Path dependent receding horizon control policies for hybrid electric vehicles

18th IEEE International Conference on Control Applications
Part of 2009 IEEE Multi-conference on Systems and Control
Saint Petersburg, Russia, July 8-10, 2009

Path Dependent Receding Horizon Control Policies for Hybrid Electric Vehicles

Georgia-Evangelia Katsargyri, Ilya Kolmanovsky, John Michelini,
Ming Kuang, Anthony Phillips, Michael Rinehart and Munther Dahleh

***Abstract*— Future Hybrid Electric Vehicles (HEVs) may use path-dependent operating policies to improve fuel economy. In our previous work, we developed a dynamic programming (DP) algorithm for prescribing the battery State of Charge (SoC) set-point, which in combination with a novel approach of route decomposition, has been shown to reduce fuel consumption over selected routes. In this paper, we propose and illustrate a receding horizon control (RHC) strategy for the on-board optimization of the fuel consumption. As compared to the DP approach, the computational requirements of the RHC strategy are lower. In addition, the RHC strategy is capable of correcting for differences between characteristics of a predicted route and a route actually traveled. Our numerical results indicate that the fuel economy potential of the RHC solution can approach that of the DP solution.**

I. INTRODUCTION

To reduce fuel consumption, the control of Hybrid Electric Vehicles (HEVs) may be tied to an expected (or to a specified by the driver) traveling route [1], [2], [3]. Utilizing route information, including road characteristics and traffic conditions, the control of the battery charging and discharging can be optimized for a specific route-to-be-traveled. The proliferation of GPS-based navigational systems and digital maps in the modern vehicles facilitates the application of the *path dependent* control methods for HEVs.

The topic of driving condition dependent HEV control has been actively researched in recent years, see e.g., [4], [5], [6], [7] and references therein. Many existing approaches utilize driving pattern recognition to determine on-line control strategy parameters. Other approaches exploit the capability of recurrent neural networks, after appropriate training, to implicitly capture driving pattern information in rendering control decisions [8]. Dynamic optimization along an anticipated vehicle route has been also considered in [9], [3], [10]. See also [11] and references therein for a discussion of approaches to predicting the route to be traveled.

In this paper, we propose a Receding Horizon Control (RHC) algorithm for prescribing the set-point for the battery State of Charge (SoC) along a route being traveled. In our approach, the original route-to-go with a large number of road segments is replaced by a virtual route with a smaller number of segments, over which the optimization of the battery SoC set-point sequence can be rapidly performed. The first element of the optimal SoC sequence from the virtual route is applied as the SoC set-point for the current segment of the actual route. The optimization is repeated at the beginning of each new route segment. Numerical results will be reported which show that a significant fraction of the maximum fuel consumption reduction benefit can be achieved with the RHC solution, yet at a fraction of the computational cost. Further, changes in the route being traveled, including road and traffic conditions, can be taken into account in the RHC solution.

This work was supported by Ford Motor Company.

G.E. Katsargyri, M. Rinehart and M. Dahleh are with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, MA 02139, USA. {gkats,mdrine,dahleh}@mit.edu

I. Kolmanovsky and J. Michelini are with Ford Research and Advanced Engineering, Ford Motor Company, Dearborn, MI 48124, USA. {ikolmano,jmichel1}@ford.com

M. Kuang and A. Phillips are with the Sustainable Mobility Technologies, Ford Motor Company, Dearborn, MI 48120, USA {mkuang,aphilli8}@ford.com

II. BACKGROUND

In our previous work [10], we discussed a Dynamic Programming (DP) algorithm, which in combination with a novel approach of route decomposition, minimizes fuel consumption over a pre-planned or predicted route. The route is decomposed into a collection of $i = 1,...,N$ road segments connected to each other in series and linking the origin (O) to the destination (D) (Fig.1). The segments may have different length, in order to provide more efficient aggregation of the relevant road conditions. The segmentation criteria used in our previous work [10] were related to changes in road and traffic conditions, e.g., changes in road grade and average traveling speed, and also to the ability of the HEV controller to track the prescribed (feasible) SoC set-points within a single route segment.

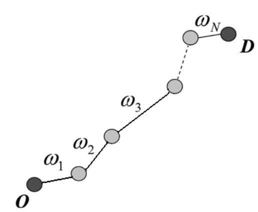

Fig. 1. Route segmentation. The ω_i denotes the fuel consumed over the ith segment of the route.

The HEV configuration that we considered in [10] is also adopted in this paper, see Fig. 2. This is a power-split (i.e., series-parallel) HEV configuration used in the Ford Escape vehicle. Our approach can also be applied to other HEV configurations or even to path-dependent fuel consumption optimization for a non-hybrid vehicle, e.g., one considered in [12].

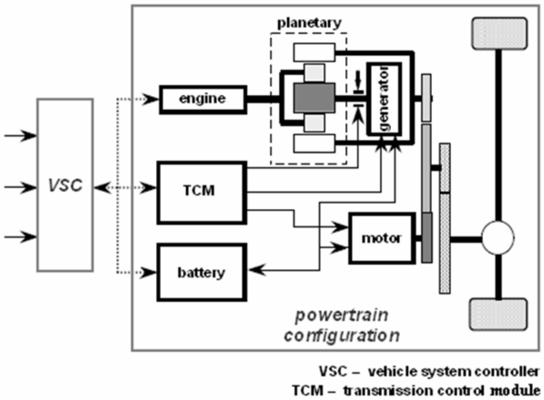

Fig. 2. Power-split HEV powertrain system which includes an internal combustion engine, a generator, an electric motor, a battery and a planetary transmission.

A hierarchical controller architecture was proposed in [10] for the path-dependent HEV control. The upper level controller prescribes the SoC set-point and the lower level controller controls power flows within HEV to satisfy the driver wheel power request and ensure that the battery SoC tracks the specified set-point.

In principle, the upper level controller can vary the SoC set-point as a function of time. While this may be feasible over a short planning horizon, over longer planning horizons the computational effort may be large and, furthermore, the road and traffic information may only be available in aggregated/averaged form over longer road portions. We thus choose to decompose the route into segments, as illustrated in Fig.1, and prescribe constant SoC set-points within each of these segments.

For the lower level controller we adopt, as in [10], a standard HEV controller available within the PSAT simulation environment of the Argonne National Laboratory. We also used PSAT environment for HEV simulations. The low level controller controls the SoC to the set-point, $SoC_d(i)$, prescribed by the upper level controller within each road segment i, while satisfying driver wheel power requests within each road segment. The upper level controller in [10] was based on a DP algorithm, which is expressed by the following relation:

$$J^*(x) = \min_{SoC_d}\{J^*\big(F(x, SoC_d)\big) + \omega(x, SoC_d)\}, \tag{1}$$

$$J^*(x_f) = 0, \tag{2}$$

where J^* is the value function in our optimization problem and $SoC_d = SOC_d(x)$ is the manipulated variable. Here, $x = [i, SoC]^{\mathrm{T}}$ is the state vector with the elements being the route segment number, i, $i = 1, \cdots, N$, and the SoC at the beginning of the segment. The vector x_f represents the final state, which is the state at the destination, and F is a nonlinear state transition function, which generates the next state from the precedent one. If the set-point, SoC_d is feasible, then for $x = [i, SoC]^{\mathrm{T}}$, it follows that $F(x, SoC_d) = [i+1, SoC_d]^{\mathrm{T}}$. Finally, $\omega(x, SoC_d)$ denotes the expected fuel consumption for the state x and the battery SoC set-point, SoC_d.

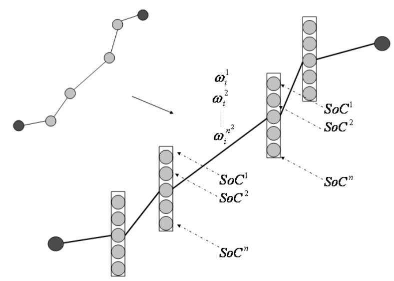

Fig. 3. SoC quantization.

For the on-line implementation, a regression model may be used to estimate the expected fuel consumption, $\omega(x, SOC_d)$, as a function of battery $SoC = x(2)$ at the beginning of the segment, battery SoC set-point, SoC_d, within the segment, and $i = x(1)$th segment properties (grade, length, features of vehicle speed trajectory such as average, maximum and minimum speeds and accelerations, etc.), The fuel consumption model may also include a dependence on the driver style (passive or aggressive) which may be inferred on-line from the variance of acceleration pedal input and vehicle speed. Either vehicle measurement data over different roads or the results of vehicle simulations may be used to develop such a fuel consumption model.

For a given SoC at the beginning of the ith segment, not all SoC_d are feasible, i.e., can be tracked within a tolerance of 0.5 percent before the end of the segment. Such infeasible values of SoC_d are eliminated from the optimization by artificially assigning high value to the fuel consumption. Thus we only consider prescribing battery SoC set-points that can be accurately tracked before the end of each road segment.

Given that SoC is a continuously-valued quantity, for the numerical implementation of DP we quantize (grid) the values of SoC. With a quantization of the form $SoC_i \in \{SoC^1, SoC^2, ..., SoC^n\}$, $SoC^1 \leq SoC^2 \leq ... \leq SoC^n$, every segment i of the route can be associated with n^2 possible pairs of initial SoC and final SoC, i.e., (SoC_i, SoC_{i+1}), and thus with n^2 possible values for the expected fuel consumption $\{\omega_i^1, \omega_i^2, ..., \omega_i^{n^2}\}$ (Fig. 3). Eleven uniformly spaced grid points $SoC_i \in \{\{44, 45, ..., 55\}$ were used in the numerical examples in this paper.

A particular route that we investigated as an example in our previous work [10] is shown in Fig. 4. For this route, O denotes the origin, D denotes the destination and $i = 1, 2, ..., 7$ are the road segments into which the route was decomposed. Length and grade information for each road segment, and a predicted vehicle speed trajectory for the entire route, that consisted of ramp-like changes and constant vehicle speed intervals, were given.

Segment	1	2	3	4	5	6	7
Length (miles)	0.87	0.68	0.74	0.68	1.02	0.59	0.42
Grade (%)	0	0	0	0	0	0	0

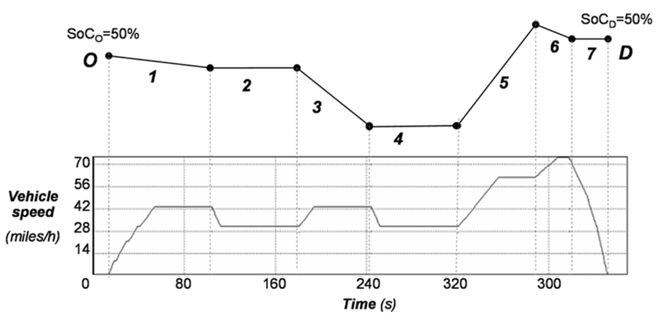

Fig. 4. Example route - Zero grade.

The SoC at the origin was assumed to be $SoC_O = 50\%$. It was desired to sustain the charge in the battery. The charge sustainability is achieved if the equivalent energy of ΔSoC between the origin and the destination is within 1% of the total fuel energy required for completing the trip. We therefore set the SoC at the destination node to be equal to $SoC_D = 50\%$.

Table I summarizes the effect that route segmentation and DP implementation had on fuel consumption. Specifically, with constant SoC set-point during the trip ("No SOC control" case), the total fuel consumption along the route was $0.37kg$. On the other hand, decomposing the route into segments and optimally prescribing their SoC set-points reduced the fuel consumption to $0.32kg$, a 13.5% total fuel consumption reduction. This fuel consumption reduction benefit is specific to the selected route. The potential benefits of varying SoC set-point within the route may be different depending on the route.

Remark 1: The form of the fuel consumption model can significantly influence the computations. For instance, using a model $\omega(x, SoC_d)$ that contains only linear and quadratic terms in $x(2) = SoC$ and SoC_d facilitates the use of standard Quadratic Programming (QP) solvers for the optimization and obviates the need to resort to the grid search of DP. The application of QP solvers will be addressed in the future publications.

TABLE I

FUEL CONSUMPTION IMPROVEMENT USING DP.

FUEL SAVINGS 13.5%	***Total fuel consumption*** *(kg)*	***SoC_d sequence*** *(%)*
No SoC control	0.37	50-50-50-50-50-50-50-50
DP SoC control	0.32	50-52-50-48-46-46-44-50

III. THE RECEDING HORIZON CONTROL APPROACH

The DP approach is helpful in delineating potential benefits of adjusting SoC set-point according to the path being traveled. In implementing this approach, two issues need to be considered. They include the computing effort and the dependence of the optimal control on the route characteristics, which may not be accurately known in advance.

Note that the computing effort of the DP depends on the number of states in the model. In our case, the model used for optimization contained only a single *vehicle* state, *SoC*. Hence, computing the optimal control on-line using DP may be feasible for routes with a small number of segments. Still, the DP approach may become computationally prohibitive for on-board applications if the route contains many segments. Since our segmentation of the route is based on changes in road and traffic conditions [10], routes over which optimization needs to be performed may contain a large number of segments.

Another issue is that we only have at our disposal a predicted route with anticipated driving characteristics along each segment, while the characteristics of the route actually traveled may turn out to be different. Variability in road and traffic conditions or in the driver's choices may easily result in such deviations. Therefore, the control policy has to either be re-computed or corrected on a regular basis to account for these changes.

Our approach to deal with these two issues relies on the use of a receding horizon implementation with virtual terminal road segment. In this approach, we assume that the vehicle travels along a route, L, with N segments, where N is large. Suppose that the vehicle is at the beginning of the kth segment of L. Then the remaining part of the route with $N-k$

segments, $L(k)$, is replaced by a virtual route, $L'(k)$, with n_c+1 segments, where $n_c << N$ segments are taken from the original route while the last, (n_c+1)st segment, is a virtual terminal segment. This virtual terminal segment is intended to represent, *on average*, the characteristics of the omitted $N-k-n_c$ segments of the original route. The set-point for the battery State of Charge, $SoC_d^*(l)$, $l=1,2,\cdots,n_c+1$ is determined using the DP and assuming that the vehicle travels along the virtual route $L'(k)$. After the DP optimization is completed, an assignment $SOC_d(k)=SoC_d^*(1)$ is made, that is the SoC set-point for the kth segment of the actual route is the SoC set-point for the first segment of the virtual route. Once the vehicle progresses to the beginning of the $(k+1)$st segment, the optimization procedure is repeated. This process continues until a segment corresponding to a sufficiently large value of k is reached for which $L(k)$ consists of just n_c+1 segments. At this point, a conventional DP is solved for this remaining part of the route once.

Note that with the receding horizon approach, the optimization is always performed over the route with n_c+1 segments and thus requires significantly less computing time and effort, compared with DP, if $n_c << N$. Further, any changes in the characteristics of the route being traveled can be accounted each time the optimization is repeated at the beginning of the new segment.

We provide a more detailed illustration of this approach next.

IV. Detailed Illustration

For the route featured in Fig. 4, we choose our horizon to be $n_c=2$ segments. Thus at every step of the procedure, we consider and optimize a 3-segment route, which consists of the current segment, the next segment and a virtual third segment that represents on average the remaining part of the original route.

More specifically, given the route of Fig. 4, we start with the first step at the beginning of the trip and we consider a virtual route consisting of segments 1, 2 and a third one which is characterized by the total length, the average speed and average grade of segments $3,4,5,6,7$. Note that the virtual segment includes acceleration and decelerations portions from inital/to final vehicle speed values. See Figure 5. On this 3-segment virtual route we apply the DP algorithm described in Section II in order to obtain an SoC sequence that minimizes fuel consumption. The DP optimization results in $50 \rightarrow 50 \rightarrow 48 \rightarrow 50$ as the optimal SoC sequence. This sequence suggests that 50% should be used as the SoC target on segment 1. We let the vehicle travel on segment 1 with that SoC target of 50% and, when the vehicle arrives at the beginning of segment 2, we proceed with the second step.

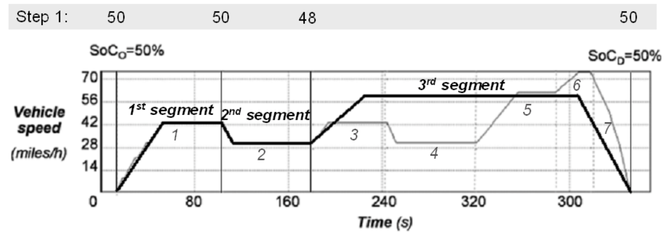

Fig. 5. Receding Horizon Control - Step 1.

Step 2 is similar to Step 1 with the difference that the route over which the optimization is performed has changed. Specifically, the beginning of the route now coincides with the beginning of segment 2. The new route consists of segments 2, 3 and a third virtual segment, that reflects on average the properties of segments $4,5,6,7$. The first element of SoC sequence that DP provides for the virtual 3-segment route at Step 2 is set as the SoC target for driving segment 2.

The procedure continues until the final Step is reached, where the vehicle has to travel just three segments before the destination so that the route-to-optimize includes the last three segments of the original route. At that point we simply compute the optimal SoC sequence for the route consisting of three terminal segments of the original route (i.e., for the remainder of the trip).

V. Results

A. Route with zero grade

In order to illustrate the RHC results and present a comparison with the DP solution, we use again the route of Fig. 4. We first consider the grade of the entire route to be zero and we initialize the SoC at the origin and the destination to be $SoC_O=SoC_D=50\%$. Table II summarizes the effect that RHC and DP implementations have on fuel consumption. Specifically, in the case that the only requirement is $SoC_D=SoC_O=50\%$ (CS requirement) if the SoC set-point is not varied during the trip, the total fuel consumption along the route is $0.37kg$. On the other hand, decomposing the route into segments and optimally prescribing their SoC set-points reduces the fuel consumption with either of the two methods (RHC, DP). With the RHC, the total fuel consumption reduces to $0.33kg$, which corresponds to a 10.8% total fuel consumption decrease compared to the case of "No SoC Control" (this case corresponds to a constant SoC set-point). With the DP method, the reduction in the total fuel consumption is $0.32kg$, which is a 13.5% total fuel consumption reduction as compared to the "No SoC Control"

TABLE II

FUEL CONSUMPTION IMPROVEMENT USING RHC AND DP ON ROUTE WITH ZERO GRADE.

	Total fuel consumption *(kg)*	*SoC_d* ***sequence*** *(%)*	***Fuel savings*** *(%)*
No SoC control	0.37	50-50-50-50-50-50-50-50	0
RHC SoC control	0.33	50-50-48-52-50-50-44-50	10.8
DP SoC control	0.32	50-52-50-48-46-46-44-50	13.5

TABLE III

FUEL CONSUMPTION IMPROVEMENT USING RHC AND DP ON ROUTE WITH NON-ZERO GRADE.

	Total fuel consumption *(kg)*	*SoC_d* ***sequence*** *(%)*	***Fuel savings*** *(%)*
No SoC control	0.37	50-50-50-50-50-50-50-50	0
RHC SoC control	0.34	50-50-48-54-52-46-44-50	8.1
DP SoC control	0.33	50-52-50-54-48-48-44-50	10.8

case. The RHC solution thus achieves a substantial fraction of the benefit of the DP solution for this particular route.

We note that the two methods (DP, RHC) produced different SoC sequences since at each step the RHC method optimizes the fuel consumption over a virtual route which is an approximation of the route-to-go, while the DP method optimizes for the actual route.

B. Non-zero grade route

Segment	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Grade (%)	0	0	0	1	0	-1	0

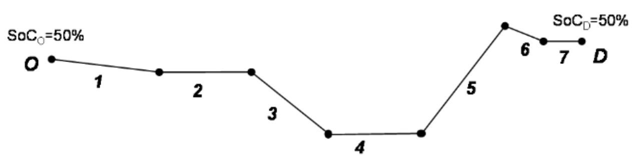

Fig. 6. Example route - Non-zero grade.

We now consider the case where road grade information is available. We insert a grade of 1% at segment 4 and -1% at segment 6 (Fig. 6). The rest of the route characteristics remain unchanged. Table III summarizes the improvement of fuel economy that RHC and DP offer when the grade is not everywhere zero. Specifically, in the case that the CS requirement ($SoC_D = SoC_O = 50\%$) is our only concern and the SoC set-point is not varied, the total fuel consumption along the route is $0.37kg$. We observe that although the grade of the route has changed, the total fuel consumption has not. This can be explained by the fact that the average grade of the entire route is slightly positive, but still very close to zero. Consequently, in the case of "No SoC Control", the change of the fuel consumption appears to be negligible. On the other hand, in RHC and DP control cases, the fuel consumption has slightly increased. Given that the two algorithms aim to minimize the fuel consumption, it is likely that the minimum fuel consumption for a route with slightly positive grade will be larger and not equal to that of the same route with zero grade. Continuing with the optimization results, the RHC method reduces the total fuel consumption to $0.34kg$, which corresponds to a 8.1% total fuel consumption decrease. With the DP, the total fuel consumption reduces to $0.32kg$, which is a 10.8% total fuel consumption decrease.

These results indicate that the RHC is able to reduce fuel consumption and attain a significant fraction of the achievable benefit, given by the DP solution. On the other hand, the RHC solution is computationally faster and is able to incorporate changes in the route information during the trip, e.g., changes in traffic information, when re-optimizing the control solution. Although certain segment characteristics, such as grade and distance, can be considered as deterministic quantities, the vehicle speed trajectories of different trips over the same segment are not. Factors such as weather and traffic conditions or even the personality and the mood of each driver may affect the speed trajectory on every trip. The ability of the RHC solution to account for these changes while providing a significant fraction of the fuel consumption reduction benefits as compared to the DP approach is thus very beneficial.

C. Longer routes with more segments

The route that we used to illustrate our results in Section V consisted of 7 segments and had a total length of 5 miles. We have also considered a longer, 50 miles route with 20 segments of 2.5 mile length each, which had varying grades (see Fig. 7). The travel speed was assumed to be constant. Over this longer route, RHC with $n_c = 2, 3, 4, 6$ has achieved, respectively, 4.82%, 5.28%, 7.02%, 7.03% of fuel consumption reduction as compared to "No SoC Control" case. The fuel consumption for $n_c = 6$ RHC solution is only slightly worse than of DP solution. See Fig. 8 which compares SoC sequences in these cases. For horizons $n_c \geq 8$, the RHC solution coincides with the DP solution.

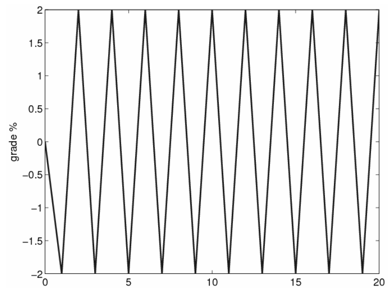

Fig. 7. Grade versus segment number.

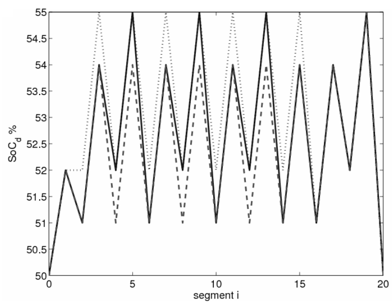

Fig. 8. Sequences of SoC_d for RHC solution with $n_c = 4$ (dotted), for RHC solution with $n_c = 6$ (dashed) and for DP solution (solid).

VI. Conclusions

We proposed a Receding Horizon Control (RHC) approach to path-dependent Hybrid Electric Vehicle (HEV) control with reduced fuel consumption. In our approach, a route with a large number of segments is replaced by a much shorter route with a few initial segments and a *virtual terminal segment*. The virtual terminal segment represents the average properties of the remainder of the original route. Each segment of the original or virtual route is characterized by its length, grade, average speed and possibly other parameters, which permit to estimate the expected fuel consumption over this segment. The optimal State of Charge set-point sequence, which minimizes the expected fuel consumption, is determined for the virtual route. The first element of the optimal SoC set-point sequence for the virtual route is applied for the current segment of the original route. During the travel along the route, the RHC solution is re-computed at the beginning of each segment of the original route until the end of the route is reached. Our simulation results have indicated the promising capability of the RHC solution to achieve the reduction in fuel consumption which is close to that maximally possible, i.e., achievable by the DP solution over the original route. At the same time, the RHC solution has advantages of faster computations and ability to incorporate changes in the route characteristics during vehicle driving.

Many aspects of this work can be further expanded on. They include approaches to compute a virtual segment which best represents on average properties of a multi-segment route and the use of Quadratic Programming solvers.

VII. Acknowledgments

The authors would like to thank Yanan Zhao of Ford Motor Company for her contributions to the fuel consumption modeling.

References

[1] J. Woestman, P. Patil, R. Stunz, and T. Pilutti, “Strategy to use an onboard navigation system for electric and hybrid electric vehicle,” *U.S. Patent 6,487,477*, 2002.

[2] A. Rajagopalan and G. Washington, “Intelligent control of hybrid electric vehicles using gps information,” *SAE Paper 2002-01-1936*, 2002.

[3] Q. Gong and Y. Li, “Trip based optimal power management of plug-in hybrid electric vehicle with advanced traffic modeling,” April 2008.

[4] S.-I. Jeon, S.-T. Jo, Y.-I. Park, and J.-M. Lee, “Multi-mode driving control of a parallel hybrid electric vehicle using driving pattern recognition,” *Journal of dynamic systems, measurement, and control*, vol. 124, 2002.

[5] R. Langari and J.-S. Won, “Intelligent energy management agent for a parallel hybrid vehicle-part i: system architecture and design of the driving situation identification process,” *Vehicular Technology, IEEE Transactions on*, vol. 54, no. 3, pp. 925–934, 2005.

[6] J.-S. Won and R. Langari, “Intelligent energy management agent for a parallel hybrid vehicle-part ii: torque distribution, charge sustenance strategies, and performance results,” *Vehicular Technology, IEEE Transactions on*, vol. 54, no. 3, pp. 935–953, 2005.

[7] Y. Murphey, Z. Chen, L. Kiliaris, J. Park, M. Kuang, A. Masrur, and A. Phillips, “Neural learning of predicting driving environment,” *Proceedings of 2008 International Conference on Neural Networks*.

[8] L. Feldkamp, M. Abou-Nasr, and I. Kolmanovsky, “Recurrent neural network training for energy management of a mild hybrid electric vehicle with an ultra-capacitor,” *Proceedings of 2009 IEEE Workshop on Computational Intelligence in Vehicles and Vehicular Systems*, Nashville, TN, to appear.

[9] E. Finkeldei and M. Back, “Implementing a mpc algorithm in a vehicle with a hybrid powertrain using telematics as a sensor for powertrain control,” *Proceedings of the 1st IFAC Symposium on Advances in Automotive Control*, Salerno, Italy, 2004.

[10] G.-E. Katsargyri, I. Kolmanovsky, J. Michelini, M. Kuang, A. Phillips, M. Rinehart, and M. Dahleh, “Optimally controlling hybrid electric vehicles using path forecasting,” *Proceedings of 2009 American Control Conference*, to appear.

[11] J. Froehlich and J. Krumm, “Route prediction from trip observations,” *Proceedings of SAE World Congress, Detroit, Michigan*, April 2008.

[12] A. Froberg, L. Nielsen, L.-G. Hedstrom, and M. Pettersson, “Controlling gear engagement and disengagement on heavy trucks for minimization of fuel consumption,” *Proceedings of the 1st IFAC Symposium on Advances in Automotive Control*, Salerno, Italy, 2004.