

The Twin-in-the-Loop approach for vehicle dynamics control

Federico Dettù, Simone Formentin and Sergio Matteo Savaresi

Abstract—In vehicle dynamics control, engineering a suitable regulator is a long and costly process. The starting point is usually the design of a nominal controller based on a simple control-oriented model and its testing on a full-fledged simulator. Then, many driving hours are required during the End-of-Line (EoL) tuning phase to calibrate the controller for the physical vehicle. Given the recent technological advances, we consider in this paper the pioneering perspective where the simulator can be run *on-board* in the electronic control unit, to calculate the nominal control action in real-time. In this way, it can be shown that, in the EoL phase, we only need to tune a simple compensator of the mismatch between the expected and the measured outputs. The resulting approach not only exploits the already available simulator and nominal controller and significantly simplifies the design process, but also outperforms the state-of-the-art in terms of tracking accuracy and robustness within a challenging active braking control case study.

I. INTRODUCTION

The use of Digital Twins (DTs) - combining software and physical connections to produce a faithful virtual replica of a given system - is revolutionizing state-of-the-art technological solutions in different fields [1].

In the automotive world, DTs are highly exploited for on-line monitoring and prognostics of vehicle components [1]. Moreover, vehicle dynamics simulators - as faithful replicas of the system - are widely employed at the mechanical design level, *i.e.*, when selecting physical components or when assessing the differences among structural choices [2]. Instead, vehicle dynamics controls are still based on simple control-oriented models, usually capturing the key features of a single maneuver [3]. Indeed, this yields well known issues when implementing the controller on the real platform, having to

This research was partially funded by company VI-Grade SrL, whose technical support is here gratefully acknowledged.

All the authors are with Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133, Milano, Italy (e-mail: federico.dettu@polimi.it, simone.formentin@polimi.it, sergio.savaresi@polimi.it).

deal with many unmodeled dynamics. In the industrial practice, this issue is overcome by finely adjusting the controller parameters, during the so-called End-of-Line (EoL) tuning process [4]. Nonetheless, the latter might be a time consuming and costly procedure, especially when considering complex industrial regulators.

In this paper, we take advantage of the most recent technological advances to show how a high-fidelity vehicle simulator can be used *in real-time* and *directly embedded* into the Electronic Control Unit (ECU) of the vehicle to enhance the closed-loop performance and significantly simplify the EoL tuning phase. We will denote the resulting architecture *Twin-in-the-loop* (TiL) *control* hereafter. Doing so, the nominal control action could be provided as the one computed on the system simulator, whereas only dynamics to be controlled is the mismatch between the vehicle and the simulator, which could be handled by a simpler compensator. We remark that, even if the possibility of running a full-fledged vehicle simulator on-board may sound far in the future, it rather fits well into the present context of an ever-increasing on-board and cloud-based [1] computing power, driven by the necessity of processing signals coming from visual sensors, such as LiDARs. Indeed, this novel approach calls for a tuning procedure for the additional compensator, which will be also treated in this paper.

TiL shares some common feature with *internal model control* (IMC) [5]. In IMC, an approximated model of the plant can be used to generate the control action necessary to obtain a desired output. In fact, if the model is invertible (*e.g.* in case of minimum phase transfer functions), the nominal control action is easily obtained; if not, the model has to be decomposed into its invertible and non-invertible parts, *e.g.* by adding high-frequency poles to the non-causal model inverse. IMC has been recently applied for spark ignition engines control [6], [7], but the applications to vehicle dynamics are not frequent in the literature, with one notable exception for the yaw-rate controller in [8].

An important problem in IMC - that can be read as a possible reason for its low popularity in spite of other approaches, like, *e.g.*, Model Predictive Control - lies in the strong requirement on the plant to be invertible or at least decomposable with an invertible part; when the plant cannot be inverted, nor its equation can be written in closed-form, IMC cannot be applied. This is the case of commercial driving simulators, which are often black-box objects, and can only be used to obtain a certain output based a specified set of inputs. TiL solves this invertibility problem for this class of models, as we show later on in this manuscript.

In order to show the potential of TiL architecture as compared to traditional vehicle controls end-of-line tuning, we consider a well-known and challenging safety-critical problem, namely *active braking control* [9]. *Active braking control* - or anti-lock braking system (ABS) - is a key piece of software in modern vehicles; its main purpose is to keep the longitudinal slip around the friction peak to maximize braking performance and thus avoiding dangerous wheel lock.

More specifically, we consider a high-performance sport car as the vehicle to be controlled, the latter being an interesting proving ground for paradigm shifting technologies, due to the extremely challenging and highly dynamic driving conditions. Model predictive control (MPC) has been recently employed in braking control [10], [11], [12], in order to explicitly account for actuator dynamics and constraints in the control law - which is impossible in simpler controls, such as PID [13]. For these reasons, we consider an MPC to build our proof-of-concept: specifically, we select the approach proposed in [10], where a specific design for high-performance vehicles is proposed. This choice is indeed not a strong assumption, as TiL can be in principle used to enhance any control algorithm and to simplify its implementation; we will discuss this in detail later on. Finally, we wish to remark here that recent work has already shown the potential of such a technology for accurate vehicle state estimation [14]; a recent development [15] also shows how to deal with the dimensionality issues when calibrating TiL estimators. Indeed, to the best of our knowledge, this is the first time that the TiL paradigm is applied to control systems.

The remainder of the paper is as follows. Section II formally states the problem, while Section III discusses the active braking case study. Section IV applies TiL control to the proposed case study, and Section IV-B shows how the algorithm parameters can be calibrated starting from a set of closed-loop experiments. Section V compares the TiL solution to the benchmark; finally, Section VI provides an Hardware-in-the-Loop real-time validation of the proposed controller on an off-the-shelf ECU, in order to prove real-time feasibility of the same. The paper is ended by some concluding remarks.

II. PROBLEM STATEMENT

The traditional framework for the development of vehicle dynamics control systems usually follows the three-step procedure depicted in Fig. 1, in which we have considered the case of reference following without loss of generality.

Specifically, the full design procedure is as follows.

- 1) A controller $C(\theta)$, fully described by a set of parameters θ , is first designed based upon a control oriented model, capturing the main dynamics of interest - e.g. a quarter-car model in case of active braking control [10]. From now on, we denote the set of parameters selected using the simple model as θ_0 .
- 2) The controller is then tested on a high-fidelity multibody simulator (a digital twin), usually accounting for unmodeled behaviours and nonlinearities. The initial values θ_0 are then adjusted until acceptable performance is reached, leading to a new optimal selection θ_s .

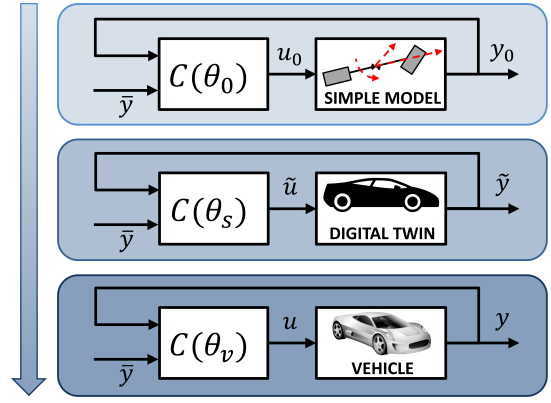


Fig. 1: Traditional vehicle controls development: the controller is first designed based on a simple control-oriented model, then it is tested on a full-fledged simulator and finally it is fine tuned on the physical vehicle.

- 3) Finally, the controller is implemented on an Electronic Control Unit (ECU) and tested on the physical vehicle along a few driving hours. This step usually requires a final refinement of the parameters, yielding the new vector θ_v , due, e.g., to neglected phenomena or measurement noises.

The key observation behind the idea of this work is that, once $C(\theta_s)$ has been tuned and tested on the simulator, an *ideal* control action \tilde{u} and the corresponding output \tilde{y} become available. The latter point opens the possibility to a *quantum leap* in the field of vehicle controls: if the simulator can be run in real-time on the ECU, the ideal \tilde{u} can be used as a nominal control action with no additional computations. If the simulator and the vehicle dynamics coincide, such an input could be directly applied to the real system. Indeed, if the simulator and the vehicle differ in some way, a second control loop needs to be designed accounting for the mismatch between \tilde{y} and the measured y . Let us denote the additional compensator as C_δ . Doing so, the control action on the physical vehicle becomes equal to $u_\delta + \tilde{u}$. We denote the above described control architecture as *Twin-in-the-Loop Control* (TiL-C), and depict its schematics in Fig. 2a.

The proposed architecture shows a set of interesting features as compared to the standard practice:

- if the simulator is a faithful replica of the vehicle, most of the system nonlinearity and complexity is managed by C , and *the nominal control action does most of the job*. The compensator C_δ would thus play the role of the controller of a linearized system in the neighborhood of an operating point;
- since C is operated on the simulator, it has access to all its states. This opens up the possibility of designing *state-feedback approaches* (for the nominal control action) even when the state is not (fully) available without the need of designing suitable observers;
- the EoL tuning procedure of Fig. 1 cannot be avoided, as C_δ needs to be tuned based on the mismatch between the vehicle and the simulator. However, *the design might be largely simplified*, as C_δ can be selected as a simple

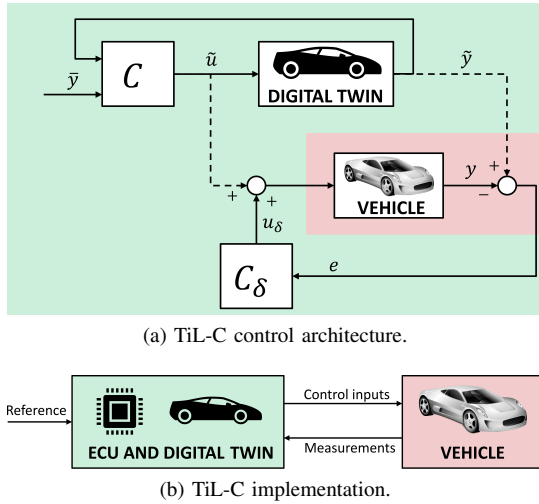


Fig. 2: TiL-C architecture (top) and its on-board implementation (bottom). The red areas denote the physical vehicle, whereas the green areas depict the software elements, namely the simulator and the controllers, running on an ECU.

controller with few parameters even when C is a complex (possibly optimization-based) controller.

- the TiL-C approach can in principle be *generalized* to any vehicle dynamics problem, classical examples being longitudinal, lateral and vertical dynamics control. In fact, many of these tasks can be formulated as reference tracking problems (see, *e.g.*, the yaw-rate tracking problem in [3] or [8]), thus the control scheme is conceptually the same of Fig. 2a.

Indeed, tuning the Digital Twin of scheme in Fig. 2a requires proper calibration. In what follows, we will assume that such calibration has been already performed by the car manufacturer according to standard practices [2], so this topic is left out of the scopes of the paper.

Remark 1. Note that the TiL scheme presented in this research might resemble an Hardware-in-the-Loop (HiL) architecture running on-board at real-time, where a faithful plant replica runs in-silico. However, there are important differences between HiL and TiL. On the one hand, the former is usually a test-bench for controller implementation [16], where some piece of hardware (*e.g.* an actuator, or an Electronic Control Unit) is inserted in the loop. On the other hand, in TiL architectures, the commands generated from the simulator are directly used to control the vehicle.

III. THE CASE STUDY: ACTIVE BRAKING CONTROL

As a case study to illustrate the potential of the proposed approach, let us consider the problem of longitudinal dynamics control during braking, *i.e.*, the design of an Anti-lock Braking System (ABS) for a high-performance car. In particular, we focus on the control of the wheel slip λ aimed to track a desired behaviour $\bar{\lambda}$ so that a certain braking force can be guaranteed, and the vehicle can stop without wheel locking. In this seminal study, we will consider two instances of the same simulator (as detailed next) to play the role of the digital

twin and the real car. The problem of selecting $\bar{\lambda}$ is instead out of the scopes of this work and will not be discussed here. Also, note that an ABS controller is generally independent on the outside road events, like traffic conditions: the ABS is activated and deactivated solely based on a request by the driver, which can stop its action at any time by releasing the brake pedal. We wish to stress that the main goal of this research is not to design a full braking control algorithm *per se*, but to simplify the EoL tuning problem by relying on the availability of the simulator on-board. As the nominal controller C in Fig. 2a, we thus start from the existing, state-of-the-art, MPC approach in [10]. The control action is computed independently for each wheel, thus $ij = fl, fr, rl, rr$, as showed in Fig. 3. In the figure, we also highlight the feedback measurements necessary for the MPC, which are the wheel ground contact point velocity $v_{x,ij}$, the actuated torque T_{ij}^{act} and the wheel longitudinal acceleration $a_{x,ij}$.




Fig. 3: MPC-based braking control [10]. The torques are computed independently for each wheel, then $ij = fl, fr, rl, rr$.

The wheel slip λ at corner ij is defined as

$$\lambda_{ij} = \frac{v_{x,ij} - \omega_{ij} R_{ij}}{\max\{v_{x,ij}, \omega_{ij} R_{ij}\}}, \quad (1)$$

where ω_{ij} is the wheel angular rate and R_{ij} is the wheel radius [9]. Note that the slip in Eq. 1 is always positive in case of braking, as the numerator is always < 0 , since the wheel decelerates.

We consider straight braking, *i.e.*, we assume almost zero camber and sideslip wheel angles. Such an assumption does not affect the scope and outcomes of the research; it is in fact well known that such variations affect the friction analogously to the vertical wheel loads, and can be neglected by suitably scheduling the slip references [9].

A. The vehicle and the simulator

Both the vehicle - a sport car - and its digital twins are here modeled in the VI-Grade CarRealTime (CRT) simulation environment [17]. The CRT simulator include a 6 degree-of-freedom (DOF) object for the chassis, a 1-DOF model for the suspensions and a 1-DOF model for the wheel dynamics. A model of the electro-hydraulic brake (EHB) is also added, so as to take into account realistic actuation limits. The most critical vehicle parameters influencing the dynamics during braking [9] are given in Table I.

M_{tot} represents the total vehicle mass, encompassing sprung and unsprung masses, whereas R_f^{nom} , R_r^{nom} represent the nominal wheel radius, for front and rear wheels respectively; note that the high-fidelity model also accounts for radius

M_{tot} [kg]	R_f^{nom} [m]	R_r^{nom} [m]	J_f^{nom} [kg · m ²]
1612	0.33	0.35	1.49
J_r^{nom} [kg · m ²]	l_f [m]	l_r [m]	h [m]
2.25	1.57	1.03	0.46

TABLE I: Braking-related parameters for the considered car.

variations due to, *e.g.*, increased wheel loads. J_f^{nom} and J_r^{nom} are the nominal spin inertia of front and rear wheels. l_f and l_r are the distances between the projection of the vehicle center-of-gravity (COG) on the ground and the front and the rear axles, respectively. h is the COG height from the ground.

Formally, the control inputs u and the set of driver commands ξ are

$$\begin{aligned} u &= [T_{fl}^{cmd} \quad T_{fr}^{cmd} \quad T_{rl}^{cmd} \quad T_{rr}^{cmd}]^t, \\ \xi &= [\phi_{throttle} \quad \phi_{brake} \quad \phi_{steer} \quad \phi_{gear}]^t \end{aligned} \quad (2)$$

where T_{ij}^{cmd} represents a torque command at wheel ij , $\phi_{throttle}$, ϕ_{brake} , ϕ_{steer} are respectively throttle, brake and steer driver request, while ϕ_{gear} is the inserted transmission gear. Indeed, when the braking controller is active, the driver brake request is bypassed. Due to the straight braking assumption, throttle and steer commands are negligible. The variables ξ simultaneously act on the simulator and on the vehicle, so the driver should be considered as an exogenous disturbance in the framework of Fig. 3.

In order to simulate a realistic difference between the vehicle and its digital twin, we will consider, in three different case studies, the effect of realistic measurement noises, the effect of a different mass distribution and of a variation of the tire-road force diagram. Specifically, since high performance cars are usually characterized by two front seats and a front trunk, we add a second concentrated mass on the passenger seat, and two unbalanced masses on the front trunk. A representation of the unmodeled loads is provided in Fig. 4, where it becomes clear how the presence of unmodelled masses changes the ratio between the distance of COG and front wheels l_f and that of COG and rear wheels l_r , thus varying the front/rear load transfer. The values of the additional masses are as in Table




Fig. 4: The considered vehicle with modeled (black) and unmodeled (light blue) masses.

II.

VI-CarRealTime models tire-road interaction via a set of Pacejka formulas, governing the longitudinal and lateral force generation at the tire - as well as more complex phenomena,




Fig. 5: Front wheel longitudinal force versus slip, in case of nominal vertical tire force. Left and right plots depict variations of the shape factor and peak friction, respectively.

m_d [kg]	m_p [kg]	m_t^l [kg]	m_t^r [kg]
75	80	90	30

TABLE II: Additional masses included in the vehicle model. Such masses are modeled as concentrated loads.

such as overturning moment. In order to further test the control robustness, we consider uncertainty in the longitudinal force generation, specifically by simulating a multiplicative mismatch μ_s on the peak friction and c_s on the shape factor. Figure 5 displays the nominal behaviour ($\mu_s = c_s = 1$), and different versions of the same model with perturbations of the two parameters.

B. Model predictive controller

The MPC in [10] is here considered as a nominal controller and is assumed to be designed before-hand, in the pre-development stage. Such a controller employs a wheel-specific predictive model of the slip dynamics

$$\dot{\lambda}_{ij}^{mpc} = \frac{1 - \lambda_{ij}^{mpc}}{v_{x,ij}} a_{x,ij} - \frac{R_{ij}^{nom}}{J_{ij}^{nom} v_{x,ij}} T_{ij}^{act}, \quad (3)$$

where λ_{ij}^{mpc} represents the prediction of the slip at each wheel. Parameters R_{ij}^{nom} , J_{ij}^{nom} are taken from Table I. The actuator dynamics is a second order system with the same model and parameters of [10]. The model (3) is used in the MPC under the assumption of constant wheel speed and acceleration during the prediction horizon T_p , namely

$$\begin{aligned} a_{x,ij}(t_0 + t) &= a_{x,ij}(t_0), \quad \forall t \in [t_0, t_0 + T_p] \\ v_{x,ij}(t_0 + t) &= v_{x,ij}(t_0), \quad \forall t \in [t_0, t_0 + T_p]. \end{aligned} \quad (4)$$

Where t_0 is the current time instant, and t is the independent time variable. Under the assumptions above - motivated by the different time scales between slip and chassis longitudinal dynamics [9] - the model in Eq. (3) becomes linear and time-invariant (LTI). Said LTI model can be then discretized and written in velocity form [18] - *i.e.*, transforming states and inputs into their instantaneous variations; the slip tracking error is then further introduced as a state. At this point, an integral action can be easily implemented in the MPC: such formulation has been shown to be robust to constant disturbances, guaranteeing zero steady-state error.

Five predictive steps are considered in the computation of the optimal control law, which is computed implicitly; with regards to other details regarding the optimal control formulation, we refer the reader to [10].

C. Sensor model

Any controller acting on the system described in Section III-A is based upon sensor measurements. For a realistic case study, we cannot neglect the effect of the noise model on the performance. In particular, we introduce the noise affecting speed, acceleration and slip as detailed next. No noise is added to the braking torque, assuming it is fully controllable and known.

Acceleration. Longitudinal acceleration measurements at each wheel ($a_{x,ij}$) are seldom available on production vehicles. Chassis acceleration is then employed when considering straight braking [10]. Such signal is obtained through a Inertial Measurement Unit (IMU), usually affected by high-frequency noise [19]. We thus consider random Gaussian noise in the acceleration measurement a_x^n

$$a_x^n = a_x + n_a, \quad n_a \sim WN(0, \sigma_a^2), \quad (5)$$

where $n_a \sim WN(\mu, \sigma^2)$ denotes a Gaussian noise with expected value μ and variance σ^2 .

Chassis speed. The COG-referenced longitudinal speed v_x cannot be directly measured without high accuracy Global Positioning System (GPS) sensors, then it is usually estimated through state observers [20], and reported at each wheel via kinematic relations. Hence, we consider low-pass filtered version n_{lp} of a white noise n_{v_x} on the speed measurement v_x^n , to mimic the state observer dynamics

$$v_x^n = v_x + n_{lp}(n_{v_x}), \quad n_{v_x} \sim WN(0, \sigma_{v_x}^2), \quad (6)$$

where $\sigma_{v_x}^2$ is the speed white noise variance.

Wheel speed. The angular rates are measured through incremental encoders: such sensors are well known to be affected by periodic noise [21], mostly due to unavoidable geometrical or misalignment errors in the sensor structure. The amplitude of such a noise increases depending on the rotational speed itself. We thus include a speed-scheduled sinusoidal error term ω_{ij}^n

$$\omega_{ij}^n = \omega_{ij} + A_\omega \sin(\omega_{ij} t), \quad A_\omega = \omega_0^n + k_\omega \omega_{ij}, \quad (7)$$

where ω_0^n and k_ω are tunable parameters.

Wheel slip. Given the noise models defined in (6) and (7), the slip in (1) is also affected by a noise term, depending on both noisy speed and wheel rate measurements. The noisy slip measurement λ_{ij}^n reads

$$\lambda_{ij}^n = \lambda_{ij} + n_\lambda^{ij}(v_{x,ij}^n, \omega_{ij}^n). \quad (8)$$

The terms σ_{v_x} , k_ω , ω_0^n appearing in equations above are tuned to achieve signal-to-noise ratios (*SNRs*) on the slip measurements which are compatible with those observed in real setups [22], namely $SNR \approx 4$. A comparison between noiseless and noisy slip signals is given in Fig. 6.

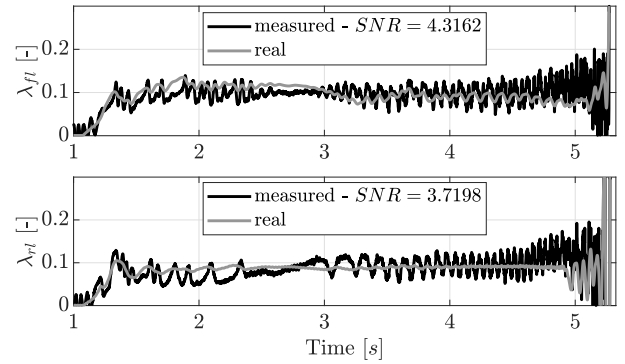


Fig. 6: Noisy and noiseless slip signals, during a closed-loop braking maneuver.

IV. TiL-C DESIGN

Given the vehicle model and the benchmark controller described in the previous section, we now discuss the TiL-C approach to the active braking problem, see Fig. 7.

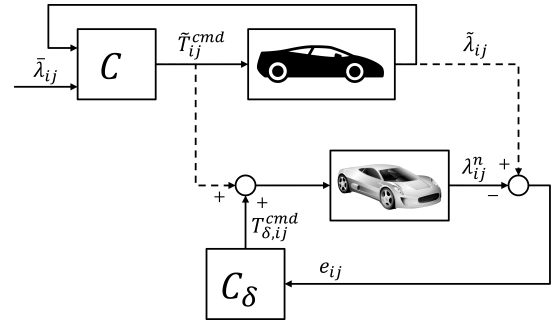


Fig. 7: TiL-C braking control. The nominal MPC runs on the simulator in the above control loop, while a second loop is closed on the physical vehicle via C_δ to compensate for unmodelled and stochastic dynamics.

The nominal MPC is designed and used on the simulator and returns a torque for each wheel \tilde{T}_{ij}^{cmd} , corresponding to a certain (ideal) slip output $\tilde{\lambda}_{ij}$. The input to the physical vehicle is given by the sum of the nominal torque \tilde{T}_{ij}^{cmd} and the outcome of the compensator C_δ , which is non-zero anytime the measured slip is different from $\tilde{\lambda}_{ij}$.

Remark 2. Let us point out that the TiL-C architecture proposed herein is absolutely general and independent on the employed nominal controller - *i.e.* the block C in Fig. 7. For this reason, the assessment of a TiL-C scheme should be made in comparison to the nominal control loop, while the absolute performance is less relevant (in fact, the purpose of the scheme is to retrieve the desired behaviour, and not to outperform any particular strategy).

Remark 3. Note that, for an actual implementation of the controller, some activation logic for the Digital Twin is necessary: if left to freely evolve over time, the simulator might eventually diverge. For this reason, the DT has to be activated whenever the controlled maneuver starts (*e.g.* when the driver starts braking). This can be simply achieved by freezing the

simulation and then re-activating it by properly setting the initial states when necessary.

A. The TiL-C block

The compensator scheme is showed in Fig. 8.




Fig. 8: The detailed C_δ architecture.

A linear regulator $R(s) = N_r(s)/D_r(s)$ processes the measurement error e_{ij} so as to obtain the control action $T_{\delta,ij}^{cmd}$. Such a regulator is implemented in an *anti-windup* fashion [23], where a suitable de-saturation function $\Gamma(s)$ is employed in the scheme. The regulator is selected as a Proportional-Integral (PI) one

$$R(s) = \frac{N_r(s)}{D_r(s)} = \frac{k_p(1 + sT_i)}{sT_i}, \quad \Gamma(s) = 1 + sT_i. \quad (9)$$

Then, it is discretized at sampling frequency $f_s = 200 \text{ Hz}$ via the Tustin approach. Notice that the parameterization of $R(s)$ might be different for front and rear wheels, *e.g.*, due to the different spin inertia or radii.

Since the wheel slip dynamics highly depend on the vehicle speed (see again (3)), we include a speed-based scheduling law, defined *a-priori*, for the controller gain [13]:

$$k_p = \begin{cases} k_p^{nom}, & v_x \geq v_x^{ub} \\ k_p^{nom} \cdot \left(k_p^{lb} + \frac{v_x - v_x^{lb}}{v_x^{ub} - v_x^{lb}} \right), & v_x^{lb} \leq v_x < v_x^{ub} \\ k_p^{nom} \cdot k_p^{lb}, & v_x < v_x^{lb}, \end{cases} \quad (10)$$

where v_x^{lb} , v_x^{ub} , k_p^{lb} are suitable tuning knobs. Considering different parameters for front and rear wheels, one obtains the following parameter vector, to be tuned

$$\theta_{\text{TiL-C}} = [k_{p,f}^{nom} \quad T_{i,f} \quad k_{p,r}^{nom} \quad T_{i,r}]. \quad (11)$$

Notice also that, in case the simulated car completes the braking maneuver before the physical vehicle *e.g.*, due to a lower simulated mass, the reference λ_{ij} goes to zero, and so does the nominal control \tilde{T}_{ij}^{cmd} .

This calls for the smooth switching structure in the lower loop of Fig. 8, in which we introduce the signal ϕ , driving the de-activation of the braking control for the simulated vehicle.

As soon as the simulator ends the braking maneuver earlier than the vehicle, the reference signal λ_{ij} is switched to a constant value $\bar{\lambda}$. Then, a second controller tracks the total torque commanded to the system $\tilde{T}_{ij}^{cmd} + T_{\delta,ij}^{cmd}$. In this way,

a smooth switching is guaranteed, from the full TiL control action to C_δ only. Such scheme *de-facto* implements a soft-insertion of the control action, which is a typical solution in PI control design [23], where $\tilde{T}_{ij}^{cmd} + T_{\delta,ij}^{cmd}$ serves as a manual-mode control action. Since the accurate tuning of such a de-activation logic is not among the scopes of this research, we assume that the TiL-C is de-activated as soon as the speed hits 10 km/h .

B. Controller tuning

As C_δ is aimed to control the dynamics of the residual between the system output and the output of the best available model, no model can be used to properly tune such C_δ in *model-based fashion*. The design of this block must therefore rely *only* on measurements collected on the plant.

Many data-driven methods for tuning PI controller parameters exist, see, *e.g.*, [24], [25], [26], [27]. However, such approaches are mainly defined for LTI systems and are not suited for the specific schemes of Fig. 8, where an additional control action is provided, coming from the nominal closed-loop. For the above reasons, we will instead employ a Bayesian Optimization (BO) rationale [28], [29]. BO deals with black-box optimization problems, where the cost function and the constraints are unknown but the values corresponding to some instances of the decision variables can be properly “measured”. Specifically, BO relies on the assumption that the cost function, here describing the closed-loop performance with certain control parameters, can be modeled as a Gaussian process (GP). Closed-loop data can then be used to estimate a GP model of such a cost function, while a suitable *acquisition function* is chosen to select the next set of parameters to evaluate, by looking for a balanced *exploration/exploitation* trade-off. For a more accurate description of the BO procedure, the interested reader can refer to, *e.g.*, [30].

In the following, we consider BO as a decision maker, in order to solve an optimization problem of the type

$$\begin{aligned} \min J(\theta) \\ \text{subject to } \theta \subseteq \Theta, \end{aligned} \quad (12)$$

whereas θ is a set of controller parameters, to be searched for within an *a-priori* defined set Θ . Since the TiL-C goal is the control of residual dynamics between ideal and real loops, J is selected as the root mean square of the average slip tracking error among the four wheels

$$J(\theta_{\text{TiL-C}}) = \left(\sum_{k=1}^{N_s} \left(\frac{\sum_{ij=fl,fr,rl,rr} e_{ij}^2(k)}{4N_s} \right) \right)^{0.5}. \quad (13)$$

where the dependence of e_{ij} , $ij = fl, fr, rl, rr$ upon $\theta_{\text{TiL-C}}$ is dropped for the sake of brevity, and N_s is the number of samples in the experiment. From a set of closed-loop experiments, one could re-calibrate the compensator parameters.

Specifically, the experiment used to train C_δ is illustrated in Fig. 9. A coasting-down phase is followed by a strong braking, yielding the braking control activation. The slip reference signal is superimposed with a pulse wave varying signal, in order to better excite the system dynamics to control.




Fig. 9: The training experiment for the BO procedure.

V. SIMULATION RESULTS

In this section, we will show the performance of the TiL-C scheme on a test braking maneuver. As a baseline for a fair assessment of the results, we will also consider a standard EoL procedure, in which the MPC is fine tuned using the same data available for the design of C_δ .

More specifically, since the aim of the EoL calibration is to obtain on the physical vehicle the same performance attained on the digital twin, the predictive model parameters are adjusted so as to minimize the distance between the ideal and the measured output in closed-loop. A suitable cost function to this scope is the root mean square of the wheel-averaged MPC prediction error

$$J(\theta_{\text{MPC}}) = \left(\sum_{k=1}^{N_s} \left(\frac{\sum_{ij=fl,fr,rl,rr} (\lambda_{ij}(k) - \lambda_{ij}^{\text{mpc}}(k))^2}{4N_s} \right)^{0.5} \right), \quad (14)$$


where $\lambda_{ij}^{\text{MPC}}$ is computed according to Eq. (3), and projected forward in time - for each time step - depending on the prediction horizon. θ_{MPC} contains the predictive model parameters, found in Eq. (3).

Three indices are employed to quantitatively assess the performance of the controllers; the first - and most important one - is the cost in Eq. (13). The second index instead represents the control effort, evaluated through the time derivative of the actuated braking torques, $\dot{T}_{ij}^{\text{act}}$ (e.g. as done in [11])


$$J_u = \left(\sum_{k=1}^{N_s} \left(\frac{\sum_{ij=fl,fr,rl,rr} (\dot{T}_{ij}^{\text{act}}(k))^2}{4N_s} \right)^{0.5} \right) \quad (15)$$

Finally, the total braking time is displayed, defined as the time passing from the braking control activation until vehicle speed hitting 10 km/h.

In what follows, we will consider three case studies. In the first one, we will assume that the only difference between the digital twin and the physical vehicle is the presence of measurement noise (as expressed in Section III-C). In the second scenario, we will neglect the effect of noise but we will investigate the case where the physical vehicle has a different mass distribution, as illustrated in Section III-A. In the third scenario, we will consider both the effect of noise and the modified mass distribution, while also introducing uncertainty on the tire-road friction



(a) Commanded (solid) and actuated (dashed) torques.



(b) Actual (solid) and reference (dashed) wheel slips.

Fig. 10: MPC and TiL-C performance in case of noisy measurements: front-left wheel (top plots) and rear-left wheel (bottom plots).

model, as described in Section III-A. Namely, the shape factor scaling c_s is set to 1.2, while the peak friction scaling is set to 0.8. This last test is indeed very close to what one could expect when driving an actual vehicle. In all tests, the braking starts after ≈ 1 s, at the speed of 196 km/h.

All the tests in this section have been performed off-line via Matlab/Simulink environments. Real-time implementation and feasibility analysis of the controller is illustrated in Section VI.

A. The case of noisy measurements

The braking experiment with noisy data is displayed in Fig. 10. Only the front-left and the rear-left wheels are illustrated, as no significant differences between left-right corners exist during straight braking. Figure 10b shows the reference tracking performance, whereas Figure 10a depicts the actuated torques. As one can note, TiL-C is able to maintain a smoother tracking of the reference, in spite of the significant measurement noise. This is due to the fact that MPC-computed torques show undesirable oscillations induced by the presence of noise (neglected in the prediction model). Figure 11 depicts the vehicle speed and deceleration profiles; the deceleration profile reflects what noted from the slip tracking, in that more accurate deceleration is achieved via TiL-C. Table III confirms what is noted in the figures: TiL-C significantly reduces both the tracking error and the control action aggressiveness, while also improving the braking time.






Fig. 11: Vehicle speed and acceleration in case of noisy measurements, for the MPC and TiL-C controllers.

Performance indices			
	J_λ [%]	J_u [Nm/s]	J_{time} [s]
TiL-C	1.06	4.17	3.88
MPC	2.85	10.82	4.00

TABLE III: Performance indices in the noisy vehicle validation experiment.



(a) Commanded (solid) and actuated (dashed) torques.



(b) Actual (solid) and reference (dashed) wheel slips.

Fig. 12: MPC and TiL-C performance in case of different masses configuration: front-left wheel (top plots) and rear-left wheel (bottom plots).




Fig. 13: Vehicle speed and acceleration in case of different masses configuration, for the MPC and TiL-C controllers.

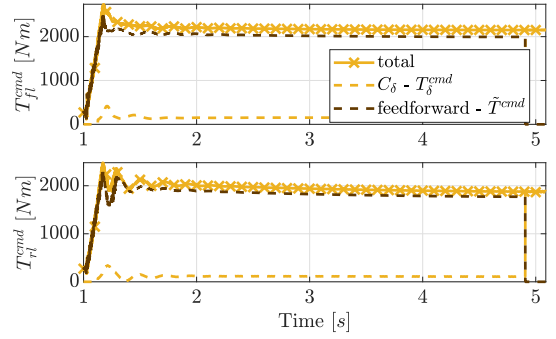


Fig. 14: Split between the nominal input and the output of C_δ , in the test with different mass configuration in the physical vehicle.

B. The case of additional unmodelled loads

Let us consider the case where the physical vehicle - with noiseless data - is equipped with additional (unmodelled) masses.

Figure 12 shows the actuated torques (upper plot) and the wheel slips (mid plot). As one can note, the MPC controlled vehicle exhibits significant slip tracking error and oscillations during the transient. This is due to the increased mass on the front trunk, almost leading to instability when coupled with the actuator nonlinearities. Instead, the TiL-C scheme




Fig. 15: MPC with piecewise slip references (dark purple) and MPC with the ideal outputs as references (light blue) for the case where the physical vehicle has a different mass configuration.

shows good performance as the reference is well tracked with a proper torque actuation. Figure 13 shows the vehicle speed and deceleration profiles; also in this case, the slip tracking oscillations noted in the MPC reflects into the deceleration profile, which shows oscillating behavior at the beginning of the braking maneuver. In Figure 14, we also show the split between the contributions given by the nominal control and the TiL-C compensator. From this additional plot, it can be noted that the most significant contribution of the control action is in fact the nominal one, while C_δ only produces a small compensation term (thus confirming the suitability of a linear compensator for small-signal regulation). One can also appreciate the importance of the switching architecture described in Fig. 8. In fact, after approximately 5 s, the nominal torque goes to zero, as the virtual vehicle is stopping due to the reduced mass. However, the overall control action is kept at the same level due to the second loop running in parallel (Fig. 8).

As a final remark, one might argue that feeding the MPC scheme in Fig. 3 with the ideal outputs $\tilde{\lambda}_{ij}$ instead of the piecewise references might increase performance, as the former signals are smoother. However, Fig. 15 shows that no visible improvements are achieved. This is due to the fact that the traditional controller is characterized by a single block that needs to be both suitable for the nominal case and robust to parameter variations. Instead, the two blocks building the TiL-C scheme play different roles and, while the MPC is aimed only to push the nominal loop at its limits, the residual dynamics is taken care of by the additional compensator. The above observations are confirmed by the performance indices reported in Table IV. For completeness, we report in

Performance indices			
	J_λ [%]	J_u [Nm/s]	J_{time} [s]
TiL-C	0.96	3.05	3.96
MPC	2.00	7.18	4.01

TABLE IV: Performance indices for the case where the physical vehicle has a different mass configuration.

Table V also a comparison between the two strategies when both measurement noises and unmodeled masses characterize the physical vehicle. As expected, the qualitative conclusions previously discussed are confirmed also in this case.

Performance indices			
	J_λ [%]	J_u [Nm/s]	J_{time} [s]
TiL-C	1.84	3.49	4.01
MPC	3.96	13.85	4.19

TABLE V: Performance indices for the case where the physical vehicle has a different mass configuration and the measurements are noisy.

C. The case of modified tire-road interaction

Finally, we consider the case where the physical vehicle has a different mass configuration and the tire-road force




Fig. 16: MPC and TiL-C slip tracking performance in case of noisy measurements, perturbed masses and tire-road interaction: front-left wheel (top plot) and rear-left wheel (bottom plot).

Performance indices			
	J_λ [%]	J_u [Nm/s]	J_{time} [s]
TiL-C	2.19	8.31	4.95
MPC	3.17	12.04	5.00

TABLE VI: Performance indices for the case where the physical vehicle has a different mass configuration, the measurements are noisy and the tire-road interaction is modified.

generation model is modified (peak friction μ_s scaled to 0.8, shape factor c_s scaled to 1.2). For the sake of brevity, we report here only the controller performance in terms of slip tracking, see Fig. 16. As one can notice, TiL-C slightly improves the performance with respect to the nominal controller - which exhibits significant oscillations, as for the case of modified mass distribution alone of Fig. 12b - and is thus robust to a low-friction test. Table VI displays the performance metrics; let us note that TiL-C yields a reduced control effort too. Finally, we wish to stress that the braking time is increased by the 20 % with respect to the other tests: this is explainable with the reduced maximum friction.

VI. HARDWARE-IN-THE-LOOP VALIDATION

So far, simulation analyses in an off-line fashion have been showed. In order to prove that the TiL-C framework is implementable at real-time on existing technology, we perform an HiL validation following the scheme in Fig. 17.

As highlighted by the figure, TiL-C is implemented in a real-time computer (RTC); more specifically, we employed AutoHawk [31], an off-the-shelf Red Hat Linux based computer which is suitable for on-board usage. The available AutoHawk is provided with six 16 GB RAM slots and a 3.20 GHz Intel Xeon 6146 processor.

AutoHawk executes both the virtual and the real control loops, in details, we implement:

- Four MPCs, one per each wheel, as described in Section III-B. The MPCs are implemented implicitly, *i.e.* solving a QP problem at each time step, with a sampling time of 5 ms;

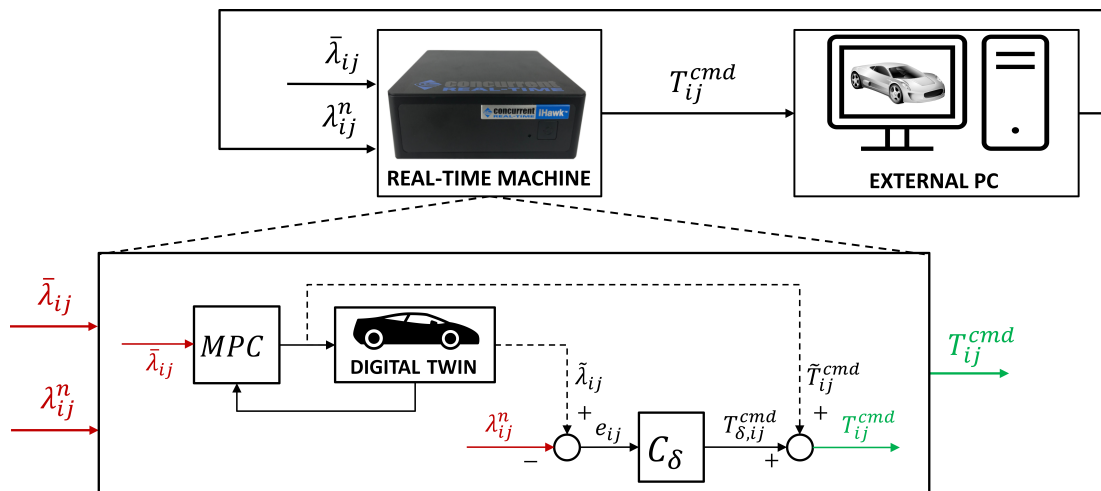


Fig. 17: Hardware-in-the-Loop experiment layout.

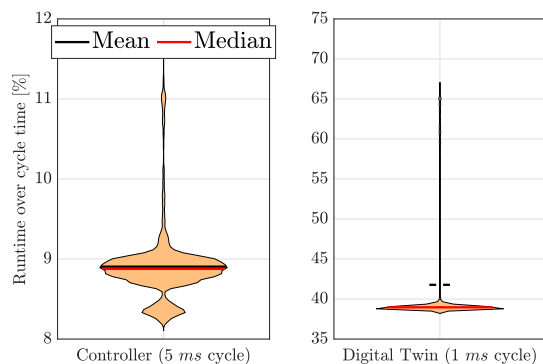


Fig. 18: Summary statistics (violin plots) of the Controller (left plot) and Digital Twin (right plot) execution runtimes. The runtimes are displayed as a percentage of the corresponding cycle time.




Fig. 19: Time domain comparison of actuated front-left braking torque in the HiL test.

- Four C_δ compensators, one per each wheel, as described in Section IV. The compensators are executed at a sampling time of 5 ms;
- The Digital Twin (modeled in CarRealTime), running at real-time with a sampling time of 1 ms.

An external laptop computer is solely employed to act as

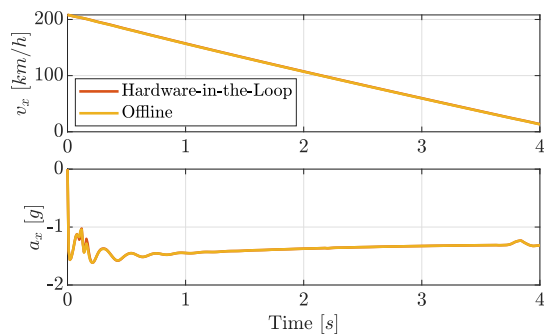


Fig. 20: Time domain comparison of speed and deceleration profiles in the HiL test.

a vehicle simulator; the vehicle is modeled in CarRealTime and executed in Simulink (via the Real-Time toolbox). The communication among the two machines is managed by User Datagram Protocol (UDP). The AutoHawk operating system features a scheduler which is in charge of executing the processes at real-time, according to the allowed cycle times, more specifically, two processes are executed, let apart the auxiliary ones (i.e. the scheduler itself, and the csv data logger): the first process is the Digital Twin one, the second one is the remaining part of the controller (featuring the MPC and C_δ). The two processes are respectively executed with cycle times of 1 ms and 5 ms.

In the HiL test, we consider a starting speed of 208 km/h, and the initial braking torque is set to its maximum value, so as to challenge the MPC and the solving time of the optimal control problem. First, a time-domain comparison of the braking torque actuated to the front-left wheel is shown in Fig. 19: as the reader can notice, the two signals are practically superimposed - only few differences are present, due to the different solver employed by the RTC. The speed and acceleration profiles in the same test are displayed in Fig. 20.

We are also interested in evaluating the scheduler task execution and the runtimes: Figure 18 shows such piece of

information (a violin plot representation is employed, given that the data are not normally distributed). The left plot highlights that the execution of the controllers (the MPC and C_δ) accounts for less than the 10% of the computational load in the 5 ms process: this means that significant margin exist for even more complex control logics. The right plot shows as the execution of the Digital Twin at 1 ms accounts for the $\approx 40\%$ of the cycle time on average, with a long tail stretching up to the 65%.

Overall, we verified that the computationally demanding proposed strategy is feasible at real-time: we are able to execute it without any overrun, at very high sample frequencies, and employing off-the-shelf hardware.

VII. CONCLUSIONS

In this paper, we have proposed a new approach for vehicle dynamics control, based on the use of a full-fledged simulator, to be run *in-the-loop* directly on the vehicle ECU in order to compute the nominal control action. The main advantage of such a configuration is that, in the EoL calibration phase, there is no need anymore to fine tune the (possibly many parameters of the) controller designed and calibrated on the simulator. Instead, as far as the simulator is an accurate software replica of the dynamics of the vehicle, a simple linear regulator can be used to compensate for the mismatch between the expected and the measured outputs. The new architecture shows promising results particularly when dealing with measurement noise and unmodeled terms and outperforms the state-of-the-art solution within an active braking control case study. Finally, we show that the proposed architecture is real-time feasible on off-the-shelf hardware.

This being a preliminary work on the topic, many research questions remain open, *e.g.*, concerning stability analysis, safe controller tuning and the generalization to different vehicle dynamics problems.

REFERENCES

- [1] G. Bhatti, H. Mohan, and R. Raja Singh, "Towards the future of smart electric vehicles: Digital twin technology," *Renewable and Sustainable Energy Reviews*, vol. 141, p. 110801, 2021.
- [2] E. Kutluay and H. Winner, "Validation of vehicle dynamics simulation models—a review," *Vehicle System Dynamics*, vol. 52, no. 2, pp. 186–200, 2014.
- [3] A. Lucchini, S. Formentin, M. Corno, D. Piga, and S. M. Savaresi, "Torque vectoring for high-performance electric vehicles: An efficient mpc calibration," *IEEE Control Systems Letters*, vol. 4, no. 3, pp. 725–730, 2020.
- [4] M. Tanelli, G. Panzani, S. M. Savaresi, and C. Pirola, "Transmission control for power-shift agricultural tractors: Design and end-of-line automatic tuning," *Mechatronics*, vol. 21, no. 1, pp. 285–297, 2011.
- [5] C. E. Garcia and M. Morari, "Internal model control. a unifying review and some new results," *Industrial & Engineering Chemistry Process Design and Development*, vol. 21, no. 2, pp. 308–323, 1982.
- [6] H. R. Ossareh, S. Wisotzki, J. Buckland Seeds, and M. Jankovic, "An internal model control-based approach for characterization and controller tuning of turbocharged gasoline engines," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 2, pp. 866–875, 2021.
- [7] H.-M. Wu and R. Tafreshi, "Observer-based internal model air–fuel ratio control of lean-burn si engines," *IFAC Journal of Systems and Control*, vol. 9, p. 100065, 2019.
- [8] M. Canale, L. Fagiano, A. Ferrara, and C. Vecchio, "Comparing internal model control and sliding-mode approaches for vehicle yaw control," *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, no. 1, pp. 31–41, 2008.
- [9] S. M. Savaresi and M. Tanelli, *Active braking control systems design for vehicles*. Springer Science & Business Media, 2010.
- [10] G. Riva, S. Formentin, M. Corno, and S. M. Savaresi, "Model predictive control of high-performance braking systems: a force-based approach," *IEEE Control Systems Letters*, pp. 1–1, 2022.
- [11] D. Tavernini, F. Vacca, M. Metzler, D. Savitski, V. Ivanov, P. Gruber, A. E. Hartavi, M. Dhaens, and A. Sorniotti, "An explicit nonlinear model predictive abs controller for electro-hydraulic braking systems," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 5, pp. 3990–4001, 2019.
- [12] M. S. Basrah, E. Siampis, E. Velenis, D. Cao, and S. Longo, "Wheel slip control with torque blending using linear and nonlinear model predictive control," *Vehicle System Dynamics*, vol. 55, no. 11, pp. 1665–1685, 2017.
- [13] T. Johansen, I. Petersen, J. Kalkkuhl, and J. Ludemann, "Gain-scheduled wheel slip control in automotive brake systems," *IEEE Transactions on Control Systems Technology*, vol. 11, no. 6, pp. 799–811, 2003.
- [14] G. Riva, S. Formentin, M. Corno, and S. M. Savaresi, "Twin-in-the-loop state estimation for vehicle dynamics control: theory and experiments," *arXiv:2209.02263*, 2022.
- [15] G. Delcaro, F. Dettù, S. Formentin, and S. M. Savaresi, "Dealing with the curse of dimensionality in twin-in-the-loop observer design," in *22nd IFAC World Congress*, 2023.
- [16] H. K. Fathy, Z. S. Filipi, J. Hagen, and J. L. Stein, "Review of hardware-in-the-loop simulation and its prospects in the automotive area," in *Modeling and simulation for military applications*, vol. 6228. SPIE, 2006, pp. 117–136.
- [17] VI-Grade, "Vi-carrealtime," <https://www.vi-grade.com/en/products/vi-carrealtime/>, 2022, [Online; accessed 23-March-2022].
- [18] L. Magni and R. Scattolini, *Advanced and multivariable control*. Pitagora editrice Bologna, 2014.
- [19] B. A. N. Campos and J. M. S. T. Motta, "Online measuring of robot positions using inertial measurement units, sensor fusion and artificial intelligence," *IEEE Access*, vol. 9, pp. 5678–5689, 2021.
- [20] E. Hashemi, A. Kasaiezadeh, S. Khosravani, A. Khajepour, N. Moshchuk, and S.-K. Chen, "Estimation of longitudinal speed robust to road conditions for ground vehicles," *Vehicle System Dynamics*, vol. 54, no. 8, pp. 1120–1146, 2016.
- [21] G. Panzani, M. Corno, and S. M. Savaresi, "On the periodic noise affecting wheel speed measurement*," *IFAC Proceedings Volumes*, vol. 45, no. 16, pp. 1695–1700, 2012, 16th IFAC Symposium on System Identification.
- [22] S. Formentin, P. De Filippi, M. Corno, M. Tanelli, and S. M. Savaresi, "Data-driven design of braking control systems," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 1, pp. 186–193, 2013.
- [23] K. J. Åström and T. Hägglund, *Advanced PID Control*. International Society of Automation, 2006.
- [24] T. Yamamoto, K. Takao, and T. Yamada, "Design of a data-driven pid controller," *IEEE Transactions on Control Systems Technology*, vol. 17, no. 1, pp. 29–39, 2008.
- [25] S. Formentin, S. M. Savaresi, and L. Del Re, "Non-iterative direct data-driven controller tuning for multivariable systems: theory and application," *IET control theory & applications*, vol. 6, no. 9, pp. 1250–1257, 2012.
- [26] S. Formentin, K. Van Heusden, and A. Karimi, "A comparison of model-based and data-driven controller tuning," *International Journal of Adaptive Control and Signal Processing*, vol. 28, no. 10, pp. 882–897, 2014.
- [27] S. Formentin, M. C. Campi, A. Carè, and S. M. Savaresi, "Deterministic continuous-time virtual reference feedback tuning (vrft) with application to pid design," *Systems & Control Letters*, vol. 127, pp. 25–34, 2019.
- [28] M. Khosravi, C. Koenig, M. Maier, R. S. Smith, J. Lygeros, and A. Rupenyan, "Safety-aware cascade controller tuning using constrained bayesian optimization," *IEEE Transactions on Industrial Electronics*, pp. 1–1, 2022.
- [29] M. Neumann-Brosig, A. Marco, D. Schwarzmann, and S. Trimpe, "Data-efficient autotuning with bayesian optimization: An industrial control study," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 3, pp. 730–740, 2020.
- [30] P. I. Frazier, "A tutorial on bayesian optimization," *arXiv preprint arXiv:1807.02811*, 2018.
- [31] VI-Grade, "Autohawk," <https://www.vi-grade.com/en/products/autohawk/>, 2022, [Online; accessed 17-October-2022].



Federico Dettù received his B.Sc. on Sep. 2017 and his M.Sc. cum laude on Oct. 2020 in Automation and Control Engineering from Politecnico di Milano. From January to September 2020, he was a Visiting Researcher at Stanford University, USA. Starting from Nov. 2020, he joined the mOve research group as a PhD student in Information Technology, Systems and Control area, at Dipartimento di Elettronica, Informazione e Bioingegneria of Politecnico di Milano. His research interests regard data-based estimation and control approaches for automotive

systems.



Simone Formentin was born in Legnano, Italy, in 1984. He received his B.Sc. and M.Sc. degrees cum laude in Automation and Control Engineering from Politecnico di Milano, Italy, in 2006 and 2008, respectively. In 2012, he obtained his Ph.D. degree cum laude in Information Technology within a joint program between Politecnico di Milano and Johannes Kepler University of Linz, Austria. After that, he held two postdoctoral appointments at the Swiss Federal Institute of Technology of Lausanne (EPFL), Switzerland and the University of Bergamo,

Italy, respectively. Since 2014, he has been with Politecnico di Milano, first as an assistant professor, then as an associate professor. He is the chair of the IEEE TC on System Identification and Adaptive Control, the social media representative of the IFAC TC on Robust Control and a member of the IFAC TC on Modelling, Identification and Signal Processing. He is an Associate Editor of Automatica and the European Journal of Control. His research interests include system identification and data-driven control with a focus on automotive and financial applications.



Sergio Matteo Savaresi received the M.Sc. in Electrical Engineering (Politecnico di Milano, 1992), the Ph.D. in Systems and Control Engineering (Politecnico di Milano, 1996), and the M.Sc. in Applied Mathematics (Catholic University, Brescia, 2000). After the Ph.D. he worked as management consultant at McKinsey&Co, Milan Office. He is Full Professor in Automatic Control at Politecnico di Milano since 2006. He is Deputy Director and Chair of the Systems&Control Section of Department of Electronics, Computer Sciences and Bioengineering

(DEIB), Politecnico di Milano. He is author of more than 500 scientific publications. His main interests are in the areas of vehicles control, automotive systems, data analysis and system identification, non-linear control theory, and control applications, with special focus on smart mobility. He has been manager and technical leader of more than 400 research projects in cooperation with private companies. He is co-founder of 8 high-tech startup companies.