Nonlinear MPC for supervisory control of hybrid electric vehicles

Johannes Buerger [a] and Mark Cannon [b] [*]

Abstract— We propose a hierarchical Model Predictive Control (MPC) strategy for energy management in plugin hybrid electric vehicles. An inner feedback loop addresses the problem of optimally tracking a given reference trajectory for the battery state of charge over a short future horizon using knowledge of the predicted driving cycle. The associated receding horizon optimization problem is solved using a projected Newton method. The controller is compared with existing approaches based on Pontryagin's Minimum Principle and the effects of imprecise knowledge of the future driving cycle are discussed. An outer feedback loop generates the state of charge reference trajectory by solving approximately the optimal control problem for the entire driving cycle. By considering averages of the driver demand over longer time intervals the required number of prediction steps is reduced such that the outer loop problem can also be efficiently solved using the proposed Newton method. Advantages over approaches that assume a linearly decreasing state of charge reference trajectory are discussed.

I. Introduction

The automotive industry has been steadily moving towards a stronger electrification of the powertrain over the recent years [1], with a particular focus currently on Plugin Hybrid Electric Vehicles (PHEVs), in which the battery is typically charged prior to the trip and steadily emptied towards the end of the trip. While the efficiency of PHEVs relative to conventional vehicles relies on many factors, it is in particular dependent on how the engine load point is shifted by utilizing the electric machine in order to minimize a strategic objective over a driving cycle (such as fuel and/or emissions) and reaching a desired energy level at the end of the driving cycle. This problem has been tackled by a variety of methods [2], but in particular methods based on optimal control formulations have been successful. These can be grouped into three classes: methods based on dynamic programing (DP) [3], [4], methods based on the Pontryagin Minimum Principle (PMP) [5], [6] and those based on Model Predictive Control (MPC) [7].

Methods based on DP result in the global optimum for a given driving cycle, but are computationally expensive, since the complexity typically scales exponentially with relevant problem dimensions (in particular with horizon length). Either the price to be paid is the burden of a high computational time for solving an online DP or the burden of high memory storage requirements for offline map-based DP approaches. Methods based on PMP, in particular the Equivalent Consumption Minimization Strategy (ECMS) [5], are suitable candidates for production vehicles since the required optimization is usually only performed at the current time

instant for a constant co-state. Unfortunately, the optimal co-state is highly sensitive to the driving cycle and therefore the method requires careful initialization and adaptation of the co-state in order to ensure that the target state can be reached [8]. In contrast, the advantage of MPC approaches is that the required tuning for ECMS methods is avoided and that information about the future driving cycle obtained from the navigation unit can be readily incorporated.

The contributions of this paper are as follows. First, a nonlinear MPC (NMPC) scheme is proposed using a specialised projected Newton method. This minimizes at each sampling instant a quadratic approximation of predicted fuel consumption, which is derived from the predicted driving cycle, subject to an end-point constraint on state of charge and pointwise-in-time bounds on electrical power. The approach provides a more general and more accurate treatment than [7], where the NMPC problem is approximated directly as a (linearly constrained) quadratic programming problem. Secondly, the NMPC method is applied to the charge sustaining and charge depleting optimal control problem for a short horizon, in which the short-horizon target state is determined from a linearly decreasing reference trajectory for the state of charge, as proposed in the context of PMP in [9] (known as "blended mode").

The third contribution is a two-level MPC framework to overcome the limitations of a linearly decreasing reference state, since in the presence of changing predicting driving conditions, such as altitude variations, a linear profile can be significantly suboptimal [10]. We propose to solve a low-sampling rate MPC problem for the entire driving cycle to obtain the long-term reference state trajectory (high level MPC) and feed this to a short-horizon MPC controller (low level MPC). Simulation results demonstrate the advantages of MPC over adaptive ECMS in particular in terms of tracking performance in charge depleting operation.

II. Modelling and Problem Formulation

This paper considers a parallel hybrid electric vehicle configuration as illustrated in Figure 1. This system consists of an electric motor and a combustion engine, both of which contribute additively to the gearbox input torque (i.e. if a clutch is present for disconnecting the engine from the gearbox, this is assumed to be permanently closed and we do not consider a pure electric driving mode). Further, the engine shaft speed ω_{eng} and electric motor shaft speed ω_{em} are identical, i.e. $\omega_{eng} = \omega_{em} = \omega$. This can be summarized in the following expression for the driver demand, P_{drv}:

$$P_{drv} = P_{eng} + P_{em} = T_{eng}\omega + T_{em}\omega \qquad (1)$$

[a]BMW Group, Munich, Germany
[b]Department of Engineering Science, University of Oxford, OX1 3PJ, UK
[*]Corresponding author: mark.cannon@eng.ox.ac.uk

where P_{eng}, T_{eng} and P_{em}, T_{em} are the power/torque demands for the engine and electric motor, respectively.

Our goal is to minimize a primary objective over a driving cycle with predefined driver power demand and speed profiles. To this end, we consider the minimization of a quasistatic map that depends on the instantaneous engine power P_{eng} and engine speed ω_{eng}. In this paper we focus on minimizing fuel consumption and hence we assume a static fuel power map, P_f, (see e.g. [7]):

$$P_f = P_f(P_{eng}, \omega) = h_{LHV} \dot{m}_f(P_{eng}, \omega) \qquad (2)$$

where $\dot{m}_f$ is the fuel flow rate and h_{LHV} denotes the lower heating value of the fuel. Secondly, the losses incurred in the conversion of electrical power, P_b, to mechanical power, P_{em}, of the electric machine are assumed to be given by a static map (e.g. [7]):

$$P_b = P_b(P_{em}, \omega). \qquad (3)$$

Both maps $P_f(P_{eng}, \omega)$ and $P_b(P_{em}, \omega)$ are assumed to be available to the controller, for example in the form of lookup tables. Furthermore the mapping between the rate of change of energy stored in the battery, P_s, and the electrical power, P_b, is determined using an equivalent electrical circuit model of the battery [2]:

$$P_s = \frac{U_{OC}^2}{2R} \left[1 - \sqrt{1 - \frac{4R}{U_{OC}^2} P_b} \right] \qquad (4)$$

where U_{OC} is the open circuit voltage of the battery and R corresponds to the internal resistance.

In line with standard optimal control approaches to the strategic power split problems in hybrid vehicles [2], the lower level powertrain dynamics are neglected and only the battery energy content E is considered as the relevant state variable. Changes in the energy level result from the battery usage P_s:

$$\dot{E} = -P_s \qquad (5)$$

The energy management problem consists of the minimization of the primary objective (2) over a driving cycle (given information about the current and future driver demand P_{drv} and the current and future longitudinal vehicle speed v). We assume that a gear selection is performed using a separate powertrain function and is therefore not considered in the energy management problem. We therefore assume that for a given vehicle speed v and driver demand P_{drv} a corresponding engine speed ω is known. Relevant constraints in the context of this problem are given by the dynamics (5), by the actuator input constraints and state path constraints and state boundary conditions on the battery energy level. We consider an optimal control formulation as the following constrained optimization problem [2], [7]:

$$\underset{P_{eng}(t),\ t \in [0,T]}{\text{minimize}} \int_0^T P_f(P_{eng}(t), \omega(t))\, \mathrm{d}t \qquad (6a)$$

$$\text{subject to} \quad \dot{E}(t) = -P_s(t) \qquad (6b)$$

$$P_{eng}^{min} \leq P_{eng}(t) \leq P_{eng}^{max} \qquad (6c)$$
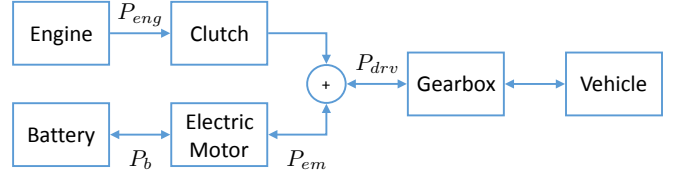

Fig. 1. Power flows in parallel hybrid vehicle configuration

$$P_{em}^{min} \leq P_{em}(t) \leq P_{em}^{max} \qquad (6d)$$

$$E_{min} \leq E(t) \leq E_{max} \qquad (6e)$$

$$E(0) = E_{init} \qquad (6f)$$

$$E(T) = E_{term} \qquad (6g)$$

which is to be performed subject to the algebraic constraints (1), (2), (3), (4), and for given trajectories $w(t)$, $P_{drv}(t)$ for $0 \leq t \leq T$, where T denotes the length of the driving cycle. Through appropriate choice of terminal conditions the above problem formulation can be employed in both charge sustaining (CS) and charge depleting (CD) modes[1] of operation, namely $E(T) = E_{term} = E_{init}$ for CS and $E(T) = E_{term} < E_{init}$ for CD, respectively.

III. CONTROL ALGORITHMS

The solution of the PHEV optimal energy management problem is likely to be affected by inaccuracies arising from discrepancies between the assumed and actual future driving cycles, and from model approximations resulting from the use of quasistatic maps to represent the fuel power map and electrical losses. To reduce the effects of these sources of uncertainty we propose a Model Predictive Control implementation that provides feedback by recomputing, repeatedly, online, the optimal predicted future power flows using current information on the state of charge of the battery and the driving cycle. This section describes the basic MPC optimization, techniques for its online solution, and a hierarchical approach for handling long prediction horizons. We begin by translating the optimal control problem (6) into an approximate discrete time MPC formulation.

For a sampling interval δ, let y_k denote the value of $P_{drv}(t + k\delta)$ predicted at time t, for $k = 0, 1, \ldots, N$, where $N\delta = T - t$, and define x_k as the k steps ahead predicted value of P_{eng} at time t. For a given engine speed profile $\omega(t)$, $t \in [0, T]$, we approximate the fuel power map (2) appearing in the objective function (6a) in terms of a time-varying quadratic function of the engine power demand: $P_f(P_{eng}(t + k\delta), \omega(t + k\delta)) \approx f_k(P_{eng}(t + k\delta)) = f_k(x_k)$, where

$$f_k(x_k) = \alpha_{2,k} x_k^2 + \alpha_{1,k} x_k + \alpha_{0,k} \qquad (7)$$

with $\alpha_{i,k} = \alpha_i(\omega(t + k\delta))$ for $i = 0, 1, 2$ and $k = 0, \ldots, N$. We also approximate the mapping (3) between the electrical power P_b and the mechanical power P_{em} in the electric machine as a time-varying quadratic function of motor power

[1]Here CD is understood to mean a strategy that aims to reach the target energy level at the end of the driving cycle.

demand: $P_b\big(P_{em}(t+k\delta), \omega(t+k\delta)\big) \approx h_k\big(P_{em}(t+k\delta)\big) = h_k(y_k - x_k)$, where

$$h_k(y_k - x_k) = \beta_{2,k}(y_k - x_k)^2 + \beta_{1,k}(y_k - x_k) + \beta_{0,k} \quad (8)$$

with $\beta_{i,k} = \beta_i\big(\omega(t+k\delta)\big)$ for $i = 0, 1, 2$ and $k = 0, \ldots, N$.

The MPC approach of [7] approximates the optimization (6) as a quadratic program through the use of a quadratic approximation of the overall map from P_f to P_s and by specifying the decision variables as the predicted samples of P_s. However the two quadratic approximations (7)-(8) and the equivalent electrical circuit model of the battery (4) provide more accurate representations of the power flows in the powertrain. In addition, by defining the variables in the optimization as the predicted samples, x_k, of P_{eng}, the current formulation allows the bounds (6c,d) on P_{eng}, P_{em} to be imposed via simple bounds on the decision variables x_k.

To simplify the presentation we neglect the bounds (6d) on P_{em} and the upper bound on P_{eng} in (6c). These constraints are rarely active in normal operation, but if needed they can be imposed as upper and lower bounds on the decision variables x_k and handled by a straightforward extension of the method proposed below for imposing the lower bounds of (6c). The lower bounds on P_{eng} in (6c) are needed in this formulation, however, since they reflect the fact (which is not implicit in the quadratic fuel map approximation (7)) that energy dissipated during engine-braking (i.e. when $P_{eng} < 0$) is not recoverable.

In the following development we also neglect the energy storage bounds (6e). Unlike the case of the bounds (6c,d) this simplification is not introduced here only to streamline presentation. Instead it is motivated by the practical consideration that the constraints on stored energy are usually inactive since the upper and lower limits of state of charge are not reached by HEVs under normal operating conditions [11].

With these considerations, a discrete time approximation of the problem (6) over the interval $[t, T]$ is given by

$$\underset{\substack{x_k \geq 0 \\ k=0,\ldots,N}}{\text{minimize}} \quad \sum_{k=0}^{N} f_k(x_k) \quad (9a)$$

$$\text{subject to} \quad \sum_{k=0}^{N} g_k(y_k - x_k) \leq \Delta E \quad (9b)$$

where $\Delta E = (E(t) - E_{term})/\delta$ and the functions g_k are defined for all k by

$$g_k(y_k - x_k) = \frac{U_{OC}^2}{2R}\left[1 - \sqrt{1 - \frac{4R}{U_{OC}^2}h_k(y_k - x_k)}\right] \quad (10)$$

Let $X_k = \{x : x \geq 0 \text{ and } h_k(y_k - x) \leq U_{OC}^2/4R\}$ denote the domain of $f_k(\cdot)$ and $g_k(y_k - \cdot)$. We make the following convexity assumptions on f_k and g_k.

Assumption 1: For all k and all $x \in X_k$, $f_k(x)$ is strictly convex and $g_k(y_k - x)$ is convex in x.

Remark 1: The definitions of f_k, h_k and g_k in (7), (8) and (10) imply that Assumption 1 holds if, for all k, $\alpha_{2,k} > 0$ and $\beta_{2,k} \geq 0$.

The problem (9) has the form of a continuous nonlinear knapsack problem, for which many solution methods are available (see e.g. [12]). To exploit convexity we propose an approach with rapid convergence properties based on the projected Newton method [13]. Before discussing this approach, we first note that the optimization is decoupled at $k = 0, \ldots, N$ if the energy constraint (9b) is inactive. In this case the optimal solution is given for all k by

$$x_k = [x_k^0]^+, \quad x_k^0 = -\tfrac{1}{2}\alpha_{1,k}/\alpha_{2,k} \quad (11)$$

where, for any real vector z, we define $[z]^+$ as the projection

$$[z]^+ = \max\{0, z\}.$$

Therefore it is relatively easy to determine whether the constraint (9b) is inactive at the solution of (9) by checking the condition:

$$\sum_{k=0}^{N} g_k\big(y_k - [x_k^0]^+\big) \leq \Delta E. \quad (12)$$

To initialize the solution when the constraint (9b) is active we consider a simplified problem in which the bounds $x_k \geq 0$ are ignored and the losses in the battery and electric motor are neglected, i.e. $P_s = P_b = P_{em}$ and hence $g_k(y_k - x_k) = y_k - x_k$. Defining the Lagrangian function

$$L(\mathbf{x}, p) = \sum_{k=0}^{N} f_k(x_k) + p\left[\sum_{k=0}^{N} g_k(y_k - x_k) - \Delta E\right], \quad (13)$$

where $\mathbf{x} = (x_0, \ldots, x_N)$ and $p \geq 0$ is the costate variable (namely the Lagrange multiplier associated with constraint (9b)), the optimality conditions $\nabla_{\mathbf{x}} L = 0$, $\nabla_p L = 0$ are equivalent to

$$f_k'(x_k) - p\, g_k'(y_k - x_k) = 0 \quad \forall k \quad (14a)$$

$$\sum_{k=0}^{N} g_k(y_k - x_k) = \Delta E. \quad (14b)$$

For the simplified problem these conditions yield

$$2\alpha_{2,k}x_k + \alpha_{1,k} - p = 0 \quad \forall k$$

$$\sum_{k=0}^{N}(y_k - x_k) = \Delta E$$

from which we obtain the closed form solution [7]:

$$\hat{x}_k = \tfrac{1}{2}(\hat{p} - \alpha_{1,k})/\alpha_{2,k} \quad \forall k \quad (15a)$$

$$\hat{p} = \frac{\sum_{k=0}^{N} y_k + \frac{1}{2}\sum_{k=0}^{N}(\alpha_{1,k}/\alpha_{2,k}) - \Delta E}{\frac{1}{2}\sum_{k=0}^{N}(1/\alpha_{2,k})} \quad (15b)$$

where $\alpha_{2,k} > 0$ for all k by Assumption 1. An initial guess of the solution of (9) and the corresponding costate is then provided by $([\hat{\mathbf{x}}]^+, \hat{p})$.

Returning to the original problem which includes the bounds $x_k \geq 0$ and incorporates electrical and mechanical losses in the model, we define a projected Newton iteration as follows. Let

$$\mathcal{A} = \{k : x_k = 0 \text{ and } f'(0) - p\, g'(y_k) > 0\}, \quad (16)$$

so that $\mathcal{A}$ denotes the set of indices at which the constraint

$x_k \geq 0$ is active, and define $(\Delta\mathbf{x}, \Delta p)$ as the solution of

$$\begin{bmatrix} D & v \\ v^T & 0 \end{bmatrix} \begin{bmatrix} \Delta\mathbf{x} \\ \Delta p \end{bmatrix} = \begin{bmatrix} \phi \\ \Delta E - \sum_{k=0}^{N} g_k(y_k - x_k) \end{bmatrix} \quad (17)$$

where

$$D = \text{diag}\{d_0, \ldots, d_N\}, \quad d_k = f_k''(x_k) + p\, g_k''(y_k - x_k) \,\forall k$$

$$v = (v_0, \ldots, v_N), \qquad v_k = \begin{cases} -g_k'(y_k - x_k), & k \notin \mathcal{A} \\ 0, & k \in \mathcal{A} \end{cases}$$

$$\phi = (\phi_0, \ldots, \phi_N), \qquad \phi_k = -f_k'(x_k) + p\, g_k'(y_k - x_k) \,\forall k$$

Then the update at each iteration is given by

$$\mathbf{x} \leftarrow [\mathbf{x} + \Delta\mathbf{x}]^+ \quad (18a)$$
$$p \leftarrow p + \Delta p \quad (18b)$$

Remark 2: The solution of (17) is given by

$$\Delta\mathbf{x} = D^{-1}(\phi - v\Delta p) \quad (19a)$$

$$\Delta p = \frac{v^T D^{-1}\phi + \sum_{k=0}^{N} g_k(y_k - x_k) - \Delta E}{v^T D^{-1} v} \quad (19b)$$

and the implied computation scales linearly with N since D is diagonal. Moreover it is only necessary to compute Δx_k for $k \notin \mathcal{A}$ since (16) ensures that $[x_k + \Delta x_k]^+ = 0$ for all $k \in \mathcal{A}$.

The following result, which is based on [13], describes the convergence properties of the iteration (18).

Proposition 1: Let $\mathbf{x}^*$ and p^* denote respectively the optimal argument of (9) and the optimal value of the costate, and assume that $p^* \in (0, \infty)$. Then $(\mathbf{x}, p)$ is a fixed point of the iteration (18) if and only if $(\mathbf{x}, p) = (\mathbf{x}^*, p^*)$. Furthermore the iteration converges superlinearly in a neighbourhood of $(\mathbf{x}^*, p^*)$ if no constraints are weakly active at the solution.

Proof: At a minimum point of (9) the solution $(\mathbf{x}^*, p^*)$ must satisfy the first order conditions

$$f_k'(x_k^*) - p^* g_k'(y_k - x_k^*) = 0 \quad \text{if } x_k^* > 0 \quad (20a)$$
$$f_k'(0) - p^* g_k'(y_k) \geq 0 \quad \text{if } x_k^* = 0, \quad (20b)$$

and, since $p^* > 0$, we must also have

$$\sum_{k=0}^{N} g_k(y_k - x_k^*) = \Delta E. \quad (20c)$$

Furthermore, if $p^* > 0$, then at least one k does not belong to the active set $\mathcal{A}^* = \{k : x_k^* = 0 \text{ and } f'(0) - p^* g'(y_k) > 0\}$ at the solution, and if p^* is finite, then $g_k'(y_k - x_k^*)$ must be non-zero for some $k \in \mathcal{A}^*$.

From (20a,b) we have $\phi_k = -f_k'(x_k^*) + p^* g_k'(y_k - x_k^*) = 0$ for all $k \notin \mathcal{A}^*$ and since $p^* \in (0, \infty)$ implies $v \neq 0$ by the preceding argument, it follows from (19b) that $\Delta p = 0$. Hence (19a) implies that $\Delta x_k = 0$ for all $k \notin \mathcal{A}^*$ so $(\mathbf{x}^*, p^*)$ must be a fixed point of (18).

Conversely, if $(\mathbf{x}, p)$ is a fixed point of (18), then $\Delta x_k = 0$ for all $k \notin \mathcal{A}$ and $\Delta p = 0$, so from (17) we can conclude that $\phi_k = -f_k'(x_k) + p\, g_k'(y_k - x_k) = 0$ for all $k \notin \mathcal{A}$ and $\sum_{k=0}^{N} g_k(y_k - x_k) = \Delta E$. Also (16) gives $x_k = 0$ and $f_k'(x_k) - p^* g_k'(y_k - x_k) > 0$ for all $x \in \mathcal{A}$. Therefore the first order conditions are satisfied and by convexity

(Assumption 1) $(\mathbf{x}, p)$ must be optimal.

To complete the proof we note that if there are no weakly active constraints at the solution of (9), then there must exist a neighbourhood of $(\mathbf{x}^*, p^*)$ within which the active set is constant, so the iteration (17)-(18) coincides with Newton's method applied to the subspace $\{\mathbf{x} : x_k = 0 \,\forall k \in \mathcal{A}\}$. A superlinear convergence rate can then be shown using the standard arguments for unconstrained optimization (see e.g. [14]). ∎

Stronger convergence results may be obtained under less restrictive assumptions if Δx and Δp are multiplied by a variable stepsize parameter α in (18). For a similar problem formulation, it is shown in [13] that α can be chosen at each iteration so as to ensure convergence to the optimal solution starting from any feasible initial point. In the current context this modification appears to be unnecessary since simulation results demonstrate consistent and rapid convergence of the iteration with a fixed stepsize $\alpha = 1$ (see Section IV). The linear constraints of (6c,d) and more general convex constraints can be incorporated in this framework [15], [12].

A summary of the MPC algorithm is as follows.

Algorithm 1: At each sampling instant t:

(i). Given the current and terminal battery energy levels $E(t), E_{term}$ and the postulated future driver power and speed demands P_{drv}, ω, determine ΔE and the quadratic approximations f_k and g_k for $k = 0, \ldots, N$.

(ii). Compute the solution (11) of problem (9) without the energy constraint (9b). If (12) is satisfied go to step (v).

(iii). Compute the initial estimate (15) of the constrained solution.

(iv). Perform the iteration (17)-(18) until $|\Delta p|$ is less than a given tolerance ϵ

(v). Implement $P_{eng}(t) = x_0$ and $P_{em}(t) = P_{drv}(t) - x_0$.

The preceding controller setup assumes that the prediction horizon of the MPC optimization (9) extends across the entire driving cycle. When applied to problems with realistic journey times this is likely to result in excessively long prediction horizons. For example the optimization (9) uses 1800 decision variables to optimize a driving cycle of 30 min with a sampling interval of 1 s; problems of this size are unlikely to be solvable within time constraints by current embedded hardware systems.

One way around this difficulty is to use a prediction horizon much shorter than the driving cycle and to assume a given reference trajectory for the battery stored energy. This is analogous to the approach that was proposed for PMP in [9], with a state of charge profile that varies linearly with time. A predefined reference trajectory of this kind can be used to determine the terminal stored energy E_{term} in the MPC optimization solved in steps (ii)-(iv) of Algorithm 1. Choosing the horizon, N, of (9) to be in the range 50–100 will then optimize performance over a future interval of 1–2 min (for which the predicted driver power and speed demands are likely to be relatively accurate) while allowing the battery stored energy to track a linearly varying profile.
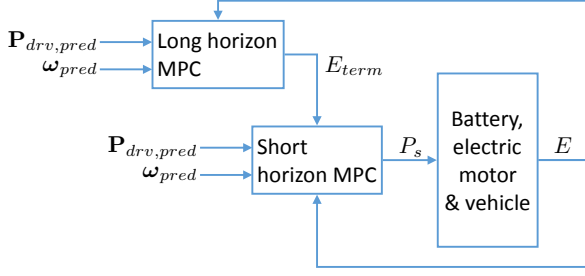
Fig. 2. Two-level MPC scheme for PHEV Energy Management; $\mathbf{P}_{drv,pred}$ and $\boldsymbol{\omega}_{pred}$ denote the sequences of predicted future driver power and speed demands

A. Two-level Hierarchical MPC

Although computationally convenient, MPC strategies based on predefined reference trajectories for the battery state of charge can be highly suboptimal when the optimal stored energy profile deviates from the assumed reference. For example, failure to account for changes in vehicle altitude in the design of the reference trajectory has a significant effect on the performance of PMP approaches (see e.g. [10]), and as discussed in Section IV, the same is true of MPC strategies. Furthermore, although it may not be possible to predict the driver power and speed demands with a high degree of accuracy over a long horizon, relevant information such as the predicted variation in altitude is likely to be available from a navigation unit.

To make use of this information we therefore propose a hierarchical MPC scheme consisting of an outer loop MPC optimization over a long prediction horizon and an inner loop short horizon MPC optimization. The prediction horizon of the outer loop covers the entire driving cycle and therefore shrinks, whereas that of the inner loop is held constant over the duration of the driving cycle. Consequently the outer and inner loop optimization problems are likely to involve similar numbers of decision variables at the start of the driving cycle.

The outer loop optimization is formulated using a long sampling interval, which is greater than that of the inner loop by a factor of e.g. $r = 50$, and with predicted power and speed demands averaged over each of these sampling intervals. The optimal battery stored energy profile computed by the outer loop is used to generate the reference target E_{term} for the inner loop (Fig. 2). Both inner and outer loop optimizations can be implemented using Algorithm 1 with the modification that, for the outer loop, step (v) should determine the target terminal state of charge for the current inner loop optimization.

IV. SIMULATIONS

This section presents the results of simulations performed using data for a $1800\,\mathrm{kg}$ passenger vehicle with a parallel hybrid electric configuration. The hybrid powertrain consists of a $100\,\mathrm{kW}$ gasoline internal combustion engine, a $50\,\mathrm{kW}$ AC electric motor and a $21.5\,\mathrm{Ah}$ lithium-ion battery. We consider the FTP-75 driving cycle with and without altitude variations. The altitude variations take the form of a positive gradient ($2°$) for the first half of the driving cycle, followed

by a negative gradient ($-2°$) for the remainder. In both cases the battery SOC is initially at 60% and the target SOC at the end of the driving cycle is set at 50%.

Three variants of the MPC algorithm described in Section III are considered: (i) a single loop controller with a shrinking prediction horizon that covers the entire driving cycle; (ii) a single loop controller with a horizon of 50 sample intervals and a linearly varying SOC reference; (iii) a two-level MPC strategy with an inner loop prediction horizon of 50 sample intervals and a shrinking horizon outer loop optimization with a sampling interval equal to 50 inner loop sample intervals. These are compared with two controllers based on Pontryagins Minimum Principle applied to the optimal control problem (6): (i) PMP with fixed costate variable p which is chosen to be optimal for the given driving cycle; and (ii) PMP with costate adapted online via proportional feedback of the error in SOC relative to a linearly varying reference. Both PMP variants are implemented by minimizing the associated Hamiltonian online (see [10] for details); the Hamiltonian is computed online directly from the fuel power map and the maps describing losses in the electric motor and the battery using interpolation.
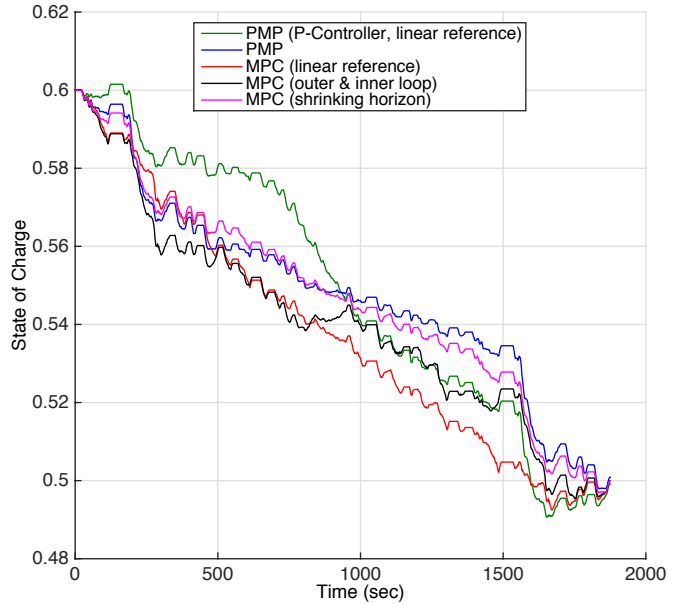

Fig. 3. Variation of battery SOC for driving cycle with zero gradient

For any given driving cycle, the PMP controller with a fixed costate optimized offline for that driving cycle necessarily achieves the best possible performance. Therefore, by comparing PMP with the response of the single loop shrinking horizon MPC strategy we obtain an indication of the effects of the approximation errors that are introduced by the quadratic maps in (7) and (8). For these two strategies, the SOC profiles remain at all times within 1.5% of each other (Fig.s 3 and 4) and the average fuel consumptions also lie within 1.5% (Table I), indicating that these approximations cause a degree of suboptimality no greater than 1.5%.

The PMP strategy with costate adaptation via proportional control and a linear SOC reference performs slightly better

than the single loop MPC with linear SOC reference for the zero gradient case (Fig. 3 and Table I). This is because the initial costate and the proportional feedback gain were tuned specifically for the zero gradient driving cycle. Although both strategies are close to optimal for a gradient of zero, they both fail to achieve the desired target SOC for the non-zero gradient. This can be attributed to the over-cautious strategy that results from the linear SOC reference while travelling uphill in the first half of the driving cycle. Choosing a large proportional gain for the costate adaptation causes instability, and the lack of accurate control of the costate adaptation makes the PMP strategy perform slightly worse than the single loop MPC in the case of non-zero gradients.
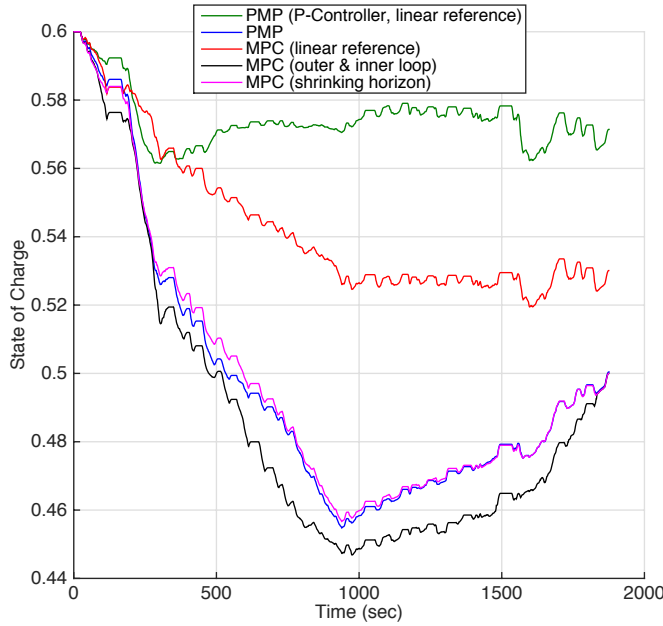

Fig. 4. Variation of battery SOC for driving with $\pm 2°$ gradient

The combined outer and inner loop MPC strategy works well for both zero and non-zero road gradients, and gets close to the optimal fuel consumption achieved by PMP with fixed costate in all cases. This near-optimal performance is obtained without the need for the extensive offline tuning that is required to find the optimal costate for the PMP approach. Furthermore the inner and outer loop MPC strategy demonstrates a degree of robustness to uncertainty in the driving cycle: perturbations of $\pm 20\%$ in the predicted driver power and speed demand result in just a few percent suboptimality.

TABLE I

Mean fuel consumption (kg/100km); null indicates target SOC not reached

	Gradient: 0°	Gradient: $\pm 2°$
PMP (P-Controller, linear reference)	3.82	–
PMP	3.80	4.26
MPC (linear reference)	3.88	–
MPC (outer & inner loop)	3.87	4.33
MPC (shrinking horizon)	3.85	4.27

The performance of the outer and inner loop MPC strategy is even more remarkable when its computational requirement is taken into account. Table II compares the online computation times required at each time step for PMP (fixed costate)

and two-level MPC; Mean and Max denote the average and greatest computation times over the whole cycle. This shows that the computational requirement of the two-level MPC strategy is comparable to that for the PMP strategy.

TABLE II

Computation times (ms) at individual time steps

	Mean	Max
PMP	24	119
MPC (inner loop)	10	76
MPC (outer loop)	6	36
MPC (outer & inner loop)	16	112

V. CONCLUSIONS

A nonlinear MPC strategy provides near-optimal energy management for plugin hybrid electric vehicles. The proposed hierarchical strategy avoids the offline tuning that is required by ECMS methods and does not suffer from the suboptimality caused by a predefined reference trajectory for the battery state of charge. Future work will investigate more general cost functions, for example incorporating emissions objectives, and more general constraints such as probabilistic constraints on energy usage when statistical information is available on the predicted driving cycle.

REFERENCES

[1] McKinsey & Company. Electric vehicles in europe: gearing up for a new phase? (http://www.mckinsey.com/insights/europe). Technical report, Amsterdam Roundtable Foundation, 2014.
[2] A. Sciarretta and L. Guzzella. Control of hybrid electric vehicles. *IEEE Control Systems Magazine*, 27(2):60–70, 2007.
[3] C. Lin, H. Peng, J.W. Grizzle, and J. Kang. Power management strategy for a parallel hybrid electric truck. *IEEE Transactions on Control Systems Technology*, 11(6):839 – 849, 2003.
[4] CC. Lin, H. Peng, and J.W. Grizzle. A stochastic control strategy for hybrid electric vehicles. In *Proc. ACC*, pages 4710 – 4715, 2004.
[5] G. Paganelli, G. Ercole, A. Brahma, Y. Guezennec, and G. Rizzoni. General supervisory control policy for the energy optimization of charge-sustaining hybrid electric vehicles. *JSAE review*, 22(4):511–518, 2001.
[6] N. Kim, S. Cha, and H. Peng. Optimal control of hybrid electric vehicles based on pontryagin's minimum principle. *IEEE Transactions on Control Systems Technology*, 19(5):1279–1287, 2011.
[7] M Koot, J.T.B.A. Kessels, B. de Jager, W.P.M.H. Heemels, P.P.J van den Bosch, and M. Steinbuch. Energy management strategies for vehicular electric power systems. *IEEE Transactions on Vehicular Technology*, 54(3):771–782, 2005.
[8] C. Musardo, G. Rizzoni, Y. Guezennec, and B. Staccia. A-ecms: An adaptive algorithm for hybrid electric vehicle energy management. *European Journal of Control*, 11(4):509–524, 2005.
[9] P. Tulpule, V. Marano, and G. Rizzoni. Energy management for plug-in hybrid electric vehicles using equivalent consumption minimisation strategy. *International Journal of Electric and Hybrid Vehicles*, 2(4):329–350, 2010.
[10] S. Onori and L. Tribioli. Adaptive Pontryagin's Minimum Principle supervisory controller design for the plug-in hybrid GM Chevrolet Volt. *Applied Energy*, 147:224–234, 2015.
[11] L. Serrao, A. Sciarretta, O. Grondin, A. Chasse, Y. Creff, D. Di Domenico, P. Pognant-Gros, C. Querel, and L. Thibault. Open issues in supervisory control of hybrid electric vehicles: A unified approach using optimal control methods. *Oil & Gas Science and Technology - Rev. IFP Energies nouvelles*, 68(1):23–33, 2013.
[12] K.M. Bretthauer and B. Shetty. The nonlinear knapsack problem – algorithms and applications. *Eur. J. Oper. Res.*, 138(3):459–472, 2002.
[13] D. Bertsekas. Projected Newton methods for optimization problems with simple constraints. *SIAM J. Control Optim.*, 20(2):221–246, 1982.
[14] R. Fletcher. *Practical Methods of Optimization*. Wiley, 2000.
[15] P.H. Calamai and J.J. Moré. Projected gradient methods for linearly constrained problems. *Math. Prog.*, 39:93–116, 1987.