Layered control strategies for hybrid electric vehicles based on optimal control

Domenico Bianchi

Department of Electrical and Information Engineering,
Center of Excellence DEWS,
University of L'Aquila,
Via Campo di Pile – Zona industriale di Pile,
67100 L'Aquila, Italy
E-mail: domenico.bianchi@univaq.it

Luciano Rolando

Dipartimento di Energetica,
Politecnico di Torino,
C.so Duca degli Abruzzi, 24,
10129 Torino, Italy
E-mail: luciano.rolando@polito.it

Lorenzo Serrao* and Simona Onori

Center for Automotive Research,
The Ohio State University,
930 Kinnear Rd., Columbus,
OH 43210, USA
E-mail: serrao.4@osu.edu
E-mail: onori.1@osu.edu
*Corresponding author

Giorgio Rizzoni

Department of Mechanical and Aerospace Engineering,
and
Center for Automotive Research,
The Ohio State University,
930 Kinnear Rd., Columbus,
OH 43210, USA
E-mail: rizzoni.1@osu.edu

Nazar Al-Khayat, Tung-Ming Hsieh and Pengju Kang

Cummins Inc.,
500 Jackson Street, Columbus,
IN 47201, USA
E-mail: nazar.al-khayat@virgin.net
E-mail: tung-ming.hsieh@cummins.com
E-mail: pengju.kang@ge.com

Abstract: Dynamic programming is known to provide the optimal solution to the energy management problem. However, it is not implementable online because it requires complete a-priori knowledge of the driving cycle and high computational requirements. This article presents a methodology to extract an implementable rule-based strategy from the dynamic programming results and thus build a near-optimal controller. The case study discussed in this paper focused on mode switching in a series/parallel hybrid vehicle, in which a clutch may be used to change the powertrain topology. Because of the complexity of the system, the controller is divided in two layers: the supervisory controller, which decides the powertrain configuration, and the energy management, which decides the power split. The process of deriving the rules from the optimal solution is described in detail. Then, the performance of the resulting rule-based strategy is studied and compared with the solution given by dynamic programming, which functions as a benchmark. Then another comparison is performed with respect to the equivalent consumption minimisation strategy (ECMS) which, if optimally tuned, can achieve optimal performance as close to DP as possible with the advantage of being implementable.

Keywords: hybrid electric vehicles; HEVs; energy management; rule-based control; dynamic programming; DP; layered control strategy; rule-based; RB; heuristic; optimal control; supervisory control.

Reference to this paper should be made as follows: Bianchi, D., Rolando, L., Serrao, L., Onori, S., Rizzoni, G., Al-Khayat, N., Hsieh, T-M. and Kang, P. (2011) 'Layered control strategies for hybrid electric vehicles based on optimal control', *Int. J. Electric and Hybrid Vehicles*, Vol. 3, No. 2, pp.191–217.

Biographical notes: Domenico Bianchi received his BS and MS in Computer Science and Automatic Control from the University of L'Aquila (Italy). After a short period in the field of information and communication technologies, he started his PhD programme in 2008 at the Department of Electrical Engineering and Computer Science of the University of L'Aquila. From 2009 to 2010, he was a Visiting Scholar at the Center for Automotive Research (CAR) of the Ohio State University in Columbus, OH (USA). His research interests are vehicle dynamics control, control of hybrid electric vehicles and applications of non-linear control.

Luciano Rolando graduated *Cum Laude* in Automotive Engineering from Politecnico di Torino (Italy) in 2008, and in 2009, he started his PhD programme in Engineering at the same university. In 2010, he spent a full semester at the Center for Automotive Research of the Ohio State University. His research activity mainly focuses on hybrid vehicles and their control strategies. He has been working on several research projects in cooperation with major automotive industries.

Lorenzo Serrao is currently a Researcher at IFP Energies Nouvelles and a Lecturer at IFP School, in Rueil-Malmaison (France). He received his MS in Mechanical Engineering from Politecnico di Torino (Italy) in 2003 and his PhD in Mechanical Engineering from the Ohio State University (OSU) in 2009. During his studies at OSU, he was affiliated with the Center for Automotive Research (CAR). His research interests include energy management of electric and hybrid vehicles, powertrain modelling and simulation, vehicle dynamics and modelling of battery aging.

Simona Onori is a Senior Research Associate at the Ohio State University Center for Automotive Research (CAR). She joined CAR in 2007 as a Postdoctoral Fellow. She received her Laurea Degree in Computer Science Engineering from the University of Rome 'Tor Vergata' (Italy), her MS in Electrical Engineering from the University of New Mexico (Albuquerque, NM, USA) and her PhD in Automation Engineering from the University of Rome 'Tor Vergata' in 2003, 2004 and 2007, respectively. She is an IEEE, ASME and SAE member. She serves as an ASME and IFAC Member Organiser for the Automotive and Transportation Technical Session and the Automotive Systems Session, respectively. Her background is in control system theory. Her research focus is on energy management control for HEV and PHEV, fault diagnosis and prognosis with application to automotive systems, aging and characterisation of advanced batteries.

Giorgio Rizzoni is the Ford Motor Company Chair in ElectroMechanical Systems and a Professor of Mechanical and Electrical Engineering at The Ohio State University. He received his BS, MS and PhD (all in Electrical and Computer Engineering) in 1980, 1982 and 1986 respectively, all from the University of Michigan. Since 1999, he has been the Director of the Ohio State University Center for Automotive Research (CAR), an interdisciplinary university research centre in the College of Engineering. His research interests are in future ground vehicle propulsion systems, including advanced engines, electric and hybrid-electric drivetrains, advanced batteries and fuel cell systems. He is a Fellow of SAE (2005), a Fellow of IEEE (2004), a recipient of the 1991 National Science Foundation Presidential Young Investigator Award, and of several other technical and teaching awards.

Nazar Al-Khayat received his MS and PhD in 1989 and 1993 respectively, from the University of Manchester (UK) and Nottingham Trent University (UK), where he worked as a Research Assistant/Fellow on condition monitoring and in-situ testing of power transformers. After leaving the academia, he worked for Danfoss Flowmetering as an Electromagnetic Design Leader, and in 1998, he joined Cummins Generator Technologies as a Chief Engineer R&D, with responsibility for advanced electrical machines and power electronics technologies. From 2000 to 2007, he led new product introduction and electronics group advancing new PM machines technologies and synchronous machines regulators. From 2007 till 2010, he moved to the USA as the Director of Advanced Powertrain at Cummins Inc., Columbus, IN, USA, to lead Cummins' effort on hybrid technologies for commercial truck and bus applications. He is a Chartered Engineer member of IET.

Tung-Ming Hsieh is currently a Technical Advisor at Cummins Inc. He received his BS and MS in Mechanical Engineering from Cheng-Kung University in Tainan, Taiwan, and PhD from the University of Wisconsin at Madison, USA. His current interests include CVT, hybrid electric vehicles and power system optimisation.

Pengju Kang is a Principle Engineer with the Electronics Systems and Controls Organization of GE Global Research. He is driving the technology development programmes in four areas: electric drive train, model-based design of safety-critical control systems, smart grid and avionics. Before joining GE, he was with the Advanced Power Systems, Cummins working on hybrid powertrain modelling, simulation and control. Before Cummins, he worked at UTC Research Center for five years and UTC Sikorsky Aircraft for three years. He received his Bachelor's and Master's from the Huazhong University of Science and Technology in 1987 and 1990 respectively, PhD from Queensland University of Technology in 2001 and MBA from Rensselaer Polytechnic Institute in 2005.

1 Introduction

Hybrid electric vehicles (HEVs) represent a powerful technology to save fuel and reduce CO_2 emissions, through the synergic use of a traditional internal combustion engine and one or more electric machines, powered by a rechargeable energy storage device (e.g., battery or supercapacitors). HEV performance strongly depends on the control strategy that shares the power demand among the engine and the electric motor(s) at each time instant, with the objective of minimising the fuel consumption over an entire driving cycle, while also maintaining the state of charge of the energy storage device in the allowable range.

The energy management strategies for HEVs can be classified according to multiple criteria. For the purposes of this paper, it is useful to consider the amount of information used and the optimisation method. On the basis of these two metrics, the following three categories can be identified: global optimisation methods, local optimisation methods, and heuristic strategies. Global optimisation methods find the optimal solution, i.e., the optimal sequence of power split, by performing a global minimisation over the entire driving cycle, which is assumed to be completely known in advance. The most common of the methods in this category is dynamic programming (DP) (Brahma et al., 2000; Sundström et al., 2009), which is not implementable online, for the need of a-priori knowledge of the driving cycle and the elevated computational requirements due to the global minimisation. Still, it can be used as a benchmark for other strategies, or as a way to assess the best possible results that a given powertrain design can provide. Local optimisation methods do not assume explicitly the knowledge of the future, relying only on information that would be available on board. The solution is generated as a sequence of local (instantaneous) minima, found at each instant. One of these methods is called equivalent consumption minimisation strategy (ECMS), being based on the conversion of the electric power into equivalent fuel consumption (Paganelli et al., 2001; Musardo et al., 2005; Serrao et al., 2009). ECMS has been proved (Sciarretta and Guzzella, 2007; Serrao et al., 2009) to give results close to the optimum given by DP, but not exactly the same, due to the lack of complete knowledge of the cycle. Local minimisation methods are implementable online, even if they tend to have high computational requirements because of the minimisation which is performed at each instant. Heuristic strategies do not perform any minimisation: instead, the control action is determined at each instant using empirical rules (He et al., 2005; Hofman et al., 2007), fuzzy logic (Salman et al., 2000; Won and Langari, 2005) and neural networks (Suzuki

et al., 2008). The advantage of heuristic strategies is the reduced computational time with respect to the minimisation techniques; the disadvantage is the fact that there is no guarantee that the solution thus generated is optimal, and the fact that they typically include numerous tuning parameters, that must be appropriately chosen in order to achieve good performance.

The strategy presented in this paper belongs to the third family: it is a rule-based (RB) strategy, thus easily implementable and with low computational requirements. However, it is derived from observation of the optimal solution obtained with DP, following a method initially proposed in Lin et al. (2003). In addition to the fact that deriving the rules in this way is a relatively fast process, this method allows to obtain results close to the optimal solution, and to reduce the number of tuning parameters. The motivation for a RB strategy is to derive a quasi-optimal solution that requires the lowest possible computational requirements, even lower than methods based on instantaneous minimisation, which are quite heavy for the application examined. The RB approach, being based on a relatively simple set of rules, does not involve minimisation or complex table lookup and therefore is very fast computationally.

The paper is organised as follows: the vehicle architecture is described in Section 2, the control problem is formulated in Section 3, and the model of the HEV is presented in Section 4. DP and ECMS solutions are outlined respectively in Sections 5 and 6. Then, Section 7 provides a description of the method used to extract the rules and Section 8 shows the simulation results for various driving cycles, comparing the resulting RB strategy to the optimal solution obtained by DP and to the ECMS.

2 Powertrain architecture

HEVs are commonly classified on the basis of their powertrain architecture, i.e., the way in which the various powertrain components are arranged and the ratio of maximum electrical to mechanical (engine) power on board the vehicle. Traditionally, two main categories have been distinguished: *series* and *parallel* configurations. In series HEVs, at least two electric machines are present: a motor and a generator. The motor is the only mean of providing power to the wheels; it receives electric power from either the battery pack or from the generator, run by an internal combustion engine. Thus, the battery and engine power are summed electrically. In parallel HEVs, on the other hand, mechanical gearings allow the engine and one or more electric machines to drive the wheels; in other words, the power of the engine and the electric machine(s) is summed mechanically at the transmission level.

The architecture considered in this article is envisioned for commercial trucks, and represents a more flexible solution than either series or parallel, with more degrees of freedom and greater opportunity for fuel economy. It is composed of two electric machines, motor and generator (called in the following *EM* and *GEN*), a diesel engine (*ICE*), a battery pack (*batt*), a clutch and an automatic gearbox (*gb*). The architecture can be described as follows. The internal combustion engine is directly connected to the generator GEN on the same shaft, while the motor is connected to the gearbox primary (input) shaft. A clutch allows to connect and disconnect the engine/generator and the motor, effectively making the three machines act on the same shaft. When the clutch is open, the vehicle behaves as a series HEV, since the engine velocity is independent from the vehicle speed and only the electric motor is able to provide torque to the wheel. When

the clutch is locked, the engine and both electric machines are connected to the powertrain, and the vehicle behaves like a parallel HEV, summing the torque of the three machines. In this configuration, the engine can drive the vehicle on its own or assisted by the electric motor. The battery is connected to both electric machines, and can be recharged using the generator while the vehicle is in series mode, or using the motor during regenerative braking.

This particular kind of hybrid powertrain is usually defined as combined series-parallel. The vehicle can operate in an all electric mode, a series hybrid configuration or a parallel hybrid configuration depending upon which is most advantageous given operating requirements. As a result of this dual drivetrain, the engine operates near optimum efficiency more often. For example, it can operate as a series vehicle at lower vehicle speed, when the low load would make inefficient to have the engine drive the wheels directly. At high vehicle speed, instead, where the engine can operate with better efficiency, the series drivetrain is less efficient for the double energy conversion; this architecture allows the engine to take over and minimise the energy losses. The increased flexibility comes at a higher initial cost than a pure parallel hybrid, for the presence of a generator, a larger battery pack, and more computing power to control the dual system. However, the series/parallel hybrid has the potential to perform better than either of the systems alone.

With the combined series-parallel architecture, the HEV controller must decide how to share the total power demand between the available machines and command the clutch to switch between various modes of operation. The formulation of the control problem is provided in Section 3, and an optimal solution is found using DP as illustrated in Section 5. However, since this algorithm requires complete knowledge of the driving cycle in advance and is extremely requiring in terms of computational loads, we illustrate in Section 7 a method to derive an implementable controller based on its results.

3 Control problem formulation

The optimal control problem in a HEV consists in finding the minimum fuel consumption during vehicle operation, while respecting the design limitations of each component and the drivability/performance specifications. The aim is to minimise a cost function defined as an integral over a finite horizon. The finite horizon typically corresponds to a complete regulatory driving cycle or a short real-world trip. The optimisation objective considered in this work is the fuel consumption during a trip and the constraints are:

- charge-sustainability: the battery SOC at the beginning and the end of the trip should be equal
- drivability constraints: at each instant, the total torque output of the powertrain should be equal to the driver's demand
- actuator limitations: at each instant, the output of each machine in the powertrain (engine, motor, and generator) cannot exceed its maximum torque/power rating; similarly, the total battery power must remain within the acceptable limits in both charge and discharge operation.

In this vehicle architecture there are several control variables: the status of the clutch (open or locked) and the status of the engine (on or off) are discrete control variables that determine the *operating mode* of the powertrain; the torques of the individual machines are continuous variables and determine how the power request is shared between components. Formally, the clutch status is represented as:

$$C = \begin{cases} 0 & \text{clutch open} \\ 1 & \text{clutch locked} \end{cases} \tag{1}$$

and, similarly, the engine status is given by:

$$E = \begin{cases} 0 & \text{engine off} \\ 1 & \text{engine on} \end{cases}. \tag{2}$$

The power split is defined by the values of the torques delivered by the two electric machines, T_{EM} (motor torque) and T_{GEN} (generator torque). The four input variables are gathered in the following control vector, defined for each time step k:

$$u_k = \left\{ T_{EM,k}, T_{GEN,k}, C_k, E_k \right\} \tag{3}$$

Figure 1 Flowchart of the powertrain controller

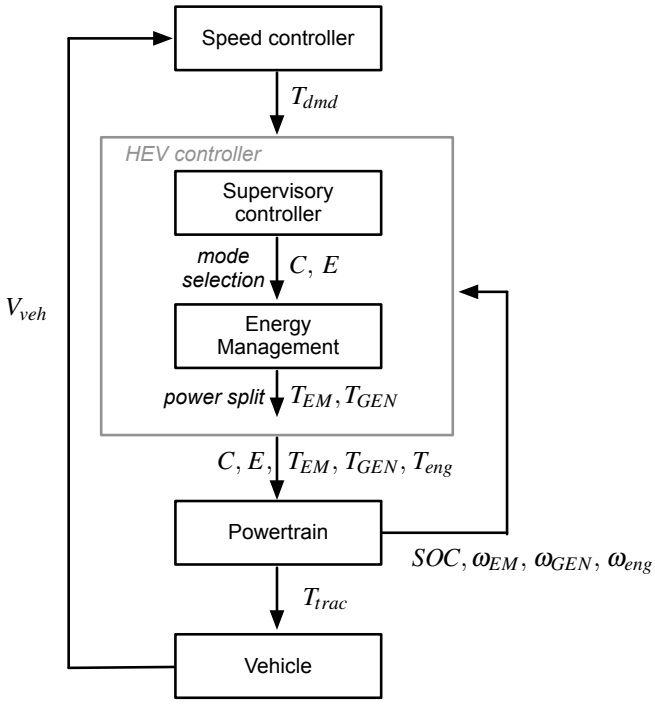

The controller of the hybrid powertrain is divided in two layers, as shown in Figure 1: the operating mode, i.e., the value of the variables C and E, is determined at the *supervisory controller* level, while T_{EM} and T_{GEN} are determined at the *energy management* level, adhering to the constraints on powertrain operation. The remaining degree of freedom of the powertrain, i.e., the transmission gear index g_{tr}, is chosen by the transmission controller, which is assumed to be external to the energy management and supervisory controller, and embedded in the transmission; therefore, the gear index is treated as an external input in this context. The vehicle velocity, the rotational speed of the two electrical machines and the engine, and the driver's torque demand are also external inputs.

The problem is formally defined as finding the control law u_k, $k = 1,\ldots,N$ that minimises the cost:

$$J = \sum_{k=0}^{N-1} m_{f,k}\left(u_k, k\right), \tag{4}$$

where $m_{f,k}$ is the mass of fuel consumed during the time step k, subject to the constraints:

$$0 \le P_{ICE,k} \le P_{ICE,\max} \quad \forall k = 0,1,\ldots,N-1 \tag{5}$$

$$P_{EM,\min} \le P_{EM,k} \le P_{EM,\max} \quad \forall k = 0,1,\ldots,N-1 \tag{6}$$

$$P_{GEN,\min} \le P_{GEN,k} \le P_{GEN,\max} \quad \forall k = 0,1,\ldots,N-1 \tag{7}$$

$$P_{batt,\min} \le P_{batt,k} \le P_{batt,\max} \quad \forall k = 0,1,\ldots,N-1 \tag{8}$$

$$SOC_{\min} \le SOC_k \le SOC_{\max} \quad \forall k = 0,1,\ldots,N-1 \tag{9}$$

where P_{ICE} is the engine mechanical power, P_{EM} is the motor electrical power, P_{GEN} is the generator electrical power, P_{batt} is the battery power. The subscripts 'max' and 'min' refer to the maximum and minimum limits of each variable. An additional constraint is the dynamic equation of the state of charge, described in Section 4.

4 Model of the hybrid electric vehicle

Two different approaches to the HEV modelling can be adopted: *backwards* or *forward* (with respect to the physical causality principles) (Pisu et al., 2005). In the forward approach the vehicle speed is a consequence of a torque delivered by the powertrain, in response to the demand generated by the driver model (usually a PID controller that compares the actual velocity with the desired value). In the backward approach, instead, no driver is necessary, since the vehicle speed is supposed known and the torque necessary to obtain it is computed by the model, as shown in Figure 2.

Figure 2 Information flow in a backward simulator

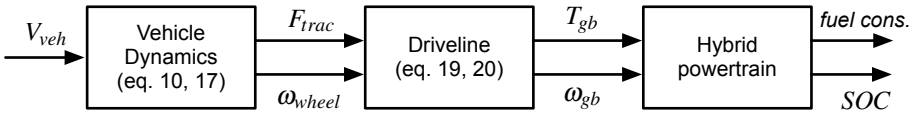

Source: Pisu et al. (2005)

In this paper, a backward, quasi static simulator is used to implement the DP algorithm, because it allows to treat the vehicle speed as an external input rather than a dynamic state. As Figure 2 shows, the vehicle speed, defined by the driving cycle, is used to calculate the tractive force; through the driveline model, the torque request upstream of the gearbox is computed. Then, through the powertrain model, both fuel consumption and battery SOC are calculated. Starting from the driving cycle inputs, it is possible to calculate the tractive force at the wheels as:

$$F_{trac} = F_{inertia} + F_{roll} + F_{aero} + F_{grade} \tag{10}$$

where $F_{inertia}$ is the inertia force, F_{roll} is the rolling resistance, F_{aero} the aerodynamic resistance, F_{grade} the force due to road slope, F_{trac} is the tractive force generated by the powertrain at wheels. Each term in the above equation is computed as follows:

- Rolling resistance

$$F_{roll} = c_{roll} \cdot M_{veh} \cdot g \cdot \cos\alpha \tag{11}$$

where g is the gravity acceleration, α is the road slope (so that $M_{veh} \cdot g \cdot \cos\alpha$ is the vertical component of the vehicle weight) and c_{roll} is the rolling resistance coefficient. In general, c_{roll} is a function of vehicle speed, tire pressure, external temperature, etc. In most cases, it can be assumed to be constant, or to be an affine function of the vehicle speed. The order of magnitude is 0.01 to 0.02, which means that the rolling resistance is roughly 1% to 2% of the vehicle weight. In this work, the approximation used is:

$$c_{roll} = c_0 + c_1 \cdot v_{veh} + c_2 \cdot v_{veh}^2 + c_3 \cdot v_{veh}^3 \tag{12}$$

where c are constant coefficients that link c_{roll} with the speed.

- Aerodynamic resistance

$$F_{aero} = \frac{1}{2} \cdot \rho_{air} \cdot A_f C_d \cdot v_{veh}^2 \tag{13}$$

where ρ_{air} is the air density, A_f the vehicle frontal area, C_d the aerodynamic drag coefficient.

- Road slope

$$F_{grade} = M_{veh} \cdot g \sin\alpha \tag{14}$$

where α is the slope of the road.

- Inertia

$$F_{inertia} = M_{eq} \cdot \frac{dv_{veh}}{dt} \tag{15}$$

where M_{eq} is the vehicle equivalent mass, which takes into account the inertia of all the driveline components and can be expressed as:

$$M_{eq} = M_{veh} + J_{wheel} \cdot \frac{1}{r^2} + J_{ICE} \cdot \frac{\tau_{gear}^2 \cdot \tau_{fd}^2}{r^2} \tag{16}$$

with J_{wheel} the wheel inertia, J_{ICE} the engine inertia, r the wheel radius, τ_{gear} the transmission gear ratio and τ_{fd} is the final drive gear ratio.

From the vehicle speed it is immediate to calculate the wheel angular velocity, with a quasi-static model neglecting wheel slip:

$$\omega_{wheel} = \frac{V_{veh}}{r}. \tag{17}$$

The tractive force F_{trac} and the wheel speed ω_{wheel} are used in the powertrain model to compute the power requirement and the energy consumption, as shown in Figure 2. Since the model is quasi-static, all the dynamics are neglected and each device is represented through stationary maps experimentally measured. The time step of the simulation is 1 s. Under this assumption, all the clutch transient phases are neglected: switching from series to parallel mode and vice versa is done instantaneously without any slip state. The engine cranking phase is also neglected, assuming that the engine can be turned on and off immediately (i.e., within one time step).

From F_{trac} and ω_{wheel}, with transmission efficiencies and ratios, it is possible to compute both the power and torque request upstream of the gearbox, as well as the speed of the gearbox input shaft:

$$P_{gb} = \frac{F_{trac} \cdot r \cdot \omega_{wheel}}{\eta_{gear} \cdot \eta_{fd}} \tag{18}$$

$$\omega_{gb} = \omega_{wheel} \cdot \tau_{fd} \cdot \tau_{gear}; \tag{19}$$

$$T_{gb} = \frac{F_{trac} \cdot r}{\tau_{fd} \cdot \tau_{gear}}. \tag{20}$$

where η_{gear} and η_{fd} are respectly the efficiency of the gearbox and the final drive.

The transmission ratio is assumed to be computed outside of the HEV powertrain controller and is considered as a parameter. Thus, the torque, speed and power at the gearbox input depend only on the driving cycle and represent the request that the hybrid propulsion system must satisfy.

The way the power request is fulfilled depends on the operating mode of the vehicle.

If the system is in the series configuration, both the engine and the generator are isolated from the gearbox since the clutch is open. The electric motor operating point is set by the gearbox since there are no additional machines that can transmit torque to the wheels. The engine torque can be written as:

$$T_{ICE} = -T_{GEN} + T_{aux,m} \tag{21}$$

where $T_{aux,m}$ is the torque requested by mechanically-driven auxiliary loads and T_{GEN} is the generator torque. The engine speed can be chosen by the energy management strategy. In particular, it was decided to keep the engine working on its best efficiency line, represented in Figure 3, in order to minimise the fuel consumption:

$$\omega_{ICE} = \omega_{ICE,opt}\left(T_{ICE}\right). \tag{22}$$

Figure 3 Engine maximum efficiency line

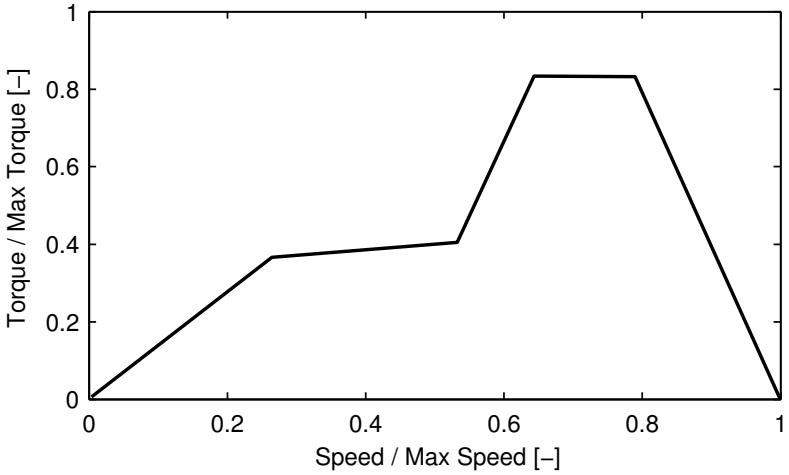

On the other hand, when the hybrid powertrain works in parallel mode, all the machines are connected to the gearbox and balance the resisting torque. In this case, the speed of the machines is determined by the external conditions, while the controller must set two of the three torques (the third is then defined by difference with the total). Using the electric machines torques as the control variables of the energy management strategy, the engine torque can be written as:

$$T_{ICE} = T_{gb} + T_{aux,m} - T_{EM} - T_{GEN}, \tag{23}$$

while the speed of all the machines is identical:

$$\omega_{gb} = \omega_{ice} = \omega_{EM} = \omega_{GEN}. \tag{24}$$

At this point, for each powertrain configuration, the operating points of all the machines are available and can be used as inputs of the efficiency maps of each component. The engine map is used to calculate the fuel consumption $\dot{m}_f$, while the electric machine maps are used to compute the electric power demand given their torque and speed.

Another significant variable is the battery power P_{batt}, which can be easily computed as:

$$P_{batt} = P_{EM,e} + P_{GEN,e} + P_{aux,e} \tag{25}$$

where $P_{EM,e}$ and $P_{GEN,e}$ are the electric power of the motor and the generator and $P_{aux,e}$ the power of the electrical auxiliary loads. The battery is normally very complex to represent. In this case, in order to obtain a reasonably simple simulator, no temperature dependency is considered, hysteresis and dynamics are neglected, and a simple circuit model (represented in Figure 4) is used to compute the state of charge variation as a function of the power at the terminals and of the circuit parameters.

Figure 4 Battery circuit model

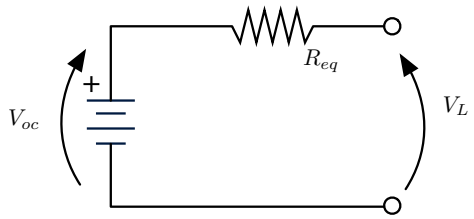

The state of charge variation represents the state equation of the energy management problem, and can be written as:

$$SOC(k+1) = SOC(k) - \alpha \cdot \frac{I(k)}{Q_{\max}} \tag{26}$$

where $I(k)$ is the current flowing through the battery (positive during discharge), $Q_{\max}$ is the battery charge capacity, and α is a correction factor that accounts for the charge losses (coulombic efficiency). In order to make this relation implementable in the framework of the model described here, it is necessary to express the current in terms of the battery power. The first step is to write the balanced equation for this equivalent circuit:

$$V_L = V_{oc} - I \cdot R_{eq}, \tag{27}$$

where $V_{oc} = V_{oc}(SOC)$ is the open circuit voltage, V_L the voltage at the battery terminals and R_{eq} is the equivalent internal resistance, composed by two terms (Szumanowski, 2000):

$$R_{eq} = R_{ohm}(SOC) + R_{pol}\left(V_{oc}\right). \tag{28}$$

R_{ohm} takes into account the ohmic losses in the accumulator and R_{pol} represents the resistance of polarisation, which can be expressed as a function of the open circuit voltage and the terminal voltage as described in Szumanowski (2000) and represented in the following equation:

$$R_{pol} = b\left(V_{oc}\right) \cdot \frac{V_{oc}(SOC)}{I}. \tag{29}$$

Here, $b(V_{oc})$ is a coefficient depending on the open circuit voltage (Szumanowski, 2000). At this point, multiplying each side of the equation (27) by the terminal current it is possible to find the battery power:

$$P_{batt} = V_L \cdot I = \left(1 - b\left(V_{oc}\right)\right) \cdot V_{oc} \cdot I + R_{ohm} \cdot I^2 \tag{30}$$

Solving equation 30 with respect to I yields:

$$I = \frac{(1-b) \cdot V_{oc} + \sqrt{\left[(1-b) \cdot V_{oc}\right]^2 - 4 \cdot P_{batt} \cdot R_{eq}}}{2 \cdot R_{eq}} \tag{31}$$

with $I > 0$ when discharging and $I < 0$ charging.

The battery power in (31) can be expressed in terms of the torque/speed of the electric machines using (25), and is given by:

$$P_{batt} = P_{aux,e} + \frac{T_{EM} \cdot \omega_{EM}}{sign\left(T_{EM}\right)} + \frac{E \cdot T_{GEN}\left[C \cdot \omega_{EM} + (1-C) \cdot \omega_{ice}\left(T_{GEN}\right)\right]}{\eta_{GEN}^{sign(T_{GEN})}} \tag{32}$$

where η_{EM} and η_{GEN} are the electric machines efficiencies, C is the clutch status ($C = 1$ for clutch closed, $C = 0$ for clutch open), and E is the engine status ($E = 1$ engine on, $E = 0$ engine off).

These are all the physical relations used in the vehicle model and they were implemented using a Matlab function.

5 Dynamic programming

DP generates a numerical solution to the optimal control problem defined in Section 3. In other words, it gives sufficient conditions for the global optimality. DP is based on Bellman's principle of optimality:

> "An optimal control policy has the property that no matter what the previous decision (i.e., controls) have been, the remaining decisions must constitute an optimal policy with regard to the state resulting from those previous decisions." (Lewis and Syrmos, 1995)

DP is capable of determining the optimal solution to the discretised problem. This solution is generally sub-optimal for the continuous problem, because of the approximation introduced with the discretisation; however, if the grid is fine enough, the approximation is negligible. The need for a backward procedure means that the solution can be obtained only off-line, for a driving cycle known a-priori, and therefore it is not possible to use DP for an online implementable solution; furthermore, the high computational load makes any DP optimisation prohibitive on typical onboard micro controllers.

To implement the DP algorithm on the hybrid architecture described, an open-source Matlab code developed at ETH-Zurich (Sundström and Guzzella, 2009) was exploited.

The DP solution is computed for several driving cycles, representative of the range of operating conditions for the vehicle considered, and provides the optimal combination of control inputs for each driving cycle. Its utilisation is represented schematically in Figure 5.

Figure 5 Flowchart of the DP controller

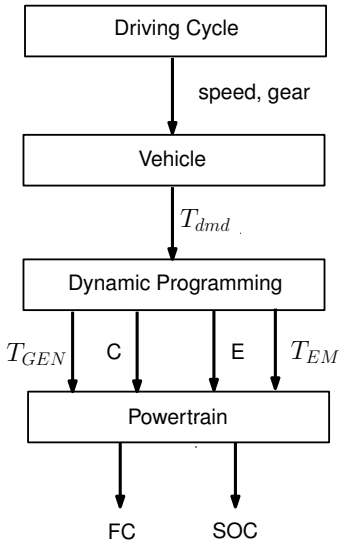

6 Equivalent consumption minimisation strategy

The equivalent consumption minimisation strategy was introduced by Paganelli et al. (2001) as a method to reduce the global minimisation problem to an instantaneous minimisation problem to be solved at each instant, without use of information regarding the future. It was later shown (Sciarretta and Guzzella, 2007; Serrao et al., 2009) that this optimisation algorithm is equivalent to the Pontryagin's minimum principle and it is able to ensure an optimal for this kind of problems. This strategy is based on the concept that, in charge-sustaining vehicles, the difference between the initial and final state of charge of the battery is negligible with respect to the total energy used. This means that the electrical energy storage is used only as an energy buffer. Since all the energy ultimately comes from fuel, the battery can be seen as an auxiliary, reversible fuel tank. The electricity used during a battery discharge phase must be replenished at a later phase using the fuel from the engine. Two cases are possible at a given operating point:

- the battery power is positive (discharge case): a recharge with the engine will require some additional fuel consumption in the future
- the battery power is negative (charge case): the stored electrical energy will be used to reduce the engine load, which implies a fuel saving.

In both cases, an equivalent fuel consumption can be associated with the use of electrical energy; the equivalent future fuel consumption can be summed to the present real fuel consumption to obtain the instantaneous equivalent fuel consumption (Serrao et al., 2009):

$$\dot{m}_{eqv} = \dot{m}_f + \dot{m}_{batt} = \dot{m}_f + s\frac{E_{batt}}{Q_{lhv}} f\left(SOC, P_{batt}\right) \tag{33}$$

where Q_{lhv} is the fuel lower heating value (energy content per unit of mass), $\dot{m}_{batt}$ is the virtual fuel consumption associated with the use of the electrical rechargeable energy storage system, E_{batt} the total battery energy capacity, and $f(SOC, P_{batt})$ is the derivative of the state of charge, described in Section 4.

The term s is called equivalence factor and is used to convert electrical power into equivalent fuel consumption; it plays an important role in the ECMS. Depending on the sign of P_{batt} (i.e., on whether the battery is charged or discharged), the virtual fuel flow rate can be either positive or negative, therefore the equivalent fuel consumption can be higher or lower than the actual fuel consumption. The instantaneous minimisation problem is computationally less demanding than the global problem solved with DP, and applicable to real-world situations since it does not rely (explicitly) on information about future driving conditions.

The concept of equivalent fuel consumption is tied with the necessity of attributing a meaningful value to the equivalence parameter s. This parameter is representative of future efficiency of the engine and the energy storage device, and its value affects both the charge sustainability and the effectiveness of the strategy: if it is too high, an excessive cost is attributed to the use of electrical energy and therefore the full hybridisation potential is not realised; if it is too low, the opposite happens and the battery is depleted too soon (loss of charge sustainability). In practical applications, it is necessary to adapt online the value of s using a feedback controller to increase stability and robustness, as proposed for example by Chasse et al. (2010) and Onori et al. (2010).

7 RB strategy

The control based on a set of empirical rules is computationally efficient for an embedded CPU, but it may generate results quite far from the optimality. Its calibration, in addition, could be quite difficult. The DP, on the contrary, provides the optimal solution on each driving cycle. Therefore, analysing its control actions, some rules can extracted that try to reproduce the optimal behaviour, and, unlike DP control signals, are implementable. This approach is known (see, e.g., Lin et al., 2003), but it is now applied to a complex architecture in which the DP is used to determine not only the hybrid power split, but also the vehicle operating mode.

The starting point for deriving a RB strategy from DP is an extensive set of simulation in which the optimal driving strategy is found for several driving cycles, covering an ample range of urban and suburban driving conditions. The results are then studied and analysed in order to find common patterns in the algorithm decisions that are then replicated by appropriate rules. The control results are represented in order to emphasise any dependency from significant input variables, such as gearbox power P_{gb}, gearbox speed ω_{gb}, and battery state of charge SOC.

As mentioned in Section 3, the powertrain controller can be divided in two parts: the *supervisory control* which decides the best operating mode and the *energy management* which shares the torque among the machines in order to satisfy the overall demand. Therefore the analysis of the DP results has to be performed at two levels: mode selection (engine and clutch status), and torque split. The extraction of rules for each of these two levels is described the following sections.

Figure 6 Hybrid mode as a function of gearbox input torque and speed

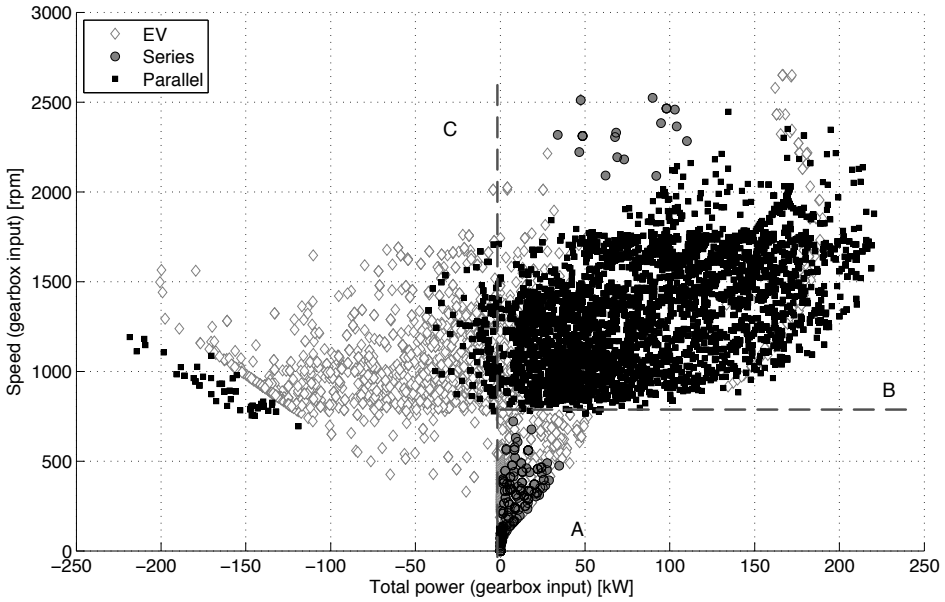

7.1 Supervisory control

To understand the behaviour of the supervisory control, the operating mode chosen by DP over all the analysed driving cycles was plotted as function of the gearbox input torque and speed, as shown in Figure 6. This plot is divided into three main areas:

A At low speed and torque, the powertrain works either in series or in pure electric (EV) mode – i.e., the clutch is open ($C = 0$) and the engine is either on or off ($E = 1$ or $E = 0$).

B This area is limited by engine idle speed and positive gearbox torque: here only the parallel configuration is present. There are also some instances of series operation, but they are only used to avoid the engine speed to exceed its maximum.

C The third area includes all the points with a negative torque: here the supervisory always switches off the engine in order to save fuel since the vehicle is decelerating. Actually in this situation the engine can also be on, but fuel injection is cut off, which in the quasi-static model is treated as engine-off.

As determined from the distribution shown in Figure 6, nevertheless, it is not possible to infer any dependence on the state of charge. In fact, it is not possible to find any clear correlation between the state of charge and the mode selection in region A. Therefore, a simple SOC threshold was established to differentiate between EV and series mode at lower speed.

The supervisory control rules are therefore implemented as follows:

A When the gearbox input torque is positive ($T_{gb} \geq 0$) and the gearbox input speed is below engine idle speed ($\omega_{gb} \leq \omega_{idle}$), the clutch is open ($C = 0$); the status of the engine is determined by the SOC value, as explained in Section 7.2.2. The vehicle is either in EV or series mode. The speed condition is not a result of the optimisation, but a physical constraint of the powertrain: the engine cannot be connected to the driveline unless the speed is above idle.

B When $T_{gb} \geq 0$ and $\omega_{gb} > \omega_{idle}$, the clutch is locked ($C = 1$) and the engine is on ($E = 1$), i.e., the vehicle is in parallel mode.

C When $T_{gb} < 0$, i.e., during regeneration, the clutch remains engaged ($C = 1$), and the engine maintains its previous state, but fuel is cut off.

7.2 Energy management

Depending on the powertrain mode decided by the supervisory controller, the power split among the two electrical machines and engine is determined in different ways.

7.2.1 Parallel mode

All the machines can directly act on the gearbox input shaft to overcome the resistance torque given by the vehicle. The energy management decides what fraction of the torque is generated by the electric machines and by the engine. The data analysis shows a linear relation between the gearbox input torque T_{gb} and the sum of the electric machine torque T_{elec}. Then, from a simple torque balance, it is possible to compute the fraction given by the engine. In order to split T_{elec} between the two machines, each of the torques computed by DP is related to T_{elec}. As a result, two linear correlations can be observed again. An example of these laws defining the parallel torque split is represented in Figure 7. In the controller implementation only one of these two will be used, while the third torque will be obtained by difference from the others.

Figure 7 Dependence of the electric machine torque from the total electric request

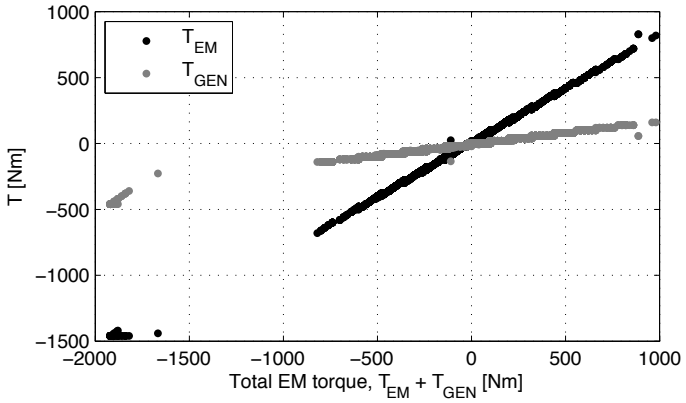

This fact allows to match the request in all cases (which would not be possible if all three torques were computed independently). The generator torque is chosen as the computed variable because preliminary tests showed better performance in comparison with the dual case in which T_{EM} was the one computed by regression.

7.2.2 Series mode

In this configuration, an approach based on the torque balance is not suitable because the engine and the generator are disconnected from the gearbox, thus the entire torque request must be satisfied by the motor. On the other hand, the electric power demand should be split between the generator and the battery. The total electric power is the sum of the motor electric power and the power demanded by the electrical accessories. Following an approach similar to the torque split, the total power request can be correlated with the battery power computed by DP, as shown in Figure 8. The power fraction that is not supplied by the battery is provided by the motor-generator group.

Figure 8 Dependence of the battery power from the total electric power

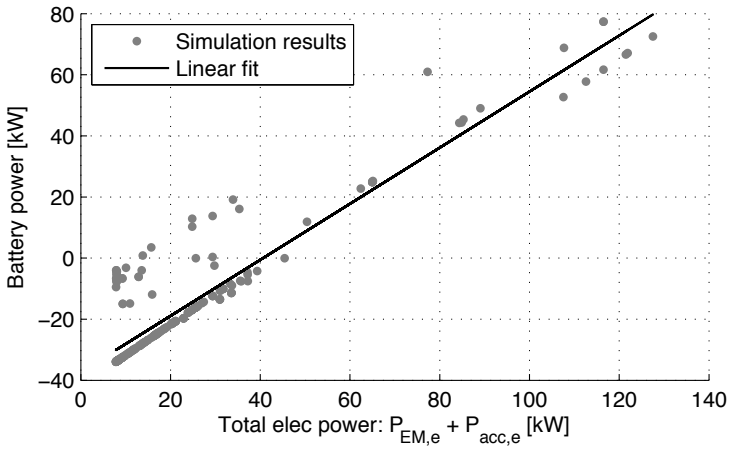

The engine torque T_{ICE} is computed by imposing a condition in which the engine operates at maximum efficiency at a given output power. Hence, engine speed will be an output rather than an input parameter.

7.3 State of charge control

The battery state of charge should affect both the supervisory decisions and the energy management. However, the effect of SOC is not present in all the empirical rules derived from DP and presented in the previous sections. Therefore, these laws need to be modified to achieve charge-sustainability. One simple way to proceed is to shift up or down the linear laws that compute the electrical loads both in parallel mode (as described in Section 7.2.1) and in series mode (as described in Section 7.2.2). To reach this target, a correction function is introduced in the linear correlations, using an additional coefficient $p(SOC)$ that multiplies the intercept of the regression lines. It is now necessary to choose

the shape of this correction function $p(SOC)$. The correction has to be minor for small deviation from the reference state of charge SOC_{ref}, and increase smoothly when the correction needs to be stronger. A cubic polynomial is a suitable function for this purpose; the correction function is thus defined as:

$$p(SOC) = -\mu \cdot x_{SOC}^3 + 1 \tag{34}$$

where x_{SOC} measures the distance of the state of charge from the reference value:

$$x_{SOC} = \frac{SOC - SOC_{ref}}{\left(SOC_{\max} + SOC_{\min}\right)/2}. \tag{35}$$

The parameter μ in (34) defines the amount of correction for the achievement of the charge sustaining condition; a higher value of μ makes the solution more robust, penalising more the variation of state of charge, but also introduces some deviation from the optimal solution.

Figure 9 Effect of the parameter μ on the penalty function

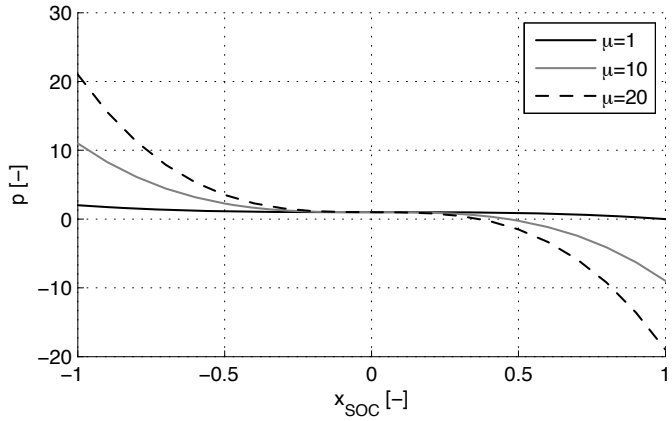

It was observed that μ depends on the driving conditions. The lowest value of μ for which the strategy achieves the charge balance is different for each cycle, and guaranteeing robustness under all conditions requires a compromise in terms of performance. This is one of the main drawbacks of a RB approach: this type of energy management requires a careful calibration of all its parameters, which have to be defined by a compromise among different driving cycles.

8 Simulation results

The rules described in Section 7 are implemented on the simplified quasi-static model described in Section 4. The resulting RB strategy is compared to the DP solution in Section 8.1, and the effect of the calibration parameter μ on the optimality of the solution is investigated in Section 8.2.

However, since the RB strategy is developed to be implemented online, a second set of simulations is run, this time using a more accurate simulator based on a forward approach, which takes into account most of the dynamics of the powertrain components. DP cannot be implemented using this simulator, because the presence of many dynamic states would make it excessively heavy computationally. Therefore, this time the reference strategy is provided by the ECMS, which has been shown to give the optimal solution when optimally tuned offline (Serrao et al., 2009). The comparison between RB and optimal ECMS is presented in Section 8.3.

The main differences between the models are summarised in Table 1.

Table 1 Comparison between model features

	Simplified model	*Complete simulator*
Modelling approach	Backwards	Forward
Vehicle dynamics	Lumped vehicle and powertrain inertia	Lumped vehicle and powertrain inertia
Engine model	Stationary fuel cons. map	First order dynamics + fuel cons. map
Engine starter	Instantaneous poweron	Electrical cranking
Electric machines	Stationary map	Stationary map
Battery model	Circuit model	Circuit model
Clutch dynamics	Instantaneous engagement	Slip dynamics

8.1 RB vs. DP comparison

The Manhattan bus cycle, representative of real-world driving conditions for an urban bus, is here considered to illustrate the performance of RB vs. DP. The parameter μ is tuned to the optimum value as described in Section 7.3. Figure 10 shows both the overall SOC profile obtained by the two strategies and the choices of the supervisory controllers. The SOC profiles are close to each other in terms of shape, which means that the power split is similar (the offset is due mainly to punctual differences), and the choices of the supervisory controller (i.e., the vehicle mode in the bottom plot) are also very similar, as expected.

Figure 11 shows the comparison of the torque split. The differences in the electric motor torque (T_{EM}) are quite small, while the generator behaviour shows that the RB tends to generate more energy to be stored in the battery. Figure 11 also shows that the torque of the electric motor obtained with the RB strategy is smoother in comparison with the DP, which is more 'nervous'. These discrepancies in the torque may also determine differences in state of charge profile and in fuel consumption. As Figure 10 demonstrates, the SOC is quite similar but the RB is characterised by a higher mean value which leads to a higher final SOC.

Figure 10 Comparison RB-DP: state of charge profile and mode selection on the Manhattan cycle

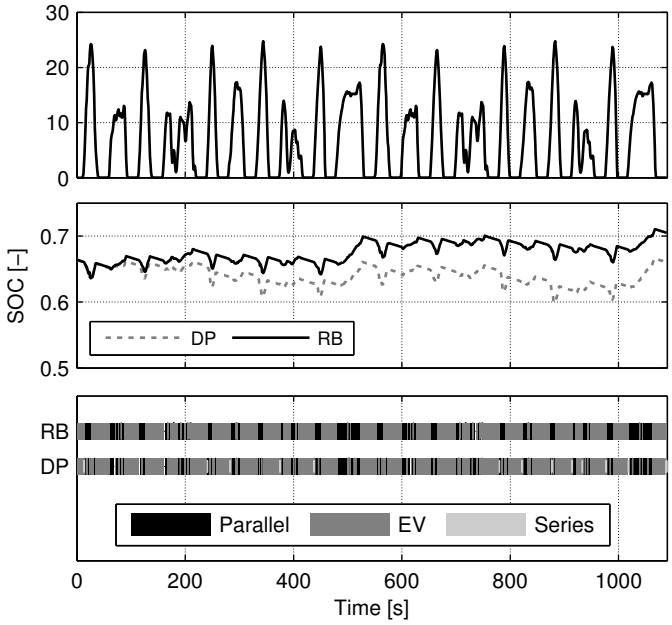

Figure 11 Comparison between RB and DP: torque split on the Manhattan cycle (detail), (a) vehicle speed (b) T_{EM} (c) T_{GEN}

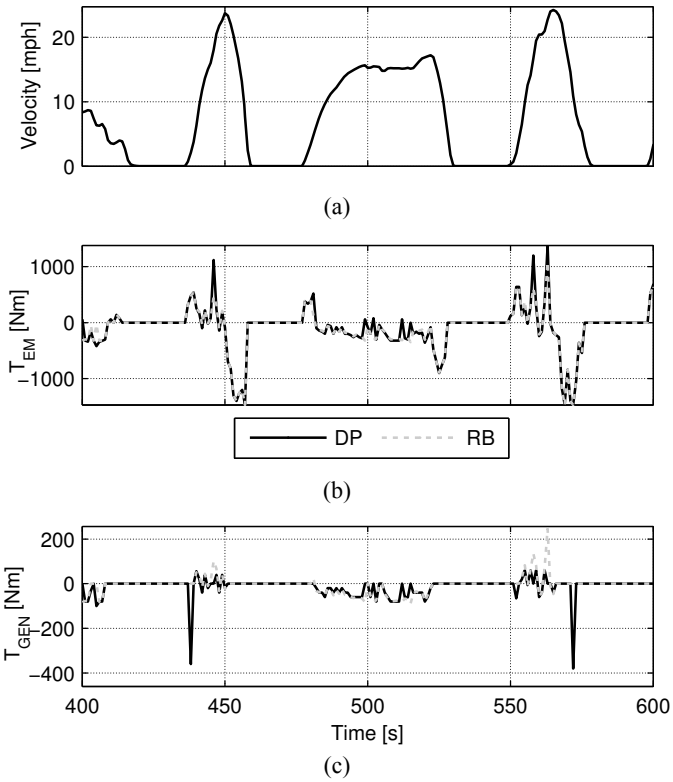

8.2 Parametric analysis

A parametric analysis is performed to evaluate the effects of the tuning variable μ (used to ensure charge-sustainability) on the vehicle performance. The parameter μ affects the overall results and should be tuned to a value that guarantees substantial charge-sustainability under a variety of driving conditions.

Figure 12 Effect of calibration parameter μ on the RB strategy, compared to DP, during the WVU-suburban cycle

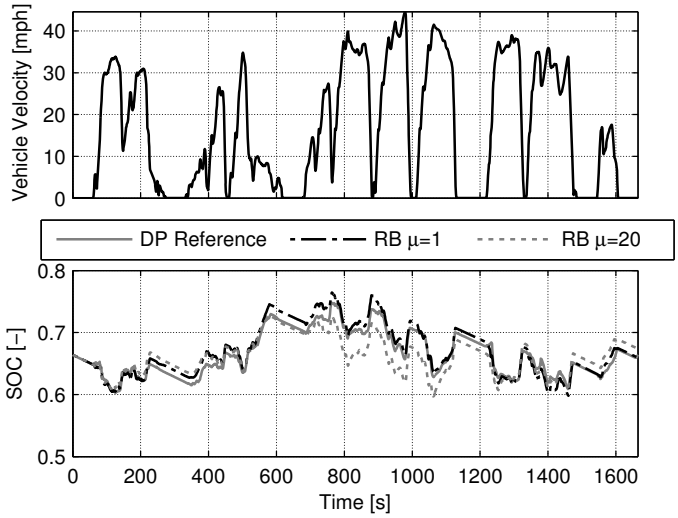

The effect of μ on the WVU-suburban driving cycle (developed by West Virginia University) is shown in Figure 12. The higher μ, the less energy stored in the battery. This is expected since the deviation of SOC from its reference value is penalised more. For this cycle, the charge-sustaining condition is perfectly achieved with $\mu = 20$, unlike the Manhattan cycle where, for the same value of μ, the final SOC is always above the reference. Table 2 represents the SOC deviation as a function of μ and shows that the controller is robust enough to be implemented: varying μ from 1 to 200 produces an effect on the fuel consumption of no more than 2%, and on ΔSOC of no more than 3%.

Similar considerations can be made on other driving cycles. A value of $\mu = 20$ is chosen as a compromise and kept constant for all cases. The results obtained on several driving cycles are in Table 3, which also includes validation cycles, i.e., cycles not used for the rule extraction.

As an example of validation cycle, the UDDS case is shown in Figure 13.

Table 2 Effect of the calibration parameter μ on the optimality and charge-sustainability of the RB strategy

	$\mu = 1$	$\mu = 20$	$\mu = 200$
Fuel consumption	100%	98%	99%
ΔSOC	–1%	1%	3%

Table 3 Comparison of RB and DP over various driving cycles

	Fuel consumption		*ΔSOC [%]*	
	DP	*RB*	*DP*	*RB*
Manhattan*	100	108	0	6
WVU-suburban*	100	102	0	1
WVU-inter*	100	102	0	3
APTA*	100	100	0	–8.3
UDDS	100	106	0	–3
HTUF	100	110	0	–6

Note: *Cycles used to extract the rules.

Figure 13 State of charge profile and mode selection on UDDS cycle

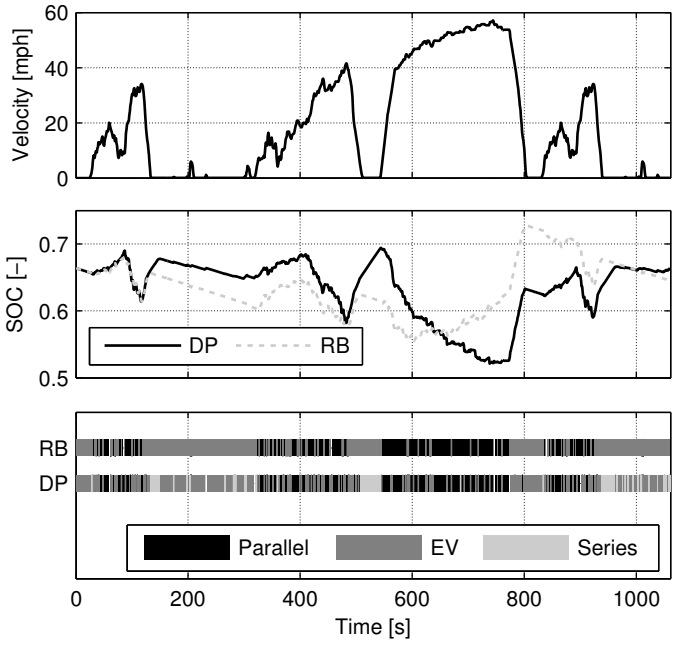

8.3 Comparison of the RB strategy with the ECMS

As a second step, the whole controller is tested in more realistic conditions. The simplified vehicle model used to implement DP and extract the rules is replaced by a more accurate simulator, which takes into account most of the dynamics of the powertrain components. The engine response is now represented by a first order dynamics (while the fuel consumption is still modelled with a stationary map). In addition, the engine cranking phase is modelled, thus, engine starting requires torque from the generator. The mode switch is no longer instantaneous, as clutch slip is also taken into account.

As mentioned earlier, DP cannot be implemented in this model, thus, the reference performance is given by the ECMS. In order to achieve results closest to the optimum, the equivalence factor s, described in Section 6, is tuned in advance in order to ensure the charge balance at the end of the driving cycle.

The rule based strategy, on the other hand, keeps the compromise value for the calibration parameter μ that was found in the static case, as described in Section 8.2.

The same supervisory controller is used for both strategies, and it implements the rules described in Section 7.1. The additional complexity of the forward simulator, however, requires a correction: a timer is added to prevent too sudden start/stop of the engine. In other words, once the engine is turned on (or off), some fixed time must elapse before it is switched off (or on) again. As in the previous case, the result on the Manhattan and WVU suburban cycle are presented. Figure 14 shows the torque split on the first cycle.

The behaviour of the electric motor (*EM*) is quite similar between the two strategies, while the generator (*GEN*) shows non-trivial differences in some portions of the cycle as shown in Figure 15.

Nevertheless, this does not affect the SOC profile. Analysing Figure 16 it is possible to see that the shape of the two curves is similar, as a consequence of the same supervisory.

Figure 14 Comparison RB-ECMS: torque split on the Manhattan cycle, (a) vehicle speed (b) T_{EM} (c) T_{GEN}

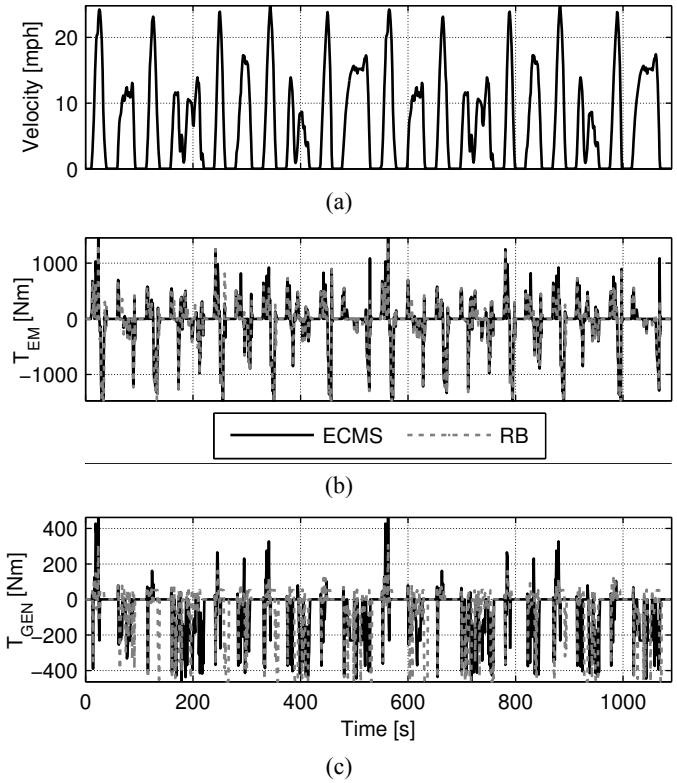

Figure 15 Comparison RB-ECMS: torque split on the Manhattan cycle (detail), (a) vehicle speed (b) T_{EM} (c) T_{GEN}

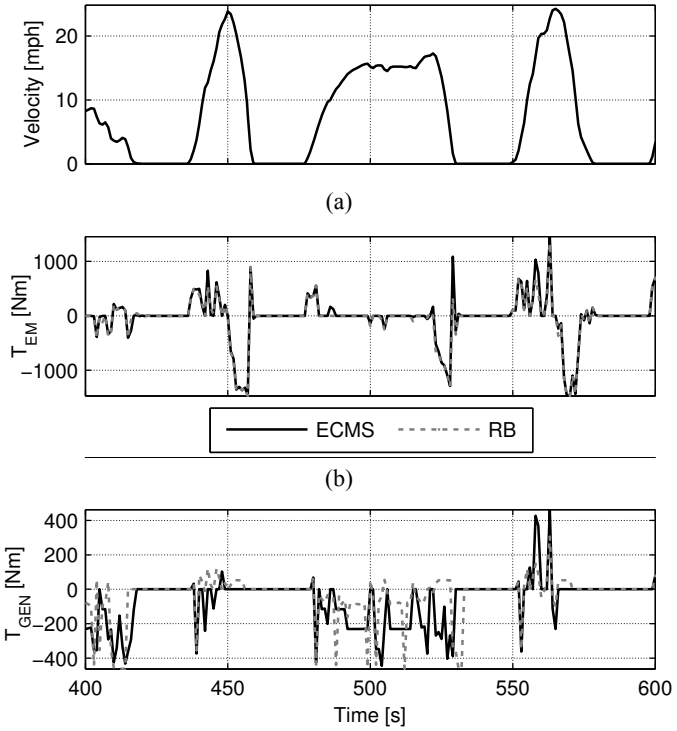

Figure 16 Comparison RB-ECMS: state of charge profile on the Manhattan cycle

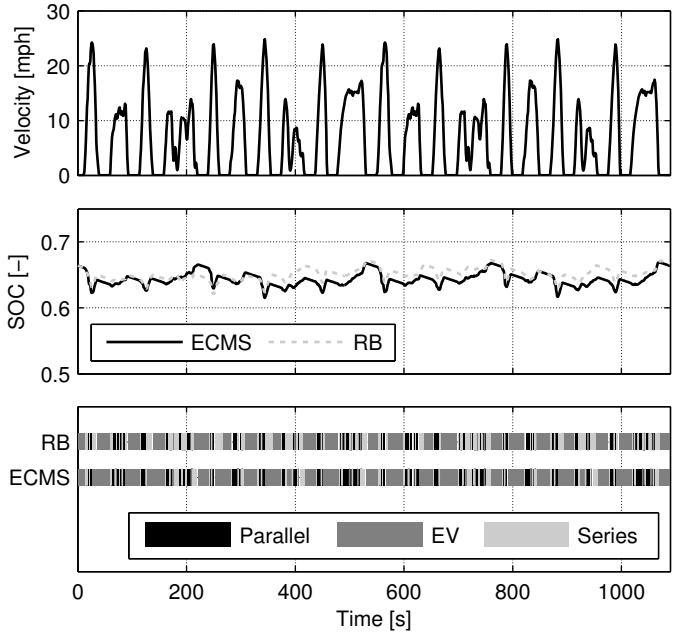

The difference in terms of fuel consumption for the various cycles is shown in Table 4; the RB strategy is remarkably similar to the quasi-optimal ECMS solution (obtained with a-priori tuning of ECMS for each cycle). Note that the fuel consumption values listed in Table 4 have been corrected to account for the SOC difference between the beginning and the end of the cycle.

Table 4 Comparison of RB and ECMS over various driving cycles

	Fuel consumption		*ΔSOC [%]*	
	DP	*RB*	*DP*	*RB*
Manhattan	100	106	0	0
WVU-suburban	100	104	0	–3
WVU-interstate	100	102	0	0
APTA	100	103	0	–8
UDDS	100	102	–8	–4

One of the strengths of the RB strategy is the smaller computational time. Since it uses only linear interpolation and algebraic operation, it is quicker than the ECMS, where a minimisation needs to be performed. On the other hand the ECMS requires less calibration and is easy to tune for different driving cycles, or to apply to different vehicle architectures.

9 Conclusions

The DP algorithm provides the optimal solution to the HEV energy management problem, and serves as a benchmark to assess the minimum fuel economy achievable along a driving mission. Both the need for a-priori knowledge of the mission profile and the high computational requirements make this strategy unrealistic to implement, since an on-board real time controller has to operate with limited computational and memory resources.

A RB strategy, on the other hand, is suitable for online implementation, due to the simple set of *if-then-else* rules. A demanding calibration phase is required though, for making the strategy charge-sustaining with respect to a wide variety of driving cycles.

In fact, RB parameters can be strongly affected by the driving conditions. The approach proposed in this paper is to study the results given by the DP in order to find common patterns in its decisions, and extract rules that can be implemented in a ‘sub-optimal’ RB controller. The resulting controller has only one calibration parameter that is tuned in order to satisfy the charge balance requirement. As shown in Section 8.2, the value of the parameter can be chosen to ensure good performance in most of the conditions. Both fuel consumption and state of charge profile are very close to the optimal, as demonstrated in Section 8.1.

The RB controller is dependent on both the powertrain components and vehicle architecture. If these change, the controller needs to be redesigned. A future improvement to the work presented consists in looking at the different decisions taken by DP as powertrain components change so as to make the RB strategy less dependent on those parameters.

References

Brahma, A., Guezennec, Y. and Rizzoni, G. (2000) 'Optimal energy management in series hybrid electric vehicles', *Proceedings of the 2000 American Control Conference*, Vol. 1, No. 6, pp.60–64.

Chasse, A., Sciarretta, A. and Chauvin, J. (2010) 'Online optimal control of a parallel hybrid with costate adaptation rule', *Proceedings of the 6th IFAC Symposium Advances in Automotive Control (AAC 2010)*.

He, X., Parten, M. and Maxwell, T. (2005) 'Energy management strategies for a hybrid electric vehicle', *Proceedings of the 2005 IEEE Vehicle Power and Propulsion Conference*, pp.536–540.

Hofman, T., Steinbuch, M., van Druten, R. and Serrarens, A. (2007) 'Rule-based energy management strategies for hybrid vehicles', *Int. J. Electric and Hybrid Vehicles*, Vol. 1, No. 1, pp.71–94.

Lewis, F. and Syrmos, V. (1995) *Optimal Control*, Wiley-Interscience.

Lin, C., Peng, H., Grizzle, J. and Kang, J. (2003) 'Power management strategy for a parallel hybrid electric truck', *IEEE Transactions on Control Systems Technology*, Vol. 11, No. 6, pp.839–849.

Musardo, C., Rizzoni, G., Guezennec, Y. and Staccia, B. (2005) 'A-ECMS: an adaptive algorithm for hybrid electric vehicle energy management', *European Journal of Control*, Vol. 11, Nos. 4–5, pp.509–524.

Onori, S., Serrao, L. and Rizzoni, G. (2010) 'Adaptive equivalent consumption minimization strategy for hybrid electric vehicles', *Proceedings of the 3rd Annual ASME Dynamic Systems and Control Conference (DSCC 2010)*.

Paganelli, G., Ercole, G., Brahma, A., Guezennec, Y. and Rizzoni, G. (2001) 'General supervisory control policy for the energy optimization of charge-sustaining hybrid electric vehicles', *JSAE Review*, Vol. 22, No. 4, pp.511–518.

Pisu, P., Hubert, C.J., Dembski, N., Rizzoni, G., Josephson, J.R., Russell, J. and Carroll, M. (2005) 'Modeling and design of heavy duty hybrid electric vehicles', *Proceedings of the 2005 ASME International Mechanical Engineering Congress and Exposition*.

Salman, M., Schouten, N. and Kheir, N. (2000) 'Control strategies for parallel hybrid vehicles', *Proceedings of the 2000 American Control Conference*, Vol. 1, No. 6, pp.524–528.

Sciarretta, A. and Guzzella, L. (2007) 'Control of hybrid electric vehicles', *IEEE Control Systems Magazine*, April, pp.60–70.

Serrao, L., Onori, S. and Rizzoni, G. (2009) 'ECMS as a realization of Pontryagin's minimum principle for HEV control', *Proceedings of the 2009 American Control Conference*.

Sundström, O. and Guzzella, L. (2009) 'A generic dynamic programming Matlab function', *Proceedings of the 18th IEEE International Conference on Control Applications*, pp.1625–1630.

Sundström, O., Ambühl, D. and Guzzella, L. (2009) 'On implementation of dynamic programming for optimal control problems with final state constraints', *Oil and Gas Science and Technology – Rev. IFP*, September.

Suzuki, M., Yamaguchi, S., Araki, T., Raksincharoensak, P., Yoshizawa, M. and Nagai, M. (2008) 'Fuel economy improvement strategy for light duty hybrid truck based on fuel consumption computational model using neural network', *Proceedings of the 17th IFAC World Congress*.

Szumanowski, A. (2000) *Fundamentals of Hybrid Vehicle Drives*, Radom, Warsaw.

Won, J. and Langari, R. (2005) 'Intelligent energy management agent for a parallel hybrid vehicle – part II: torque distribution, charge sustenance strategies, and performance results', *IEEE Transactions on Vehicular Technology*, Vol. 54, No.3, pp.935–953.