Proceeding of the 2004 American Control Conference
Boston, Massachusetts June 30 - July 2, 2004

A Stochastic Control Strategy for Hybrid Electric Vehicles

Chan-Chiao Lin[1], Huei Peng[1], and J.W. Grizzle[2]

[1] Dept. of Mechanical Engineering, University of Michigan, MI 48109-2125
chancl@umich.edu, hpeng@umich.edu
[2] Dept. of Electrical Engineering and Computer Science, University of Michigan, MI 48109-2122
grizzle@umich.edu

Abstract

The supervisory control strategy of a hybrid vehicle coordinates the operation of vehicle sub-systems to achieve performance targets such as maximizing fuel economy and reducing exhaust emissions. This high-level control problem is commonly referred as the power management problem. In the past, many supervisory control strategies were developed on the basis of a few pre-defined driving cycles, using intuition and heuristics. The resulting control strategy was often inherently cycle-beating and lacked a guaranteed level of optimality. In this study, the power management problem is tackled from a stochastic viewpoint. An infinite-horizon stochastic dynamic optimization problem is formulated. The power demand from the driver is modeled as a random Markov process. The optimal control strategy is then obtained by using Stochastic Dynamic Programming (SDP). The obtained control law is in the form of a stationary full-state feedback and can be directly implemented. Simulation results over standard driving cycles and random driving cycles are presented to demonstrate the effectiveness of the proposed stochastic approach. It was found that the obtained SDP control algorithm outperforms a sub-optimal rule-based control strategy trained from deterministic DP results.

1. Introduction

Hybrid vehicle powertrains have been widely studied recently because of their potential to significantly improve fuel economy and reduce emissions of ground vehicles. Due to the multiple-power-source nature and the complex configuration, the control strategy of a hybrid vehicle is more complicated than that of an engine-only vehicle. The main function of the control strategy is power management, i.e., the design of the high-level control algorithm that determines the proper power split between the motor and the engine to minimize fuel consumption and emissions, while satisfying constraints such as drivability, charge sustenance and component reliability.

Many existing power management strategies employ heuristic control techniques such as rules/fuzzy logic for the control algorithm development [1, 2]. The idea of this approach is commonly based on the concept of "load-leveling", which attempts to operate the engine in an efficient region and uses the battery as a load-leveling device. Another approach is based on static optimization method, whereby the proper split between the two energy sources is decided by minimizing the total equivalent consumption cost [3]. In an earlier work [4], we proposed a design procedure that uses deterministic dynamic programming (DDP) to find the optimal solution and then extracts implementable rules to form the control strategy. Even though the control laws we obtained have performed well in a real hybrid electric vehicle [5], there are two drawbacks to this approach. First, this approach optimizes with respect to a specific driving cycle and might be neither optimal nor charge-sustaining under other cycles; secondly, the feedback solution to the DDP is not directly implementable and the rule extraction process can be time-consuming. To overcome these drawbacks, a design procedure based on stochastic dynamic optimization techniques is proposed in this paper. This alternative approach assumes that there is an underlying Markov process to represent the power demand from the driver. A similar approach to an automotive powertrain control problem can be found in [6]. Instead of being optimized over a given driving cycle, the power management strategy is optimized over a family of random driving cycles in an average sense. In order to obtain a time-invariant control strategy, an infinite-horizon optimization problem is formulated and solved by using stochastic dynamic programming (SDP). The control law derived from SDP is a time-invariant state variable feedback and can be directly used in real-time implementation. More importantly, the goal of charge sustaining can be incorporated in the control design process through this SDP approach.

The paper is organized as follows: In Section 2, the configuration of a hybrid electric truck is described, followed by the stochastic modeling of power demands in Section 3. The stochastic dynamic programming approach is introduced in Section 4. Section 5 presents the simulation results using the SDP optimal policy. Conclusions are presented in Section 6.

2. Parallel Hybrid Electric Vehicle Model

2.1 Vehicle Configuration

The hybrid vehicle studied in this paper was modified from an International 4700 series truck, a 4X2 Class VI diesel truck produced by Navistar [7]. The hybrid version of the truck has a parallel configuration with the electric motor positioned after the transmission. The original diesel engine was downsized from V8 (7.3L) to V6 (5.5L) and a 49 KW electric motor was added. A schematic of the vehicle is given in Figure 1. Basic vehicle information is given in Table 1.

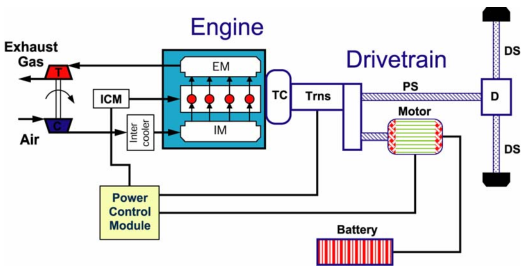

Figure 1: Schematic diagram of the hybrid electric truck

Table 1: Basic vehicle specification

DI Diesel Engine	V6, 5.475L, 157HP/2400rpm
DC Motor	49kW
Lead-acid Battery	Capacity: 18Ah, Number: 25
Automatic Transmission	4 speed, GR: 3.45/2.24/1.41/1.0
Vehicle	Total mass: 7504 kg

2.2 System Equations

Because the computation time of dynamic programming problems increases exponentially with the number of states, a minimum number of state variables are kept. Due to the fact that the system-level performance is the main concern, dynamics that are much faster than 1Hz are ignored. It was determined that only two state variables need to be kept: the battery SOC and the wheel speed.

The discrete-time state equation of the battery SOC can be expressed as

$$SOC_{k+1} = SOC_k - \frac{i_{b,k}(SOC_k, P_{b,k})}{Q_b}, \quad (1)$$

where k is the time index (time step is assumed to be one second), Q_b is the maximum battery charge, and i_b is the battery current which is calculated from the battery power, P_b, by using a static equivalent circuit model. By convention, positive P_b indicates discharging and negative P_b indicates charging. The battery power is derived from $P_b = \tau_m \cdot \omega_m \cdot \eta_m^{-\mathrm{sgn}(\tau_m)}$, where the motor efficiency $\eta_m(\tau_m, \omega_m)$ is a function of motor torque and speed.

The engine dynamics are ignored based on the quasi-static assumption [8]. The fuel consumption rate $W_{fuel}(\omega_e, \tau_e)$, engine-out NOx emission $W_{NOx}(\omega_e, \tau_e)$ and PM emission $W_{PM}(\omega_e, \tau_e)$ are assumed to be static functions of engine speed and engine torque. We assume the engine is fully warmed-up and effect of the engine temperature is not considered. The driveline components are fast and thus are represented by static models. Nonlinear maps are used to calculate the engine speed, $\omega_e(\omega_t, \tau_e)$, and the torque converter output torque, $\tau_t(\omega_t, \tau_e)$, where ω_t is the torque converter output speed. The transmission output torque is obtained from $\tau_x = R_x \eta_x (\tau_t - \tau_{x,l})$, where R_x is the transmission gear ratio, $\eta_x(g_x)$ is the gear efficiency, which varies with the gear position, g_x, and $\tau_{x,l}(\omega_t, g_x)$ is the torque loss due to friction and churning loss. Note that the transmission inefficiency is modeled by reducing the torque multiplication while keeping the speed reduction the same (i.e., $\omega_x = \omega_t / R_x$). The gear position is determined by the shift logic $g_x(\tau_e, \omega_x)$ generated in [4].

The torques from the engine and the motor are summed after the transmission by using a torque coupling device that has a gear reduction ratio R_c and efficiency η_c. The final propulsion torque after the differential is calculated by $\tau_d = R_d \eta_d \left(\tau_x + R_c \tau_m \eta_c - \tau_{d,l}\right)$, where R_d and η_d are the gear ratio and gear efficiency, and $T_{d,l}(\omega_x)$ is the torque loss of the differential.

The vehicle is modeled as a point-mass, and the wheel speed is calculated from the state equation

$$\omega_{wh,k+1} = \omega_{wh,k} + \frac{1}{M_r r_d^2}\left(\tau_{wh,k} - B_{wh}\omega_{wh,k} - r_d\left(F_r + F_a\right)\right), \quad (2)$$

where $\tau_{wh} = \tau_d - \tau_{fb}$ is the net wheel torque (τ_{fb} is the friction braking torque), r_d is the dynamic tire radius, B_{wh} is the viscous damping, F_r and F_a are the rolling resistance force and the aerodynamic drag force, $M_r = M_v + J_r / r_d^2$ is the effective mass of the vehicle, and J_r is the equivalent moment of inertia of the rotating components in the vehicle.

3 Stochastic Modeling of Driver Power Demand

The power demand is an input to the power management controller. The driver's throttle and brake pedal commands are interpreted as a power demand to be satisfied by the powertrain. Unlike other work that treats the power demand as a-priori information (e.g., a known power demand path to follow a given driving cycle), we propose to model the driver's demand as a discrete-time stochastic dynamic process. A stationary Markov chain is used to generate the power demand from the driver, P_{dem}, which is assumed to take on a finite number of values

$$P_{dem} \in \left\{P_{dem}^1, P_{dem}^2, \ldots, P_{dem}^{N_p}\right\}. \quad (3)$$

The wheel speed (directly related to the vehicle speed) is also discretized into a finite number of values

$$\omega_{wh} = \left\{\omega_{wh}^1, \omega_{wh}^2, \ldots, \omega_{wh}^{N_\omega}\right\}. \quad (4)$$

The dynamics of P_{dem} is assumed to be

$$P_{dem,k+1} = w_k \,. \tag{5}$$

where the probability distribution of w_k is assumed to be

$$\begin{aligned} &\Pr\left\{ w = P_{dem}^{j} \,\middle|\, P_{dem} = P_{dem}^{i}, \omega_{wh} = \omega_{wh}^{l} \right\} = p_{il,j} \\ &i, j = 1, 2, ..., N_p, \quad l = 1, 2, ..., N_\omega \end{aligned} \,. \tag{6}$$

where $\sum_{j=1}^{N_p} p_{il,j} = 1$, and $p_{il,j}$ represents the one-step transition probability that the system is in P_{dem}^{j} at time $k+1$ given the system is in P_{dem}^{i} and ω_{wh}^{l} at time k. It should be noted that $p_{il,j}$ is simplified from a two-dimensional Markov chain $p_{il,jm}$ with $m = 1, 2, ..., N_\omega$, due to the fact that the transition of ω_{wh} is a deterministic process.

It follows that specifying driving-cycle characteristics is equivalent to specifying the transition probabilities, $p_{il,j}$. A natural way to determine values for the transition probabilities is to estimate them on the basis of observed sample paths, such as past driving records, or standard driving cycles. We used standard driving cycles in this study to determine the transition probabilities as follows: various standard driving cycles were selected to represent mixed city, suburban, and highway driving. From the speed profile, P_{dem} and ω_{wh} could be calculated by the vehicle model. Using nearest-neighbor quantization, the sequence of observations (P_{dem}, ω_{wh}) was mapped into a sequence of quantized states $(P_{dem}^{i}, \omega_{wh}^{l})$. The transition probability could then be estimated by the maximum likelihood estimator, which counts the observation data as

$$\hat{p}_{il,j} = \frac{m_{il,j}}{m_{il}} \qquad \text{if } m_{il} \neq 0 \,, \tag{7}$$

where $m_{il,j}$ is the number of times that the transition from P_{dem}^{i} to P_{dem}^{j} has occurred given the wheel speed state was ω_{wh}^{l}, and $m_{il} = \sum_{j=1}^{n} m_{il,j}$ is the total number of times that P_{dem}^{i} has occurred at wheel speed ω_{wh}^{l}. However, it is possible a training set may not be rich enough to cover the whole state space, and, in some cases, m_{il} may be zero for some i and l. Therefore, a smoothing technique was used to the estimated parameters [9].

After the Markov model is built, we constructed a stochastic hybrid vehicle model as shown in Figure 2. The model includes three state variables: SOC, wheel speed, and driver power demand. The Markov driver model is used to determine the probability distribution of future power demands and to generate a sequence of synthetic power demands, in order words, a random driving cycle. A power management controller that is optimized on the basis of this stochastic model will be described in the next section.

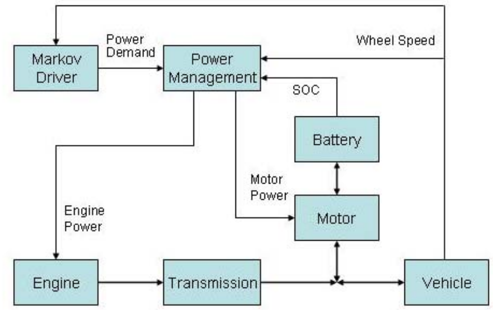

Figure 2: Block diagram of the hybrid electric vehicle

4. Stochastic Dynamic Optimization

The section determines an optimal policy for the power split between two power sources (engine and motor) so that the fuel consumption and emissions are minimized. In the meantime, vehicle drivability and the desire to maintain battery SOC have to be satisfied.

4.1 Problem Formulation

We start by formulating an infinite horizon problem, which is proper because the system dynamics and the cost are time-invariant and we do not have a final time or terminal constraint defined. A key benefit of the infinite horizon problem is that the generated control policy is time-invariant and thus can be easily implemented.

Our objective is to find an optimal control policy $u = \pi(x)$ that maps observed states to control decisions so as to minimize the expected total cost over an infinite horizon:

$$J_\pi(x_0) = \lim_{N \to \infty} \underset{w_k}{E} \left\{ \sum_{k=0}^{N-1} \gamma^k g\left(x_k, \pi(x_k)\right) \right\}, \tag{8}$$

where g is the instantaneous cost incurred, $0 < \gamma < 1$ is the discount factor, and $J_\pi(x_0)$ indicates the resulting expected cost when the system starts at state x_0 and follows the policy π thereafter. For our hybrid vehicle control problem, the control signal is engine power P_e. When we impose a drivability requirement, the motor power P_m becomes a dependent variable due to the power balance requirement imposed at each time step,

$$P_{e,k} + P_{m,k} = P_{dem,k} \,. \tag{9}$$

The instantaneous cost is the weighted sum of the fuel consumption, NOx emission, PM emission, and a penalty for SOC deviation; it takes the form

$$g = W_{fuel} + \mu \cdot W_{NOx} + \nu \cdot W_{PM} + \alpha \cdot M \,, \tag{10}$$

where μ and ν are the weighting factors on the emissions, and M penalizes the SOC deviation, which is measured by a quadratic distance between the current SOC value and the SOC reference point

$$M = \left(SOC - SOC_{ref}\right)^2 . \quad (11)$$

The integral constraint on charge sustenance is moved into the cost function so that the SOC depletion can also be minimized. To satisfy the charge-sustaining constraint, the weighting factor α is chosen to ensure the SOC is operated in an allowable region. The optimization is subject to equality constraints (system equations) described in Section 2, and the following inequality constraints:

$$\begin{aligned} &\omega_{e\_\min} \le \omega_{e,k} \le \omega_{e\_\max} \\ &\tau_{e\_\min}\left(\omega_{e,k}\right) \le \tau_{e,k} \le \tau_{e\_\max}\left(\omega_{e,k}\right) \\ &\tau_{m\_\min}\left(\omega_{m,k}, SOC_k\right) \le \tau_{m,k} \le \tau_{m\_\max}\left(\omega_{m,k}, SOC_k\right) \\ &SOC_{\min} \le SOC_k \le SOC_{\max} \end{aligned} . \quad (12)$$

4.2 Stochastic Dynamic Programming Approach

Stochastic dynamic programming has been extensively studied in the literature. It can handle constrained nonlinear optimization problems under uncertainties [10]. In this study, an approximate policy iteration algorithm is used.

For our hybrid electric vehicle problem, the state vector $x = (SOC, \omega_{wh}, P_{dem})$ forms a three-dimensional state space, where SOC and ω_{wh} originally take on continuous values, and P_{dem} has finite values as shown in (3). To solve the SDP problem, we discretize SOC as $SOC \in \{SOC^1, SOC^2, ..., SOC^{N_s}\}$, and the wheel speed as (4). The total state space is then represented by finite grids $\{x^i, i = 1,2,..., N_s N_\omega N_p\}$. The control variable $u = P_e$ is also discretized into $\{P_e^1, P_e^2, ..., P_e^{N_u}\}$.

Based on the Bellman's optimality equation, the SDP problem is solved by a policy iteration algorithm. The policy iteration conducts a policy evaluation step and a policy improvement step in an iterative manner until the optimal cost function converges. In the policy evaluation step, given a control policy π, we calculate the corresponding cost function $J_\pi(x)$ by iteratively updating the Bellman equation

$$J_\pi^{s+1}(x^i) = g(x^i, \pi(x^i)) + \underset{w}{E}\left\{\gamma J_\pi^s\left(x'\right)\right\} \quad (13)$$

for all i, where s is the iteration number, and x' is the new state, i.e., $x' = f\left(x^i, \pi(x^i), w\right)$, given by (1), (2), and (5). However, the first two components of state x', (SOC, ω_{wh}), do not necessarily fall exactly on the state grid. In this case, a linear interpolation of the cost function along the first two dimensions is used. In order to accelerate the computations, only a fixed number of iterations are performed regardless of the convergence of the estimated cost function. This truncated policy-evaluation method has been shown to reduce the computation time effectively [11].

In the policy improvement step, the improved policy is found through the following equation

$$\pi'(x^i) = \underset{u \in U(x^i)}{\operatorname{argmin}}\left[g(x^i, u) + \underset{w}{E}\left\{\gamma J_\pi(x')\right\}\right] \quad (14)$$

for all i, where J_π is the approximate cost function obtained from the policy evaluation step. Each of the minimizations is performed subject to the constraints shown in (12). After the new policy is obtained, we go back to the policy evaluation step to update the cost function by using the new policy. This iterative process is repeated, until J_π converges within a selected tolerance level.

5. Optimization and Simulation Results

5.1 SDP Results

The stochastic optimization procedure is summarized in Figure 3. Four representative driving cycles are chosen to construct the observation samples of the stochastic system. The Markov model is built by estimating the transition probabilities from the sample cycle data as described in Section 2. An example of the estimated transition probabilities is shown in Figure 4. Since the power demand has 30 discrete states, the transition probability for a given wheel speed is a 30×30 matrix that maps the current power demand to the next power demand. It can be seen that the power demand is highly correlated as the probabilities are centered around the diagonal axis.

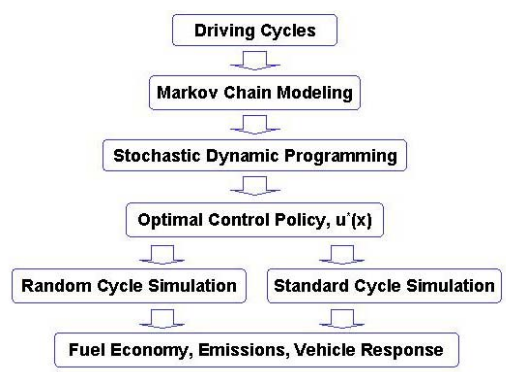

Figure 3: Stochastic dynamic optimization process

In the SDP optimization, the weighting factors $\mu = 40$ and $\nu = 800$ are selected based on the tradeoff analysis conducted in [4]. The discount factor is $\gamma = 0.95$ and α is chosen as 600. The policy iteration procedure is terminated when the maximum absolute value of the difference of two successive cost functions is less than 0.01. The resulting policy takes the form of a look-up table, $\pi^*(SOC, \omega_{wh}, P_{dem})$, i.e., the optimal engine power P_e^* is a function (look-up table) of SOC, wheel speed, and driver power demand. The desired motor power can then be calculated from $P_m^* = P_{dem} - P_e^*$. To illustrate the control policy more clearly, the optimal engine power P_e^* is transformed to an optimal power-split-ratio (PSR) defined as $PSR^* = P_e^* / P_{dem}$. The PSR shows four possible operating modes: motor-only ($PSR = 0$), engine-only ($PSR = 1$), power-assist ($0 < PSR < 1$), and recharging ($PSR > 1$). For better visualization, we plot PSR against P_{dem} and SOC in Figure 5. The figure shows that the optimal policy is to use motor-only mode at extremely low P_{dem} , the recharging mode when the power demand is low/medium and the power-assist mode when the power demand is high. In addition, note that PSR increases as SOC decreases, in order to ensure charge sustenance.

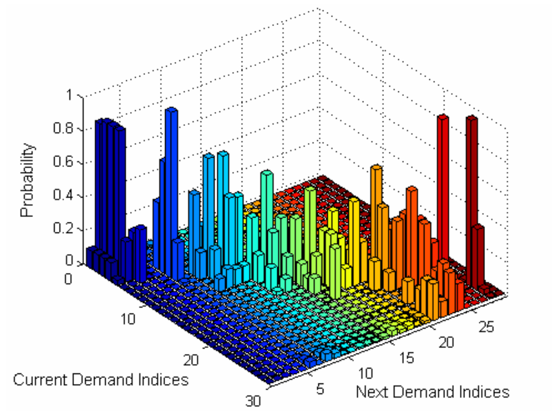

Figure 4: Transition probability of power demand at ω_{wh} =39 (rad/s)

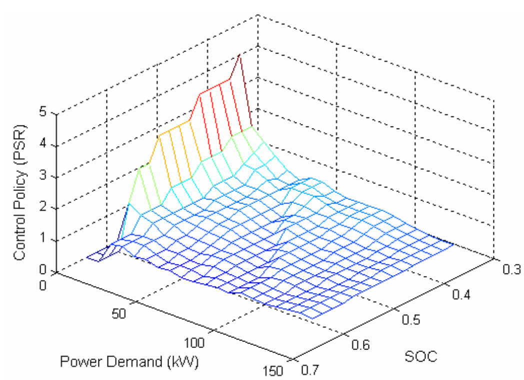

Figure 5: Optimal SDP policy at ω_{wh} =44 (rad/s)

5.2 Simulation Results

After the optimal policy is obtained, we first perform random cycle simulations by using the simplified model shown in Figure 2. A random power demand trajectory can be generated by the Markov model. More specifically, at time k , the current power demand and wheel speed (P_{dem}, ω_{wh}) are interpolated to the nearest grid point $(P_{dem}^i, \omega_{wh}^l)$. The corresponding probability of the next power demand $\Pr\left\{P_{dem}^j\right\}$ is obtained. Using a uniform random number generator, the power demand at time $k+1$ is then determined. The discrete-time simulation updates the system dynamics and generates the control based on SDP policy every sampling time. Since the power demand is stochastic, the wheel speed (vehicle speed) also becomes stochastic and can emulate diverse driving scenarios. Figure 6 shows an example of this random cycle simulation. The simulation time is 1500 seconds and the initial conditions are $SOC_0 = 0.55$, $\omega_{wh,0} = 0$, and $P_{dem,0} = 0$. From Figure 6, it can be seen that the simulated power demands from the Markov chain lead to mixed highway, suburban, and urban driving patterns in terms of wheel speed. This demonstrates that the Markov stochastic modeling can generate diverse driving cycles which resemble real driving conditions. Our control policy is optimized over all the possible random cycles in an average sense.

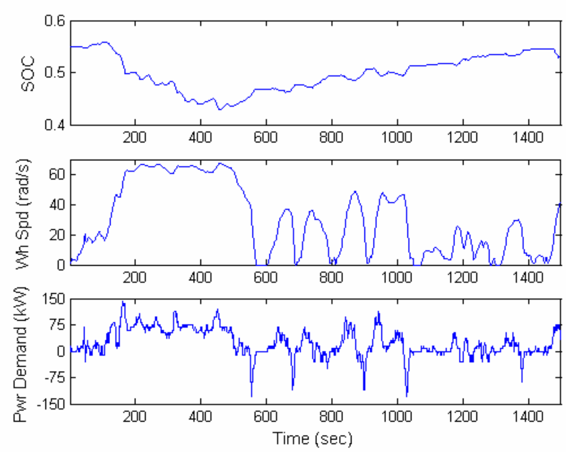

Figure 6: Random cycle simulation

When a vehicle is officially evaluated, it is driven over pre-determined cycles. To illustrate the effectiveness of the proposed control design approach, a continuous-time standard driving cycle simulation is performed. The control policy derived from SDP is integrated in the original, detailed HEV model [7]. The model is implemented in the SIMULINK environment, as presented in Figure 7. Instead of using a Markov driver, a PI (proportional-integral) controller replaces the driver and is used to ensure the specified velocity profile is followed. The driver model outputs the acceleration pedal and braking pedal signals, which are subsequently interpreted as a power demand for the power management controller. This simulation set-up allows us to study the performance of control algorithms under realistic driving conditions. Through the continuous-time simulation, the performance of the control policy from SDP is compared with our prior work over different driving cycles as given in Table 2. The "Rule-Based (DDP)" refers to a rule-based control strategy trained based on the results of deterministic dynamic programming [4]. Compared with the rule-based results, the SDP control strategy achieves better performance over most of the test cycles. Since the rule-based controller was trained on the UDDSHDV cycle, the relative performance of SDP is slightly worse on that cycle. For driving cycles that were not used in the training of the Markov chain, SDP also shows better performance over rule-based control.

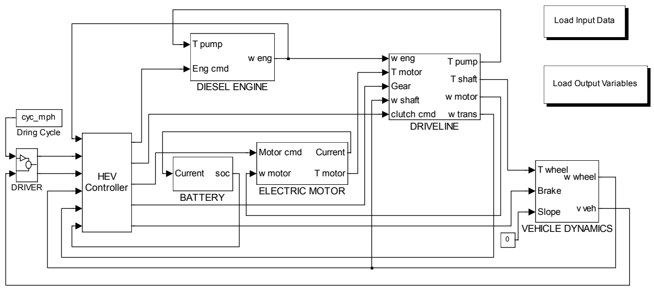

Figure 7: Hybrid-electric vehicle simulation in SIMULINK

Table 2: Simulation results of SDP control policy ($fuel + 40 \cdot NOx + 800 \cdot PM$ [g/mile])

Driving cycle	Rule-Based (DDP)	SDP
UDDSHDV * †	793.2	796.2
WVUINTER *	896.0	885.4
WVUSUB *	582.2	587.3
WVUCITY *	494.1	477.9
NYCCOMP	401.2	393.5
HWFET	958.7	941.0
SC03	1011.8	993.5

* Cycles that used as training sets for learning the Markov model.

† Cycle that used to apply DDP and train the rule-based control.

6. Conclusions

A design method for the power management control algorithm for hybrid electric vehicles is developed by using Markov chain modeling and stochastic dynamic programming techniques. The driver power demand is modeled as a Markov process to represent the future uncertainty of the driver power request under diverse driving conditions. As opposed to deterministic optimization over a given driving cycle, the stochastic approach optimizes the control policy over a family of diverse driving patterns. The infinite-horizon SDP solution generates a time-invariant state-dependent power split strategy, which governs the engine and motor operations. The algorithm can be directly implemented in simulations and vehicle testing. Simulation results indicate that the SDP control strategy achieves improved performance in most of testing scenarios over the sub-optimal rule-based control strategy which is trained based on deterministic DP-results. Furthermore, the proposed approach provides a directly implementable control design path, which is highly desirable because of its potential for a fully integrated optimal design and control process.

Acknowledgments

This research is supported by the U.S. Army TARDEC under the contract DAAE07-98-C-R-L008.

References

[1] B. M. Baumann, G. N. Washington, B. C. Glenn, and G. Rizzoni, "Mechatronic Design and Control of Hybrid Electric Vehicles," *IEEE/ASME Transactions on Mechatronics*, vol. 5, pp. 58-72, 2000.

[2] N. Schouten, M. Salman, and N. Kheir, "Fuzzy Logic Control for Parallel Hybrid Vehicles," *IEEE Transactions on Control Systems Technology*, vol. 10, pp. 460-468, 2002.

[3] G. Paganelli, G. Ercole, A. Brahma, Y. Guezennec, and G. Rizzoni, "General Supervisory Control Policy for the Energy Optimization of Charge-Sustaining Hybrid Electric Vehicles," *JSAE Review*, vol. 22, pp. 511-518, 2001.

[4] C. C. Lin, H. Peng, J. W. Grizzle, and J. Kang, "Power Management Strategy for a Parallel Hybrid Electric Truck," *IEEE Transactions on Control Systems Technology*, vol. 11, pp. 839-849, 2003.

[5] C. C. Lin, H. Peng, J. W. Grizzle, J. Liu, and M. Busdiecker, "Control System Development of an Advanced-Technology Medium-Duty Hybrid Electric Truck," *SAE Paper 2003-01-3369*, 2003.

[6] I. Kolmanovsky, I. Siverguina, and B. Lygoe, "Optimization of Powertrain Operation Policy for Feasibility Assessment and Calibration: Stochastic Dynamic Programming Approach," presented at American Control Conference, Anchorage, Alaska, 2002.

[7] C. C. Lin, Z. S. Filipi, Y. Wang, L. S. Louca, H. Peng, D. N. Assanis, and J. L. Stein, "Integrated, Feed-Forward Hybrid Electric Vehicle Simulation in SIMULINK and its Use for Power Management Studies," *SAE Paper No. 2001-01-1334*, 2001.

[8] I. Kolmanovsky, M. Nieuwstadt, and J. Sun, "Optimization of Complex Powertrain Systems for Fuel Economy and Emissions," presented at IEEE International Conference on Control Applications, Hawaii, 1999.

[9] J. Dai, "Isolated word recognition using Markov chain models," *IEEE Transactions on Speech and Audio Processing*, vol. 3, pp. 458 - 463, 1995.

[10]D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, Mass.: Athena Scientific, 1995.

[11]M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York: J. Wiley, 1994.