

Minimizing Battery Stress during Hybrid Electric Vehicle Control Design

Real World Considerations for Model-Based Control Development

Christopher Vagg, Chris J. Brace, Sam Akehurst
Department of Mechanical Engineering
University of Bath
Bath, U.K.
C.R.M.Vagg@bath.ac.uk

Lloyd Ash
Research and Development Department
Ashwoods Automotive Ltd.
Exeter, U.K.
Lloyd.Ash@Ashwoods.org

Abstract—In a mild hybrid electric vehicle (HEV) aggressive use of the electrical powertrain is desired to maximize the benefits from hybridizing the vehicle, however this has negative consequences for battery management, battery state of health, and motor temperature. In this paper a control strategy cost function is presented which can minimize these negative effects without significantly affecting the achievable reduction in fuel consumption, and without requiring a detailed battery model or a motor thermal model. This concept is demonstrated on a retrofit HEV unit developed by Ashwoods Automotive, with a model validated using chassis dynamometer test data. Dynamic Programming (DP) is used to optimize the controller, and some limitations of DP which are not often recognized are discussed.

Keywords— control; battery health; battery management system (BMS); battery stress; dynamic programming; hybrid electric vehicle (HEV); real world; retrofit; simulation

I. INTRODUCTION

Optimal control of Hybrid Electric Vehicles (HEVs) has been the subject of much study, and numerous ingenious algorithms have been proposed to maximize the benefits of having a hybrid powertrain. The problem is not trivial, because the best use of electricity often depends on the future actions of the driver, which are not known. The potential of new algorithms is commonly demonstrated in a simulation environment without any physical testing, and this often leads to various practical considerations being neglected. As a result of extensive practical development work the authors have come to realize that two factors commonly neglected in simulation often play a dominant role in the effectiveness of a HEV: battery cell imbalance, and motor temperature.

A. The Dilemma of High Power Operation

It has been noted by Plett [1] that the HEV environment is particularly harsh for batteries, as we frequently want to draw and return energy at extremely high rates. This makes sense because batteries are heavy and expensive, so we must ensure that we obtain maximum benefit from carrying them. Unfortunately high power use of the battery cells is a well-known stress factor, accelerating capacity fade and reducing the useable life of the battery [2, 3]. The problem is worse in

battery packs made up of numerous cells in series, because these tend to develop State of Charge (SOC) imbalances between individual cells. These imbalances mean that the Battery Management System (BMS) is forced to cut back the power demand in order to ensure no permanent damage is done to any of the cells, and to try to re-balance the pack.

Operating the electrical systems at high powers also has consequences for the Electric Machine (EM), because ohmic losses in the motor windings cause the EM to become hot. If the EM is consistently operated at high powers and the rate of cooling is not sufficient, the system may be forced into thermal cutback.

In the following sections the authors aim to show that the control strategy can be very simply enhanced to mitigate both of these practical considerations, without need for detailed battery pack or motor thermal models.

B. Hybrid Vehicle Description

Ashwoods Automotive have developed a hybrid-electric powertrain system which can be retrofitted into an existing conventional Internal Combustion Engine (ICE) vehicle with

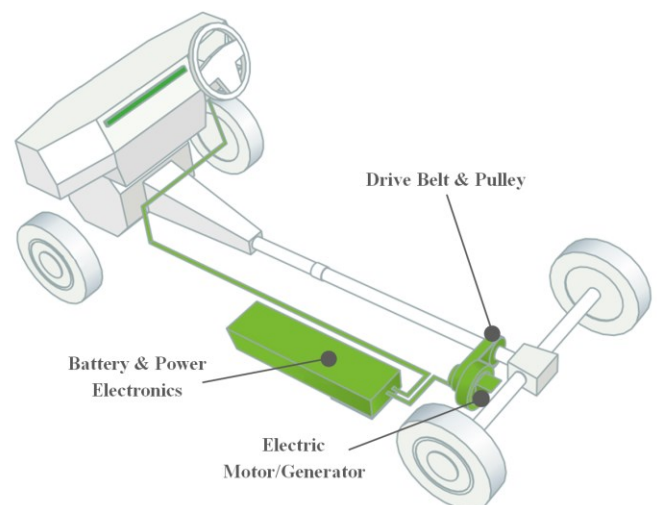


Fig. 1. Architecture of the hybrid electric vehicle, showing how the retrofit system is integrated into the standard vehicle.

relative ease. This retrofit system, described here for the first time, is novel in itself and is primarily aimed at the Light Commercial Vehicle (LCV) market. The authors see retrofit technology as an important step in improving the fuel economy of vehicles; since none of the leading suppliers of LCVs are currently mass producing hybrid or electric versions at a market competitive price, it seems likely that even a decade from now the majority of LCVs on the road will not be OEM produced hybrids. This observation suggests that there is a gap in the present research to examine good control of retrofit HEVs, whilst the findings remain relevant to all hybrid vehicles.

At the most fundamental level the system consists of an EM and a battery pack which enable energy usually lost during braking to be captured in batteries through regenerative braking. This energy is then used to assist the engine at other times, reducing the vehicle's fuel consumption and emissions. This is often called a Kinetic Energy Recovery System, and the configuration adopted is commonly categorized as a parallel torque assist (Fig. 1). The EM adds tractive power through a belt and pulley, which is sandwiched between the propeller shaft and the final drive (differential). Although there are many other configurations which could be adopted, possibly with greater technical scope, the simplicity of this architecture means that most manufacturers will still honor their standard warranty, since no standard components have actually been modified – this is a vital issue when selling into the commercial sector. Furthermore the system is entirely fail-safe, since if the system were to malfunction the vehicle continues to be drivable as though the hybrid system were not present. Even in the most dramatic failure mode, where for some reason the EM becomes seized, the drive belt simply breaks, preventing any damage. The hybrid system adds tractive force, and so as a result the driver uses less throttle pedal activation – there is no direct intervention in the throttle pedal or engine control signals. The battery pack is able to store around 0.6 kWh of energy at a nominal voltage of 79 V, at the EM has a peak power of 7 kW.

The system described here is a “mild hybrid” because the electrical powertrain is much less powerful than the standard ICE. Some of the functionality which is common in full hybrid vehicles, such as the ability to turn off the ICE and operate in “electric only” mode, is not available. In mild hybrid vehicles the problem of high power use (Section I-A) is particularly apparent because the capacity of the battery pack is comparatively small, and the electrical system is typically operating near its full capacity much of the time. In contrast, full Battery Electric Vehicles require much larger battery packs to achieve satisfactory vehicle range, and this added capacity also tends to yield a higher power capability (even if the batteries in the two vehicles have the same energy and power densities), meaning that the powertrain is rarely operating at its limit.

II. HYBRID VEHICLE CONTROL AND COST FUNCTIONS

Most HEV control strategies being developed employ a cost function to define when a control strategy is optimal [4, 5]; the optimal strategy is the one which minimizes the total cost over a drive cycle. In its most simple form this cost function is

probably equal to fuel consumption, however efforts have been made to include several other parameters to achieve multi-objective optimization; examples include equivalent cost of electricity used [6-9], NO_x and PM emissions [10, 11], and gear shift busyness [12].

As outlined in Section I-A the aim is to integrate battery stress considerations into the cost function. To this end it is noted that ohmic heating (I^2R loss) in individual pack cells plays a pivotal role, since this is the primary cause of heat generation in the pack, and temperature is among the most widely documented stress factors. Furthermore, it is common that cell manufacturing tolerances mean that cells in a pack have a spread of internal resistance, which means that some cells generate more ohmic heating than others. As a consequence of this spread packs can develop thermal gradients within them during normal use, which in turn can cause cell potentials to spread due to voltage-temperature dependencies. Since ohmic heating is proportional to the square of current, it seems logical to apply a cost proportional to the square of current, to target reductions in high power operation. The concept of our cost function is therefore

$$cost = fuel + \alpha \cdot I_{DC}^2 \quad (1)$$

where α is a weighting factor and I_{DC} is the DC battery current.

Ohmic heating is also the primary heat source within the EM, usually followed by bearing friction losses. The cost proposed does not work to penalize motor current (and therefore heat) in exactly the same way, because the DC current is not directly linked to the AC current in the windings. However since battery voltage is relatively steady, the proposed cost does penalize the square of power, which will inevitably reduce heat generation in the EM and make thermal cutback scenarios less likely.

III. HYBRID VEHICLE MODEL

A fully functional vehicle equipped with the hybrid-electric retrofit system was tested on the chassis dynamometer facilities at the University of Bath. Engine maps were obtained experimentally by mapping fuel consumption against vehicle speed and tractive force in each gear, and these data were used to develop a bespoke vehicle model in the Matlab/Simulink environment. Data for the batteries were obtained partly from manufacturer's data sheets and partly from experimental work conducted using battery test facilities at Ashwoods Automotive. A model of the power electronics and electric motor was developed from data collected using the motor test facilities at the University of Bath. These tests were conducted using comparable but not identical components to those on the vehicle, as the inverter-motor combination used on the vehicle was unavailable for testing. Nevertheless the results are broadly representative of the system used on-vehicle within its normal operating window.

Model validation was achieved by comparing the model predicted fuel consumption against measured drive cycle fuel consumption on the chassis dynamometer for eight scenarios, where fuel consumption is measured using the industry standard “bag analysis” method [13] and in accordance with

ISO 16183:2002 [14]. These scenarios were generated by changing three variables, each of which had two possible options, resulting in $2^3 = 8$ variations. The three variables and their options are shown in Table I.

When modeling the behavior of the hybrid system, the recorded trace of measured battery current was fed directly into the model so the control strategy did not need to be represented in the model.

As can be seen from Fig. 2, the simulation is capable of predicting the vehicle fuel consumption to a good degree of accuracy. It should be noted that the Coefficient of Variance of this chassis dynamometer facility has previously been shown to be in the region of 1%, and so chasing a degree of accuracy beyond this quickly becomes futile.

IV. DYNAMIC PROGRAMMING

Dynamic Programming (DP) is a technique which can be used to find the optimal control strategy over a known drive cycle. The result is essentially no different to direct enumeration (calculating the cost of every possible strategy) but the algorithm is able to dramatically reduce the computational time required by observing that the top-level problem can be broken down into a series of far more trivial problems. For example, in the common shortest path problem we may need to find the shortest path from $B \rightarrow C$. If we are then asked to find the shortest route from $A \rightarrow C$ via B ($A \rightarrow B \rightarrow C$) we only need to find the solution to $A \rightarrow B$, because the solution to $B \rightarrow C$ has previously been found. The same logic is applied to our problem, where the goal is changed to minimizing the cost function over a drive cycle. For this purpose we require a state vector describing the vehicle at each time, defined as

$$x = x(v, a, g, SOC) \quad (2)$$

where v is vehicle speed, a is acceleration, g is gear selected and SOC is the battery State of Charge. Note that this number of state variables is more than is commonly found in the literature, and the computational procedure is therefore a little more complex than is usual; this is in fact necessitated by the comparatively simple vehicle architecture. In most other

TABLE I. SIMULATION VALIDATION TEST VARIATIONS

Variable	Options	Details
Drive Cycle	Urban	Half of the urban phase of the NEDC, i.e. 2 repetitions of the ECE-15
	Extra-Urban	EUDC phase of the NEDC
Gear shifts	STD	Standard shift points described by the NEDC
	GSI	Shifts effected as advised by an on-board gear shift indicator (GSI)
Hybrid System	OFF	Hybrid system electrically deactivated
	ON	Hybrid system active, following a simple heuristic control strategy

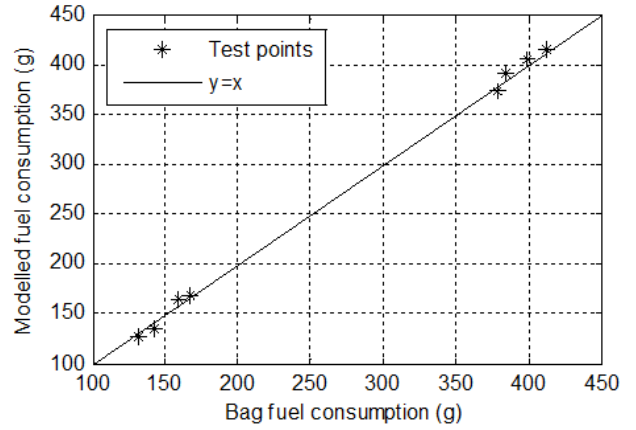


Fig. 2. Correlation between simulated fuel consumption and that measured on the chassis dynamometer.

studies the hybrid system is assumed to be heavily integrated into the vehicle, and so the control decision dictates the required engine power, which is often assumed to be delivered by the Engine Control Unit (ECU) at the most efficient engine operating point and gear, and as a result v , a and g may be collapsed into a single *Power* variable. In this case however, the gear selection is controlled by the driver and the engine may not be assumed to be at its most efficient point. Therefore in order to get a realistic representation of the fuel consumption all three states are required.

The computational procedure for DP is helpfully described by Larson and Casti [15] with worked examples. Essentially, we require two matrices to be populated: an instantaneous cost matrix, and a state transition matrix. The instantaneous cost matrix contains the cost for each combination of vehicle state and control decision (u), and so is four dimensional in this case:

$$cost = L(v, a, g, u). \quad (3)$$

Note that the instantaneous cost is not a function of SOC, despite this forming part of the vehicle state vector, because fuel consumption is affected only by the electrical power (resulting from the control decision, u) and not by the battery SOC.

The state transition matrix describes the SOC of the vehicle in the next timestep, given the vehicle state and control decision at the current timestep (k):

$$SOC(k+1) = f(SOC(k), v(k), u(k)). \quad (4)$$

Again, note that the full state vector is not required because SOC is not a function of a or g .

Given these two matrices, the computational procedure starts from the penultimate step, $k = N-1$, where N is the number of time steps, and works backwards in time towards the first step, $k = 1$. The control decision at $k = N$ is null, because the simulation stops. At each timestep the following steps are made:

1. The current v , a and g are used to isolate the appropriate instantaneous cost matrix L_k at the current vehicle state, which in this case leaves a row vector, $L_k(u_k)$.
2. For every $[SOC_k, v_k, u_k]$ the next SOC (SOC_{k+1}) is calculated from (4).
3. For each SOC_{k+1} the minimum possible future cost, $J_{k+1}(SOC_{k+1})$, is found by interrogating the table generated at the previous iteration. If SOC_{k+1} does not lie on the SOC grid then interpolation is required.
4. $L_k(u_k)$ is expanded to have the same dimensions as $J_{k+1}(SOC_k, u_k)$, and the two are summed to find the total cost-to-go, which is a 2-D matrix describing the minimum cost, J_k , for progressing from each present state $[SOC_k, u_k]$ to the end of the drive cycle.
5. The procedure is repeated for the next timestep (reducing k) using the J_k calculated in step 4 as the future cost matrix in step 3.

Having tabulated the problem at each timestep, as described above, the optimality problem may be easily solved by tracing the optimal solution forward through the table, starting at $k = 1$, selecting the control decision with the minimum future cost, and moving to the resulting SOC at $k+1$. More formally, the solution is obtained by solving the recursive equation:

$$J_k^*(x_k) = \min [L_k + J_{k+1}^*(x_{k+1})] \quad (5)$$

where J_k^* represents the total cost of the optimal control trajectory u^* which starts from state x_k .

The instantaneous cost at each step is a function of fuel consumption and DC current, as per (1), but the DC current has been exchanged for the instantaneous battery C-rate (which is essentially a measure of current normalized to the battery capacity), to make the results more transferable. Therefore the cost function implemented is:

$$L(x, u) = F(x, u) + \alpha [C(x, u)]^2 \quad (6)$$

where F is the ICE fuel consumption (g/s), C is the C-rate (h^{-1}) at which the battery is being operated and α is a weighting factor which may be tuned.

V. RESULTS

Using the DP method optimal control trajectories were found which minimized the cost function over a NEDC test cycle for a range of α . As would be expected, increasing the value of α increased the cost of electricity use – particularly at the highest powers – resulting in reduced C-rates. The reduction in electrical energy use unavoidably increased the fuel consumption. However, the linearity of this relationship is of great significance, and it can be seen from Fig. 3 that the relationship is highly non-linear.

Fig. 3 indicates that increasing the C-rate yields diminishing returns in fuel saving – a finding which is common

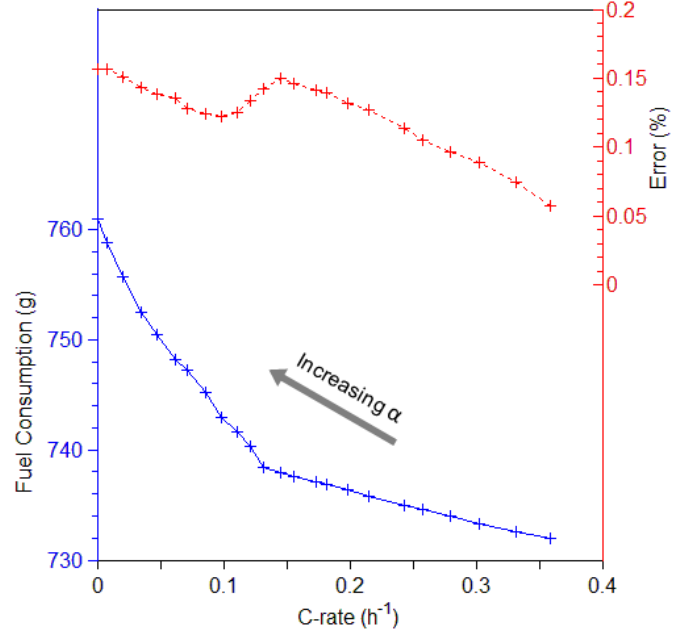


Fig. 3. Correlation between the average battery C-rate over an NEDC and simulated fuel consumption (—) and DP accumulated error (- -).

when trading-off one optimization parameter against another. The shape of this curve reveals that by increasing α the average C-rate can be reduced from 0.35 to around 0.25 (a reduction of about 32%) while only sacrificing about 10% of the fuel consumption benefit.

The discontinuity observed at approximately 0.13 C corresponds to an important change in control strategy behavior. As α increases through this point it becomes more costly to charge the batteries through high power regeneration during braking than at a lower power during cruise phases.

VI. CAUTIONARY NOTES ON DYNAMIC PROGRAMMING

Dynamic Programming is frequently used by researchers in HEV control to assess the potential of various hybrid vehicle architectures, and as a baseline against which to measure novel algorithms. It is often taken for granted that the DP solution is the ultimate benchmark, but the details of its implementation and the inherent error accumulated are seldom mentioned. In this work the authors have become aware of numerous decisions and trade-offs that must be made during the implementation of the algorithm. Some of these are briefly discussed in the following sections, in the hope that this may be of use to others, and perhaps start a discussion on “good practice” implementation.

A. Accuracy

Practical implementation of DP necessitates that the state space is represented as a discretized grid. When a state transition leads to a next state which is not exactly on the grid interpolation of the future cost is required. Even with an interpolation regime there is an error, and this is compounded at each timestep. As a result, the tabulated value for future cost starting from a given state at $k = 1$ will not exactly equal the

actual cost incurred when the optimal control trajectory is followed. Larson and Casti [15] noted that “This discrepancy shows that for [a] highly nonlinear system ... it is sometimes necessary either to use a finer quantization interval and/or higher order interpolation to obtain precise results.”

Note that this is *not* a simulation error, which could be combated with a high fidelity model – it is an error which is inherent in the process of tabulating and interpolating the DP cost function. The error incurred in this implementation is shown in Fig. 3, which is generally less than 0.15%. This level of error was tolerated in exchange for computational time, however it appears to the authors that there is some need of discussion as to what degree of error is acceptable in this application. Once again, this does not imply that the fuel consumption data displayed in Fig. 3 necessarily carry these errors, but that the cost function interpolated along the optimal control trajectory does.

During the development of the DP scripts an earlier version was implemented which ran with a timestep of 1 second and used nearest neighbor interpolation for vehicle speed, rather than linear interpolation. For comparison the plot for fuel consumption and DP accumulated error has been included (Fig. 4). Note that the timestep only affects the fidelity of the simulation to the real world – that is the model’s ability to reproduce transient behavior, and how accurate the modeled fuel consumption is. Conversely, the interpolation regime(s) affect the DP error accumulated – that is the discrepancy between the tabulated minimum cost and the cost actually incurred when the optimal control trajectory is traced through the tables. In this implementation the DP accumulated error was generally in the region 1-2%. In fact it can be seen in Fig. 4 that when $\alpha = 0$ and the C-rate is at its highest, the DP algorithm has not properly minimized the fuel consumption. At this point the fuel consumption is the only component of the cost function, and so should be at its lowest: it should not be possible to achieve lower fuel consumption by increasing α . This erroneous result highlights the need to understand some of the algorithm’s weaknesses, and ensure that the implementation is fit for purpose.

Key information about the final DP implementation is given in Table II. Further investigations are required to establish the sensitivity of the algorithm to each parameter, and how best to maximize its performance.

B. Sensitivity

At each timestep the algorithm selects the control decision which minimizes the cost function. In some circumstances this can lead to the selection of some peculiar behavior, where a much more intuitive solution is only marginally less optimal. This issue is closely connected to the question of how much we trust our model. In this case the authors noted that the model seems to perform very well over an NEDC, however it does not perform so well over an Artemis cycle, presumably because the engine model is based on steady-state behavior and the Artemis cycle is highly dynamic. Developers should be aware that the accuracy of the model is crucial to the trustworthiness of the DP solution, and that there often remains a need for sanity-checking. Of course this problem is not limited to the DP

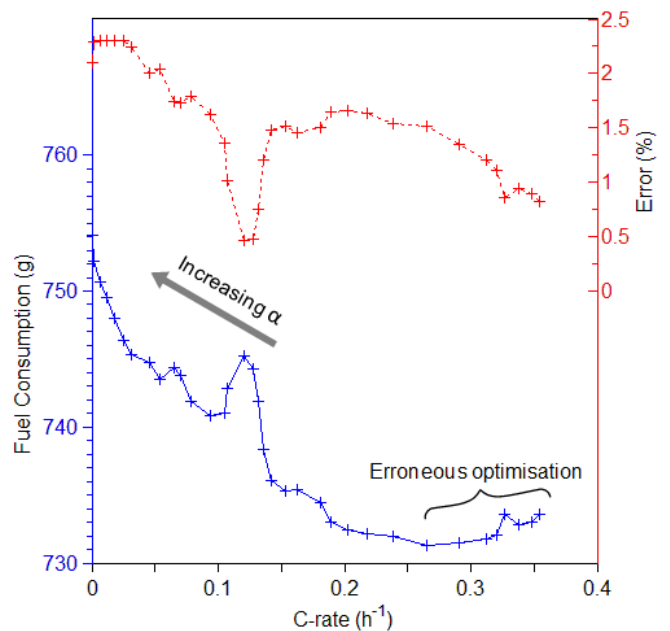


Fig. 4. Correlation between average battery C-rate over an NEDC and simulated fuel consumption (—) and DP accumulated error (---). Results of an early implementation of the DP algorithm (timestep = 1 s, nearest neighbor interpolation for v state).

algorithm, but is an issue for all model-based-design techniques.

C. Control bifurcations

Although seemingly unlikely, it is possible that given a vehicle state, two possible control decisions will have exactly the same future cost. When using the earlier (1 Hz) implementation it was found that this could occur as many as 60 times in 1180 timesteps, though this was reduced in the final implementation when nearest neighbor interpolation was replaced with linear interpolation for the vehicle speed. Bifurcations are also far less likely for complex cost functions, and more likely when α is small. The vast majority of these are inconsequential, and occur because the vehicle state is constant for several timesteps (e.g. cruising) and a sequence of decisions could be taken in a variety of orders with the same result (i.e. the cost is the same for [A, B, C] as for [C, A, B]). However in some cases there is a genuine bifurcation in control trajectories, and some method of selecting one is required. The approach

TABLE II. DYNAMIC PROGRAMMING PARAMETERS

Parameter	Quantization interval	Interpolation method
Speed (v)	2 km/h	Linear
Acceleration (a)	0.1 m/s ²	Nearest
Gear (g)	1	Exact
State of Charge (SOC)	0.005 %	Linear
Demand (u)	± 4 % ^a	-
Timestep (Δt)	0.2 s	-

^a. The control signal can vary from -100% to $+100\%$, corresponding to regeneration and assist.

taken here was to select the control decision closest to the previous one, with the logic that this will reduce control noise and busyness.

VII. CONCLUSION

We have demonstrated in this paper that the average C-rate to which the battery is exposed can be reduced considerably without significantly affecting the fuel savings achieved. This was accomplished by adding a term into the controller cost function which adds a penalty proportional to the square of the C-rate. The effect of this is to protect the battery against potentially high stress situations, which should help to increase its effective life. Furthermore this technique reduces the likelihood of the BMS needing to scale back the control demands, and also reduces the likelihood of the system entering thermal cutback due to a high EM temperature. All of this is achieved without need for detailed modeling of the battery or EM thermal properties.

This paper has also discussed several issues in the implementation of the Dynamic Programming algorithm which users should be aware of, but are rarely discussed.

ACKNOWLEDGMENT

The authors would like to thank Ashwoods Automotive Limited for supply of the hybrid electric vehicle, and assistance with its modeling.

REFERENCES

- [1] G. L. Plett, "Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 1. Background," *Journal of Power Sources*, vol. 134, pp. 252-261, 2004.
- [2] Z. Li, L. Lu, M. Ouyang, and Y. Xiao, "Modeling the capacity degradation of LiFePO₄/graphite batteries based on stress coupling analysis," *Journal of Power Sources*, vol. 196, pp. 9757-9766, 2011.
- [3] B. Y. Liaw and M. Dubarry, "From driving cycle analysis to understanding battery performance in real-life electric hybrid vehicle operation," *Journal of Power Sources*, vol. 174, pp. 76-88, 2007.
- [4] L. Johannesson, M. Åsbogård, and B. Egardt, "Assessing the Potential of Predictive Control for Hybrid Vehicle Powertrains Using Stochastic Dynamic Programming," *Intelligent Transportation Systems, IEEE Transactions on*, vol. 8, pp. 71-83, 2007.
- [5] J. Liu and H. Peng, "Modeling and Control of a Power-Split Hybrid Vehicle," *Control Systems Technology, IEEE Transactions on*, vol. 16, pp. 1242-1251, 2008.
- [6] G. Paganelli, S. Delprat, T. M. Guerra, J. Rimaux, and J. J. Santin, "Equivalent consumption minimization strategy for parallel hybrid powertrains," in *Vehicular Technology Conference, 2002. IEEE 55th*, Birmingham, AL, USA, 2002, pp. 2076-2081 vol.4.
- [7] G. Paganelli, M. Tateno, A. Brahma, G. Rizzoni, and Y. Guezennec, "Control development for a hybrid-electric sport-utility vehicle: Strategy, implementation and field test results," in *2001 American Control Conference*, Arlington, VA, United states, 2001, pp. 5064-5069.
- [8] T. S. Kim, C. Manzie, and R. Sharma, "Model predictive control of velocity and torque split in a parallel hybrid vehicle," in *2009 IEEE International Conference on Systems, Man and Cybernetics (SMC 2009)*, San Antonio, TX, United states, 2009, pp. 2014-2019.
- [9] S. J. Moura, H. K. Fathy, D. S. Callaway, and J. L. Stein, "A Stochastic Optimal Control Approach for Power Management in Plug-In Hybrid Electric Vehicles," *Control Systems Technology, IEEE Transactions on*, vol. 19, pp. 545-555, 2011.
- [10] C.-C. Lin, H. Peng, and J. W. Grizzle, "A stochastic control strategy for hybrid electric vehicles," in *2004 American Control Conference*, Boston, MA, United states, 2004, pp. 4710-4715.
- [11] C.-C. Lin, H. Peng, J. W. Grizzle, and J.-M. Kang, "Power management strategy for a parallel hybrid electric truck," *Control Systems Technology, IEEE Transactions on*, vol. 11, pp. 839-849, 2003.
- [12] C.-C. Lin, H. Peng, J. W. Grizzle, J. Liu, and M. Busdiecker, "Control System Development for an Advanced-Technology Medium-Duty Hybrid Electric Truck," presented at the 2003 SAE International Truck and Bus Meeting and Exhibition, Fort Worth, Texas, 2003.
- [13] *Uniform provisions concerning the approval of vehicles with regard to the emission of pollutants according to engine fuel requirements*, UN/ECE Regulation No 83 Rev.4, 2011.
- [14] *Heavy-duty engines — Measurement of gaseous emissions from raw exhaust gas and of particulate emissions using partial flow dilution systems under transient test conditions*, ISO Standard 16183:2002.
- [15] R. E. Larson and J. L. Casti, *Principles of Dynamic Programming: Part 1 - Basic Analytic and Computational Methods* vol. 7. New York: Marcel Dekker, 1978.