Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 54-10 (2021) 350–355

Energy-efficient powertrain control of an automated and connected power-split HEV in an urban environment

Lucas Koch * **Xiaonan Klingbeil** * **Jakob Andert** *

* *Teaching and Research Area for Mechatronics in Mobile Propulsion, RWTH Aachen University, Germany (E-Mail: koch_luc/klingbeil/andert_j@vka.rwth-aachen.de)*

Abstract: The ongoing trend toward powertrain electrification and vehicle automation enables the exploitation of additional energy saving potential through the joint optimization of the driving trajectory and the hybrid management. To tackle this complex control task within the E-COSM 2021 Benchmark Challenge, we combine a GLOSA algorithm generating a uniform speed profile with an optimization-based hybrid powertrain controller employing an equivalent consumption minimization strategy.

Keywords: eco-driving, connected vehicles, GLOSA, ECMS, automated driving, hybrid vehicles

1. INTRODUCTION

Reducing emissions of greenhouse gases in order to mitigate man-made climate change poses an immense challenge to the mobility sector, which is currently still heavily dependent on fossil fuels. In addition to the progressive change in propulsion technology toward electrification of the powertrain, automated and connected driving is becoming the focus of development due to previously untapped energy-saving potentials resulting from the embedding of additional environmental information in the control of vehicle dynamics (Ma et al. (2019); Wegener et al. (2020)). However, the combination of the high degree of freedom over the powertrain operation of a power-split Hybrid Electric Vehicles (HEVs) with its longitudinal trajectory make the powertrain control task highly complex.

To tackle these challenging issues of next generation powertrain control, a competition of the benchmark problem "optimal control of HEVs under V2X connected environment" is organized within the conference program of IFAC E-COSM 2021. Within this benchmark challenge, a simulation model of a passenger HEV with power-split hybrid system and a connected traffic environment are provided from the organization to validate the developed control strategy.

In our institute for Mechatronics in Mobile Propulsion from the RWTH Aachen University, we developed a unique approach to control the hybrid powertrain using V2X-information with the aim of increasing total vehicle energy efficiency. We focus on the generation of a speed profile with minimal acceleration and deceleration in combination with an optimization-based hybrid management showing that already a relatively simple control logic can achieve significant energy savings when the trajectory is planned in a smart way. The purpose of this is to emphasize the relevance of an efficient driving profile versus an optimized energy management, to put the necessity of their common consideration up for discussion and to share our methodology with researchers in the field of control strategies for hybrid vehicles.

* Research reported in this article was conducted in part in the context of the Hy-Nets4all project (Grant No. EFRE-0801698), supported by the European Regional Development Fund (ERDF). Moreover, the work was partially performed in the Center for Mobile Propulsion (CMP) funded by the German Research Foundation and the German Council of Science and Humanities.

2. RELATED WORK

For the purpose of analyzing energy-saving potentials of HEVs, the entire energy conversion chain of the powertrain from energy storage to vehicle movement can be divided into a portion determined by the driving style ("wheel-to-distance") and a portion determined by the mode of operation of the powertrain components ("tank-to-wheel") (Vahidi and Sciarretta (2018)). Therefore, also operation strategies can be divided into two categories: driving strategy and energy management. Driving strategies include all functions that influence vehicle velocity as well as ensure safety and drivability. Since a hybrid vehicle has more than one energy source, this requires an energy management system to distribute the power required to follow the planned trajectory among the different powertrain components. The overall operating strategy can, according to the operating goals of driving and energy management, select the optimal operating actions, i.e. speed and torque of the power units depending on the hybrid topology, and thereby control the vehicle's longitudinal dynamics.

Different methodologies were investigated for the operation strategies. Heuristic or rule-based operation strategies have been used most widely in the past due to their ability to achieve decent results with a moderate implementation effort. For driving functions, the rule-based (Cooperative) Adaptive Cruise Control ((C)ACC) with a three layer cascaded control system is already integrated in the serial

10.1016/j.ifacol.2021.10.187

production of many OEMs (Pananurak et al. (2009)). The rule-based energy management strategies attempt to reduce the fuel consumption by calibrating the parameters of operation modes and torque split. The operating modes and torque split are typically decided by the requested wheel torque, actual vehicle velocity and battery charge states (Ye et al. (2019)).

Dynamic Programming (DP) is a common method for solving optimal control problems. It provides a globally optimal solution to nonlinear, nonconvex or mixed-integer optimal control problems. As a method of vehicle velocity optimization, for an a priori known route both the energy consumption and the duration of the journey will be considered in a cost function and can be minimized during the solution of the optimization problem. The state-space of the vehicle will be discretized and decision costs will be minimized (Gausemeier et al. (2010)). As energy management system, if the velocity and the altitude profile of a route are known a priori, the globally optimal control and state trajectories for achieving the operating objective (reduction of fuel consumption) can be determined by the DP (Larsson et al. (2014)). In addition to the control variables, the internal state variables of the system and the battery state of charge must also be investigated over the time interval (Uebel et al. (2018)). Apart from the required a priori knowledge, limitations of DP arise from the computational effort. Because of the discretization the state-space of the solution is large and the optimization is computationally expensive.

Besides DP, methods based on Pontyagin's Maximum Principle (PMP) can also be used to solve the optimal control problem of hybrid vehicles energy management through the a priori known route information. Operating strategies based on PMP are called equivalent-based strategies. Here, the sum of the actual fuel consumption and the equivalent fuel consumption of the consumed electric energy of the battery are represented as an objective function. Since fuel energy and electric energy are not directly comparable, the equivalence between the two forms of energy is represented by an equivalence factor (Zhang et al. (2016); Han et al. (2017)). If a global solution of the optimisation is not possible due to missing information about future traffic scenarios, the Equivalent Consumption Minimization Strategy (ECMS) method derived from PMP can determine the equivalence factor depending on the difference between actual and target battery State of Charge (SOC) and solve the optimisation problem online, although this produces a suboptimal result (Onori and Serrao (2011)).

Recently, also data-driven algorithms have entered the focus of research on operating strategies of HEVs by addressing the optimal control problems with Reinforcement Learning (RL). Here, the control task is modeled as a Markov Decision Process consisting of states, actions, rewards and a set of transition probabilities. The RL agent derives its control strategy, called policy, through learning from interaction with its environment. This class of algorithms can potentially achieve results close to the global optimum. Moreover, neither a priori knowledge nor a prediction model is needed. However, disadvantages remain in the high training duration and the strong dependence on hyperparameters (Sutton and Barto (2018)). So far, researchers have utilized RL for different HEV topologies either to determine the torque or power split for a given power request (Lian et al. (2020)) or as a stand-alone control function that directly maps environment information to powertrain actions (Zhu et al. (2021)).

3. METHODOLOGY

The goal is to derive an energy-efficient driving and hybrid management strategy for an automated power-split HEV in a connected, urban traffic scenario utilizing the E-COSM Benchmark Challenge Simulation Framework proposed by Xu et al. (2020). The developed controller determines the torque requests of the powertrain based on information about the ego vehicle's state, the route, the traffic lights ahead and the preceding vehicle. For a detailed description of the modeling of the powertrain and the connected traffic scenario, reference is made to the aforementioned source.

Due to their exponential increase with the velocity, low energy demands can be achieved by driving as constantly and slowly as possible. Furthermore, the avoidance of unnecessary braking events has the effect of minimized energy losses since braking, whether with the hydraulic brake, with cut-off of the combustion engine or recuperation, always involves a dissipation of energy. The "tank-to-wheel" efficiency is depending on the operating mode of the powertrain components, mainly the engine and the two electric motors. While in conventional and electric vehicles the driving trajectory defines the operating points of the engine or motor respectively, the combined hybrid powertrain gives additional degrees of freedom regarding the choice of speeds and torques. However, the components must always deliver the requested power at the wheel. Our hypotheses based on previous studies by Plum et al. (2019) is that the energy saving potentials of the driving trajectory that result in a low power demand at the wheel are higher than the savings from an optimized operation of the engine and motors. Therefore, our highest priority is to generate a driving trajectory that exhibits as little acceleration and deceleration as possible to reach the destination as energy-efficient as possible with a minor increase in travel time while adhering the given traffic lights. Since the traffic lights are completely deterministic, the only cause that requires braking is a slowing preceding vehicle. Hence, we decouple the driving trajectory as far as possible from the preceding vehicle. The operating mode of the powertrain components is not considered in the trajectory planing and is individually optimized subsequently. The resulting functional architecture of the proposed powertrain controller is shown in Fig. 1. It consists of the two main parts, which are the Green Light Optimal Speed Advisory (GLOSA) responsible for the driving trajectory and the powertrain controller that converts the target acceleration into the desired torque of the motor, generator and combustion engine at the output of the function.

3.1 GLOSA Algorithm

The GLOSA algorithm itself can again be split into two functions executed serially. The task of the green phase scheduler is to compute the targeted time at which the upcoming traffic light shall be passed by the ego

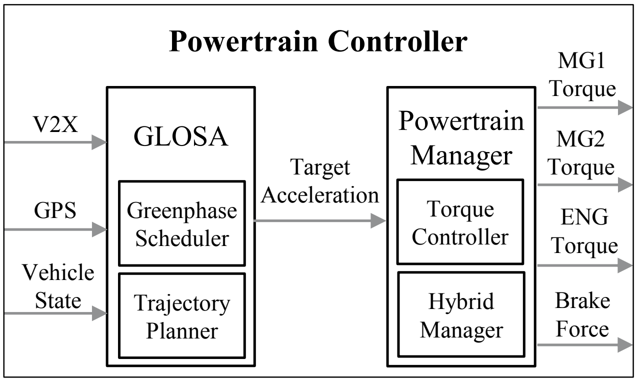

Fig. 1. Functional architecture of proposed powertrain controller

vehicle and by that provide the target time of arrival at the upcoming traffic light. Subsequently, the trajectory planner generates a trajectory that satisfies this desired passing time.

The functionality of the green phase scheduler is visualised in Fig. 2 where the ego vehicle in black has just crossed a traffic light and the algorithm calculates the next target time. Firstly, the green phase of the preceding vehicle is either predicted or, in case the upcoming traffic light has already been passed, sampled from its past trajectory. For the prediction, a constant acceleration until 60 km/h followed by a phase of constant velocity are assumed. If the traffic light, unlike the example in Fig. 2, is red when the preceding vehicle is predicted to arrive, it is assumed that the traffic light is passed at the beginning of the next green phase. This rather simple way of prediction matches the synthetic, longitudinal trajectory of the preceding vehicle accurately for free travel. However, it does not account for other traffic participants in front of the preceding vehicle or queues in front of traffic lights that can cause the preceding vehicle to decelerate before the stop line of the upcoming traffic light since vehicles in front of the preceding one cannot be sensed. Once the passing time and thereby the green phase of the preceding vehicle are predicted or sampled, the target passing time of the upcoming traffic light is calculated for the ego vehicle. To make the trajectory independent from the preceding vehicle in order to minimize the necessity to react to braking events of the preceding vehicle while satisfying the maximum delay criteria, a minimum time-gap of 20 s is considered for setting the target for the ego vehicle as visualised by the red area in Fig. 2. The passing timing is raised iteratively until the average velocity required to pass the upcoming traffic light (cf. the black dotted lines in Fig. 2) deceeds a calibratable velocity threshold lower than the legal speed of 60 km/h. Thereby, the first and last 2 s of each green phase are left out to account for possible controller deviations that might lead to red light violations. The function also constantly monitors the preceding vehicle and once the predicted green phase changes, e.g. due to a queue in front of the traffic light, also the target time for the ego vehicle is recalculated accordingly.

After the scheduling is completed, the trajectory is calculated in order to pass the upcoming traffic light at the predefined time. For this purpose a constant acceleration or deceleration is applied until the velocity is sufficient to pass the traffic light at the stated time with a defined tolerance. Since the calculation of the target does not take into account the distance to the preceding vehicle, it does not yet ensure a collision-free ride. Therefore, the target acceleration is constantly monitored by the intelligent driver model by Treiber et al. (2000). This model ensures a collision-free ride by applying a deceleration when the distance or headway to the preceding vehicle becomes too small.

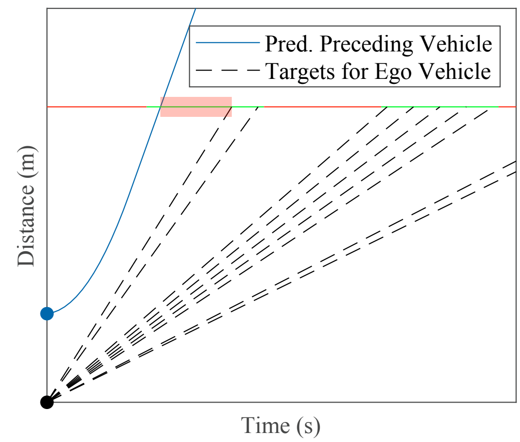

Fig. 2. Example for generation of target traffic light passing time

3.2 Powertrain Manager

The second part of the proposed controller incorporates the two-stage transmission from the target acceleration into the motor and engine torques. At first, the target acceleration is converted into a torque demand at the wheel by considering the driving resistances from rolling resistance, air drag, slope and acceleration. After the calculation of the required wheel torque, it is divided between the required engine, generator and motor in the second part of the powertrain manager, the hybrid manager with ECMS methodology. Since the equivalent factor between fuel consumption and changes of SOC are known, the torque split can be calculated energetically optimal online with the ECMS.

To be able to use the ECMS, the combustion engine fuel consumption and e-motor efficiency maps are required to estimate the fuel equivalent. These data were generated with the simulation model, which is provided from the E-COSM benchmark challenge.

During the optimization, the whole range of the combustion engine speed $n_{ICE,i}$ and torque $tq_{ICE,i}$ is discretized firstly. The speed and torque of the generator $n_{MG1,i}$, $tq_{MG1,i}$ and motor $n_{MG2,i}$, $tq_{MG2,i}$ will be calculated with the the kinematic equations of the planetary gearbox based on the combustion engine operation point and the wheel speed n_{Whl} and torque tq_{Whl}:

$$tq_{MG1,i} = \frac{tq_{ICE,i}}{(i_0 - 1)} \tag{1}$$

$$n_{MG1,i} = \frac{n_{Whl} \cdot i_{Diff} \cdot i_r - n_{ICE,i} \cdot (1 - i_0)}{i_0} \quad (2)$$

$$n_{MG2,i} = n_{Whl} \cdot i_{Diff} \cdot i_m \quad (3)$$

$$tq_{MG2,i} = \frac{\frac{tq_{Whl}}{\eta_{Gear} \cdot i_{Diff}} - \frac{tq_{ICE,i} \cdot \eta_{Gear}) \cdot i_0}{(i_0 - 1) \cdot i_r}}{i_m}. \quad (4)$$

Hereby, i_{Diff}, i_0, i_r and i_m are the gear ratios of the differential, planetary gear, counter gear and reduction gear. η_{Gear} is the estimated efficiency of the differential.

The fuel consumption $m_{fuel,i}$ and the changes of the SOC ΔSoC_i can be calculated utilizing every combination of ICE speed and torque. The equivalent fuel consumption $m_{fuel,equi,i}$ will be calculated with the equivalent factor λ:

$$m_{fuel,equi,i} = m_{fuel,i} + \lambda \cdot \Delta SoC_i \quad (5)$$

where the ΔSoC_i can be calculated with the power of MG1 $P_{MG1,i}$ and MG2 $P_{MG2,i}$ with the auxiliary power P_{AUX} and battery power losses $P_{Batt,loss}$ divided by the battery capacity Cap_{Batt}:

$$\Delta SoC_i = \frac{P_{MG1,i} + P_{MG2,i} + P_{AUX} + P_{Batt,loss}}{Cap_{Batt}} \quad (6)$$

The equivalent fuel consumption is thus calculated for each possible speed and torque combination and it's minimal value decides the optimal control of the vehicle. The operation points of the combustion engine, generator and motor with the minimal equivalent fuel consumption will be transmitted to the vehicle model. Since hydraulic braking always involves energy loss that could have been partly recuperated, the hydraulic brake is not used at all.

4. SIMULATION RESULTS

In the following, the results and the control strategy obtained with the presented control scheme are presented. The E-COSM simulation model provides ten different traces of the same 16.002 km route with varying preceding vehicle and traffic light timings. In Tab. 1 the simulation results for all 10 simulation scenarios are listed. Each simulation run starts with a balanced SOC of 50 % and meets the constraints of maximum speed, distance, headway, red light violation and travel time. With the given equivalent factor between electric and chemical energy of 25, the total fuel consumption is composed of the actual fuel and the weighted difference between initial and final SOC (cf. formula 5). The results reveal a clear trend toward charge increasing through load point shifting since the relatively high equivalent factor favors electric energy stored in the battery. Further, despite using the same overall length, traffic light distances and height profile, a significant variance in total fuel consumption and travel time can be noted without a clear correlation between these two quantities. While the two fastest simulation runs (case 1 and 3) yield the highest total fuel consumption, case 2 and 6 which are only slightly slower are the most efficient ones. Also slow travel times can exhibit both relatively high (case 5) and low (case 4) consumption which underlines the great influence of the traffic situation on the travel time and fuel economy. Overall, the total fuel consumption ranges from 261.8 to 459.1 g, corresponding to 2.15 liter/100km and 3.77 liter/100km.

Table 1. Simulation results in E-COSM Benchmark Challenge Simulator

Case	Fuel (g)	Final SoC (%)	Total Fuel (g)	Travel Time (s)
1	986.8	71.66	445.3	1966
2	726.9	68.6	261.8	1993
3	1428.1	89.77	433.9	1951
4	687	66.09	284.7	2287
5	561.5	55.4	426.4	2397
6	579	61.52	291	1973
7	844.2	69.31	361.6	2359
8	1434.3	89.01	459.1	2190
9	627.7	61.53	339.3	2235
10	694	66.53	280.8	2181

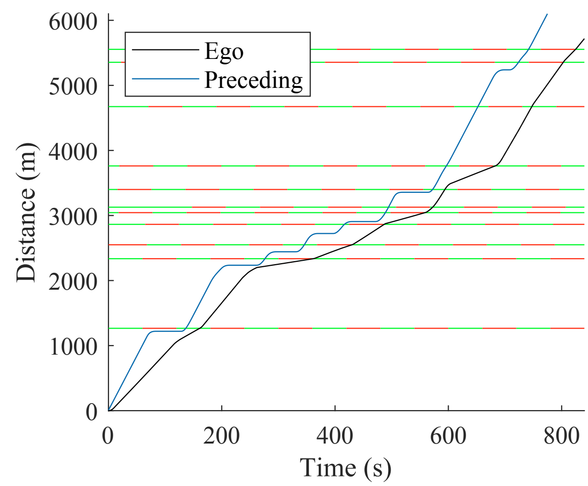

Fig. 3. Vehicle distance over time for simulation case 2

For a deeper analysis of the driving strategy simulation case 2 was selected for visualization. Fig. 3 shows the driven distance of the ego and preceding vehicle over time together with the location and signal phases of the traffic lights. This illustrates the mayor differences in the driving style between the preceding vehicle and the controlled ego vehicle. While the preceding vehicle accelerates to its maximum velocity until it has to stop in front of almost every traffic light, the velocity of the ego vehicle usually changes only once after each traffic light. The time gap in the calculation of the target passing time (cf. Sec. 3) successfully prevents unnecessary acceleration and braking maneuvers. The only exception can be observed before the second traffic light. Here, a queue in front of the traffic light forces the preceding vehicle to stop and wait around 100 m before the stop line. Since the prediction does not consider those cases where the preceding vehicle stops at any distance before the traffic light due to the inability to track the traffic in front of the preceding vehicle, the target passing time constantly stays in the middle of the corresponding green phase. Only when the headway to the preceding vehicle drops below the monitoring limit, the IDM forces the ego vehicle to brake and a new target is calculated. These cases, however, only occur rarely as the time gap is usually sufficient for the preceding vehicle to start accelerating before the ego vehicle arrives at the queue.

Fig. 4 shows the most important quantities of the powertrain for simulation case 2. The velocity and acceleration profiles in the upper plot illustrate the sections of constant velocity between two traffic lights which was already observed in Fig. 3. Apart from some minor control deviations, the vehicle only accelerates or decelerates shortly after passing a traffic light to reach the new target passing time. MG1 regulates the operating point of the combustion engine and charges the battery if the ICE is turned on. Therefore, when the ICE is off, also torque and speed of MG1 are zero. The optimal hybrid management only turns on the combustion engine when a certain wheel power is exceeded and chooses a constant operating point of the ICE at 90 Nm and 1000 1/min. The energy generated by the ICE exceeds the required wheel power most of the time which can be seen by the negative torque demands of MG2. Thus, the remaining energy is stored in the battery and eventually released from the battery back to the wheel later. Despite the energy loss due to double energy conversion, the relatively high, SOC independent equivalent factor favors this kind of powertrain operation that leads to an increasing battery SOC. MG2 compensates for any deviation between the constant power output of the planetary gearbox and the required power at the wheels. During pure electric drive mode, MG2 provides the whole power output.

The lower plot of Fig. 4 visualizes the actual fuel consumption of the ICE, the battery SOC and the total fuel consumption calculated with the equivalent factor. At the beginning, the total fuel becomes negative although the ICE is active. The consumed fuel is partly converted into electrochemical energy in the battery and also a constant downward slope enabled stronger recuperation. In the following section with the combustion engine switched off (around 350 to 800s), the requested power is so low that the SOC level barely decreases. When the ICE is switched on afterwards, the effect of the shifted load point becomes visible most clearly. Between 850 and 1200s, the SOC drops despite the active ICE which means that the power request at the wheel exceeds the power provided by the ICE. Here, the total fuel consumption rises most rapidly. Afterwards, the SOC rises all the way until the vehicle has reached its final destination. Hence, the actual fuel consumed by the ICE rises faster than the total fuel.

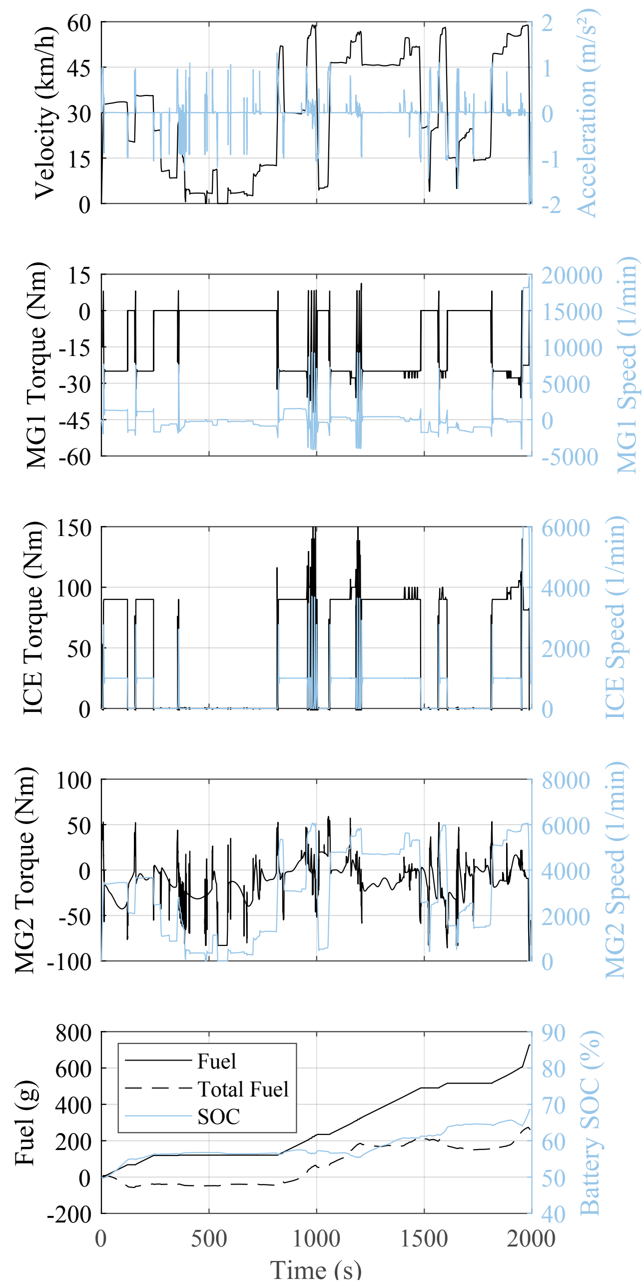

Fig. 4. Simulation results for scenario 2

5. CONCLUSION AND OUTLOOK

In this paper, we have combined a rule-based GLOSA Algorithm with an optimal hybrid management based on ECMS for an energetically optimal operation of a power-split HEV in a connected, urban traffic environment. The simulation results indicate that is is possible to decouple the ego vehicle's trajectory from the preceding vehicle to keep a constant velocity during driving and thereby minimize the requested power at the wheel while achieving an acceptable travel time. With the given equivalent factor between chemical and electrical energy, the optimal hybrid management tends to charging the battery via load point shifting. The combustion engine is only active at speeds above around 30 kilometer/hour and is operated for the most part at a constant operating point.

Future work involves an extension of the prediction algorithm to estimate and react on queues in front of traffic lights. We also target the implementation of a controller based on reinforcement learning for solving the optimal problem directly without the necessity of splitting trajectory generation and hybrid management. Here, the designed controller can serve as a benchmark for any controllers developed in future research.

REFERENCES

Gausemeier, S., Jäker, K.P., and Trächtler, A. (2010). Multi-objective optimization of a vehicle velocity profile by means of dynamic programming. *IFAC Proceedings Volumes*, 43(7), 366–371. doi: https://doi.org/10.3182/20100712-3-DE-2013.00167. 6th IFAC Symposium on Advances in Automotive Control.

Han, J., Kum, D., and Park, Y. (2017). Synthesis of predictive equivalent consumption minimization strategy for hybrid electric vehicles based on closed-form solution of optimal equivalence factor. *IEEE Transactions on Vehicular Technology*, 66(7), 5604–5616. doi: 10.1109/TVT.2017.2660764.

Larsson, V., Johannesson, L., and Egardt, B. (2014). Cubic spline approximations of the dynamic programming cost-to-go in hev energy management problems. In *2014 European Control Conference (ECC)*, 1699–1704. doi: 10.1109/ECC.2014.6862404.

Lian, R., Tan, H., Peng, J., Li, Q., and Wu, Y. (2020). Cross-type transfer for deep reinforcement learning based hybrid electric vehicle energy management. *IEEE Transactions on Vehicular Technology*, 69(8), 8367–8380. doi:10.1109/TVT.2020.2999263.

Ma, F., Yang, Y., Wang, J., Liu, Z., Li, J., Nie, J., Shen, Y., and Wu, L. (2019). Predictive energy-saving optimization based on nonlinear model predictive control for cooperative connected vehicles platoon with v2v communication. *Energy*, 189, 116120. doi: https://doi.org/10.1016/j.energy.2019.116120.

Onori, S. and Serrao, L. (2011). On adaptive-ecms strategies for hybrid electric vehicles. *Int. Sci. Conf. Hybrid Electr. Veh.*, 1–7.

Pananurak, W., Thanok, S., and Parnichkun, M. (2009). Adaptive cruise control for an intelligent vehicle. 1794 – 1799. doi:10.1109/ROBIO.2009.4913274.

Plum, T., Wegener, M., Eisenbarth, M., Ye, Z., Etzold, K., Pischinger, S., and Andert, J. (2019). A simulation-based case study for powertrain efficiency improvement by automated driving functions. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 233.

Sutton, R.S. and Barto, A.G. (2018). *Reinforcement Learning: An Introduction*. The MIT Press, second edition. URL `http://incompleteideas.net/book/the-book-2nd.html`.

Treiber, M., Hennecke, A., and Helbing, D. (2000). Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62, 1805–1824. doi:10.1103/PhysRevE.62.1805.

Uebel, S., Murgovski, N., Tempelhahn, C., and Bäker, B. (2018). Optimal energy management and velocity control of hybrid electric vehicles. *IEEE Transactions on Vehicular Technology*, 67(1), 327–337. doi: 10.1109/TVT.2017.2727680.

Vahidi, A. and Sciarretta, A. (2018). Energy saving potentials of connected and automated vehicles. *Transportation Research Part C: Emerging Technologies*, 95, 822–843.

Wegener, M., Plum, T., Eisenbarth, M., and Andert, J. (2020). Energy saving potentials of modern powertrains utilizing predictive driving algorithms in different traffic scenarios. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 234(4), 992–1005. doi:10.1177/0954407019867172. URL `https://doi.org/10.1177/0954407019867172`.

Xu, F., Tsunogawa, H., Kako, J., Hu, X., Li, S.E., Shen, T., Eriksson, L., and Guardiola, C. (2020). Real-time energy optimization of hevs under connected environment: Ecosm 2021 benchmark problem and a case study.

Ye, Z., Li, K., Stapelbroek, M., Savelsberg, R., Günther, M., and Pischinger, S. (2019). Variable step-size discrete dynamic programming for vehicle speed trajectory optimization. *IEEE Transactions on Intelligent Transportation Systems*, 20(2), 476–484. doi: 10.1109/TITS.2018.2812921.

Zhang, F., Xi, J., and Langari, R. (2016). An adaptive equivalent consumption minimization strategy for parallel hybrid electric vehicle based on fuzzy pi. In *2016 IEEE Intelligent Vehicles Symposium (IV)*, 460–465. doi:10.1109/IVS.2016.7535426.

Zhu, Z., Gupta, S., Gupta, A., and Canova, M. (2021). A deep reinforcement learning framework for eco-driving in connected and automated hybrid electric vehicles.