Chapter 1
Multi-horizon MPC and Its Application to the Integrated Power and Thermal Management of Electrified Vehicles

Qiuhao Hu, Mohammad Reza Amini, Ilya Kolmanovsky, and Jing Sun

Abstract The chapter provides an overview of several recent applications of model predictive control (MPC) to coordinated power and thermal management in connected and automated electrified vehicles. Such applications of MPC face significant challenges due to the need to accommodate different timescales of faster power and slower thermal dynamics, interactive constraints, and the need to generate and accommodate long-term vehicle speed and traction power forecasts in the optimization. It will be shown that a multi-horizon formulation of MPC consisting of a shorter receding horizon phase and a longer shrinking horizon phase can help address these challenges. Perspectives on creating forecasts and exploiting multi-horizon MPC for energy efficiency optimization based on several case studies are summarized.

Keywords Model predictive control · Multi-horizon model predictive control · Real-time optimization · Energy management · Thermal management · Integrated power and thermal management · Hybrid electric vehicles · Electric vehicles · Battery charging · Fast charging · Connected vehicles · Long term forecasting · Long term preview · Long prediction horizon · Uncertainty mitigation · Computational complexity reduction

Q. Hu · M. R. Amini · J. Sun
Department of Naval Architecture and Marine Engineering, University of Michigan, Ann Arbor, MI, USA
e-mail: qhhu@umich.edu

M. R. Amini
e-mail: mamini@umich.edu

J. Sun
e-mail: jingsun@umich.edu

I. Kolmanovsky (✉)
Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI, USA
e-mail: ilya@umich.edu

A. Daniilidis et al. (eds.), *Model Predictive Control*, Dynamic Modeling and Econometrics in Economics and Finance 31,
https://doi.org/10.1007/978-3-031-85256-5_1

1.1 Introduction

Over the past few decades, Model Predictive Control (MPC) [22, 41] has emerged as a vibrant discipline for control of constrained linear and nonlinear systems with numerous applications, including mobility and transportation. Despite many advances, the application of MPC to control *integrated* systems, where heterogeneous physical entities are combined to form a highly functional union, still remains one of the important research frontiers. In such integrated systems MPC must exploit complementary characteristics of integrated components and account for the component constraints and their operating environment in achieving safe and optimal performance. Such desired performance can be achieved by leveraging the tight physical coupling among the multiple components involved, calling for integrated optimization solutions over multiple domains and timescales.

Many integrated systems have dynamics evolving over different timescales. Examples of such systems are abound in chemical process industry [47], control of electrical power systems and microgrids [14], in mining industry [11], in electrified ground vehicles [3, 4, 33], and in aircraft [16, 34]. When MPC is utilized for control of such a multi-timescale system, the slower dynamics often dictate the controller design, as they need to be optimized over a relatively long prediction horizon to achieve the best performance. At the same time, due to the presence of faster dynamics and the need to handle the transients, the sampling period has to be kept short. Consequently, the computational footprint of the resulting MPC solutions can be very large. Furthermore, an accurate preview of operating conditions needs to be available over a long horizon, which represents another formidable challenge for implementing MPC.

The singular perturbation theory [35] can be utilized for controlling systems with dynamics evolving over multiple timescales. In such an approach, a system with explicit timescale separation in the dynamics is decomposed into two reduced-order subsystems with different timescales, see [6, 12, 13, 17]. Then "fast" and "slow" MPCs are designed and applied to control fast and slow dynamics, respectively. However, a standard requirement in the singular perturbation theory is for the system to have an isolated equilibrium manifold for the fast dynamics, which may not always be the case. Moreover, the timescale decomposition strategy may not be applicable to systems with constraints which are dynamically changing at different (slow and fast) rates.

An alternative approach for dealing with multi-timescale systems is the hierarchical MPC (H-MPC) [3, 4, 10, 18, 34, 49]. H-MPC computes the optimal reference values for the "slow" dynamics over a relatively long prediction horizon. Then, a tracking problem is solved over a much shorter prediction horizon with fast sampling to track the planned references and compute the control commands. Unlike the decentralized fast-slow MPC approach based on the singular perturbation theory, H-MPC typically requires and leverages the communication between the slow and fast MPCs to deal with the coupled constraints.

In [3, 4], an H-MPC solution was developed that exploited two receding horizon MPCs. In the upper (scheduling) layer, an economic MPC was used over a long prediction horizon with a slow sampling rate. In the lower (piloting) layer, a tracking MPC was used over a much shorter prediction horizon and with a fast update rate. The proposed H-MPC was applied to the integrated battery and cabin thermal management in electrified connected and automated vehicles (CAVs), demonstrating its effectiveness in handling coupled constraints.

While the developed H-MPC framework in [3, 4] allows for optimizing multi-timescale dynamics over different receding horizons, it may not be the best solution for "mission-centric" applications in which a system needs to accomplish a specific mission/task (e.g., reach a destination) while having access to a limited onboard energy source and minimizing energy use. Electrified vehicles, aircraft, and ships are a few examples that fall into this category. For mission-centric applications, the controller is required to not only optimize the performance over a finite receding horizon but also accurately account for the "cost-to-go" beyond the receding horizon until the end of the mission in order to achieve the overall mission objectives and optimal performance.

In particular, terminal constraints, such as the battery charge sustainability conditions for hybrid electric vehicles at the end of the trip, often need to be enforced by the power and thermal management systems. In conventional economic MPC design with a receding horizon, this charge sustainability constraint is often enforced by adding a terminal penalty term to the stage cost of MPC to penalize the battery state-of-charge (*SOC*) deviation from its reference value—which is often set to be a constant value—by the end of the prediction horizon [8, 15]. Given the finite horizon of the MPC, such a terminal penalty term in the stage cost limits the electric battery *SOC* operating range and may force the Hybrid Electric Vehicle (HEV) powertrain to operate in a narrow and less efficient region [25]. The effectiveness of this approach in enforcing the charge sustainability highly depends on the weighting of the terminal cost, which has to be calibrated. To alleviate this issue, one can pre-compute the optimal *SOC* references offline, for example, using Dynamic Programming (DP) [50]), and then design a finite-horizon tracking MPC to follow the optimal SOC trajectory. Such an approach requires prior knowledge about the entire vehicle speed profile, which is not always available and is subject to uncertainties.

In this chapter, a multi-horizon MPC (MH-MPC) control strategy is described which addresses the aforementioned challenges and facilitates the design of MPC for mission-centric operation and for systems that exhibit multi-timescale dynamics. The approach considers the prediction horizon from the current time instant to the end of the mission duration and subdivides it into two phases: (i) a shorter receding horizon phase and (ii) a longer shrinking horizon phase. The preview information over these two horizons can have different levels of accuracy, different fidelity models may be employed, and sampling rates can also be different. Typically, a more accurate preview and faster sampling rate to capture fast dynamics are necessary over the short receding horizon, while a less accurate preview and slower sampling are used over the longer shrinking horizon phase. In effect, the optimization over the shrinking horizon phase provides estimates of the cost-to-go, thereby informing the terminal

cost for the receding horizon phase to account for the overall mission cost. We note an earlier work related to MH-MPC strategy that replaces the end of the trip road segments with a single virtual terminal segment in HEV energy management [28]. There is also a connection of MH-MPC to MPC approaches which exploit adaptive time meshing, see, e.g., [40].

In the forthcoming sections of this chapter, we will describe two practical applications that illustrate the use of the proposed MH-MPC. The first application centers on the integrated power and thermal management in HEVs and follows [24]. Through this application, we aim to elucidate the rationale behind adopting multiple prediction horizons and highlight how estimating the "cost-to-go" over the shrinking horizon phase can enhance optimal performance. The second application pertains to the fast charging of electric vehicles, and the treatment follows [26]. We will exemplify the utilization of MH-MPC with weight adaptation in this setting, even when the precise end of the charging duration is not predetermined.

1.2 Integrated Power and Thermal Management (iPTM) System of Hybrid Electrified Vehicles

The efficient power and thermal management (PTM) of electrified vehicles, including that of the internal combustion engine, exhaust aftertreatment, battery, electric motors, and cabin temperature, can significantly impact the overall fuel economy and driving range, especially in extreme cold and hot ambient conditions [32, 33]. This is because (i) the actuators used for thermal management (e.g., compressor, pumps, and fans) can consume a considerable amount of power by more than 2.5 kW [30], and (ii) the efficiencies of the vehicular power and thermal systems are often temperature-dependent and are degraded when the temperature falls outside the optimum range. For HEVs, with the engine and electric battery as the two main sources of energy, thermal management priority varies as the ambient temperature changes. Under cold ambient conditions, engine thermal management and cabin heating are the main thermal loads. In hot summer weather, on the other hand, engine thermal management, as well as cabin and battery cooling, dominate thermal loads.

Figure 1.1 shows the power and thermal loops of a common power-split HEV. As shown in Fig. 1.1, a part of the generated heat during the combustion process (the rest of the fuel energy is converted to either mechanical work or wasted through the exhaust gases) is absorbed by the engine coolant. If there is any demand for cabin heating, the stored thermal energy in the coolant is then utilized to heat the cabin air at the heater cores. This means cabin heating reduces the coolant temperature, forcing the engine to idle to generate heat once the coolant temperature drops to below a certain threshold (e.g., 40–50 °C [21])—even if there is no demand for the drive. Figure 1.1 also shows the air conditioning (A/C) refrigerant loop, including the A/C compressor, expansion valve, and evaporator. The electric-driven A/C compressor, radiator fan, coolant pump, valves, and A/C blower are the actuators of the engine and

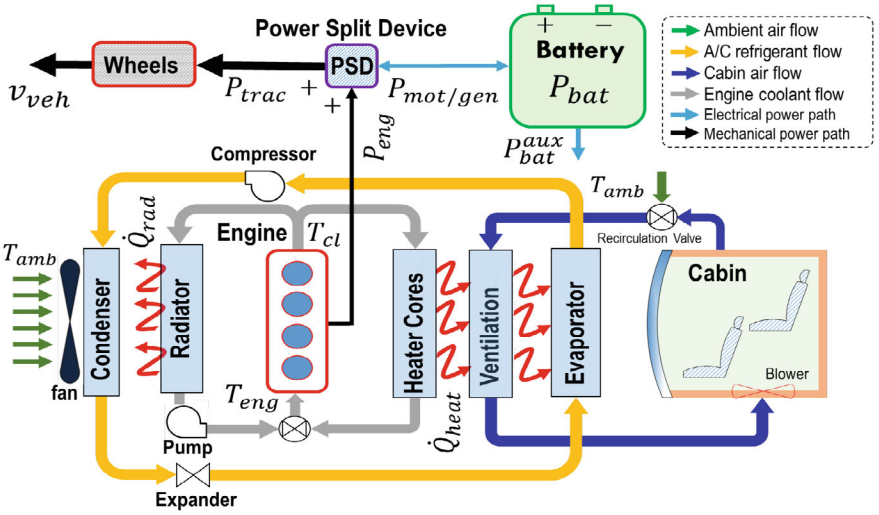

Fig. 1.1 The schematic of a power-split HEV with power and thermal loops. Reprinted with permission from [24]

cabin thermal management system, consuming battery power for their operation. The operation of these actuators is considered as the main auxiliary load on the battery in an HEV.

There have been numerous studies dedicated to the development of energy management strategies (EMS) for HEVs using both (i) offline approaches, e.g., Dynamic Programming (DP) [9, 27, 36], Pontryagin's Minimum Principle (PMP) [31, 43–45], and (ii) online approaches, e.g., equivalent consumption minimization strategy (ECMS) [19, 29, 46] and MPC [7, 8, 20]. These approaches have been extended to PTM with consideration of thermal dynamics and constraints. In [21], multi-dimensional DP was adopted to study the impact of engine efficiency degradation and cabin heating requirements in cold winter on the energy efficiency of an HEV. The results presented in [21] showed that such thermal loads significantly influence the optimal energy flow within the vehicle PTM system. Similarly, in [51], an integrated PTM (iPTM) strategy was developed based on DP to improve the fuel economy during the cold-start phase. Although DP provides a globally optimal solution, it is infeasible for real-time implementation due to the high computational complexity of multi-dimensional optimization.

The emergence of connected and automated vehicles (CAVs) technologies in recent years has provided unique opportunities for situational awareness [48], allowing for vehicle speed forecasting. Exploiting these forecasts brought in by the CAVs technologies for energy flow optimization has made MPC a natural choice for real-time iPTM of HEVs [23]. Two main challenges we face in applying MPC to iPTM to fully leverage the traffic condition and speed trajectory forecasting:

- the coupling between power and thermal dynamics and the associated different timescales,
- uncertainties in vehicle speed forecast can affect the MPC performance unfavorably.

Figure 1.2 illustrates the separation of timescales between the power and thermal system dynamics in an HEV, based on data collected from a test vehicle. As shown, the engine power (Fig. 1.2b) responds relatively fast, typically on a timescale of $1 - 3\ s$. In contrast, the engine coolant temperature (Fig. 1.2c) changes slower, over a timescale of approximately 30 s. Despite these differing response times, the power and thermal dynamics are strongly interdependent. As Fig. 1.2 demonstrates, there is a close correlation between engine power generation and engine coolant temperature; as the engine produces traction power, the engine coolant temperature rises concurrently.

The multi-timescale dynamics and mission-based nature of HEV operation make MH-MPC a suitable approach for iPTM. In this section, we explore the design and performance of MH-MPC for iPTM in HEVs, focusing on a winter scenario with cold ambient temperatures where cabin heating is required. The primary objective of iPTM in this setting is to minimize fuel consumption while meeting power and thermal constraints. However, as previously noted, uncertainties in vehicle speed forecast can adversely affect MPC performance. To assess the impact of these uncertainties and the robustness of the MH-MPC controller, we conduct a comprehensive statistical

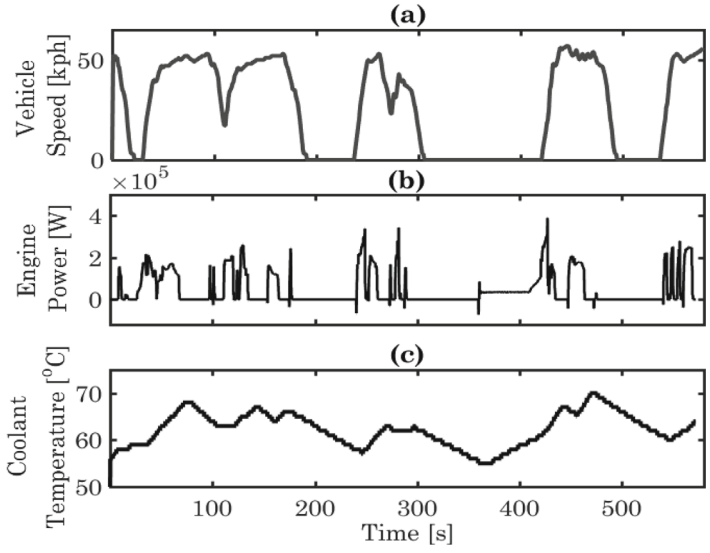

Fig. 1.2 Timescale separation between power and thermal subsystems in an HEV: **a** vehicle speed, **b** engine power, and **c** engine coolant temperature. Data collected from a test HEV. Reprinted with permission from [24]

and sensitivity analysis using data from real-world driving cycles. This analysis helps identify critical traffic events and information whose accurate prediction can significantly enhance the efficiency of iPTM.

1.2.1 Power and Thermal Models of an HEV

We use the iPTM problem in HEVs as a motivating example to highlight the unique controller design requirements for integrated systems with multi-timescale dynamics, as well as the need for a novel MPC approach, namely MH-MPC, for such systems. Figure 1.1 presents the schematics of the power and thermal loops in a power-split HEV. In the thermal loop, the engine coolant temperature (T_{cl}) represents the engine's thermal dynamics. The following subsections introduce the physics-based models of the battery state-of-charge (SOC) and engine coolant temperature (T_{cl}) dynamics, which are the primary dynamics representing the power and thermal subsystems, respectively.

1.2.1.1 Battery Power-Balance Model

For a power-split HEV, the power provided by the battery (P_{bat}^{trac}) and internal combustion engine (P_{eng}) are blended to meet the traction power demand (P_{trac}) for driving:

$$P_{trac} = P_{bat}^{trac} + P_{eng}. \tag{1.1}$$

Additionally, the battery provides the power for auxiliary systems (P_{bat}^{aux}), e.g., the electric coolant pump, and the vehicle climate control system:

$$P_{bat} = P_{bat}^{trac} + P_{bat}^{aux}, \tag{1.2}$$

where P_{bat} is the total battery power. The evolution of the battery SOC is represented by an equivalent circuit model:

$$\dot{SOC}(t) = f_{SOC}(t) = \frac{U_{oc} - \sqrt{U_{oc}^2 - 4R_{int}P_{bat}}}{2R_{int}C_{bat}}, \tag{1.3}$$

where C_{bat}, R_{int}, U_{oc}, and t are the capacity, internal resistance, and open-circuit voltage of the battery and time, respectively.

1.2.1.2 Engine Coolant Temperature Model

At cold ambient temperatures, the main thermal loads are due to the thermal management of the engine and cabin heating. The temperature dynamics of the engine coolant are represented by the following equation [32]:

$$\dot{T}_{cl}(t) = f_{T_{cl}}(t) = \frac{1}{M_{eng}C_{eng}}(\dot{Q}_{fuel} - P_{eng} - \dot{Q}_{exh} - \dot{Q}_{air} - \dot{Q}_{heat}), \tag{1.4}$$

where M_{eng} and C_{eng} are the equivalent thermal mass and capacity of the engine cooling system, respectively. The variables $\dot{Q}_{fuel}$, $\dot{Q}_{exh}$, $\dot{Q}_{air}$ and $\dot{Q}_{heat}$ represent the heat rates released from the combustion process, exchanged through exhaust gases, dissipated by air convection, and delivered for cabin heating, respectively. In particular, $\dot{Q}_{fuel}$ is computed using the fuel consumption rate and lower heating value (LHV) of the fuel (gasoline):

$$\dot{Q}_{fuel} = LHV \cdot \dot{m}_{fuel}, \tag{1.5}$$

where $\dot{m}_{fuel}$ is the fuel consumption rate calculated as a function of engine speed (ω_e), torque (τ_e), and T_{cl}:

$$\dot{m}_{fuel}(\omega_e, \tau_e, T_{cl}) = \alpha(T_{cl}) \cdot f_{fuel}(\omega_e, \tau_e), \tag{1.6}$$

where, $f_{fuel}(\omega_e, \tau_e)$ is the nominal fuel consumption rate and $\alpha(T_{cl})$ is a multiplier reflecting the fuel consumption sensitivity to the coolant temperature. These functions can be found in the library of the Autonomie[1] software for a power-split HEV, also see [21] for more details. When T_{cl} is higher than 60 °C, $\alpha = 1$, as the engine has been fully warmed. When T_{cl} is less than 60 °C, α increases, reflecting the engine efficiency degradation caused by low coolant and ambient temperatures. The experimental validation of the control-oriented models in (1.3) and (1.4) can be found in [2, 21].

1.2.2 MPC-Based iPTM of HEVs

The goal of iPTM is to minimize fuel consumption while enforcing constraints on the battery state-of-charge (SOC) and engine coolant temperature (T_{cl}) and while responding to traction and cabin heating demands. In what follows, a scenario is considered in which the vehicle is operating in winter conditions, with an ambient temperature of -10 °C and a constant cabin heating power demand of $\dot{Q}_{heat} = 1.5$ kW. In this section, we first assess a conventional MPC approach with a short

[1] Autonomie® is a MATLAB®/Simulink®-based system simulation tool for vehicle energy consumption and performance analysis developed by Argonne National Laboratory (ANL) [33].

receding horizon as the baseline. Motivated by the limitations of conventional MPC, we then apply the MH-MPC approach. For all evaluations, the battery *SOC* and engine coolant temperature (T_{cl}) are treated as system states, while battery power (P_{bat}) is optimized as the control input.

1.2.2.1 Conventional MPC

The conventional MPC-based strategy for energy management in HEVs, following in [8, 50], is designed to minimize fuel consumption over a finite prediction horizon while penalizing deviations of the battery *SOC* from a reference value to ensure charge sustainability. To achieve this, a finite-horizon economic MPC law is defined by solving the following discrete-time optimal control problem:

$$\begin{aligned} \underset{P_{bat}(i)}{\arg\min} \; & \sum_{i=t}^{t+H-1} \left\{ \dot{m}_{fuel}(i)\delta t \right\} + \lambda (SOC(t+H) - SOC_r)^2 \\ \text{s.t.} \quad & T_{cl}(i+1) = T_{cl}(i) + \delta t \cdot f_{T_{cl}}(i), \\ & SOC(i+1) = SOC(i) + \delta t \cdot f_{SOC}(i), \\ & 0.4 \leq SOC(i) \leq 0.8, \\ & 40^{\circ}C \leq T_{cl}(i) \leq 90^{\circ}\, C, \\ & T_{cl}(0) = T_{cl,init}, \;\; SOC(0) = SOC_{init}, \end{aligned} \tag{1.7}$$

where H is the prediction horizon, t is the current time, $\delta t = 1\ s$ is the discrete-time step and state constraints are imposed on *SOC* and T_{cl}. The initial conditions are $SOC_{init} = 0.6$ (60%), and $T_{cl,init} = 50\,^{\circ}$C, meaning that the engine is warmed-up at the beginning of the trip. The weight $\lambda = 3$ is multiplying the quadratic penalty on *SOC* deviation from the reference value for the battery state-of-charge (SOC_r). The reference value SOC_r is set to be the same as SOC_{init} to promote charge sustainability. Furthermore, f_{SOC} and $f_{T_{cl}}$ are functions describing the dynamics of *SOC* and T_{cl}, respectively, that were previously introduced in (1.3) and (1.4). The optimal battery power (P_{bat}) at each time instant t is determined by solving (1.7). Then, the required engine power (P_{eng}) is computed to match the traction power request (1.1). Additionally, it is assumed that the engine operates with the minimum brake-specific fuel consumption (BSFC) value for the given engine power, which enables the computation of the engine speed and engine torque. The MPC problem in (1.7) is solved using MPCTools package [42], which exploits the IPOPT solver [5] and CasADi for numerical optimization. The simulations are performed on a desktop computer with an Intel® E-2136@3.30 GHz processor.

The New European Driving Cycle (NEDC) shown in Fig. 1.3a is considered when evaluating the performance of the MPC in (1.7). It includes a combination of city (with multiple stop-and-go) and highway driving. Note that in this section, it is assumed that the accurate prediction of the vehicle speed is available over the

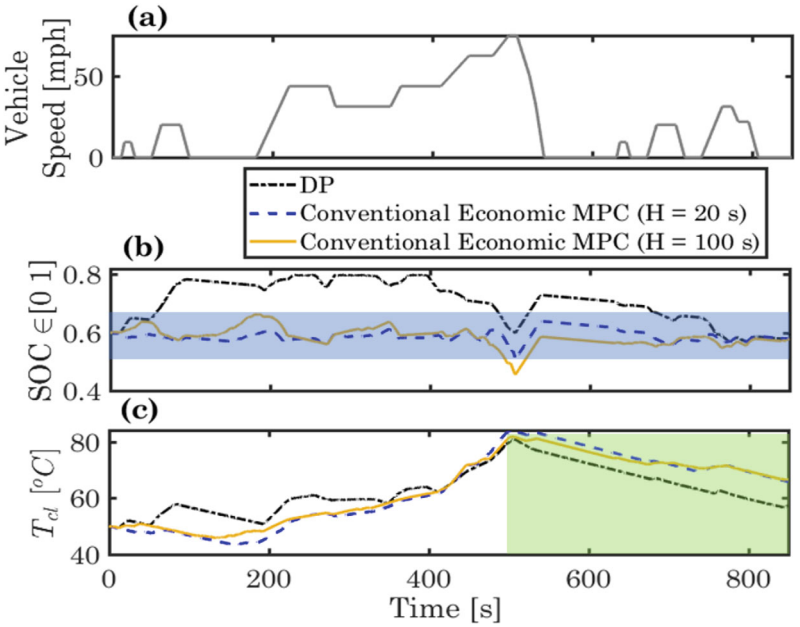

Fig. 1.3 State trajectories (SOC, T_{cl}) of conventional MPC with different prediction horizons and DP. Reprinted with permission from [24]

prediction horizon, without considering uncertainties associated with vehicle speed prediction.

As MPC performance is often influenced by the length of the prediction horizon [38], the MPC strategy defined in (1.7) was simulated with varying prediction horizons, and the resulting fuel consumption is summarized in Fig. 1.4a. The global optimal solution provided by Dynamic Programming (DP) is also shown for comparison in Fig. 1.4. As can be observed, although fuel consumption with the conventional MPC decreases as the prediction horizon extends from 20 s (20 s with $\delta t = 1\ s$) to 100 s, the results still significantly deviate from the global optimal solution by more than 3%. Figures 1.3b, c display the state trajectories with DP and conventional MPC. Notably, the SOC under the MPC solution varies within a limited range of approximately 10%, and extending the prediction horizon has little impact on widening this range. This is because conventional MPC is only aware of future vehicle speed within the prediction horizon and includes a quadratic term in its cost function that penalizes deviations of the SOC from its terminal reference value. In contrast, DP has full knowledge of the entire driving cycle beforehand, allowing it to utilize the battery more efficiently by expanding the operational range of the SOC, as shown in Fig. 1.3b. Furthermore, as shown in Fig. 1.3c, the engine coolant temperature (T_{cl}) response under conventional MPC differs significantly from that of the DP solution. Knowing the occurrence of a long stop (around $t = 100$ to 180 s) before entering the highway, DP increases T_{cl} in advance, thereby preventing

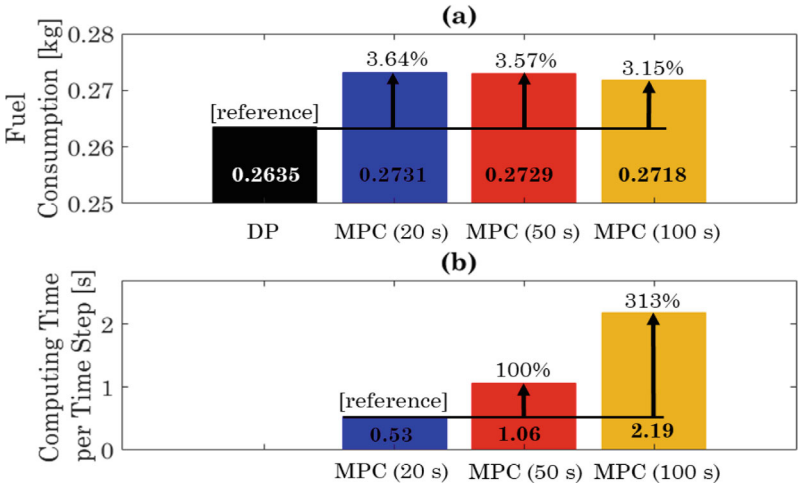

Fig. 1.4 Simulation results of DP and conventional MPCs with different prediction horizons. **a** fuel consumption, **b** average computing time per time step. Reprinted with permission from [24]

the coolant temperature from dropping below its lower limit, which would otherwise trigger engine idling. In contrast, conventional MPC is not able to deliver this optimal response, resulting in the engine operating within an inefficient coolant temperature range (i.e., below $<50\,^{\circ}\mathrm{C}$).

Toward the end of the trip, as highlighted in Fig. 1.3c, the vehicle exits the highway and starts a city driving period with low traction power demand and multiple stop-and-go. Since DP uses accurate information about the trip, it manages to store enough electrical energy in the battery (for traction) and thermal energy in the coolant (for cabin heating) before exiting the highway at around $t = 500$ s, so that the vehicle can reach the end of the trip mainly in the full electric mode. For the MPCs, on the other hand, since not enough electrical energy has been stored in the battery (Fig. 1.3b), the engine is being used more often after $t = 500$ s, resulting in unnecessarily higher coolant temperatures by the end of the trip.

Although the fuel consumption can be reduced with a longer prediction horizon, the MPC computational footprint grows with the length of the prediction horizon as shown in Fig. 1.4b. As the prediction horizon increases from 20 s to 100 s, the computing time increases by 313% on average, which hinders real-time implementation.

Based on the results shown in Figs. 1.4 and 1.3, the shortcomings of the conventional MPC (1.7) can be summarized as follows:

- The quadratic penalty term in the MPC cost function incorporated to promote the battery charge sustainability limits the operating range of SOC, resulting in inefficient use of the battery. This is the main reason for not observing significant improvements in fuel consumption as the MPC prediction horizon is increased.

- While the fuel consumption is reduced by extending the MPC prediction horizon, the computational footprint grows to unaffordable levels.
- Although the vehicle speed over the prediction horizon is assumed to be known a priori, in the real-world traffic environment, any long-term prediction of the vehicle speed is subject to uncertainties. Given that the conventional MPC (1.7) falls short of DP performance in the absence of uncertainties, the energy efficiency may degrade even further with the uncertainties.

To address these challenges, a novel multi-horizon MPC is described in the next subsection to address the trade-off between energy efficiency, and computational footprint.

1.2.2.2 Multi-Horizon MPC (MH-MPC)

The concept of the MH-MPC is shown in Fig. 1.5. The prediction horizon covering the entire trip is divided into two sub-horizons, (i) a short receding horizon (red window) and (ii) a long shrinking horizon (green window). The vehicle speed preview is assumed to be accurate over the short receding horizon.

- **Remark 1**: For connected vehicles, a high-accuracy short-term prediction of the vehicle speed can be obtained, often using vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications. See [39] for an example.

An "approximate" vehicle speed preview is assumed to be available over the long shrinking horizon. This long-term preview does not require to have a detailed second-by-second profile of the vehicle's speed. Instead, it only captures the main traffic events, e.g., acceleration and deceleration at signalized intersections and the cruise speed between the intersections.

- **Remark 2**: For a specific road segment, a long-term vehicle speed preview can be obtained through a big-data analysis of the historic traffic data collected from connected vehicles driving through the same corridor, see [1] for more details.

The MH-MPC is formulated as follows, subject to the same state constraints of conventional MPC (1.7):

$$
\begin{aligned}
&\arg\min_{P_{bat}(i)} \quad \sum_{i=t}^{t+N-1} \left\{\dot{m}_{fuel}(i)\Delta t_1\right\} + \sum_{i=t+N}^{t_{end}} \left\{\dot{m}_{fuel}(i)\Delta t_2\right\},\\
&\text{s.t.} \quad T_{cl}(i+1) = T_{cl}(i) + \Delta t \cdot f_{T_{cl}}(i), \quad j \in \{1,2\},\\
&\qquad SOC(i+1) = SOC(i) + \Delta t_j \cdot f_{SOC}(i), \quad j \in \{1,2\},\\
&\qquad 0.4 \le SOC(i) \le 0.8,\\
&\qquad 40^{\circ}C \le T_{cl}(i) \le 90^{\circ}C,\\
&\qquad 0.99 \times SOC(0) \le SOC(t_{end}) \le SOC(0) \times 1.01,\\
&\qquad T_{cl}(0) = T_{cl,init}, \;\; SOC(0) = SOC_{init},
\end{aligned}
\tag{1.8}
$$

where N is the short receding horizon, t_{end} is the end time of the trip, and Δt_1 and Δt_2 are the sampling times over the receding and shrinking horizons, respectively. The variable $j \in \{1, 2\}$ is determined as follows:

$$j = \begin{cases} 1, & \textbf{if } i \leq t + N - 1, \\ 2, & \textbf{if } i \geq t + N, \end{cases} \tag{1.9}$$

The MH-MPC cost function has two terms:

- the fuel consumption over the short receding horizon calculated based on accurate vehicle speed preview,
- an estimate of the fuel consumption over the long shrinking horizon representing the "cost-to-go" toward the end of the trip beyond the receding horizon.

The cost function of MH-MPC is designed to account for fuel consumption over the entire trip, unlike the conventional MPC (1.7), which only minimizes the fuel consumption over the receding horizon. Additionally, the MH-MPC cost function does not require a quadratic terminal penalty term for the *SOC* deviation; instead, battery charge sustainability is ensured through the terminal constraint. To maintain the feasibility of the optimization problem, the final state-of-charge ($SOC(t_{end})$) is permitted to deviate by up to 1% from the initial SOC_{init}. Furthermore, there is no need for a pre-computed reference trajectory (SOC_r) for MH-MPC, simplifying the implementation.

To reduce the computational footprint of MH-MPC over the long shrinking horizon, both the data and prediction model used in this horizon are sampled at

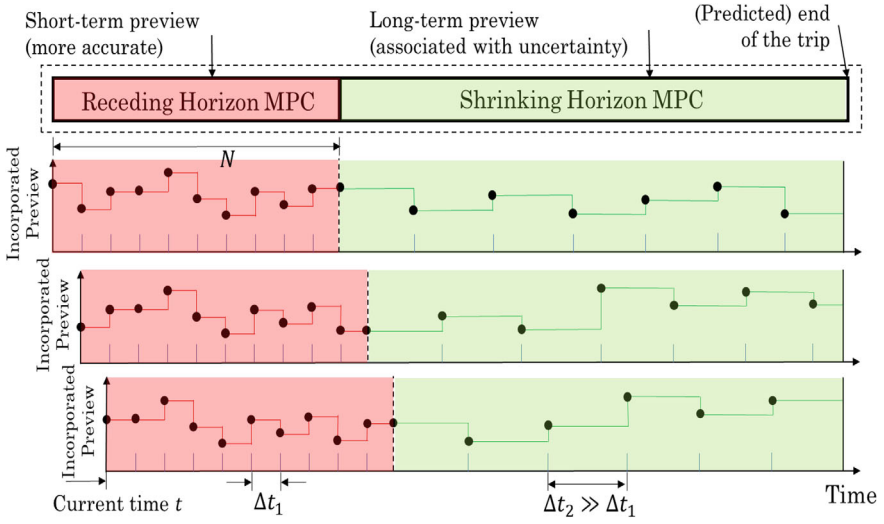

Fig. 1.5 The concept of Multi-horizon MPC with short receding horizon and long shrinking horizon. Reprinted with permission from [24]

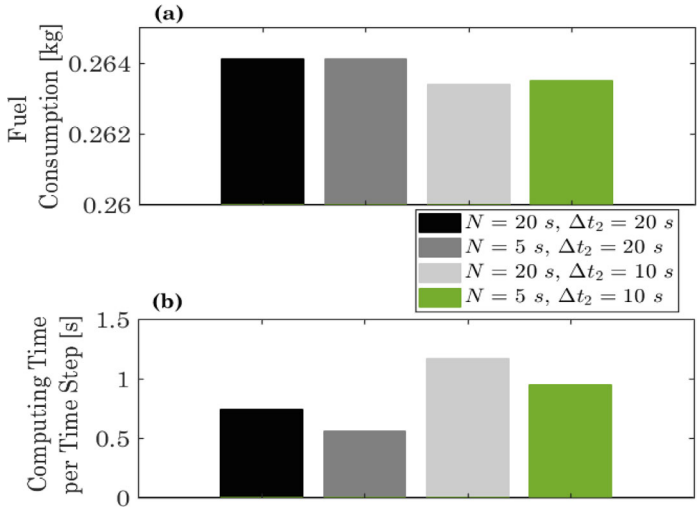

Fig. 1.6 Fuel consumption and average computing time with different short receding horizons and sampling times over the long shrinking horizon. Reprinted with permission from [24]

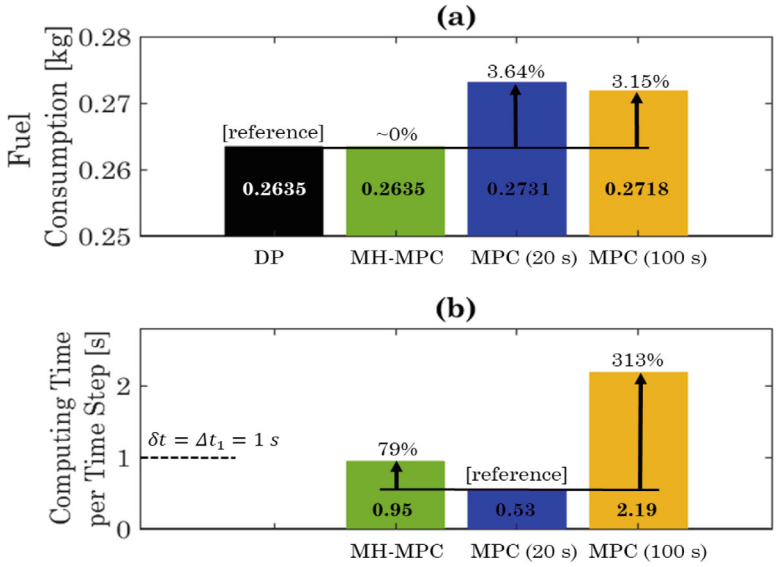

Fig. 1.7 Simulation results of DP, MH-MPC, and conventional MPC with different prediction horizons. **a** fuel consumption, and **b** average computing time per time step. Reprinted with permission from [24]

a slower rate compared to the short receding horizon, see Fig. 1.5. Specifically, a sampling time of Δt_1 is used for the short receding horizon, whereas a much larger sampling time, $\Delta t_2 \gg \Delta t_1$, is employed for the long shrinking horizon. The MH-MPC optimization problem is solved every $\Delta t_1 = 1\ s$, and the first element of the computed control input is applied to the system. Subsequently, the receding horizon shifts forward by Δt_1, and the shrinking horizon is shortened by Δt_1 seconds. Note that when the remaining trip time is shorter than the receding horizon length, the multi-horizon approach is no longer necessary, and the cost function is evaluated solely over the shrinking horizon.

To analyze the sensitivity of the MH-MPC to changes in the receding horizon length (N) and the resolution (i.e., sampling time, Δt_2) of the long shrinking horizon, simulations were conducted using various parameter configurations of the MH-MPC over the same driving cycle shown in Fig. 1.3a. The results for fuel consumption and the average computation time per optimization iteration are presented in Fig. 1.6. Figure 1.6a shows that reducing Δt_2 from 20 s to 10 s slightly lowers fuel consumption while increasing N has a marginal effect on fuel consumption. On the other hand, Fig. 1.6b shows how different horizon lengths and sampling times affect the computations. Balancing fuel consumption impact with computational requirements, the parameters $N = 5$ and $\Delta t_2 = 10\ s$ seconds were selected. Furthermore, the selection of $N = 5$ renders the assumption of accurate speed preview being available over the receding horizon quite reasonable [37]. Note that, for the results presented in this section, it is still assumed that long-term preview is known a priori. However, since this preview is sampled at a slow rate of every Δt_2 seconds over the shrinking horizon, imprecision is introduced due to data down-sampling.

Figure 1.7 presents a comparison of MH-MPC with DP and conventional MPC. As shown in Fig. 1.7a, MH-MPC achieves a fuel consumption reduction of over 3% compared to conventional MPC, with performance closely matching that of DP. In terms of computational effort, Fig. 1.7b indicates that the computational cost of MH-MPC is comparable to that of conventional MPC with a shorter horizon of 20 s and is significantly lower than that of conventional MPC with a longer horizon of 100 s. This computational affordability of MH-MPC is largely due to the slower sampling rate applied over the long shrinking horizon, which reduces the overall computational demand.

The powertrain trajectories with MH-MPC are shown in Fig. 1.8, and are compared with those of DP and conventional MPC (with a horizon length of $H = 20\ s$). Unlike conventional MPC, the SOC trajectory under MH-MPC varies over a wide range exceeding 20%, and closely following the trend observed with DP. This behavior is attributed to MH-MPC's ability to account for the entire driving cycle and incorporate an approximation of the cost-to-go beyond the receding horizon. During the initial phase of the trip, when the engine coolant temperature is relatively low, MH-MPC prioritizes the use of the engine to provide traction power while simultaneously increasing the coolant temperature. Consequently, the battery is also being charged during this period using the extra power available from the engine output. Later, when the vehicle exits the highway around (around $t = 540\ s$), the vehicle predominantly operates in electric mode until the end of the trip ($t = 540$ to 850 s),

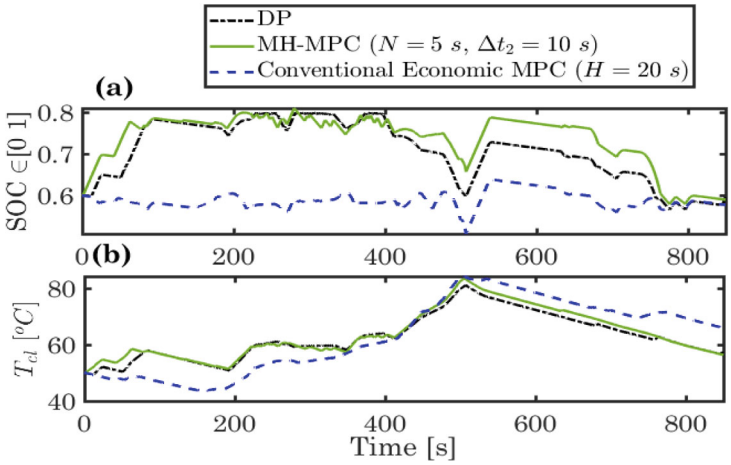

Fig. 1.8 State trajectories (*SOC*, T_{cl}) of DP, MH-MPC, and conventional MPC. Reprinted with permission from [24]

which is similar to DP. To achieve this, MH-MPC was able to strategically store sufficient thermal energy in the coolant, enabling it to meet cabin heating requirements throughout the remainder of the trip. As these results indicate, MH-MPC is able to effectively utilize the engine coolant as a form of thermal energy storage, adding another layer of flexibility to HEV energy management alongside the battery, which serves as electrical energy storage. Thus MH-MPC is able to exploit both electric energy storage and thermal energy storage to optimize energy flows and enhance overall HEV efficiency.

1.3 Electric Vehicle Enhanced Fast Charging Enabled by Battery Thermal Management

In this section, the MH-MPC is applied to the power and thermal management in Electric Vehicles (EVs) during the fast-charging process. We consider a commercial EV operating in hot ambient temperature and with a low initial *SOC*, requiring the battery to be charged at a nearby fast-charging station. The objectives of the iPTM for the EV in this study are threefold. The first objective is to ensure that the final *SOC* after charging is above a prescribed threshold. The second objective is to maintain the total charging time within the prescribed duration. The third objective is to minimize the energy consumption for the battery thermal management system while enforcing power and thermal constraints. In a real-world application, the requirements on the final *SOC* and charging time are typically determined by the type of the trip and mission of the commercial vehicle after the charging event.

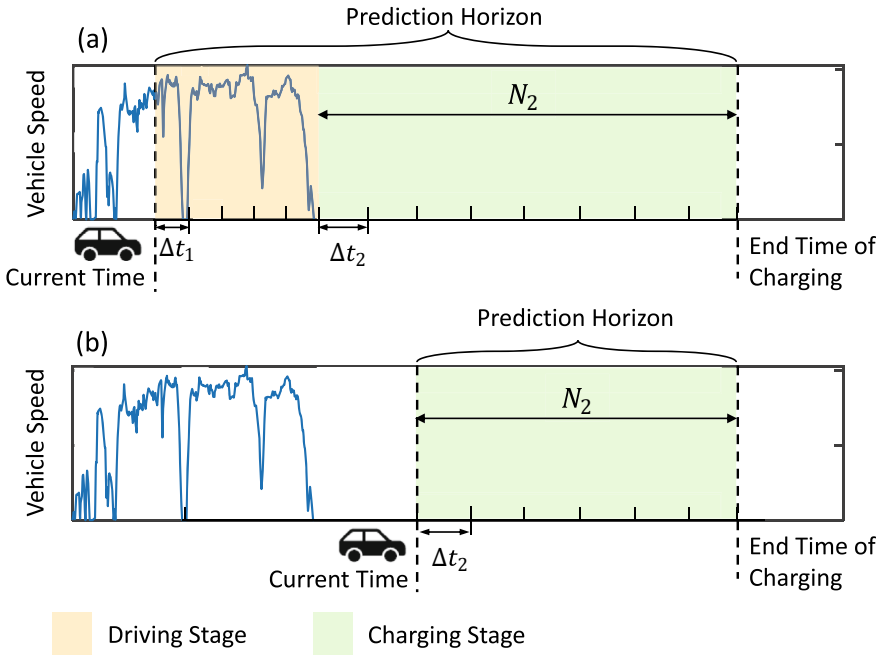

Fig. 1.9 The implementation concept of the MPC-based iPTM strategy with two different scenarios: **a** before the charging starts, when the vehicle moves toward the charging station, and **b** after the charging starts, when the vehicle stays in the charging station. Reprinted with permission from [26]

1.3.1 MH-MPC Formulation for Enhanced Fast Charging

The MH-MPC is now applied to address the aforementioned objectives. Figure 1.9 illustrates that two different scenarios, before and after the start of the charging event, need to be accounted for. Before the start of the charging event, the vehicle moves toward the charging station, and the prediction horizon of the MH-MPC extends from the current time instant until the projected end time of the charging event. This horizon and vehicle operation are then naturally subdivided into two stages, a driving stage and a charging stage. After the charging starts, the vehicle stays at the charging station, and the prediction horizon only has one stage from the current time to the end of the charging event. As in Sect. 1.2, the prediction horizon always extends from the current time instant to the end time of charging. Consistently with MH-MPC approach, the sampling time used can be different in different stages. The MH-MPC formulation is based on solving numerically the following discrete-time optimal control problem:

$$\begin{aligned}
&\min_{\substack{\dot{Q}_{cl}(i),\\ P_{chg}(i),\,\Delta t_2(i)}} \quad \sum_{i=t}^{t+N_1-1}\left(\frac{\dot{Q}_{cl}(i)}{COP(i)}\Delta t_1\right)^2 + \sum_{i=t+N_1}^{t+N_1+N_2-1}\left\{\begin{array}{l}\left(\dfrac{\dot{Q}_{cl}(i)}{COP(i)}\Delta t_2(i)\right)^2\\ +\alpha(\Delta t_2(i))^2+\beta\epsilon^2\end{array}\right\},\\
&\text{s.t.} \quad SOC(i+1) = SOC(i) + f_{SOC}(i)\Delta t_j, \quad j\in\{1,2\}\\
&\qquad T_{bat}(i+1) = T_{bat}(i) + f_{bat}(i)\Delta t_j, \quad j\in\{1,2\}\\
&\qquad SOC(t+N_1+N_2) = SOC_{\text{targ}},\\
&\qquad SOC_{\min} \le SOC(i) \le SOC_{\max},\\
&\qquad T_{bat,\min} \le T_{bat}(i) \le T_{bat,\max}+\epsilon,\\
&\qquad -\dot{Q}_{cl,\max} \le \dot{Q}_{cl}(i) \le 0,\\
&\qquad 0 \le P_{chg}(i) \le P_{chg,\max},\\
&\qquad 0 \le \Delta t_2(i) \le \Delta t_{2,\max},
\end{aligned} \tag{1.10}$$

where f_{SOC} is the same as previously introduced in (1.3), while f_{bat} describes the battery temperature dynamics, which are expressed using the following equation:

$$\dot{T}_{bat} = f_{bat}(t) = \frac{1}{m_{bat,thm}C_{bat,thm}}(\dot{Q}_{gen} - \dot{Q}_{amb} - \dot{Q}_{cl}), \tag{1.11}$$

where $m_{bat,thm}$, $C_{bat,thm}$ are the battery thermal mass and specific heat capacity, respectively. The $\dot{Q}_{gen}$ is the irreversible battery heat generation rate attributed to the internal resistance and is expressed as $\dot{Q}_{gen} = I_{bat}^2 R_{int}$. The $\dot{Q}_{amb}$ is the rate of the heat dissipation to ambient through air convection, which is driven by the temperature difference between ambient and battery.

Moreover, in (1.10), Δt_1 and Δt_2 are the sampling periods over the driving and charging stages, respectively, and N_1 and N_2 correspond to the numbers of sampling points over these two stages. Therefore, the prediction horizon length is $\Delta t_1 N_1 + \Delta t_2 N_2$ where $\Delta t_1 N_1$ is the remaining time for the vehicle to arrive at the charging station while $\Delta t_2 N_2$ is the total predicted time spent at the charging station. Note that $N_1 = 0$ once the driving stage is completed and the charging stage has started. The index $j \in \{1, 2\}$ is determined as follows:

$$j = \begin{cases} 1, & \textbf{if } i \le t + N_1 - 1,\\ 2, & \textbf{if } i \ge t + N_1. \end{cases} \tag{1.12}$$

It can be seen from (1.10) that the control/decision variables are $\dot{Q}_{cl}$, P_{chg}, and Δt_2, and the cost function consists of four terms. The first and second terms are the accumulated energy consumption of the battery cooling system over the driving and charging stages. The third term is the square of Δt_2, which penalizes the total charging time. The last term is used to relax the constraint on the battery temperature by introducing a slack variable (ϵ) and treating the battery temperature limit as a soft

constraint. The use of this last term avoids infeasibility in solving the discrete-time optimization problem.

As the prediction horizon always extends from the current time instant to the end of the charging event, the horizon length of the driving stage shrinks over time while the horizon length of the charging stage is not predetermined and it depends on the computed solution of the problem (1.10) with Δt_2 being an optimization variable. This distinguishes the problem formulation in (1.10) from the conventional receding or shrinking horizon MPC. To this end, a novel sampling strategy with varying sampling time and sampling points over different optimization horizons is further elaborated next.

Specifically, during the driving stage, the sampling time Δt_1 is fixed, and N_1 is calculated based on the remaining time before the vehicle arrives at the charging station. However, at the charging stage, the sampling time Δt_2 is one of the adjustable variables to be determined by the solution of the optimization problem (1.10). Therefore, it is no longer fixed. In order to solve the optimization problem (1.10) numerically with a fixed dimension, we fix the number of samples, N_2. Then, by applying the first element of the computed control sequence to the vehicle and repeating the optimization with updated initial conditions, a feedback law is formed as in the conventional receding horizon MPC scheme.

1.3.2 Simulation Results and Discussion

In this section, we report simulation results of applying MH-MPC to a commercial EV model. The ambient temperature is set to be 38 °C, and the initial battery *SOC* is 0.3 (30%). In the simulated scenario, the vehicle follows through an urban route to arrive at the fast-charging station, where the battery needs to be charged to $SOC = 0.6$. The maximum battery cooling power and charging power of the station are 5 kW and 80 kW, respectively. Moreover, Δt_1 and N_2 are set to 10 s and 40, respectively.

In this study, the battery temperature must be maintained within an operating range of 15–35 °C, during both the driving and charging stages. To discourage violations of the soft temperature constraint, we set the slack variable to $\beta = 10^8$. This value has been chosen by trial and error so that it is sufficiently large to avoid constraint violation in simulation scenarios.

1.3.2.1 Trade-Off Between BTM Energy Use and Charging Time

We first assume that the vehicle speed profile over the urban route is known a priori, based on which the arrival time at the charging station can also be accurately predicted. To investigate the impact of the weight (α) on the charging time, a sensitivity analysis has been conducted, and the results are summarized in Fig. 1.10.

A trade-off can be observed from Fig. 1.10 between the energy consumed for the battery cooling and the battery charging time. As α increases, the battery charging

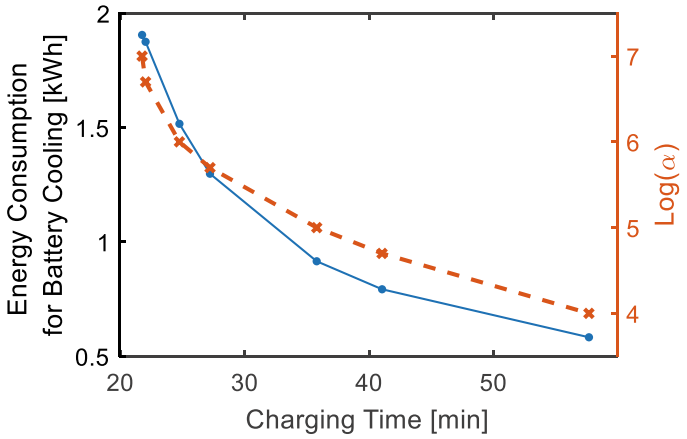

Fig. 1.10 The energy consumption for battery cooling and battery charging time results with different α. Reprinted with permission from [26]

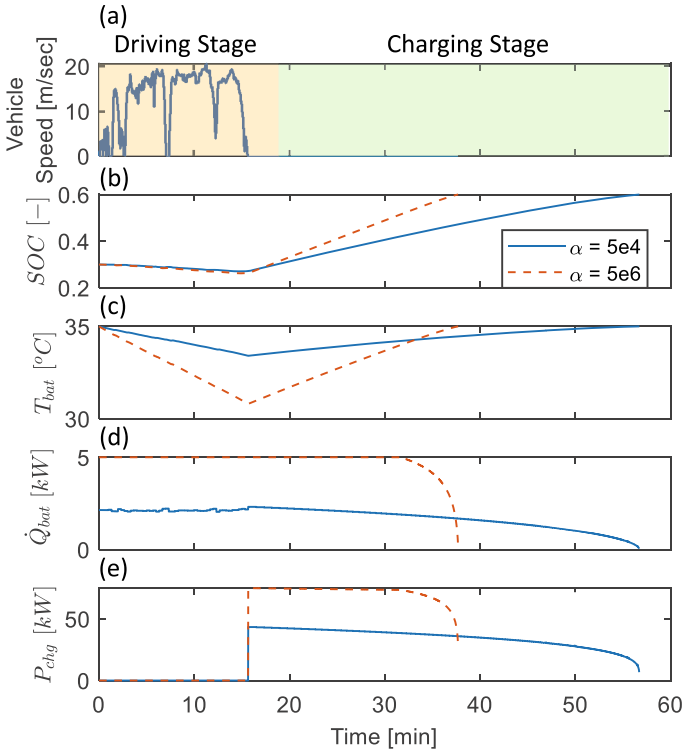

Fig. 1.11 State and input trajectories with different α values: **a** vehicle speed, **b** *SOC*, **c** battery temperature, **d** battery cooling power, and **d** battery charging power. Reprinted with permission from [26]

time decreases while the energy consumed for the battery cooling increases. To account for the impact of different weights, the state and input trajectories for two different values of α are presented in Fig. 1.11. It can be seen that in both cases, the battery temperature is maintained within the prescribed range at all times. By comparing the two cases, it can be seen that, with a larger α, the controller tends to draw a larger battery charging power during the charging stage to reduce the charging time, as more focus is put on the penalty term for the charging time in the cost function. At the same time, a larger battery charging power requires a larger battery cooling power to avoid raising the temperature above the prescribed limit. Moreover, with accurate knowledge of the arrival time and charging event timing, battery pre-cooling is effected to create room for the rise of battery temperature, thereby enabling faster charging with larger charging power. The above observations explain the trade-off between energy consumed for battery cooling and total charging time, as well as the emergent pre-cooling behavior associated with larger α.

1.3.2.2 A Strategy for Weight Adaptation

The results shown in Sect. 1.3.2.1 demonstrate that in order to balance the charging time and BTM energy consumption, α needs to be properly tuned to enforce the constraint on charging time while minimizing the energy consumption. There are several challenges in scheduling α in real-time. Note that α is determined by the arrival time at the charging station and by the required charging time. It needs to be adjusted during the trip to deal with uncertainties in vehicle speed preview, availability of charging, etc. In this section, an adaptive strategy to schedule α in real-time is described, which can be summarized as follows:

- **Step 1**: An initial value of α is selected from the range shown in Fig. 1.10, at $t = 0\ s$.
- **Step 2**: The optimization problem (1.10) is solved using the current α_t at time t, and an estimated charging time ($t_{chg,est}$) is calculated as $\Delta t_2 N_2$.
- **Step 3**: α is updated by the following adaptive law:

$$log(\alpha_{t+\Delta t_j}) = log(\alpha_t) + \lambda(t_{chg,est} - t_{chg,req}). \tag{1.13}$$

- **Step 4**: Repeat Steps 2 and 3 till the end of the charging stage.

In (1.13), λ is the adaptation rate, and Δt_j is the sampling time defined in (1.10). It can been seen that the adaptive law in Step 3 leverages the relationship presented in Fig. 1.10, and α is updated by comparing the estimated charging time ($t_{chg,est}$) with the required charging time ($t_{chg,req}$). The use of logarithm in the adaptation law (1.13) is motivated by the shape of the dependence in Fig. 1.10, and it enables larger steps in the parameter value and faster adaptation as compared to conventional linear adaptation.

To demonstrate the effectiveness of the adaptive strategy for updating α, the required charging time is set to be within 30 *min*, and different λ values are applied

in (1.13). The simulation results are presented in Fig. 1.12, and the actual charging times of three cases are summarized in Table. 1.1. It can be seen that for these three cases, the initial guess for α is too small, which leads to an estimated battery charging time of almost 50 *min*. With the α updated by the adaptive law, the estimated charging time converges to the value near $t_{chg,req}$. Moreover, as λ increases, the convergence rate increases, and the steady-state error in the final charging time decreases. Based on these results, $\lambda = 0.1$ has been chosen.

It can be seen from Figs. 1.11 and 1.12 that MH-MPC incorporates the knowledge of the upcoming charging event to pre-cool the battery, allowing larger battery charging power to reduce the charging time. To quantify the benefits of leveraging

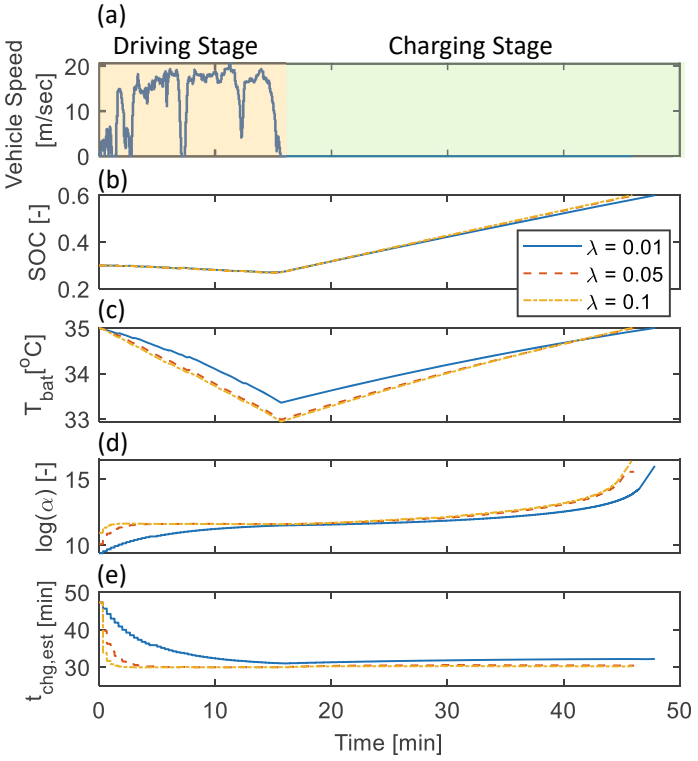

Fig. 1.12 State and input trajectories with different adaptation rates λ: **a** vehicle speed, **b** *SOC*, **c** battery temperature, **d** weight on charging time, and **e** estimated charging time. Reprinted with permission from [26]

Table 1.1 The actual charging time with different adaptation rates

λ	0.01	0.05	0.1
t_{chg} [*min*]	32.2	30.4	30.2
BTM energy consumption [kWh]	1.00	1.06	1.07

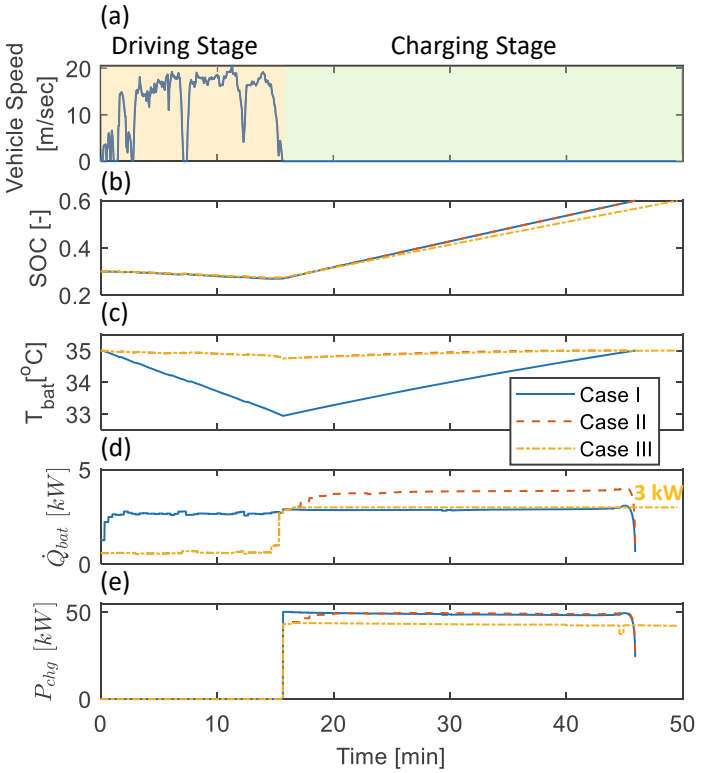

Fig. 1.13 State and input trajectories with different α for Cases I, II, and III: **a** vehicle speed, **b** *SOC*, **c** battery temperature, **d** battery cooling power, and **d** battery charging power. Reprinted with permission from [26]

the preview information and characterize the conditions under which the pre-cooling is beneficial, the following three cases are considered for comparison:

- **Case I**: The charging event is predicted accurately over the prediction horizon.
- **Case II**: The charging event is not predicted until the vehicle arrives at the charging station, and the maximum cooling power is 5 kW.
- **Case III**: The charging event is not predicted until the vehicle arrives at the charging station, and the maximum cooling power is 3 kW.

For Case I, it is assumed that the preview information is known a priori, and the MPC-based iPTM is applied at the driving and charging stages. For Cases II and Case III, before the vehicle arrives at the station, as the charging event is not predicted over the prediction horizon, the cost function only has the first term in (1.10) to minimize the BTM energy consumption during the driving stage. After the vehicle arrives at the station and starts charging, the same optimization problem in (1.10) is solved by MH-MPC. The only difference between Cases II and III is the maximum battery

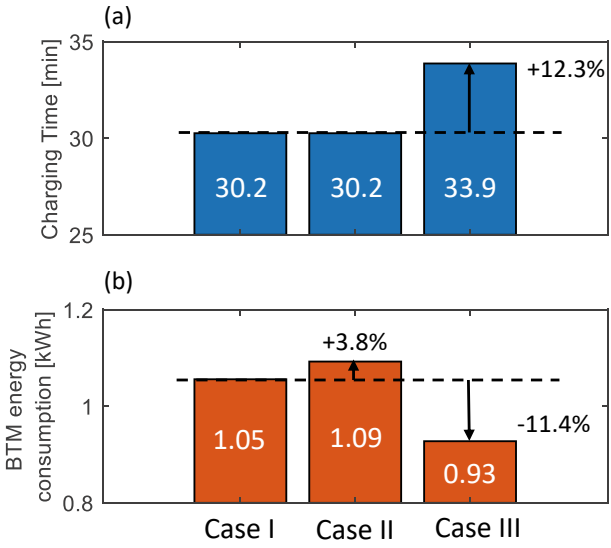

Fig. 1.14 Simulations results of Cases I, II, and III: **a** charging time, **b** BTM energy consumption. Reprinted with permission from [26]

cooling power. Figure 1.13 presents the state and input trajectories of these three cases and the charging time, while BTM energy consumption results are compared in Fig. 1.14.

It can be seen from Fig. 1.13 that pre-cooling is only performed in Case I, while the battery temperature in Cases II and III follows the upper bound constraint before arriving at the charging station to minimize the cooling loads. While the initial battery temperatures at the beginning of the charging stage are different, charging times in Cases I and II are similar. This is attributed to the battery charging power in these two cases being almost the same, as shown in Fig. 1.13e. Compared to Case I, which has a lower initial battery temperature at the start of the charging stage, Case II needs to increase the battery cooling power to avoid constraint violation when applying the same charging power. It can be seen from Fig. 1.13d that the battery cooling power for Case I is always below 3 kW, while the cooling power for Case II is above 4 kW during the charging stage, leading to a slight increase in the BTM energy consumption, as shown in Fig. 1.14b.

As a result, when the maximum cooling power is reduced to 3 kW in Case III, the same charging power is unattainable without pre-cooling, and the charging time is increased by 12.3%. Note that because the upper bound constraint of cooling power in Case I is inactive, the optimal performance would remain the same when the maximum cooling power is reduced from 5 kW to 3 kW.

1.4 Summary

This chapter summarized a multi-horizon MPC (MH-MPC) strategy, which was developed for control of integrated systems with dynamics responding over multiple timescales. The MH-MPC exploits multi-range prediction and performs optimization of the cost function defined over a shorter receding horizon phase followed by a longer shrinking horizon phase. The discretization time step, models, and preview information accuracy can be different in the two phases. In effect, the MH-MPC estimates the "cost-to-go" over the longer shrinking horizon beyond the conventional receding horizon phase. This makes it an appealing choice for mission-centric applications in which the objective is to accomplish a mission with a limited onboard energy resource while satisfying state and control constraints. For such systems, the MH-MPC relaxes the requirement for the inclusion of a terminal penalty term in the optimization stage cost, allowing energy-related states to operate on or close to their admissible boundary, thereby improving the performance.

In this chapter, we first described the use of MH-MPC for integrated power and thermal management (iPTM) of HEVs operating in a connected traffic environment. In such a scenario, short- and long-term predictions of the vehicle's speed can be obtained using connectivity-enabled data and technologies (e.g., V2V/V2I) and incorporated over the receding and shrinking horizons of the MH-MPC, respectively. The simulation results of applying the MH-MPC to experimentally validated models of a power-split HEV illustrated performance improvements with MH-MPC as compared to more traditional MPC, which uses terminal penalties to promote charge sustainability. Notably, in the absence of uncertainties in the vehicle speed forecast, the MH-MPC provided results close to that of Dynamic Programming (with a deviation of 1%), at a much lower and more affordable computational cost.

Furthermore, we showed how MH-MPC can be used to minimize charging time and energy consumption for battery thermal management of a commercial electric vehicle (EV). In this case, MH-MPC is applied to achieve a target battery state-of-charge (SOC) within the prescribed time while enforcing the power and thermal constraints for the battery system. The simulation results showed that, by leveraging the preview information, MH-MPC is able to reduce the charging time; it achieves this by pre-cooling the battery before the start of the charging event. Moreover, an adaptive strategy was described for adjusting the weight of the charging time in the MPC stage cost to manage the trade-off between charging time and battery thermal management (BTM) energy consumption.

In a broader sense, this chapter illustrated how MPC techniques can be creatively extended and applied in challenging real-world engineering applications. The use of multiple horizons, stages, and adaptation of the parameters in the cost function can be exploited to improve economic performance and safety while lowering computational costs.

References

1. Amini, M., Feng, Y., Yang, Z., Kolmanovsky, I., Sun, J.: Long-term vehicle speed prediction via historical traffic data analysis for improved energy efficiency of connected electric vehicles. In: TRR, vol. 2674, pp. 17–29 (2020)
2. Amini, M., Gong, X., Wang, H., Feng, Y., Kolmanovsky, I., Sun, J.: Sequential optimization of speed, thermal load, and power split in connected HEVs. In: ACC. Philadelphia, PA, USA (2019)
3. Amini, M., Kolmanovsky, I., Sun, J.: Cabin and battery thermal management of connected and automated HEVs for improved energy efficiency using hierarchical model predictive control. IEEE Trans. Control Syst. Technol. **28**(5), 1711–1726 (2020)
4. Amini, M., Kolmanovsky, I., Sun, J.: Hierarchical MPC for robust eco-cooling of connected and automated vehicles and its application to electric vehicle battery thermal management. IEEE Trans. Control Syst. Technol. **29**(1), 316–328 (2021)
5. Andersson, J., Gillis, J., Horn, G., Rawlings, J., Diehl, M.: CasADi: a software framework for nonlinear optimization and optimal control. Math. Program. Comput. **11**(1), 1–36 (2019)
6. Baldea, M., Daoutidis, P., Nagy, Z.: Nonlinear model predictive control of integrated process systems. IFAC Proc. Vol. **43**(14), 1040–1045 (2010)
7. Bichi, M., Ripaccioli, G., Di Cairano, S., Bernardini, D., Bemporad, A., Kolmanovsky, I.: Stochastic model predictive control with driver behavior learning for improved powertrain control. In: CDC, Atlanta, GA, USA (2010)
8. Borhan, H., Vahidi, A., Phillips, A., Kuang, M., Kolmanovsky, I., Di Cairano, S.: MPC-based energy management of a power-split hybrid electric vehicle. IEEE Trans. Control Syst. Technol. **20**(3), 593–603 (2011)
9. Brahma, A., Guezennec, Y., Rizzoni, G.: Optimal energy management in series hybrid electric vehicles. In: ACC, Chicago, IL, USA (2000)
10. Brdys, M., Grochowski, M., Gminski, T., Konarczak, K., Drewa, M.: Hierarchical predictive control of integrated wastewater treatment systems. Control Eng. Pract. **16**(6), 751–767 (2008)
11. Bürger, R., Damasceno, J., Karlsen, K.: A mathematical model for batch and continuous thickening of flocculated suspensions in vessels with varying cross-section. Int. J. Miner. Process. **73**(2–4), 183–208 (2004)
12. Chen, X., Heidarinejad, M., Liu, J., Christofides, P.: Composite fast-slow MPC design for nonlinear singularly perturbed systems. AIChE J. **58**(6), 1802–1811 (2012)
13. Chen, X., Heidarinejad, M., Liu, J., de la Peña, D., Christofides, P.: Model predictive control of nonlinear singularly perturbed systems: application to a large-scale process network. J. Process Control **21**(9), 1296–1305 (2011)
14. Clarke, W., Manzie, C., Brear, M.: Hierarchical economic MPC for systems with storage states. Automatica **94**, 487–507 (2018)
15. Debert, M., Yhamaillard, G., Ketfi-herifellicaud, G.: Predictive energy management for hybrid electric vehicles-prediction horizon and battery capacity sensitivity. IFAC Proc. Vol. **43**(7), 270–275 (2010)
16. Dunham, W., Hencey, B., Girard, A., Kolmanovsky, I.: Distributed model predictive control for more electric aircraft subsystems operating at multiple time scales. IEEE Trans. Control Syst. Technol. **28**(6), 2177–2190 (2019)
17. Ellis, M., Heidarinejad, M., Christofides, P.: Economic model predictive control of nonlinear singularly perturbed systems. J. Process Control **23**(5), 743–754 (2013)
18. Farina, M., Zhang, X., Scattolini, R.: A hierarchical multi-rate MPC scheme for interconnected systems. Automatica **90**, 38–46 (2018)
19. Gong, Q., Li, Y., Peng, Z.: Trip-based optimal power management of plug-in hybrid electric vehicles. IEEE Trans. Veh. Technol. **57**(6), 3393–3401 (2008)
20. Gong, Q., Li, Y., Peng, Z.: Stochastic MPC with learning for driver-predictive vehicle control and its application to HEV energy management. IEEE Trans. Control Syst. Technol. **22**(3), 1018–1031 (2013)

21. Gong, X., Wang, H., Amini, M., Kolmanovsky, I., Sun, J.: Integrated optimization of power split, engine thermal management, and cabin heating for hybrid electric vehicles. In: CCTA, Hong Kong, China (2019)
22. Grune, L., Pannek, J.: Nonlinear Model Predictive Control. Springer (2017)
23. Guanetti, J., Kim, Y., Borrelli, F.: Control of connected and automated vehicles: state of the art and future challenges. Annu. Rev. Control **45**, 18–40 (2018)
24. Hu, Q., Amini, M., Kolmanovsky, I., Sun, J., Wiese, A., Buckland-Seeds, J.: Multihorizon model predictive control: an application to integrated power and thermal management of connected hybrid electric vehicles. IEEE Trans. Control Syst. Technol. **30**(3), 1052–1064 (2022)
25. Hu, Q., Amini, M., Wang, H., Kolmanovsky, I., Sun, J.: Integrated power and thermal management of connected HEVs via multi-horizon MPC. In: ACC, Denver, CO, USA (2020)
26. Hu, Q., Amini, R., Wiese, A., Buckland-Seeds, J., Kolmanovsky, I., Sun, J.: Electric vehicle enhanced fast charging enabled by battery thermal management and Model Predictive Control. In: 22nd IFAC World Congress, Yokohama, Japan (2023)
27. Johannesson, L., Asbogard, M., Egardt, B.: Assessing the potential of predictive control for hybrid vehicle powertrains using stochastic dynamic programming. IEEE Trans. Intell. Transp. Syst. **8**(1), 71–83 (2007)
28. Katsargyri, G., Kolmanovsky, I., Michelini, J., Kuang, M., Phillips, A., Rinehart, M., Dahleh, M.: Path dependent receding horizon control policies for hybrid electric vehicles. In: IEEE CCTA. St. Petersburg Russia (2009)
29. Kessels, J., Koot, M., Van Den Bosch, P., Kok, D.: Online energy management for hybrid electric vehicles. IEEE Trans. Veh. Technol. **57**(6), 3428–3440 (2008)
30. Khoury, G., Clodic, D.: Method of test and measurements of fuel consumption due to air conditioning operation on the new Prius II hybrid vehicle. SAE Trans. **114**(16), 2563–2571 (2005)
31. Kim, N., Cha, S., Peng, H.: Optimal control of hybrid electric vehicles based on Pontryagin's minimum principle. IEEE Trans. Control Syst. Technol. **19**(5), 1279–1287 (2010)
32. Kim, N., Rousseau, A.: Thermal impact on the control and the efficiency of the 2010 Toyota Prius hybrid electric vehicle. Proc. Inst. Mech. Eng. Part D: J. Automob. Eng. **230**(1), 82–92 (2016)
33. Kim, N., Rousseau, A., Lee, D., Lohse-Busch, H.: Thermal model development and validation for 2010 Toyota Prius. SAE Technical Paper (2014-01-1784) (2014)
34. Koeln, J., Pangborn, H., Williams, M., Kawamura, M., Alleyne, A.: Hierarchical control of aircraft electro-thermal systems. IEEE Trans. Control Syst. Technol. **28**(4), 1218–1232 (2020)
35. Kokotović, P., Khalil, H., O'reilly, J.: Singular Perturbation Methods in Control: Analysis and Design. SIAM (1999)
36. Liu, J., Peng, H.: Modeling and control of a power-split hybrid vehicle. IEEE Trans. Control Syst. Technol. **16**(6), 1242–1251 (2008)
37. Liu, K., Asher, Z., Gong, X., Huang, M., Kolmanovsky, I.: Vehicle velocity prediction and energy management strategy part 1: deterministic and stochastic vehicle velocity prediction using machine learning. In: SAE Technical Paper 2019-01-1051 (2019)
38. Morari, M., Lee, J.: Model predictive control: past, present and future. Comput. Chem. Eng. **23**(4–6), 667–682 (1999)
39. Moser, D., Waschl, H., Schmied, R., Efendic, H., Del Re, L.: Short term prediction of a vehicle's velocity trajectory using ITS. SAE Int. J. Passeng. Cars: Electron. Electr. Syst. **8**, 364–370 (2015)
40. Paiva, L., Fontes, F.A.: Sampled-data using adaptive time-mesh refinement algorithms. In: 12th Portuguese Conference on Automatic Control (2016)
41. Rawlings, J., Mayne, D., Diehl, M.: Model Predictive Control: Theory, Computation, and Design, 2nd edn. Nob Hill Publishing, Madison, Wisconsin (2017)
42. Risbeck, M., Rawlings, J.: MPCTools: Nonlinear Model Predictive Control tools for CasADi (2019). https://bitbucket.org/rawlings-group/octave-mpctools

43. Rousseau, G., Sinoquet, D., Rouchon, P.: Constrained optimization of energy management for a mild-hybrid vehicle. Oil Gas Sci. Technol.-Rev. IFP **62**(4), 623–634 (2007)
44. Sciarretta, A., Back, M., Guzzella, L.: Optimal control of parallel hybrid electric vehicles. IEEE Trans. Control Syst. Technol. **12**(3), 352–363 (2004)
45. Sciarretta, A., Guzzella, L.: Control of hybrid electric vehicles. IEEE Control Syst. Mag. **27**(2), 60–70 (2007)
46. Serrao, L., Onori, S., Rizzoni, G.: ECMS as a realization of Pontryagin's minimum principle for HEV control. In: ACC. St. Louis, MO, USA (2009)
47. Skogestad, S.: Plantwide control: the search for the self-optimizing control structure. J. Process Control **10**(5), 487–507 (2000)
48. Vahidi, A., Sciarretta, A.: Energy saving potentials of connected and automated vehicles. Transp. Res. Part C: Emerg. Technol. **95**, 822–843 (2018)
49. Van Henten, E., Bontsema, J.: Time-scale decomposition of an optimal control problem in greenhouse climate management. Control Eng. Pract. **17**(1), 88–96 (2009)
50. Wang, H., Huang, Y., Khajepour, A., Song, Q.: Model predictive control-based energy management strategy for a series hybrid electric tracked vehicle. Appl. Energy **182**, 105–114 (2016)
51. Wei, C., Hofman, T., Caarls, E., van Iperen, R.: Optimal control of an integrated energy and thermal management system for electrified powertrains. In: ACC, Philadelphia, PA, USA (2019)