

Fast dual loop nonlinear receding horizon control for energy management in hybrid electric vehicles

Johannes Buerger, Sebastian East, and Mark Cannon

Abstract—This paper proposes a receding horizon optimization strategy for the problem of energy management in plugin hybrid electric vehicles. The approach employs a dual loop Model Predictive Control (MPC) strategy. An inner feedback loop addresses the problem of optimally tracking a given reference trajectory for the battery state of energy over a short future horizon using knowledge of the predicted driving cycle. An outer feedback loop generates the battery state of energy reference trajectory by solving approximately the optimal energy management problem for the entire driving cycle. The receding horizon optimization problems associated with both inner and outer loops are solved using a specialized 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. The paper contains a detailed simulation study: first, this assesses the optimality of the associated uncertainty-free approach and its computational load. Secondly, the effects of imprecise knowledge of the future driving cycle are illustrated.

Index Terms—Automotive controls, Energy management, Receding horizon control, Nonlinear programming, Hybrid electric vehicles.

I. INTRODUCTION

The automotive industry has been steadily moving towards a stronger electrification of the powertrain over recent years [1]. A particular focus in industry is on Plugin Hybrid Electric Vehicles (PHEVs), in which the battery is usually charged prior to the trip and gradually emptied towards the end of the trip. While the improvement in 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 using the electric machine to minimize a strategic objective, such as fuel and/or emissions, over a driving cycle. Typically this optimization must be performed subject to the constraint that a desired energy level is reached at the end of the driving cycle. This problem has been tackled by a variety of approaches [2], and methods based on optimal control formulations have been particularly successful. These can be grouped into three

classes: methods based on dynamic programming (DP) [3], [4], methods based on Pontryagin's Minimum Principle (PMP) [5], [6] and those based on Model Predictive Control (MPC) [7].

In principle, DP-based methods can determine the optimal strategy for a given driving cycle. However the approach is computationally expensive, and its complexity typically scales exponentially with relevant problem dimensions (in particular with the horizon length). The price to be paid for this is either the high computational demand of solving an online DP or the high memory storage requirements of offline map-based DP approaches. On the other hand, methods based on PMP, in particular the Equivalent Consumption Minimization Strategy (ECMS) [5], are suitable candidates for production vehicles since the required optimization can be performed at each time instant given knowledge of the battery state of energy and the optimal costate variable. Unfortunately, the optimal costate is highly sensitive to the driving cycle and the method requires careful initialization and adaptation of the costate in order to ensure that the target state of energy is reached [8]. In contrast, MPC approaches have the advantage that the tuning required for ECMS methods is avoided and information about the future driving cycle obtained from the navigation unit can be systematically included in the receding horizon MPC optimization. This allows the predicted power and speed to be updated during the journey, resulting in a greater degree of inherent robustness to uncertainty in the driving cycle.

Robustness to uncertainty in the predicted drive cycle is a key aspect of controller performance. Control strategies that account explicitly for uncertainty in future driver demand have previously been proposed, and methods based on stochastic MPC (e.g. [9], [10]) are potentially useful in this context. However these approaches can be very computationally intensive and may not be suitable for applications with limited computational resources. Specifically the online optimization in stochastic MPC algorithms typically involves large numbers of decision variables and constraints because stochastic models of driver demand are used either to enumerate scenario trees [9] or to generate scenarios for use with sampled convex programming techniques [10]. As a result these approaches have so far been limited to linear models obtained by linearising around operating points for the PHEV powertrain components. The emphasis of the current paper is on developing a computationally simple nonlinear MPC strategy for the optimal energy management problem without resorting to linear models of engine efficiency and electrical losses in the

This paper was first submitted on 7 May 2017, revised on 27 October 2017, and received in final form on XXXX 2018.

Johannes Buerger is with BMW Group, Munich, Germany (e-mail: johannes.albert.buerger@gmail.com).

Sebastian East and Mark Cannon are with the Department of Engineering Science, University of Oxford, OX1 3PJ, UK (e-mail: sebastian.east@eng.ox.ac.uk, mark.cannon@eng.ox.ac.uk).

hybrid powertrain. Rather than incorporate explicit knowledge of uncertainty in the driver demand into the control strategy, the paper demonstrates instead the implicit robustness of the approach to various types of model and drive cycle uncertainty.

The contributions of this paper are as follows. First, a non-linear MPC scheme is proposed using a specialised projected Newton method. This minimizes at each sampling instant the predicted fuel consumption, which is derived from the predicted driving cycle, subject to an end-point constraint on battery state of energy and pointwise-in-time bounds on electrical and engine power. The approach is able to account explicitly for electrical losses, and it provides a more general and more accurate treatment than [7], where the MPC problem is approximated directly as a linearly constrained quadratic programming problem. Secondly, the MPC method is applied to a charge depleting optimal control problem, in which a “short-horizon” target state of energy is determined from a reference trajectory that decreases linearly over time, as proposed in the context of PMP in [11] (known as “blended mode”).

The third contribution is a dual loop MPC framework to overcome the limitations of a linearly decreasing reference state, since in the presence of changing predicted driving conditions such as altitude variations caused by varying road gradients, a linear profile can be significantly suboptimal [12]. We propose to solve a low-sampling rate MPC problem for the entire predicted 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). Preliminary results on this approach were presented in [13]. Relative to this contribution, the current paper employs a more realistic powertrain model, includes additional constraints in the optimal power management problem, and provides a detailed simulation study demonstrating the practical usefulness of the dual loop method. This analysis focuses on optimality, robustness and computational aspects of the approach.

II. MODELLING AND PROBLEM FORMULATION

We consider the parallel hybrid electric vehicle configuration illustrated in Figure 1. This system consists of an electric motor, a combustion engine, a clutch, and a gearbox. When the clutch is disengaged, only the electric motor contributes to the powertrain; when it is engaged, the engine and motor contribute additively. The motor shaft speed ω_{em} is calculated from the road speed, v , of the vehicle assuming a simple gear selection strategy (in i th gear, $\omega_{em} = r_i v$ and $v_i \leq v \leq v_{i+1}$, for given gear ratio r_i and threshold speeds v_i, v_{i+1}). The engine speed, ω_{eng} , is identical to the motor speed when the clutch is engaged, i.e. $\omega_{eng} = \omega_{em}$, and equal to some minimum value, $\underline{\omega}_{eng}$, when the clutch is disengaged. The clutch is assumed to be disengaged when the motor speed falls below the engine idle speed $\underline{\omega}_{eng}$.

A fixed fraction, $\gamma_b \in (0, 1)$, of the available braking power is assumed to be recovered (via the electric motor-generator) during braking, and we assume that the motor and engine do not provide power when the vehicle is stationary. We further assume that while the power demand is positive and

the clutch is disengaged, the engine is idle and the vehicle is operating in an all-electric mode. Therefore, a control strategy for load balance is only required when the power demand, P_{drv} , is positive and the motor speed is greater than the engine stall speed. These statements are summarized in the following conditions

$$P_{drv} > 0, \omega_{em} \geq \underline{\omega}_{eng} \Rightarrow P_{eng} + P_{em} = P_{drv} \quad (1a)$$

$$P_{drv} > 0, \omega_{em} < \underline{\omega}_{eng} \Rightarrow P_{eng} = 0, P_{em} = P_{drv} \quad (1b)$$

$$P_{drv} \leq 0 \Rightarrow P_{eng} = 0, P_{em} = \gamma_b P_{drv} \quad (1c)$$

where P_{eng} and P_{em} are, respectively, the power supplied to the drivetrain by the engine and by the motor (Fig. 1).

Our goal is to minimize a primary objective over a driving cycle with predefined 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]):

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

where \dot{m}_f is the fuel flow rate and h_{LHV} denotes the lower heating value of the fuel. (Note that this model accounts for the fuel consumed while the vehicle is braking or the engine is idling, since P_f may be non-zero when $P_{eng} = 0$.) 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]):

$$P_b = P_b(P_{em}, \omega_{em}). \quad (3)$$

Both maps $P_f(P_{eng}, \omega_{eng})$ and $P_b(P_{em}, \omega_{em})$ 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]:

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

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

In line with standard optimal control approaches to 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 (i.e. the battery state of energy). Changes in the energy level result from the battery usage P_s :

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

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 this is not therefore considered in the energy management problem. Relevant constraints in the




Fig. 1. Power flows in parallel hybrid vehicle configuration.

context of this problem are given by the dynamics (5), by the actuator input constraints and by 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]:

$$\underset{P_{eng}(t), t \in [0, T]}{\text{minimize}} \int_0^T P_f(P_{eng}(t), \omega_{eng}(t)) dt \quad (6a)$$

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

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

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

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

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

$$E(T) \geq E_{term} \quad (6g)$$

which is to be performed subject to the algebraic constraints (1), (2), (3), (4), and for given trajectories $\omega_{eng}(t)$, $P_{driv}(t)$ 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 of operation, namely $E(T) = E_{term} = E_{init}$ for CS and $E(T) = E_{term} < E_{init}$ for CD, respectively.

III. OPTIMIZATION ALGORITHM

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, a Model Predictive Control implementation provides feedback by recomputing, repeatedly, online, the optimal predicted future power flows using current information on the state of energy of the battery and the driving cycle. The MPC strategy proposed in this paper is based on a specialized Newton method for solving approximately the underlying optimization problem (6); this is the main topic of the current section. The online implementation of the associated control law, and a hierarchical approach for handling long prediction horizons based on the proposed Newton method are treated in Section IV.

We begin by translating the optimal control problem (6) into an approximate discrete time MPC formulation. For a given sampling interval $\delta_k = t_{k+1} - t_k$ where t_0 is the current time, let y_k , $\omega_{em,k}$, and $\omega_{eng,k}$ respectively denote the values of $P_{driv}(t_k)$, $\omega_{em}(t_k)$, and $\omega_{eng}(t_k)$, predicted at time t_k for $k = 1, \dots, N$, where $\sum_{k=0}^N \delta_k = T - t_0$ and T is the predicted

end time. Also define x_k as the k steps-ahead predicted value of P_{eng} at time t and let z_k denote the k steps-ahead predicted value of P_{em} . The values of y_0 , $\omega_{em,0}$, and $\omega_{eng,0}$ are not taken from the predicted drive cycle as they are known as the current driver power demand, motor state, and engine state respectively. Therefore, at each interval the calculated values of x_0 and z_0 are implemented as the P_{eng} and P_{em} control inputs during simulation. The indices k are grouped into sets based on the conditions (1a-c)

$$\mathcal{P} = \{k : y_k > 0 \text{ and } \omega_{em,k} \geq \underline{\omega}_{eng}\}$$

$$\mathcal{C} = \{k : y_k > 0 \text{ and } \omega_{em,k} < \underline{\omega}_{eng}\}$$

$$\mathcal{B} = \{k : y_k \leq 0\}.$$

Note that each $k \in \{0, \dots, N\}$ lies in either \mathcal{P} , \mathcal{C} , or \mathcal{B} .

For a given engine speed profile $\omega_{eng}(t)$, $t \in [0, T]$, 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 output: $P_f(x_k, \omega_{eng,k}) \approx \alpha_{2,k}x_k^2 + \alpha_{1,k}x_k + \alpha_{0,k}$. Therefore, the predicted fuel consumption during each interval is given by $f_k(x_k)$, where

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

with $\alpha_{i,k} = \alpha_i(\omega_{eng,k})$ for $i = 0, 1, 2$ and $k = 0, \dots, N$. We likewise 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 demand: $P_b(z_k, \omega_{em,k}) \approx \beta_{2,k}z_k^2 + \beta_{1,k}z_k + \beta_{0,k}$. By also considering the electrical circuit battery losses (4), the predicted loss in battery energy is given at each interval by $g_k(z_k)$, where

$$h_k(z_k) = \beta_{2,k}z_k^2 + \beta_{1,k}z_k + \beta_{0,k} \quad (8a)$$

$$g_k(z_k) = \delta_k \frac{U_{OC}^2}{2R} \left[1 - \sqrt{1 - \frac{4R}{U_{OC}^2} h_k(z_k)} \right] \quad (8b)$$

with $\beta_{i,k} = \beta_i(\omega_{em,k})$ for $i = 0, 1, 2$ and $k = 0, \dots, N$. 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 in (7) and (8a), combined with the equivalent electrical circuit model of the battery (4), provide more accurate representations of the power flows in the powertrain.

The current formulation allows the bounds (6c,d) on P_{eng} and P_{em} to be imposed via simple bounds on the decision variables x_k , given by:

$$\bar{x}_k = \min\{\bar{T}_{eng}\omega_{eng,k}, y_k - \underline{T}_{em}\omega_{em,k}\}$$

$$\underline{x}_k = \max\{\underline{T}_{eng}\omega_{eng,k}, y_k - \bar{T}_{em}\omega_{em,k}\}$$

where \bar{T} and \underline{T} denote the upper and lower limits on the engine and motor torque. However we neglect the energy storage bounds (6e), which constitute state constraints in the optimal control problem (6). This is motivated by the observation that the upper and lower limits of the battery state of energy are not reached by PHEVs under normal operating conditions (see e.g. [14]), since these vehicles typically have batteries with

large capacity and short-term violations of state constraints are therefore unlikely.

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

$$\underset{x_0, \dots, x_N, z_0, \dots, z_N}{\text{minimize}} \quad \sum_{k=0}^N f_k(x_k) \quad (9a)$$

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

$$\left. \begin{aligned} z_k + x_k &= y_k \\ \underline{x}_k &\leq x_k \leq \bar{x}_k \end{aligned} \right\} \quad k \in \mathcal{P} \quad (9c)$$

$$x_k = 0 \quad k \notin \mathcal{P} \quad (9d)$$

$$z_k = y_k \quad k \in \mathcal{C} \quad (9e)$$

$$z_k = \gamma_b y_k \quad k \in \mathcal{B} \quad (9f)$$

where $\Delta E = E(t_0) - E_{term}$. For given y_k , the variables x_k and z_k in the minimization (9) are fully specified at all times $k \notin \mathcal{P}$. Hence (9) is equivalent to the following optimization over $x_k, k \in \mathcal{P}$,

$$\underset{x_k, k \in \mathcal{P}}{\text{minimize}} \quad \sum_{k \in \mathcal{P}} f_k(x_k) \quad (10a)$$

$$\text{subject to} \quad \sum_{k \in \mathcal{P}} g_k(y_k - x_k) \leq \Delta E_{\mathcal{P}} \quad (10b)$$

$$\underline{x}_k \leq x_k \leq \bar{x}_k \quad (10c)$$

where

$$\Delta E_{\mathcal{P}} = \Delta E - \sum_{k \in \mathcal{C}} g_k(y_k) - \sum_{k \in \mathcal{B}} g_k(\gamma_b y_k).$$

and clearly no optimisation is required if $\mathcal{P} = \emptyset$.

Let $X_k = \{x \in [\underline{x}_k, \bar{x}_k] : h_k(y_k - x) \leq U_{OC}^2/4R\}$ denote the domain of $f_k(\cdot)$ and $g_k(y_k - \cdot)$ for $k \in \mathcal{P}$. We make the following convexity assumptions on f_k and g_k .

Assumption 1: For $k \in \mathcal{P}$ 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), (8a) and (8b) imply that Assumption 1 holds if $\alpha_{2,k} > 0$ and $\beta_{2,k} \geq 0$ for $k \in \mathcal{P}$.

The optimization (10) has the form of a continuous non-linear knapsack problem, for which various solution methods are available (see e.g. [15]). To exploit convexity we propose an approach with rapid convergence properties based on the projected Newton method [16]. Before discussing this approach, we first note that the optimization is decoupled at $k = 0, \dots, N$ if the energy constraint (10b) is inactive. In this case the optimal solution is given for $k \in \mathcal{P}$ by

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

where, for any real vector u and elementwise bounds \underline{u}, \bar{u} , we define $[u]^\dagger$ as the elementwise saturated version of u :

$$[u]^\dagger = \min\{\bar{u}, \max\{\underline{u}, u\}\}.$$

It is therefore relatively easy to determine whether constraint (10b) is inactive at the solution of Problem (10) by

checking the condition:

$$\sum_{k \in \mathcal{P}} g_k(y_k - [x_k^0]^\dagger) \leq \Delta E_{\mathcal{P}}. \quad (12)$$

To initialize the solution when the constraint (10b) is active we consider a simplified problem in which the bounds $\underline{x}_k \leq x_k \leq \bar{x}_k$ are ignored and the losses in the battery and electric motor are neglected, i.e. $P_s = P_b = P_{em}$ and hence $g_k(z) = \delta_k z$. Defining the Lagrangian function

$$L(\mathbf{x}, p) = \sum_{k \in \mathcal{P}} f_k(x_k) + p \left[\sum_{k \in \mathcal{P}} g_k(y_k - x_k) - \Delta E_{\mathcal{P}} \right], \quad (13)$$

where $\mathbf{x} = \{x_k, k \in \mathcal{P}\}$, and $p \geq 0$ is the costate (namely the Lagrange multiplier associated with constraints (9b) and (10b)), the optimality conditions $\nabla_{\mathbf{x}} L = 0, \nabla_p L = 0$ are equivalent to

$$f'_k(x_k) - p g'_k(y_k - x_k) = 0, \quad k \in \mathcal{P} \quad (14a)$$

$$\sum_{k \in \mathcal{P}} g_k(y_k - x_k) = \Delta E_{\mathcal{P}}. \quad (14b)$$

For the simplified problem with $g_k(z) = \delta_k z$, these yield

$$2\delta_k \alpha_{2,k} x_k + \delta_k \alpha_{1,k} - p \delta_k = 0, \quad k \in \mathcal{P}$$

$$\sum_{k \in \mathcal{P}} \delta_k (y_k - x_k) = \Delta E_{\mathcal{P}}$$

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

$$\hat{p} = \frac{\sum_{k \in \mathcal{P}} \delta_k y_k + \frac{1}{2} \sum_{k \in \mathcal{P}} (\delta_k \alpha_{1,k} / \alpha_{2,k}) - \Delta E_{\mathcal{P}}}{\frac{1}{2} \sum_{k \in \mathcal{P}} (\delta_k / \alpha_{2,k})} \quad (15a)$$

$$\hat{x}_k = \frac{1}{2}(\hat{p} - \alpha_{1,k}) / \alpha_{2,k}, \quad k \in \mathcal{P} \quad (15b)$$

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 $([\hat{\mathbf{x}}]^\dagger, \hat{p})$.

Returning to the original problem, which includes the bounds $\underline{x}_k \leq x_k \leq \bar{x}_k$ and incorporates electrical and mechanical losses in the model, we define a projected Newton iteration as follows. Let \mathcal{A} be the subset of \mathcal{P} defined by

$$\mathcal{A} = \left\{ k \in \mathcal{P} : \left\{ \begin{aligned} &x_k = \underline{x}_k \text{ and } f'(\underline{x}_k) - p g'(y_k - \underline{x}_k) > 0 \\ &\text{or } x_k = \bar{x}_k \text{ and } f'(\bar{x}_k) - p g'(y_k - \bar{x}_k) < 0 \end{aligned} \right\} \right\}, \quad (16)$$

so that \mathcal{A} denotes the set of indices $k \in \mathcal{P}$ for which the constraint $\underline{x}_k \leq x_k \leq \bar{x}_k$ is active, and define $(\Delta \mathbf{x}, \Delta p)$ via

$$\Delta p = \frac{\sum_{k \in \mathcal{P}} g_k(y_k - x_k) + \sum_{k \in \mathcal{P} \setminus \mathcal{A}} v_k \phi_k / d_k - \Delta E_{\mathcal{P}}}{\sum_{k \in \mathcal{P} \setminus \mathcal{A}} v_k^2 / d_k} \quad (17a)$$

$$\Delta x_k = \begin{cases} (\phi_k - v_k \Delta p) / d_k, & k \in \mathcal{P} \setminus \mathcal{A} \\ 0 & k \in \mathcal{A} \end{cases} \quad (17b)$$

where

$$\begin{aligned} d_k &= f''_k(x_k) + p g''_k(y_k - x_k) \\ v_k &= -g'_k(y_k - x_k) \\ \phi_k &= -f'_k(x_k) + p g'_k(y_k - x_k). \end{aligned}$$

Then the update at each iteration is given by

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

$$p \leftarrow p + \Delta p. \quad (18b)$$

Remark 2: The computation required to perform the projected Newton iteration (16)-(18) scales linearly with N .

The following result, based on [16] and [17], describes the convergence properties of the iteration (18).

Proposition 1: Let \mathbf{x}^* and p^* denote respectively the optimal argument of (10) and the optimal value of the costate in (13), 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 (10) the solution (\mathbf{x}^*, p^*) must satisfy the first order optimality conditions:

$$f'_k(x_k^*) - p^* g'_k(y_k - x_k^*) = 0 \quad \text{if } x_k^* \in (\underline{x}_k, \bar{x}_k) \quad (19a)$$

$$f'_k(x_k^*) - p^* g'_k(y_k - x_k^*) \geq 0 \quad \text{if } x_k^* = \underline{x}_k \quad (19b)$$

$$f'_k(x_k^*) - p^* g'_k(y_k - x_k^*) \leq 0 \quad \text{if } x_k^* = \bar{x}_k \quad (19c)$$

and since $p^* > 0$ is assumed, (10b) must be active, i.e.

$$\sum_{k \in \mathcal{P}} g_k(y_k - x_k^*) = \Delta E_{\mathcal{P}}. \quad (19d)$$

Let \mathcal{A}^* denote the active set at the solution (\mathbf{x}^*, p^*) , then from the optimality conditions (19) and the definition of \mathcal{A} in (16) it follows that: (i) $\mathcal{P} \setminus \mathcal{A}^*$ must be non-empty (since constraints (10c) cannot be active for all $k \in \mathcal{P}$ if (19d) holds because of the assumption that no constraint is weakly active at the solution), (ii) $\phi_k = -f'_k(x_k^*) + p^* g'_k(y_k - x_k^*) = 0$ for all $k \in \mathcal{P} \setminus \mathcal{A}^*$, and (iii) $v_k = -g'_k(y_k - x_k^*) \neq 0$ for some $k \in \mathcal{P} \setminus \mathcal{A}^*$ (since p^* is by assumption finite). Therefore (17a) implies $\Delta p = 0$ while (17b) implies $\Delta x_k = 0$ for all $k \in \mathcal{P}$, and hence (\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 $k \in \mathcal{P}$ and $\Delta p = 0$. In this case, from (17b) we have that $\phi_k = -f'_k(x_k) + p g'_k(y_k - x_k) = 0$ for all $k \in \mathcal{P} \setminus \mathcal{A}$ so that (17a) implies $\sum_{k \in \mathcal{P}} g_k(y_k - x_k) = \Delta E_{\mathcal{P}}$. From these conditions and the definition of \mathcal{A} in (16) it follows that the first order conditions (19a-d) are satisfied, and, by convexity (Assumption 1), (\mathbf{x}, p) must therefore be optimal for (10).

To complete the proof we note that if there are no weakly active constraints at the solution of (10), then there must exist a neighbourhood of (\mathbf{x}^*, p^*) within which the active set is constant, so the iteration (16)-(18) coincides with Newton's method applied to a subspace (namely $\{\mathbf{x} : x_k = \underline{x}_k \text{ or } x_k = \bar{x}_k \text{ for all } k \in \mathcal{A}\}$). A superlinear convergence rate can then be shown using the standard arguments for unconstrained optimization (see e.g. [18]). ■

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 [16] 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 V).

IV. CONTROL ALGORITHM

This section considers the definition of MPC algorithms based on the receding horizon implementation of the optimization method proposed in Section III. We start by listing the major steps of the MPC algorithm:

Algorithm 1: At each sampling instant t :

- (i). Given the current and minimum end-point values of battery state of energy, $E(t)$, E_{term} , and the postulated future power and speed demands P_{drv} , ω_{em} , determine ΔE and quadratic approximations f_k , g_k , $k = 0, \dots, N$.
- (ii). Compute the solution (11) of Problem (10) without the energy constraint (10b). If (12) is satisfied go to step (v).
- (iii). Compute an initial solution estimate using (15a,b).
- (iv). Perform the iteration (16)-(18) until either: (i) $|\Delta p|$ is less than a given tolerance ϵ , or (ii) $\mathcal{A} = \mathcal{P}$ (in this case the problem is infeasible).
- (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 (10) 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 (10) 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 using the computing hardware that is currently available onboard a typical production vehicle. To address this computational challenge, two alternative solutions are discussed: first by pre-defining a linearly decreasing long-term battery energy reference, and secondly by computing an optimal long-term battery energy reference.

A. MPC with linear reference

In order to avoid high computational loads associated with long prediction horizons, it could be advantageous 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 [11], with a battery state of energy profile that varies linearly with time. A predefined reference trajectory of this kind can be used to determine the end-point 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 reference profile. This setup is illustrated in Figure 2.

B. Dual loop MPC

Although computationally convenient, MPC strategies based on predefined reference trajectories for the battery state of energy can be highly suboptimal when the optimal stored




Fig. 2. Two strategies for defining the reference for inner loop MPC terminal conditions: a pre-defined linear reference and an optimized reference.




Fig. 3. Dual loop MPC scheme for PHEV Energy Management; $P_{drv,pred}$ and ω_{pred} denote the sequences of predicted future driver power and speed demands.

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. [12]), and, as discussed in Section V, 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 an onboard 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 over time, whereas that of the inner loop is held constant over the duration of the driving cycle. The outer and inner loop optimization problems are expected 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 defining a window size of e.g. $win = 50$, 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. 3). 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. This setup is also illustrated in Figure 2.

V. SIMULATIONS

This section presents the results of simulations performed using data for a 1800 kg passenger vehicle with a parallel hybrid electric configuration. The hybrid powertrain consists of a 100 kW gasoline internal combustion engine, a 50 kW AC electric motor and a 21.5 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 or vice versa. In all cases the battery state of energy is initially at 60% and the target state at the end of the driving cycle is set at 50%. The fraction of available braking power that is recovered by regenerative braking is assumed to be fixed at 40%, so that $\gamma_b = 0.4$.

A. Optimality assessment and computational load for uncertainty-free conditions

Three variants of the MPC algorithm described in Section IV are considered in this section: (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 $N = 50$ sample intervals and a linearly varying battery state of energy reference; (iii) a dual loop MPC strategy with an inner loop prediction horizon of $N = 50$ sample intervals and a shrinking horizon outer loop optimization with a sampling interval equal to $win = 50$ inner loop sample intervals.

These are compared with a controller based on Pontryagin's Minimum Principle applied to the optimal control problem (6), implemented with a fixed costate to minimize the associated Hamiltonian (see [12] for details). The Hamiltonian is minimized online using the battery loss function (4) and quadratic approximations of the fuel power map (2) and the electric motor power map (3). We note that the use of quadratic approximations (7) and (8a) in the PMP-based strategy results in negligible performance degradation relative to a PMP approach in which the Hamiltonian is minimized by directly interpolating on the fuel power map and the electric motor power and battery loss maps.




Fig. 4. Variation of battery state of energy 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 (8a). For these two strategies, the state of energy profiles remain at all times within 0.6% of each other (Fig.s 4 and 5) and the average fuel consumptions also lie within 0.6% (Table I), indicating that these approximations cause a degree of suboptimality no greater than 0.6%.

For the zero gradient case, the performance of MPC with a linearly varying reference for the battery state of charge is similar to the other methods (Fig. 4 and Table I) and the fuel consumption is only marginally increased, to 3.70 kg/100 km. However, when the gradient is non-zero, the fuel consumption for MPC with a linear state of charge reference increases dramatically to 4.45 kg/100 km. This is due to the highly sub-optimal state trajectory (Fig. 5), which, unlike the optimal trajectories for PMP and MPC with shrinking horizon and with dual loop optimization, reaches a final state of charge greater than 50% due to its over-cautious strategy during the first half of the driving cycle.




Fig. 5. Variation of battery energy state for driving with $\pm 2^\circ$ gradient.

The dual loop MPC strategy works well for both zero and non-zero road gradients, as can be seen in Table I, achieving a fuel consumption within 1% of the optimal value that is obtained with PMP with a fixed optimized costate in all cases. We note that 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.

TABLE I
MEAN FUEL CONSUMPTION (kg/100km)

	Gradient: 0°	Gradient: $\pm 2^\circ$
PMP	3.67	4.11
MPC (linear reference)	3.70	4.45
MPC (dual loop)	3.69	4.14
MPC (shrinking horizon)	3.67	4.12

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 (with fixed costate) and dual-loop MPC, with each algorithm implemented in Matlab on a 2.6 GHz Intel Core i7 processor. Mean and Max denote the average and greatest computation times over the whole cycle. This shows that dual-loop MPC provides an order of magnitude improvement in computational time against the shrinking horizon MPC formulation. Dual-loop MPC remains an order of magnitude slower than PMP. However this increase in computation is the price to be paid in order to avoid the offline tuning required for the PMP method; as we show in Section V-C, this results in a degree of robustness that is not achievable using PMP.

TABLE II
COMPUTATION TIMES (ms) AT INDIVIDUAL TIME STEPS

	Mean	Max
PMP	0.80	2.36
MPC (inner loop)	7.75	19.52
MPC (outer loop)	3.28	20.28
MPC (dual loop)	11.03	30.67
MPC (shrinking horizon)	158.39	604.35

B. Optimality assessment for variations in electrical loss model

This section investigates the degree of sub-optimality caused by neglecting aspects of electrical losses in the prediction model. This is shown for two cases: neglecting the electrical losses in the battery (by setting $g_k(z_k) = h_k(z_k)$ in the prediction model), and neglecting the electrical losses from in the battery and electric motor (by setting $g_k(z_k) = z_k$ in the prediction model). In each case, all electrical losses are retained in the powertrain model that is used to simulate closed loop behaviour. It can be seen in Figure 6 that neglecting both losses causes a far greater deviation from the optimal MPC state trajectory than the battery losses alone, which implies that the motor losses are more important in establishing an optimal control formulation. Despite the deviations from the nominal MPC method, the final SOC constraint is still met in both cases.

Furthermore, table III shows that the optimization neglecting all electrical losses results in the fuel consumption increasing from 4.12 kg/100 km to 4.16 kg/100 km, whereas neglecting the circuit losses reduces the fuel consumption to 4.11 kg/100 km. This slight improvement is caused by the state trajectory for MPC without battery losses following a trajectory closer to the optimum PMP curve shown in Figure 5, and is due to a coincidence of the sub-optimality from both the MPC formulation and the model reduction cancelling out, rather than due to any systematic improvement. It would be expected that during a longer or higher power cycle that the effect of this sub-optimality would increase, and therefore increase the fuel consumption compared to the nominal MPC. The results do suggest, however, that optimal control formulations that neglect the battery losses can produce near optimal results.




Fig. 6. Variation of battery energy state for driving with $\pm 2^\circ$ gradient.

TABLE III
MEAN FUEL CONSUMPTION (kg/100 km)

MPC (shrinking horizon)	4.12
MPC (shrinking horizon no battery losses)	4.11
MPC (shrinking horizon no electrical losses)	4.16

C. Optimality assessment for variations in dual loop MPC parameters

This section considers the effects of varying two fundamental parameters in the dual loop MPC setup, the inner loop horizon length N and the outer loop sampling window size win for a driving cycle with non-zero road gradient. We first compare outer loop window sizes of $win = 10$ and $win = 100$ with the inner loop horizon length fixed at $N = 100$. Figure 7 illustrates that with the smaller value of win , the dual loop solution more closely resembles the shrinking horizon MPC solution and, as can be seen in Table IV, achieves a marginally better average fuel consumption.




Fig. 7. Optimality assessment for varying outer loop sampling window.

Secondly, the outer loop window size is fixed at $win = 100$ and the inner loop horizon length is set to $N = 100$ and $N = 200$, respectively. Figure 8 does not indicate a clear effect of this variation. The most likely explanation is that the overall dual loop performance is mostly dominated by the outer loop

state of energy reference and to a lesser degree by the inner loop horizon length. However, a longer inner loop horizon allows more scope for a short term optimization without an overly stringent terminal condition on the state. This explains the small improvement in average fuel consumption that is obtained with larger N (Table IV).

TABLE IV
MEAN FUEL CONSUMPTION (kg/100 km)

	Gradient: $\pm 2^\circ$
PMP	4.10
MPC (shrinking horizon)	4.12
MPC (dual loop, $N = 100$, $win = 10$)	4.12
MPC (dual loop, $N = 100$, $win = 100$)	4.13
MPC (dual loop, $N = 200$, $win = 100$)	4.12




Fig. 8. Optimality assessment for varying inner loop horizon length N .

D. Optimality assessment in the presence of uncertainty

The dual loop MPC strategy demonstrates a degree of robustness to uncertainty in predicted driver demand, which is investigated in this section. We consider the driving cycle with an upward/downward altitude profile and dual loop MPC with $N = 50$ and $win = 50$. First we consider the effect of introducing discrepancies between the predicted and actual values for the power demand and speed demand. This is done by perturbing the predicted power and predicted speed demand for the outer loop by factors of 0.5 and 1.5 respectively (Fig. 9), while keeping the actual driver power and speed demand unchanged. For the inner loop optimization the actual power and speed demands are assumed to be known. This models the situation in which a large amount of uncertainty is present in the predicted drive cycle for the outer loop MPC (with a horizon of up to 30 min), whereas the predicted power and speed demands for the inner loop (with a horizon of 50 s) are accurate. These perturbation factors are enforced globally by uniformly scaling the predicted power and speed demand sequences. For both types of uncertainty, MPC achieves the target end-point battery energy level to within 0.25%; however discrepancies are found in the corresponding fuel consumption. The fuel consumption for MPC deteriorates from

4.14 kg/100 km to 4.17 kg/100 km for the case of scaled power, and to 4.15 kg/100 km for the case of scaled speed.

The effects on the PMP-based strategy of the same scaling perturbations applied to predicted driving cycles are much more severe (Fig. 10). In each case the costate is optimized for the predicted driving cycle, which includes a scaling factor on either the power or speed demand, and the PMP approach is applied to the actual driving cycle without scaling. Clearly PMP is very sensitive to uncertainty in the initial costate value (Table V).

TABLE V

OPTIMAL COSTATE VALUES FOR PERTURBED DRIVING CYCLES

	Costate
Nominal power and nominal speed	1.64
Scaled power ($\times 1.5$) and nominal speed	1.70
Scaled speed ($\times 0.5$) and nominal power	1.82




Fig. 9. Robustness of dual loop MPC to global scaling perturbations of driver demand and speed.




Fig. 10. Effect of global scaling perturbations of driver demand and speed on the PMP-based strategy.

Figure 9 shows that the battery energy state trajectories for the nominal case and the case of scaled speed demand are highly correlated, however there is a large discrepancy in the state trajectory for the case in which the predicted power

demand is scaled. This is because the power consumption during the uphill phase of the cycle is then lower than what is predicted by the controller, hence less electrical energy is consumed and the battery energy level is generally higher than in the nominal case. As it approaches the mid-point of the simulation the controller determines that additional battery energy can be consumed without violating the terminal constraint on battery stored energy, and after the mid-point of the simulation, the trajectories are again correlated. Both perturbations in the predicted driving cycle result in increased fuel consumption as they push the battery energy state away from the optimal trajectory.

We next consider robustness against final time uncertainty. To model this type of uncertainty the duration of the predicted drive cycle is scaled so that the predicted horizon is of length $T - \Delta T$ while the length of the actual drive cycle remains equal to T . To account for the error ΔT thus introduced in the final time of the predicted driving cycle, the magnitude of the predicted speed demand is scaled to maintain the original total distance (Fig. 11). We also adjust the prediction time at which the road gradient is expected to switch between uphill and downhill to $(T - \Delta T)/2$.




Fig. 11. FTP75 drive cycle scaled for ΔT of 500 seconds.

The scaled drive cycle is simulated using PMP and MPC using two different methods. For the strategy based on PMP, the unscaled drive cycle is used for simulation but the costate is optimised using the scaled drive cycle. For dual-loop MPC, the unscaled drive cycle is also used for simulation, but at each time step the predictions are generated by re-scaling the remaining drive cycle in order to model the effect of improving information on the drive cycle duration towards the end of the journey. Therefore the effective ΔT is constant for PMP, whereas it reduces linearly to zero during the simulation for MPC. This is a fair comparison as feedback is available for MPC but not for PMP, and the error in estimated arrival time would be expected to reduce to zero during a journey (although not necessarily linearly). Both the outer-loop window and inner-loop horizon for the dual-loop MPC implementation are set to 50 s.

Figures 12 and 13 show the battery energy state trajectories for PMP and MPC with the initial ΔT varying from 0 s to




Fig. 12. Robustness of PMP to predicted time of arrival error.




Fig. 13. Robustness of MPC to predicted time of arrival error.

800 s in steps of 100 s. It can be seen that MPC robustly satisfies the inequality constraint (6g) on the final battery energy level for all initial values of ΔT , whereas the PMP method does not meet the constraint on the final state of energy if there is any uncertainty in the arrival time. For MPC, the final battery energy state remains at the lower limit of 50% for all ΔT values between 0 and 500 s, but for $\Delta T > 500$ the initial conservatism in the predicted battery energy state results in the later predictions using the unconstrained optimum solution to problem (10). This is illustrated by the predicted trajectories shown in Figure 14 for $\Delta T = 600$ s: the end-point battery energy levels predicted at time steps $k < 600$ are at 50%, but they rise above 50% for $k \geq 600$.

VI. CONCLUSIONS AND FURTHER WORK

A nonlinear MPC strategy which provides near-optimal energy management for plugin hybrid electric vehicles is presented. The proposed dual loop strategy avoids the offline tuning that is required by popular PMP methods such as ECMS and does not suffer from the suboptimality caused by a predefined reference trajectory for the battery state of energy. This paper demonstrates the advantages of the method in terms of its computational convenience, performance and robustness. Therefore we believe it is of particular practical interest.




Fig. 14. Prediction of state of charge made at several intervals during simulation with an initial ΔT of 600 seconds.

Further work will investigate extensions of the dual loop approach in order to exploit the similarities between PMP with constant costate p and solutions based on MPC. Specifically, a hybrid variant of the dual loop control scheme could determine the optimal costate using the outer loop MPC optimization, and use this to construct a Hamiltonian function which is optimized at each inner loop sampling instant instead of solving an inner loop MPC optimization.

Other future research directions are the inclusion of more general state constraints (in addition to constraints on the initial and final battery state of energy) in the receding horizon optimization problem; consideration of optimal strategies for purely electric driving modes; and incorporation of measures to account explicitly for stochastic uncertainty in the predicted driving cycle and in estimates of the battery state of energy. Obviously, these modifications will make the optimization problem more challenging, but they could significantly improve the performance and applicability of this control strategy.

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] C. 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] S. Di Cairano, D. Bernardini, A. Bemporad, and I.V. Kolmanovsky. Stochastic MPC with learning for driver-predictive vehicle control and its application to HEV energy management. *IEEE Transactions on Control Systems Technology*, 22(3):1018–1031, 2014.
- [10] M. Joševski, A. Katriniok, and D. Abel. Scenario MPC for fuel economy optimization of hybrid electric powertrains on real-world driving cycles. In *Proc. American Control Conference*, pages 5629–5635, 2017.
- [11] 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.
- [12] 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.
- [13] J. Buerger and M. Cannon. Nonlinear MPC for supervisory control of hybrid electric vehicles. In *European Control Conference*, pages 135–140, 2016.
- [14] 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.
- [15] K.M. Bretthauer and B. Shetty. The nonlinear knapsack problem – algorithms and applications. *Eur. J. Oper. Res.*, 138(3):459–472, 2002.
- [16] D. Bertsekas. Projected Newton methods for optimization problems with simple constraints. *SIAM J. Control Optim.*, 20(2):221–246, 1982.
- [17] P.H. Calamai and J.J. Moré. Projected gradient methods for linearly constrained problems. *Math. Prog.*, 39:93–116, 1987.
- [18] R. Fletcher. *Practical Methods of Optimization*. Wiley, 2000.



Johannes Buerger received the M.Eng. degree in Engineering, Economics and Management from the University of Oxford in 2008. In 2013 he obtained the D.Phil. degree in Engineering Science, also from the University of Oxford, for his work on efficient optimization techniques for Model Predictive Control. Since 2013 has worked as a powertrain control engineer at the BMW Group in Munich. His research interests include optimal and robust control, predictive control and optimization. In particular, he is currently

interested in applying these methods in the automotive domain.



Sebastian East obtained the M.Eng. degree in Mechanical Engineering from the University of Bath in 2015 and he is currently a doctoral student in at the University of Oxford. His doctoral research is concerned with model predictive control and optimisation, with applications to energy management problems in hybrid electric vehicles.



Mark Cannon received the M.Eng. degree in Engineering Science in 1993 and the D.Phil. degree in Control Engineering in 1998, both from the University of Oxford. He also received the M.S. degree in Mechanical Engineering in 1995 from Massachusetts Institute of Technology. He is currently Associate Professor of Engineering Science, Oxford University and a Fellow of St Johns College, Oxford. His research interests are in robust and optimal control for constrained and uncertain systems, and optimization, in particular for receding horizon control involving robust constraints and stochastic uncertainty.

particular for receding horizon control involving robust constraints and stochastic uncertainty.