	0.37	0.29	0.37

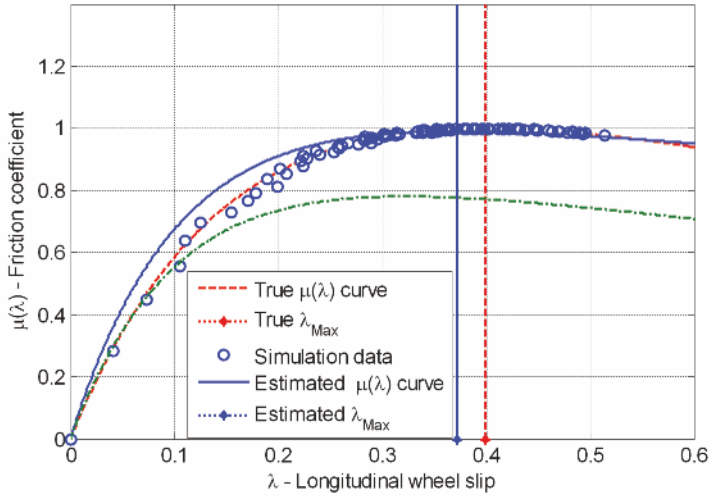


Figure 8.5 Curve fitting results with RML on cobblestone: simulated data (*dots*), theoretical $\mu(\lambda)$ curve (*dashed line*), initial estimation (*dash-dotted line*), estimated $\mu(\lambda)$ curve (*solid line*), true λ_{Max} (*dashed vertical line*) and estimated λ_{Max} (*solid vertical line*)




Figure 8.6 Results of $\hat{\lambda}_{Max}$ estimation with RML on cobblestone: true value (*dashed line*), estimated value (*solid line*), initialisation period (*dark grey*) and 5% error band (*light grey*)

Let us consider in more detail the cobblestone case, for which simulation results are depicted in Figures 8.5 and 8.6. Note that, among all those avail-

able, this particular type of road condition has been selected for discussion as it represents a particularly hard case for identification purposes, being almost flat in the neighbourhood of the maximum. Notice that data are not available for high slip values, so it should not be surprising that unsatisfactory accuracy is experienced in that range. Recall, however, that the ultimate objective of the estimation problem is the peak abscissa evaluation, which can be satisfactorily addressed, as long as the available data encompass the maximum point of the curve. In the case reported in Figure 8.5, the curve fit is only approximately correct in the operating range, resulting in an under-estimation of λ_{Max} . What is more, while the estimation of λ_{Max} quickly approaches the correct value, after an initialisation period of 20 samples, it converges extremely slowly, revealing that the RML method inefficiently exploits the information in the data. Also, accurate curve fitting is obtained only where the samples are more numerous, that is close to the peak point. In any case, a remarkably low 10% error in the estimation of λ_{Max} is achieved after only 30 samples, *i.e.*, with sufficient advance with respect to the critical point in the braking process. This information, as will be further discussed in Section 8.3.3, can be valuably exploited by ABS and wheel slip control systems.

Table 8.2 reports the results of the estimation of λ_{Max} with both the batch and the recursive ML method, based on the respective estimated friction curves. Apparently, RML provides generally worse peak estimation, an exception being the snow case.

The difference in performance between the ML and RML methods on the snow case is due to the fact (see also Figure 2.3) that the friction curve in this case has an almost flat right portion and this fools the ML method – as the great majority of the available data are not descriptive of the peak dynamics – while better results are achieved with the RML approach. This is because the *ad hoc* initialisation procedure is particularly effective on this surface. As a matter of fact, the λ_{Max} value being very low, the interesting part of the curve is well captured by the initialisation samples (see also Equation 8.11). In the RML case, then, the recursive nature of the method can better exploit a good initialisation than the batch version, which tries to fit, in the ML sense, the whole dataset in one shot.

Some general remarks concerning RML are in order. Apparently, due to numerical ill-conditioning of the gradient term of the parameter tuning equation, parameters ϑ_{r_2} and ϑ_{r_3} are only slightly moved from the initial values. Given the better performance of the batch version of the algorithm, this can only be ascribed to the fact that RML combines the iterative nature of the ML approach with data recursion. To deal with this problem, it is convenient to modify the update equations for ϑ_{r_2} and ϑ_{r_3} inserting an additional gain factor to speed up convergence. If computationally viable, a few batch iterations could also be performed as soon as RML convergence does not significantly improve. This last variation, denoted in the following as RML+ML, has yielded significantly better results in simulation and its performance on experimental data will be analysed in Section 8.3.3.




Figure 8.7 Curve fitting results with RLS on cobblestone: simulated data (dots), theoretical $\mu(\lambda)$ curve (dashed line), initial estimation (dash-dotted line), estimated $\mu(\lambda)$ curve (solid line), true λ_{Max} (dashed vertical line) and estimated λ_{Max} (solid vertical line)

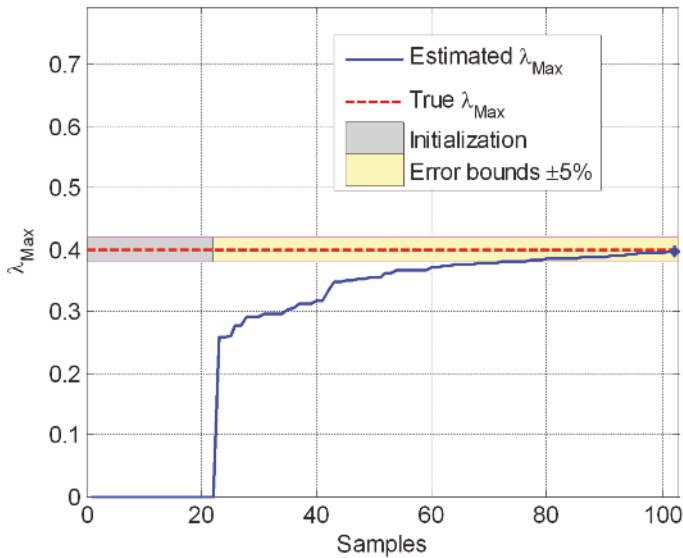


Figure 8.8 $\hat{\lambda}_{Max}$ estimation results with RLS on cobblestone: true value (dashed line), estimated value (solid line), initialisation period (dark grey) and 5% error band (light grey)

Table 8.2 Summary of the λ_{Max} estimation performances obtained with ML and RML on noisy simulation data

		Dry asphalt	Wet asphalt	Cobblestone	Snow
True values	λ_{Max}	0.1700	0.1307	0.3995	0.0605
	μ_{Max}	1.1699	0.8039	1.0014	0.1857
ML	$\hat{\lambda}_{\text{Max}}$	0.1737	0.1350	0.3989	0.0656
	$\hat{\mu}_{\text{Max}}$	1.1700	0.8035	0.9998	0.1848
	$\epsilon_{\lambda}\%$	2.17%	3.28%	0.15%	8.42%
RML	$\hat{\lambda}_{\text{Max}}$	0.1641	0.1242	0.3723	0.0594
	$\hat{\mu}_{\text{Max}}$	1.1679	0.8048	0.9906	0.1929
	$\epsilon_{\lambda}\%$	3.47%	4.97%	6.8%	1.81%

Table 8.3 Summary of the λ_{Max} estimation performances obtained with LS and RLS on noisy simulation data

		Dry asphalt	Wet asphalt	Cobblestone	Snow
True values	λ_{Max}	0.1700	0.1307	0.3995	0.0605
	μ_{Max}	1.1699	0.8039	1.0014	0.1857
LS and RLS	$\hat{\lambda}_{\text{Max}}$	0.1786	0.1370	0.3966	0.0658
	$\hat{\mu}_{\text{Max}}$	1.1696	0.8024	1.0019	0.1849
	$\epsilon_{\lambda}\%$	5.05%	4.84%	0.73%	8.76%

Let us now move to investigate the performance of the LS and the RLS algorithms. In the simulations, the RLS algorithm was initialised using the batch LS method on the first 20 available data. The estimation performances for λ_{Max} are reported in Table 8.3. Notice that 100 data are sufficient for the RLS to converge exactly to the estimates obtained with the LS method. Overall, the LS approach is slightly less accurate than ML, but no significant deterioration resulting from the recursive version of the algorithm is experienced.

Also for the RLS method, the snow case experiences the worst estimation performance in terms of relative error (note that the absolute one is still comparable to the other surfaces). This is mainly because, as observed analysing the ML approach, λ_{Max} takes on a very low value which makes relative error an unfair metric for this specific case. Moreover, because of the large difference between the numerical values of all the snow parameters ϑ_r (all of them differ of more than an order of magnitude from each other), one would need specific scalings in the RLS gains to significantly improve performance. However, as the results are acceptable also on this kind of surface, scaling has been avoided, as it would be surface dependent and therefore not applicable in practice.

For comparison purposes, consider again the cobblestone case, for which simulation results are represented in Figures 8.7 and 8.8. In Figure 8.7 notice the significant estimation error in the initial part of the curve, which is a result

of using an empirical model as opposed to the Burckhardt one. Nevertheless, a remarkable accuracy is achieved near the peak, so that λ_{Max} can be estimated almost exactly.

Concerning the convergence properties of λ_{Max} estimation (see Figure 8.8), one can observe that with RLS the initialisation process is not so efficient in approaching the correct value quickly as in the RML case, but convergence is faster, so that 10% λ_{Max} estimation error is achieved only slightly later (after 40 samples) and 5% error after 70 samples.




Figure 8.9 Example of data (dots) measured on the front wheel on high-grip asphalt road and the theoretical $\mu(\lambda)$ curve (red dashed line)

8.3.3 Experimental Results

Before analysing the estimation results obtained with the identification approaches on measured data, some remarks are due. Specifically, refer to Figure 8.9, where an example of the samples of $\mu(\lambda; t)$ measured on the front left wheel on high-grip asphalt road is shown together with the theoretical tyre–road friction curve related to the road condition used in the tests. Note that the theoretical curve in Figure 8.9 has been derived considering also the increased vertical load experienced by the front wheel during braking. As is apparent by inspection of Figure 8.9, experimental data are significantly dif-

ferent from those obtained in simulation, and this is due to several concurring phenomena:

1. First of all, the $\mu(\lambda)$ curve is intrinsically a static description of the tyre–road friction condition, in that it represents the collection of the steady-state friction coefficient values associated with the corresponding value of the wheel slip. Note that this fact does not depend on the analytical model chosen (the Burckhardt one in this case); the same would have held true if, for example, the Pacejka model had been used (see Section 2.2). When looking at real data, instead, several dynamic phenomena, mainly suspension elasticity and tyre relaxation, cause the real $\mu(\lambda)$ values to exhibit the overshoot clearly visible in Figure 8.9, which also accounts for dynamic load transfer effects.
2. Secondly, as the test was carried out with a real, albeit professional, driver, it is very unlikely that the measured data go significantly beyond the peak value of the $\mu(\lambda)$ curve. When this happens, in fact, the car experiences a significant loss of driveability, and the driver is forced to release the brake pedal to avoid losing control of the vehicle.
3. Finally, the initialisation phase for both algorithms on experimental data has been linked to the wheel slip behaviour. Specifically, both RLS and RML+ML are initialised when the wheel slip is be such that $\lambda > \lambda_{\text{init}}$. The value of $\lambda_{\text{init}} = 0.07$ was chosen based on the analysis of the measured data.

Notice that even though these dynamic phenomena apparently change the shape of the $\mu(\lambda)$ curve, they cause a very small change (if any) of the value of λ_{Max} , that is the abscissa of its peak. Therefore, as the main aim of the identification procedure is that of estimating such a value so that it can be used as input for braking control systems, the identification algorithms do not lose their capability of pursuing this objective.

8.3.3.1 Wheel Slip Measurements Available

We first discuss the estimation results obtained when the vehicle speed is exactly measurable, *i.e.*, by employing the speed signal obtained with an optical sensor, and thus the wheel slip can be directly measured.

Note that here the main aim is to verify that upon controller activation, the proposed estimation algorithm is able to provide a consistent estimate $\hat{\lambda}_{\text{Max}}$ of the value of the abscissa of the peak point of the friction curve. To this end, we assume that the ABS is switched on when the current wheel slip value exceeds a predefined threshold, which has been set to $\lambda = 0.15$. Note that on real ABS systems, see Section 3.7, it is possible to have a dynamic activation threshold selection, which tunes the activation slip value also based on the braking intensity. Nonetheless, as the aim of the experiments is that of testing the identification performance obtainable with a limited number

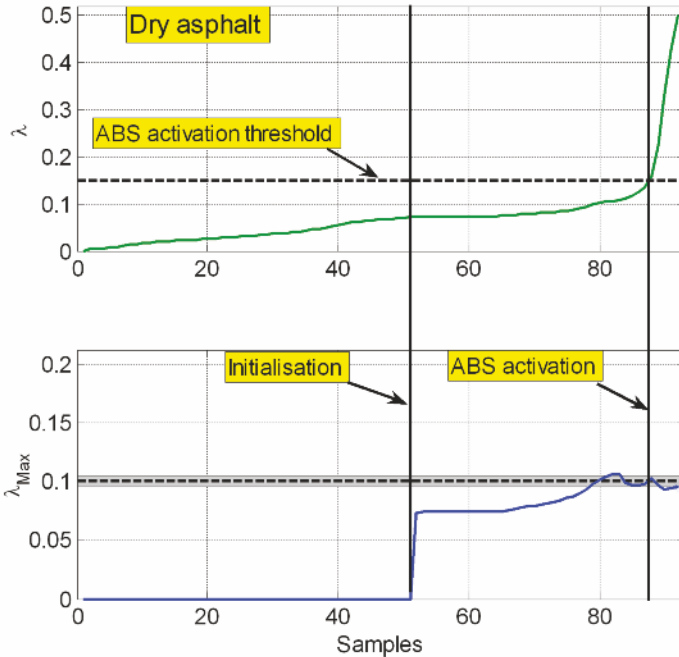


Figure 8.10 Plot of the slip λ as function of samples (*top, solid line*) and the ABS activation threshold (*top, dashed line*) and $\hat{\lambda}_{Max}$ estimation results with RLS on high-grip asphalt road (*bottom*): theoretical value (*dashed line*), estimated value (*solid line*), initialisation period (*grey*) and 5% error band (*grey*)

of data a constant threshold is appropriate. Accordingly, besides evaluating the final results of the estimation when all the available 100 samples are processed, we also monitor at which sample the ABS would virtually be activated and how accurate our estimation is at that time instant.

The estimation results obtained on dry asphalt road with the RLS method are shown in Figure 8.10, while those obtained on low-grip off-road with the RML+ML method are shown in Figure 8.11. Both these figures also show the time history of the wheel slip and highlight the ABS activation threshold. As can be seen, the estimation of λ_{Max} is quite satisfactory and, most importantly, it is reliable also at the time instant at which the ABS would be activated, even if this happens when less than 100 samples have been processed.

A quantitative summary of the overall results obtained with RLS and RML+ML with measured vehicle speed for both friction conditions is provided in Table 8.4. As can be seen, while on dry asphalt the performances of the two algorithms are comparable, RLS shows its better features with respect to RML+ML on low-grip off-roads, where the availability of fewer




Figure 8.11 Plot of the slip λ as function of samples (*top, solid line*) and the ABS activation threshold (*top, dashed line*) and $\hat{\lambda}_{Max}$ estimation results with RML+ML on low-grip off-road (*bottom*): theoretical value (*dashed line*), estimated value (*solid line*), initialisation period (*grey*) and 5% error band (*grey*)




Figure 8.12 Plot of the slip λ as function of samples (*top, solid line*) and the ABS activation threshold (*top, dashed line*) and $\hat{\lambda}_{Max}$ estimation results with RLS and speed estimation on high-grip asphalt road (*bottom*): theoretical value (*dashed line*), estimated value (*solid line*), initialisation period (*grey*) and 5% error band (*grey*)

Table 8.4 Summary of the λ_{Max} estimation performance obtained with RLS and RML+ML on experimental data with measured vehicle speed

v measured	Dry – RLS	Dry – RML+ML	Off-road – RLS	Off-road – RML+ML
True λ_{Max}	0.1	0.1	0.13	0.13
# samples	92 [0.46 s]	92 [0.46 s]	52 [0.26 s]	52 [0.26 s]
$\hat{\lambda}_{\text{Max}}$	0.0955	0.0942	0.1226	0.1063
ϵ_{λ} %	-4.5%	-5.8%	-5.69%	-18.23%
# samples at ABS activation	88 [0.44 s]	88 [0.44 s]	44 [0.22 s]	44 [0.22 s]
$\epsilon_{\lambda_{\text{ACT}}}$ %	1.9%	4.7%	-1.69%	-19.46%

samples makes the initialisation phase crucial, thereby confirming the theoretical analysis of the two algorithms.

8.3.3.2 Wheel Slip Measurements not Available

The estimation results obtained when non-exact wheel slip measurements are available are now discussed. This analysis reflects the practical case in which the tyre-road friction estimation is in fact implemented on a passenger vehicle equipped with standard ABS/ESC sensors (*i.e.*, wheel encoders and a longitudinal accelerometer), in which the wheel slip is estimated. Specifically, the estimation algorithm described in Chapter 5 has been used.

Due to the superior properties of the RLS algorithm observed in the case of measured vehicle speed, we only present results obtained with this algorithm in combination with estimated vehicle speed. The results obtained with the RLS method on a dry asphalt road and on a low-grip off-road are shown in Figures 8.12 and 8.13, respectively. Again, these figures also show the time history of the wheel slip and highlight the ABS activation threshold. As can be seen, the reliability of the estimation remains unchanged (both looking at the ABS activation time instant and at the overall estimation results) if the estimated value of the vehicle speed is provided as input to the algorithm, thereby confirming its practical applicability. Similar results with respect to the consistency between the case of measured and estimated vehicle speed have been obtained also for the RML+ML algorithm. A quantitative summary of the overall results on experimental data obtained with RLS and estimated vehicle speed is provided in Table 8.5.

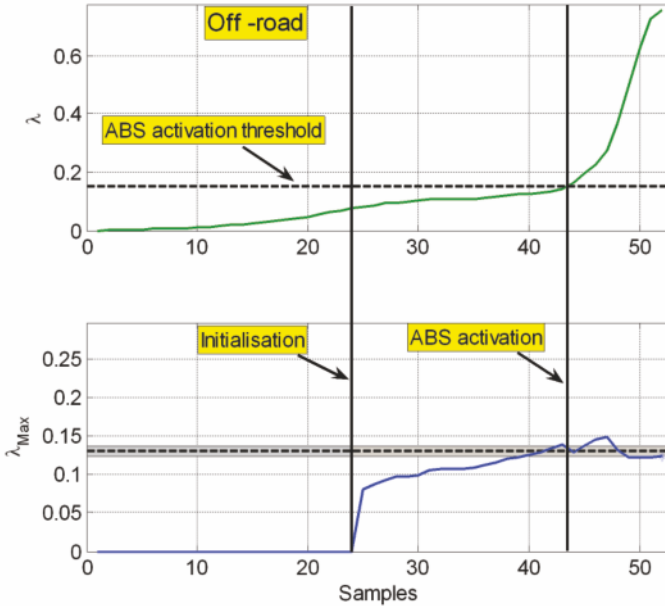


Figure 8.13 Plot of the slip λ as function of samples (*top*, *solid line*) and the ABS activation threshold (*top*, *dashed line*) and $\hat{\lambda}_{\text{Max}}$ estimation results with RLS and speed estimation on low-grip off-road (*bottom*): theoretical value (*dashed line*), estimated value (*solid line*), initialisation period (*grey*) and 5% error band (*grey*)

Table 8.5 Summary of the λ_{Max} estimation performance obtained with RLS on experimental data with estimated vehicle speed

v estimated	Dry – RLS	Off road – RLS
True λ_{Max}	0.1	0.13
# samples	92 [0.46 s]	52 [0.26 s]
$\hat{\lambda}_{\text{Max}}$	0.0953	0.123
ϵ_{λ} %	-4.7%	-5.3%
# samples at ABS activation	88 [0.44 s]	44 [0.22 s]
$\epsilon_{\lambda_{\text{ACT}}}$ %	0.2%	-1.38%

8.4 Direct Estimation of Contact Forces *via* In-tyre Sensors

Traditionally, tyre–road contact forces are indirectly estimated from vehicle-dynamics measurements (*e.g.*, chassis accelerations, yaw and roll rates, suspension deflections). The emerging of the *smart-tyre* concept (tyre with embedded sensors and digital-computing capability) has made possible, in principle, a more direct estimation of contact forces. In this field, which is still in

its infancy, a basic and fundamental problem is the choice of the sensor(s) and of the regressor(s) that are most appropriate to be used for force estimation.

The objective of this section is to illustrate a sensor-regressor choice tailored to accomplish this task, and to provide some experimental results to discuss the validity of this choice. The idea is to use a wheel encoder and an accelerometer mounted directly in the tyre (see Figure 8.14). The measurement of the in-tyre acceleration is transmitted through a wireless channel. The key concept is to use the *phase shift* between the wheel encoder and the pulse-like signals provided by the accelerometer as the main regressor for force estimation.




Figure 8.14 Detail of the in-tyre accelerometer

8.4.1 Introduction

In the field of smart-tyres, *i.e.*, tyres equipped with electronic devices that make them active components of the vehicle, one of the main challenges is the direct real-time estimation of the tyre–road contact forces *via* in-tyre sensors. This research is still actively ongoing and a number of non-trivial issues are still open. Among others: the choice of in-tyre sensors, the in-tyre preprocessing of the signal, the wireless transmission, the post-processing, the regressor choice and the estimation algorithm. The problem is made even more complicated by technological and industrial issues like durability, cost and energy consumption of the in-tyre electronic devices.

The great interest in a smart-tyre capable of providing a real-time estimation of the tyre–road contact forces is easily explained by its huge potential benefits: the direct measurement of tyre–road contact forces can stimulate the development of a new generation of traction, braking and stability control systems which may outperform the existing ones in terms of safety, driving satisfaction and also energy consumption.



Figure 8.15 Detail of the battery and wireless transmitter installed on the wheel rim

At this stage, no commercial products that can estimate the friction forces directly from on-vehicle measurements are available on the market, but many research groups are exploring and testing different solutions. The two key issues (strictly interwoven), which are still open, can be summarised as follows:

1. What is the best tyre-embedded sensor choice?
2. Given the sensor configuration, what is the best set of regressors to be used for force estimation?

This section presents a possible solution to these two open issues and provides some experimental results as a preliminary validation. Specifically, this method is tailored to estimate the tyre–road vertical and longitudinal forces and it makes use of two sensors in each wheel; namely

- a standard wheel encoder typically used by ABS systems for the wheel speed measurement (see Appendix B); and
- a one-axis accelerometer mounted directly *in the tyre*, which measures the acceleration experienced by the tyre in the radial direction.

Based on these signals, we aim at evaluating the phase shift between the wheel hub and the tyre, based on the wheel encoder and the accelerometer signals, respectively, by detecting when the tyre encounters and leaves the tyre–road contact patch. In fact, such a phase shift appears to be strongly correlated with the longitudinal and vertical tyre deformation; henceforth it can be suitably employed for the direct identification of the contact forces.

8.4.2 *Experimental Set Up*

The test-car used in the experiments is a rear-wheel-driven BMW, equipped with the following sensors:

- four inductive 48-teeth Hall-effect encoders (see also Appendix B) that measure the wheel rotational speed, whose output is a sinusoidal-like signal, with amplitude and frequency proportional to the rotational wheel speed $\omega(t)$; and
- a one-axis, ± 500 g piezoresistive low-mass linear accelerometer, mounted (glued) inside the tyre (see Figure 8.14), which measures the in-tyre radial acceleration $a_{\text{tyre}}(t)$; the accelerometer has been installed alternatively on the front-left tyre and on the rear-left tyre. Its bandwidth is of approximately 3 kHz.

The wireless data transmission of the in-tyre acceleration signal is made via a transmitter, which is mounted on the wheel rim together with its battery (see Figure 8.15); the receiver antenna is placed on the car roof. All signals are sampled at 10 kHz, with a resolution of 16 bit.

The driving tests were all made on the same road surface, a high-grip flat dry-asphalt road. Two main types of tests were carried out, namely

- quasi-static tests: very slow decelerations and accelerations (with no gear shifts) in the range 10-25 m/s; and
- dynamic tests: strong braking and acceleration manoeuvres, interleaved with constant-speed intervals.

All the tests were performed on a straight road. Two tyre-pressure conditions were tested: 2.0 bar (nominal condition) and 1.6 bar (low-pressure condition).

8.4.3 *Main Concept*

Consider the signal detected by the in-tyre accelerometer. It measures the instantaneous acceleration experienced by the tyre at the installation point in the radial direction. An example of this signal (the raw signal, without any kind of pre-processing) over a 0.7 s time window is displayed in Figure 8.16. The analysis of this signal reveals that:

- The radial acceleration is approximately constant over a short time-window (it is the centripetal acceleration, proportional to the wheel speed), but when the accelerometer passes through the contact patch the acceleration is characterised by two impulse-like signals. This twin-spike is obviously repeated every wheel revolution.
- The two main *fronts* of the acceleration signal around the contact patch (the first is descending, the second is ascending) can be used to detect

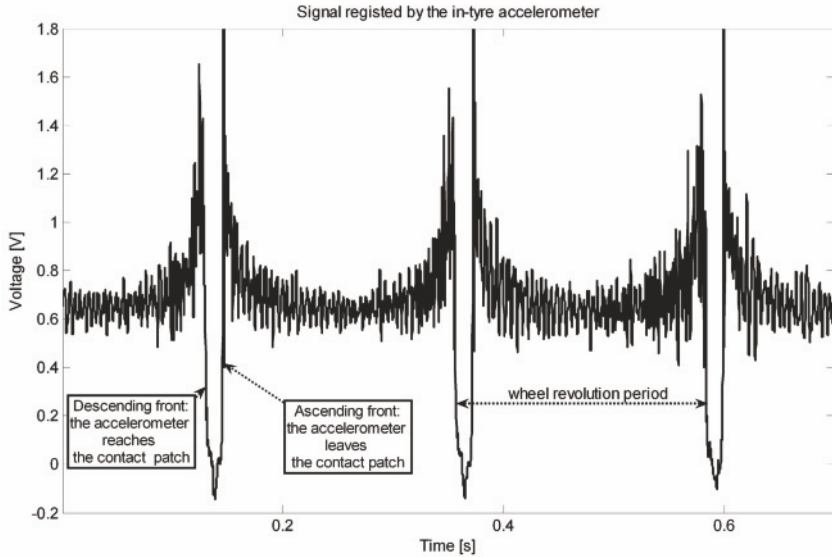


Figure 8.16 Raw signal detected by the radial in-tyre accelerometer

the time instants when the accelerometer reaches and leaves the contact patch, respectively.

Based on this signal and on the wheel encoders measurements, the aim is to estimate the vertical and the longitudinal force acting on the tyre by exploiting the information contained in the phase shift between the wheel hub and the tyre.

In order to obtain a more visual description of this idea, consider Figure 8.17(a), where a schematic picture of a wheel and its contact patch is displayed. If we consider the two acceleration spikes, and their absolute position in the angular reference frame α given by wheel encoder (which is not subject to elastic deformation), three quantities can be computed at each wheel revolution:

1. the angular position α_1 of the initial position of the contact patch;
2. the angular position α_2 of the final position of the contact patch; and
3. the length $\Delta\phi = \alpha_2 - \alpha_1$ of the contact patch.

In Figures 8.17(b) and 8.17(c) the phase shift phenomenon is pictorially described. More specifically:

- Due to the elastic properties of the tyre in the radial direction (see *e.g.*, [7, 8, 71]), the vertical tyre–road contact force F_z is assumed to be strictly correlated with the length

$$\Delta\phi = \alpha_2 - \alpha_1 \quad (8.13)$$




Figure 8.17 A schematic representation of a wheel and its contact patch (a), the effect of a vertical force variation (b) and the effect of a longitudinal force variation (c)

of the contact patch (see Figure 8.17(b)); the idea hence is to use the measured angle $\Delta\phi$ to estimate F_z , *i.e.*,

$$F_z = f_z(\Delta\phi). \quad (8.14)$$

- Due to the elastic properties of the tyre in the longitudinal (tangential) direction (see again [7, 8, 71]), the longitudinal tyre–road contact force F_x is assumed to be strictly correlated to the phase shift of the centre of the contact patch (see Figure 8.17(c)), given by

$$\delta\phi = (\alpha_1 + \alpha_2)/2. \quad (8.15)$$

The idea is to use the measured angle $\delta\phi$ to estimate F_x , *i.e.*,

$$F_x = f_x(\delta\phi). \quad (8.16)$$

In particular, note that $\delta\phi$ is expected to be negative during braking and positive during acceleration; in other words, this means that the centre of the contact patch is assumed to rotate backwards or forwards (with respect to a conventional zero position), respectively. Also notice that, in general, the real centre of the contact patch is not perpendicular to the wheel hub. Hence, (8.15) provides a *conventional* centre of the contact patch. This, however, does not affect the estimation procedure, since it is insensitive to the conventional choice of the contact patch centre.


8.4.4 Signal Processing

The whole method is based on the measurement of two signals for each wheel: the wheel rotational speed ω and the radial acceleration a_{tyre} experienced by a point in the tyre.


The wheel encoder is the standard 48-teeth encoder used by the ABS and ESC control systems (see also Appendix B). The original signal coming from the encoder is a sinusoidal-like voltage signal. An example of this signal is displayed in Figure 8.18(a), over a time-window of 100 ms. The wheel rotational speed is computed from the sinusoidal wheel encoder *via* a frequency tracking algorithm, see [83], and the wheel radius is calibrated as detailed in Section B.2.2.

In Figure 8.18(b), an example of the calibrated linear wheel speed signal (for the front-left wheel) is displayed. Notice that the estimation of the wheel radius is not particularly critical for this application, since the considered regressors are based on phase shifts, which are not strongly correlated with the actual wheel radius.


The last step of the pre-processing of the wheel encoder signal, specific to the considered estimation problem, is to estimate the instantaneous angular



(a)



(b)



(c)

Figure 8.18 Examples of the measured wheel encoder signal (a), the calibrated wheel speed at the contact point (b) and the estimated angular position α of the wheel from the wheel encoder (c)

position α of the wheel. This estimation is fundamental since it represents the baseline for the identification of the phase shift. The basic idea for the estimation of the angular position is simply to make an incremental step counter, which resets every wheel revolution (or 360°). Since in a 48-teeth encoder the peak-to-peak distance of one period corresponds to 7.5° of wheel rotation, the whole sinusoidal profile of the signal (Figure 8.18(a)) must be used to improve the accuracy of the angular position. The angular resolution obtained can be estimated to be about 0.1° , which is suitable for this application. An example of the estimated wheel angular position is displayed in Figure 8.18(c).




Figure 8.19 Example of output of the front-detection algorithm

As has already been mentioned, in-tyre accelerometer is a one-axis sensor (oriented in the radial direction) linear accelerometer glued inside the tyre (see Figure 8.14). The output signal is a voltage signal, characterised by a bump (or a twin-impulse) around the tyre–road contact patch. The main signal processing issue is to detect the time instant when each impulse takes place, namely to detect the angular position α_1 and α_2 of the two fronts (one descending and one ascending) of the bump.

In Figure 8.19 an example of the output of the front-detection algorithm is displayed (dotted line). The implemented front-detection algorithm essentially performs a simple zero-crossing search in the neighbourhood of the two main fronts; the zero-crossing algorithm is applied to the normalised and detrended signal. Note that in Figure 8.19 the signals are plotted as a function of the wheel's absolute angular position α , not as a function of time. This is

a key step that allows us to eliminate most of the dependency of the phase shift from the wheel rotational frequency.

Starting from α_1 and α_2 , the estimation of the length of the contact patch $\Delta\phi$ is straightforward, as $\Delta\phi = \alpha_2 - \alpha_1$. For the computation of the phase shift of the centre of the contact patch Ω_0 the following procedure was used. The first part of every test drive is always characterised by a low, constant-speed (lasting approximately 10 ms) condition. It is conventionally assumed that the phase shift in that condition is zero, namely

$$\delta\phi = \frac{\alpha_1 + \alpha_2}{2} + \Omega_0 = 0.$$

The calibration offset Ω_0 is then added to the absolute angular position of the wheel for the entire experiment. In other words, it is assumed that at the beginning of each experiment the phase shift is zero, and all the phase shifts of the rest of the experiment are referred to that conventional zero-condition. As has already been remarked, the choice of this conventional *zero* does not affect the quality of the estimation.




Figure 8.20 Example of a complete dynamic test: vehicle speed (a), estimated phase shift (b) and estimated length of the contact patch (c). The phase shift and length of the contact patch refer to the rear-left wheel

The results of the above described pre-processing of the two main signals can be appreciated in Figure 8.20, where the vehicle speed, the phase shift of

the mid-point of the contact patch $\delta\phi$ and the contact patch length $\Delta\phi$ are displayed, for a 1 min long dynamic experiment. From the behaviour of the vehicle speed, note that experiment is constituted by five main parts: a first part of constant low-speed, an acceleration, a new long constant-speed window, a strong braking manoeuvre and a final constant low-speed condition.

By inspecting the estimated $\delta\phi$ and $\Delta\phi$ in Figure 8.20 it is immediately apparent that their behaviour is highly correlated with the vehicle speed. Unfortunately, this speed-dependency phenomenon almost completely hides the important part of the relationships between the pair $(\delta\phi, \Delta\phi)$ and the contact forces (F_x, F_z) .

The removal of the speed-dependent trends in $\delta\phi$ and $\Delta\phi$ hence is mandatory. Note that this effect was somehow expected and it is mainly due to the effects of the aerodynamic forces and (at mid/low speed values) of the rolling resistance.

In order to remove this speed dependence, a simple quasi-static experiment was performed: the car was slowly accelerated (without gear-shift) from 10 m/s to 25 m/s; the same experiment was repeated in deceleration (coasting down).

Since the acceleration/deceleration ramp is extremely slow, in this experiment the dynamic effects can be neglected. At each wheel revolution, the pairs $(\omega r, \delta\phi)$ and $(\omega r, \Delta\phi)$ have been computed. The results are plotted in Figures 8.21(a) and 8.21(b).

As the experiment has been made in a quasi-static setting, the relationships depicted in Figures 8.21(a) and 8.21(b) are static and can be easily fitted with one-dimensional nonlinear functions. In particular, both maps were fitted with simple second-order polynomials. A unique map was used both for the front and the rear tyres. The estimated maps are displayed in Figures 8.21(a) and 8.21(b).

Using the estimated maps, the speed-dependent trends have been removed from the dynamic experiments. The results are displayed in Figures 8.22 and 8.23, for both the rear and the front wheels. Note that after the trend-removal, the phase shift and the length of the contact patch have the same value (0° and 25° , respectively) in every constant-speed condition.

By carefully inspecting Figure 8.23 (front tyre), another residual spurious effect can be observed. As a matter of fact, notice that – during the acceleration phase – no significant longitudinal force F_x should be developed by the front wheels; as a consequence, we know *a priori* that the phase shift of the contact patch on a front wheel in that condition should be zero. This condition is not perfectly met by the data displayed in Figure 8.23 (see the time interval $t \in [15, 20]$ s in the middle plot). This phenomenon has a simple and intuitive explanation: there is a slight dependency (or *cross-talk*) between the phase shift and the vertical force. Hence, the phase shift must be subject to an additional correction as follows:

$$\delta\phi = \delta\tilde{\phi} + f_F(F_z), \quad (8.17)$$

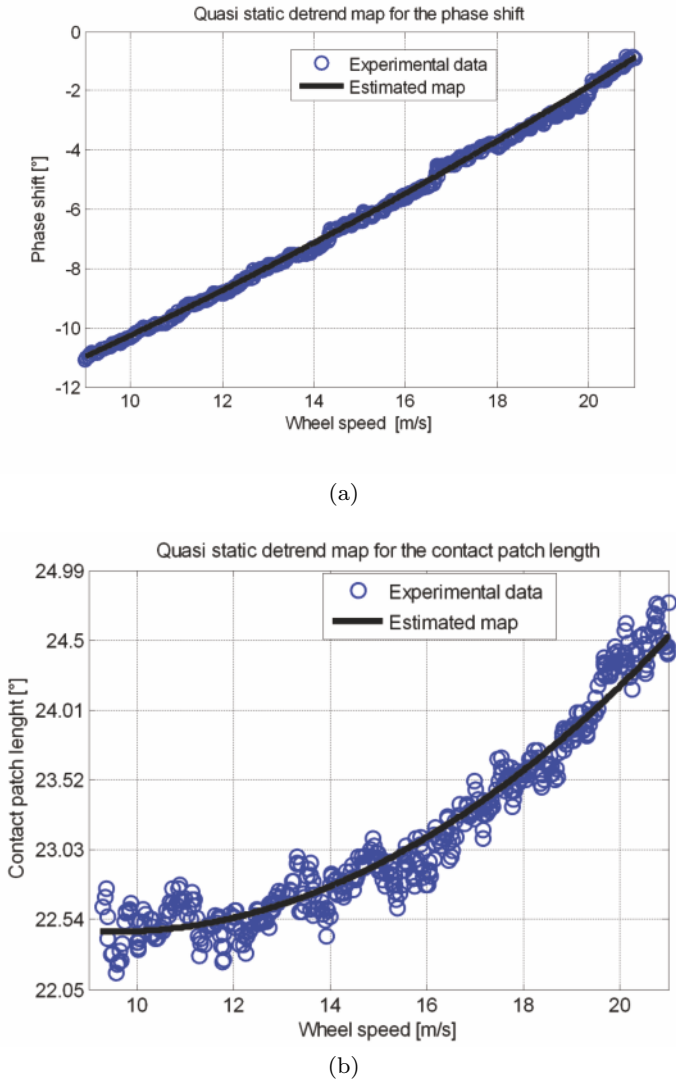


Figure 8.21 Quasi-static maps of the phase shift (a) and contact patch length (b) as functions of wheel speed

where $\delta\tilde{\phi}$ is the phase shift without correction and $f_F(F_z)$ is the correction term (to be estimated). Unfortunately, F_z is not directly known. However, since we have assumed a direct static relationship between F_z and $\Delta\phi$, we can approximate Equation 8.17 with the following equation:

$$\delta\phi = \delta\tilde{\phi} + f_F(\Delta\phi). \tag{8.18}$$




Figure 8.22 Example of a complete dynamic experiment: vehicle speed (a), estimated phase shift (b) and estimated length of the contact patch (c). Rear-left wheel; nominal pressure

For simplicity, we have assumed a linear dependency in the correction term, namely

$$\delta\phi = \delta\tilde{\phi} + \beta(\Delta\phi - \Delta\phi_0). \quad (8.19)$$

The term $(\Delta\phi - \Delta\phi_0)$ in Equation 8.19 is the dynamic variation of the contact patch length ($\Delta\phi_0$ is the average value of the contact patch, at constant speed). The only unknown term in Equation 8.19 is the coefficient β . The optimal value of β has been identified by numerical optimisation from data, in order to guarantee no phase shift on the front tyre during acceleration in every working condition.

In Figure 8.24 the detail of the acceleration phase for the front tyre before and after the correction with (8.19) is shown. Using the estimated value of β , the phase shift signals of both the front and the rear wheels have been modified according to Equation 8.19. After the single (for the contact patch length) and the double (for the phase shift) trend removal, all the main spurious effects are eliminated, and data can be employed to extract the phase shift and contact patch length information.




Figure 8.23 Example of a complete dynamic experiment: vehicle speed (a), estimated phase shift (b) and estimated length of the contact patch (c). Front-left wheel; nominal pressure

8.4.5 Experimental Results

The experimental results are based on a set of test drives performed on a flat dry-asphalt surface, when driving in a straight line at low/mid-range speed (0-80) km/h. Two tyre pressure settings have been tested, in order to analyse the sensitivity of the method with respect to this critical parameter. Specifically, a nominal pressure of 2.0 bar and a reduced pressure of 1.6 bar were considered.

Figures 8.25 and 8.26 show the results for the nominal pressure at the rear and front wheels, respectively. By carefully inspecting these figures, the following observations can be made.

- *Acceleration phase – Longitudinal force*

During this phase no longitudinal force is applied at the front wheel: $F_{x\text{Front}} = 0$, whereas a positive force is applied at the rear (drive) wheel: $F_{x\text{Rear}} > 0$.

Accordingly, since we have assumed a direct monotone static relationship $F_x = f_x(\delta\phi)$ between the longitudinal force F_x and the phase shift of




Figure 8.24 Detail of the acceleration phase (front wheel): vehicle speed (a) and estimated phase shift before (b) and after (c) the removal of the load-transfer parasitic effect

the contact patch $\delta\phi$, the phase shift of the front wheel should be zero ($\delta\phi_{\text{Front}} = 0$).

Conversely, the phase shift at the rear wheel should be positive ($\delta\phi_{\text{Rear}} > 0$). From Figures 8.25 and 8.26 it is easy to see that both these conditions are met.

- *Acceleration phase – Vertical force*

During this phase the front wheel should experience a decrease of the vertical force, *i.e.*, $F_{z_{\text{Front}}} < \bar{F}_{z_{\text{Front}}}$, where $\bar{F}_{z_{\text{Front}}}$ is the static load at the front wheel, whereas the vertical force at the rear wheel should increase, *i.e.*, $F_{z_{\text{Rear}}} > \bar{F}_{z_{\text{Rear}}}$, where $\bar{F}_{z_{\text{Rear}}}$ is the static load at the rear wheel.

Accordingly, since we have assumed a direct monotone static relationship $F_z = f_z(\Delta\phi)$ between the vertical force F_z and the length of the contact patch $\Delta\phi$, the length at the front wheel should decrease, *i.e.*, $\Delta\phi_{\text{Front}} < \Delta\phi_0$, whereas that at the rear wheel should increase *i.e.*, $\Delta\phi_{\text{Rear}} > \Delta\phi_0$. From Figures 8.25 and 8.26 it is easy to see that also both these conditions are met.

- *Braking phase – Longitudinal force*

During this phase, a negative longitudinal force is applied at both wheels *i.e.*, $F_{x_{\text{Front}}} < 0$, $F_{x_{\text{Rear}}} < 0$. Accordingly, we expect that $\delta\phi_{\text{Rear}} < 0$ and $\delta\phi_{\text{Front}} < 0$.

From Figures 8.25 and 8.26 it is easy to see that these conditions are met.

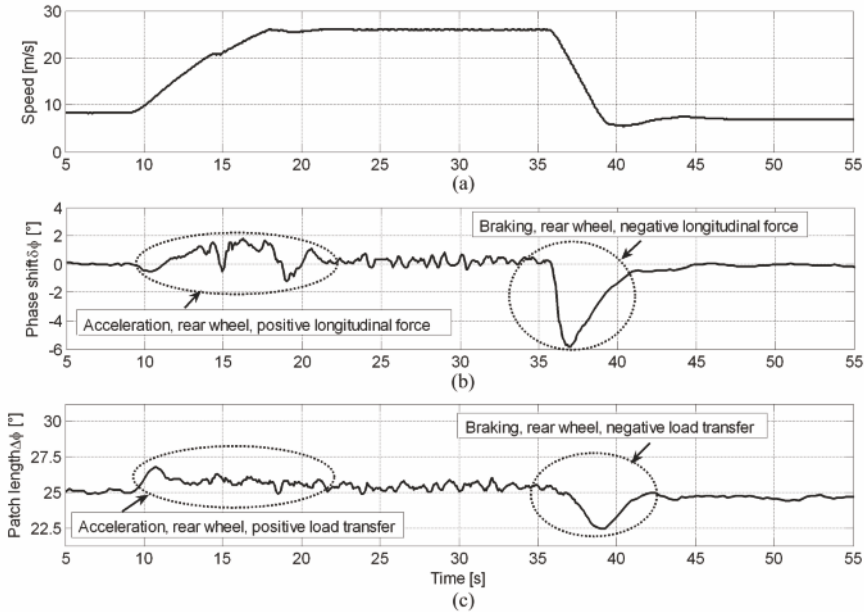


Figure 8.25 Final results on a complete dynamic experiment: vehicle speed (a), estimated phase shift (b) and estimated length of the contact patch (c). Rear-left wheel; nominal pressure

- *Braking phase – Vertical force*

During this phase, the front wheel should experience an increase of the vertical force, *i.e.*, $F_{z_{Front}} > \bar{F}_{z_{Front}}$, whereas the vertical force at the rear wheel should decrease, *i.e.*, $F_{z_{Rear}} < \bar{F}_{z_{Rear}}$.

Accordingly, we expect that $\Delta\phi_{Front} > \Delta\phi_0$ and $\Delta\phi_{Rear} < \Delta\phi_0$.

From Figures 8.25 and 8.26 it is easy to see that both these conditions are met.

In Figures 8.27 and 8.28 the results of the experiments performed with low-pressure tyres on rear and front wheels, respectively, are displayed. Notice that the results displayed in Figures 8.27 and 8.28 were obtained without recalibration; they were computed using the same pre-processing and calibrations as for the 2.0 bar case in order to test the robustness of the method. All the considerations made in the nominal-pressure case still hold; this fact is encouraging as the method shows a good robustness with respect to pressure variations.

By carefully comparing the results at nominal and low pressure, one can notice that, as expected, the contact patch length is slightly larger in the case of low-pressure tyres; the variations in the phase shift instead are very similar to the nominal-pressure setting.




Figure 8.26 Final results on a complete dynamic experiment: vehicle speed (a), estimated phase shift (b) and estimated length of the contact patch (c). Front-left wheel, nominal pressure

8.5 Summary

This chapter provided some approaches to estimate online the tyre–road friction characteristics. As friction information can be extremely valuable for all active vehicle control systems, much research activity has focused on the estimation and monitoring of tyre–road friction characteristics (see, *e.g.*, the detailed review in [65]), resulting in the proposal of many different approaches, which differ both regarding the estimation technique and the required sensor equipment.

Some interesting approaches are tyre-oriented. For example, in [12] an acoustic sensor is used to gain information on road-condition by registering the acoustic waves emitted by the tyres. The drawback of such an approach is the high noise level in the acoustic signal, which makes it very hard to extract the real effect of friction changes on the measured signal. Another tyre-oriented approach is documented in [12] and [24], where tyre-tread deformation sensors are employed. Such an approach, besides suffering from the same drawbacks as the former, is also quite expensive and of difficult implementation, as the tyre-tread sensors have to be embedded in the tyre with specific vulcanisation processes.




Figure 8.27 Final results on a complete dynamic experiment: vehicle speed (a), estimated phase shift (b) and estimated length of the contact patch (c). Rear-left wheel, reduced pressure

Another approach for the estimation of tyre–road friction is the so-called *slip-based* approach, which appears particularly appealing as it needs only standard ABS-ESC sensor equipment. Slip-based estimation is mostly addressed during braking manoeuvres, where sufficiently large slip levels are encountered, although estimation techniques based on low slip measurements, as available during traction, are also documented in the literature (see, *e.g.*, [65]). In [29], an adaptive estimation method is proposed based on a linear approximation of the tyre–road friction description. In the last few years, the interest has also shifted towards the estimation of the dynamic behaviour of tyre friction forces (see, *e.g.*, [12, 24, 29, 85, 86, 93, 129]). Another example of this field of research is given in [69], where the authors estimate the tyre extended braking stiffness – *i.e.*, the derivative of the longitudinal friction force with respect to the wheel slip, which indicates the residual longitudinal friction force available to the driver. Such an estimate serves as additional information for the design of an ABS system.

Further, developments in tyre materials, structure and manufacturing techniques have been enormous in the last decades. However, as has already been remarked, the tyre has up to now essentially remained a *passive* object.

In the last few years a new trend has emerged, whose aim is to equip the tyre with embedded sensors and digital-computing capability; the measure-




Figure 8.28 Final results on a complete dynamic experiment: vehicle speed (a), estimated phase shift (b) and estimated length of the contact patch (c). Front-left wheel, reduced pressure

ment and low-bandwidth transmission of the tyre pressure is an industrial reality (tyre pressure monitoring systems, see, *e.g.*, [31]) and new in-tyre sensors and electronics are currently under research [63, 75]. This trend is mainly driven by information and communication technology methods and devices and it represents a sort of revolution in tyre manufacturing. The term *smart-tyre* is frequently used to label this new generation of tyres.

In the open scientific literature, little has been published so far on the topic of direct estimation of tyre–road contact forces by in-tyre sensors.

Most of the current research activity on this topic has up to now been described in industrial patents (see, *e.g.*, [13, 79]) or in oral presentations (see, *e.g.*, [63]).