35th International Electric Vehicle Symposium and Exhibition (EVS35)
Oslo, Norway, June 11-15, 2022

Optimizing Torque Delivery for an Energy Limited Electric Race Car Using Model Predictive Control

Thomas Maull[1], Adriano Schommer[2]
[1] *Williams Advanced Engineering, Grove, Wantage, OX12 0DQ, UK, tom.maull@wae.com*
[2] *Oxford Brookes University, Wheatley Campus, Oxford, OX33 1HX, adriano.schommer@brookes.ac.uk*

Executive Summary

This paper deals with an energy optimized torque controller for the powertrain of an energy-limited electric Formula Student car. Limited battery capacity within electric race car designs requires energy management solutions to minimize race time while simultaneously controlling and managing overall energy consumption to finish the race. An energy-managing torque control algorithm was developed in order to use the finite onboard energy in a manner that optimized lap time performance and decreased energy consumption when energy deficits occurred. The longitudinal dynamics of the vehicle was represented by a linearized first-principles model. The plant was validated against a parameterized electric Formula Student race car model in a commercial lap time simulation software. A Simulink-based Model Predictive Controller (MPC) controller architecture was created to balance energy use requirements with optimum lap time. This controller was then tested against a hardware-limited system as well as a torque-limited system in a constant torque request and varying torque request scenario. The controller decreased the elapsed time to complete a 150m straight line acceleration by 11.4 % over the torque-limited solution, and 13.5 % in a 150m representative Formula Student maneuver.

Keywords: MPC, Formula Student, torque-limited, energy-limited

1 Introduction

In the development of an electric race car, the battery pack sizing often reflects a compromise between range and mass, and therefore ultimate performance. Most electric motorsport is limited by the energy capacity of the battery, rather than the total power capability of the drivetrain. This applies especially to Formula Student (FS) battery systems, which can have capacities of less than 7 kWh. The car is limited to 80kW power by FS regulations, however preliminary lap time simulations indicate that an energy capacity of approximately 10 kWh would be required to finish the endurance event at the maximum regulatory power level. Thus, the optimal use of this energy can produce significant lap time improvements over a constant power or torque limitation. While highly skilled drivers can approximate an optimal strategy by using lift and coast techniques, Formula Student relies on amateur drivers and automated torque control and energy saving methods are preferable.

The need to optimize energy consumption in internal combustion engine (ICE) powered motorsport has been long understood, with associated trade-offs between fuel consumption, fuel mass, range, and available power [1]. This has been analysed in numerous ways with most real-time applications of power optimization independent of driving path [2, 3, 4]. In order to minimize lap time, or maximize average speed, and simultaneously reduce energy consumption, the controller must use energy when it contributes most to acceleration. As drag squares with speed, this occurs at lower velocities [5]. This is exacerbated by the drag created by downforce generating aerodynamic packages commonly used in motorsport competition.

A dynamic optimization problem was created to control this Single Input Multi Output plant. The Model Predictive Controller (MPC) calculates a minimal cost over a receding prediction horizon using a defined plant model at a repeating time step as shown in Figure 1 [6].

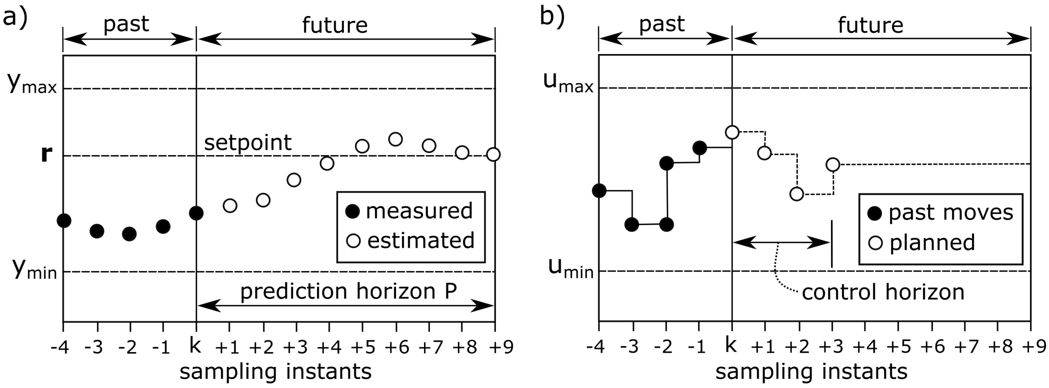

Figure 1: Prediction Horizons within Boundary Limits. Figure adapted from [6].

The manipulated variables u can be held within established limits, as can measured outputs y, allowing MPC to consider the physical constraints of the system. As the optimal output is recomputed at every time step k, MPC tends to incur relatively high computational costs compared to other Proportional, Integral, Derivative (PID) control and Linear Quadratic Regulation (LQR) methods. However, as real-time controllers gain computational power, MPCs are now able to be used in high frequency real-time applications [2, 7]. These methods are of great interest to driverless systems' researchers as they are often used for motion and route planning in driverless motorsport [8, 9].

MPC uses two horizon values in its calculation. The first, P, is the prediction horizon, or number of steps into the future the model predicts and optimizes for. This allows the controller to avoid constraints as well as make decisions while understanding the long term implications of the control action, such as plants with a negative short term response and positive long term response, or in plants with a long delay period [6]. The second horizon is the control horizon, which predicts the next Mcontrol moves in the future. Only the first of these control moves is implemented, the rest are discarded. This is typically a small number, between two and five [6, 7].

Online MPC control is of significant interest to driverless researchers, and has recently been shown to work in driverless systems at $10 - 50\,ms$ sample rates, indicating that this control methodology is possible in real time onboard systems [8, 9]. This paper will introduce a novel feedback control methodology for energy-limited race vehicles, that uses a model based control algorithm to reduce energy use in an optimal fashion when consumption is greater than targeted. This real-time algorithm is adaptive to unforeseen changes in race conditions, such as a safety car or driver change, as it will repeatedly re-target the optimal solution [10]. Similar control methodologies are employed in powertrain and energy management control strategies implemented in Formula 1, but a gap in the research exists when applied to a purely electric race vehicles [3, 11, 12].

In this paper, a motorsport specific, model based predictive energy managing torque control system is described. A similar predictive algorithm applied to energy management in electric motorsport could not be identified in the scientific literature. Compared to a torque-limited approach with the same energy consumption, this methodology provides a $11.4\,\%$ improvement in elapsed time to complete the course of an 150 meters straight line acceleration. A Formula Student representative 150m maneuver was also investigated – the maneuver was extract from a real endurance event at the Michigan – USA competition, which is a course of around 1km where vehicles compete for the minimum elapsed time over 22km (22 laps). The the MPC approach reduced the elapsed time by $13.5\,\%$, as well as responding to energy deficit conditions by reducing torque in an optimal manner.

In the following section, a state-space linearized vehicle model is developed and compared to a transient model using the commercial software AVL VSM™. The MPC is presented in Section 3 alongside design considerations for the number of prediction and control horizons and the definition of setpoints and signal weighting. Results are presented in Section 4 for three scenarios: straight line acceleration, Formula Student maneuver and a reduced prediction and control horizons scenario to verify the algorithm stability. Finally, this work is summarized and concluded in Section 5.

2 Vehicle Model

Designed to reduce computational cost of the controller, the plant model is based on a simplified vehicle longitudinal dynamics model using first principles (Equation 1). It express a lumped mass being accelerated by a thrust force – torque from the electric motors, drive ratio and tire rolling radius; which is resisted by aerodynamic drag, rolling resistance and road gradient forces. Equation 2 express the power output based on the battery pack voltage v and torque of the motors.

$$\frac{dV}{dt} = \frac{1}{em}\left[\frac{\eta_t T_r i_\theta}{r} - \left\{\frac{1}{2}\rho C_d A V^2 + \left(mgcos(\alpha) - \frac{1}{2}\rho C_L A V^2\right) f + mgsin(\alpha)\right\}\right] \quad (1)$$

$$\frac{dE}{dt} = v\left(-T_r(k_t)\right) \quad (2)$$

Where, V is the vehicle velocity, E is the energy available on the battery pack and T_r is the total motor torque – from the four in-hub motors. The remaining parameters are described in Table 1 alongside the values used to populate the model.

Table 1: Vehicle parameters

Parameter	Value	Units	Description
η_t	0.9	-	Transmission efficiency
e	1.4	-	Rotational mass factor
m	300	$[kg]$	Vehicle mass
ρ	1.225	$[kg/m^3]$	Air density
A	2.2	$[m^2]$	Vehicle frontal area
g	9.81	$[m/s^2]$	Gravitational acceleration
r	0.203	$[m]$	Tire rolling radius
k_t	0.26	$[A/N.m]$	Torque constant
i_θ	15.55	-	Final drive ratio
C_L		-	Coefficient of lift
C_D	0.40	-	Coefficient of drag
α	0	$[deg]$	Road inclination angle
f	0.015	-	Rolling resistance coefficient

To verify the fidelity of the simplified model in comparison to the real dynamics of the system, a transient model was created using AVL VSM™ software. The energy consumption was then compared to the simplified model in which a Root Mean Squared Error (RMSE) of $0.008\,kWh$ was found (Figure 2). The comparison was based on a Formula Student representative maneuver over a series of medium speed corners with continuous positive torque request – no braking or coasting. This data comes from a Formula Student endurance event in Michigan, USA.

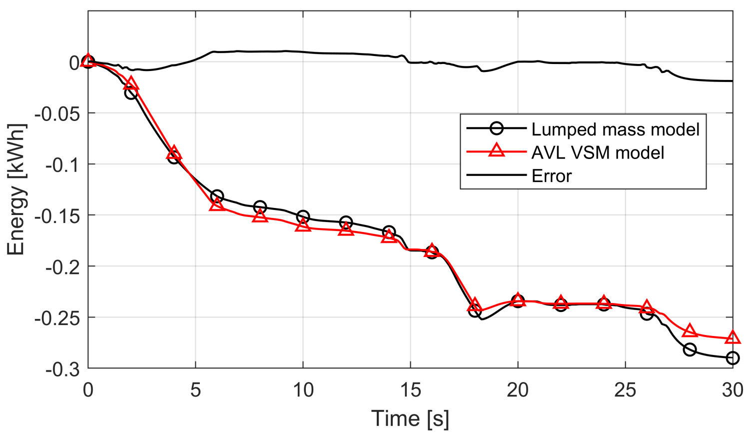

Figure 2: Energy consumption comparison between a transient model developed using AVL VSM software and a simplified lumped mass model.

2.1 Identified Linear Time-Invariant (IDLTI) Model

MPC based control methodologies generally rely on linear time-invariant (LTI) state-space representations of systems [13]. However, real world dynamical systems are very rarely purely linear [10] and in fact, the lumped mass vehicle model used in this paper is time invariant and nonlinear: as drag forces square with velocity, the outputs remain independent of absolute time. The representation of a LTI assumes the state-space format of Equations 3 and 4.

$$x(k+1) = \boldsymbol{A}x(k) + \boldsymbol{B}u(k) \tag{3}$$

$$y(k) = \boldsymbol{C}x(k) + \boldsymbol{D}u(k) \tag{4}$$

In the first equation, x is the state vector of the system, $\boldsymbol{A}$ is the state matrix, $\boldsymbol{B}$ is the input matrix and u is the input vector of the plant. In the second equation, the output vector y is a function of the output matrix $\boldsymbol{C}$, the feed-through matrix $\boldsymbol{D}$, and the state and input vectors. The $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$ and $\boldsymbol{D}$ matrices are function of the plant dynamics and were determined through the process of system identification.

With the help of the MathWorks System Identification Toolbox™, a third and a fourth-order linearized plant were developed. The process consists of evaluating the system response to a given set of input and identifying the set of parameters of the state-space model that better approximate the system response compared to the nonlinear plant. Hence, these models are called Identified Linear Time-Invariant (IDLTI) models [14]. In this work, three step inputs of magnitudes similar to those expected in real-world conditions were used for the system identification. Figure 3 compares both models with the nonlinear plant.

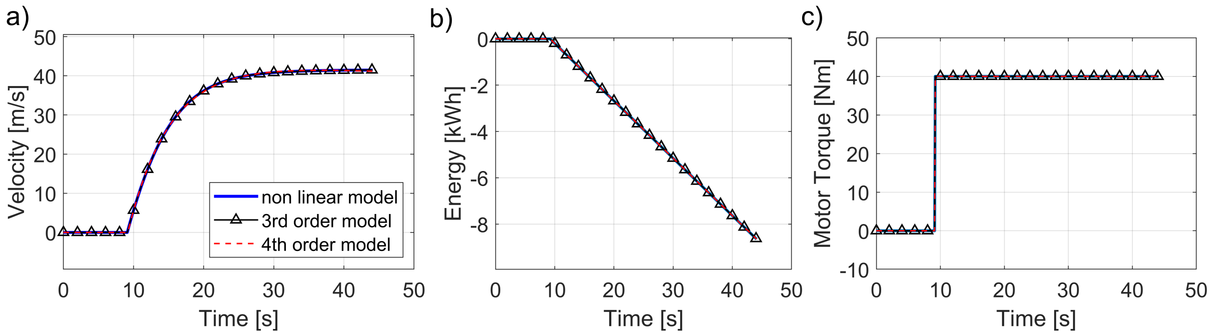

Figure 3: Comparison between a 3rd and 4th order linearized models with the nonlinear plant. NRMSE for all three inputs were greater than 98% for the 3rd order model and greater 99% for the forth order model.

Because the second derivative of Equation 1 is relatively small, this vehicle model is considered only slightly nonlinear [10]. Even then, the linearization process introduces inaccuracies to the model and thus, a comparison between the linearized model against the nonlinear plant needs to be evaluated. This was accomplished by a Normalized Root Mean Squared Error (NRMSE) cost function (Equation 5):

$$fit = 100\left(1 - \frac{\|y - \hat{y}\|}{\|y - mean(y)\|}\right) \tag{5}$$

Where y is the validation data output, $\hat{y}$ is the output of the model, and fit is a percentage. Values of 100% indicate a perfect fit, whereas a fit of at least greater than 75% is recommended for stable control [7]. The 3rd order model achieved a NRMSE greater than 98% whilst the 4th order linearized model had 99% fit with the corresponding results of the non linear plant. Because of the negligible difference of the model fit, the 3rd order model was selected to reduce computational costs.

3 MPC Design

The MPC controller design is illustrated in Figure 4. The set point block diagram gives the controller the hard limits of the motor torque, u_{max} and the target velocity set point, r. The MPC then optimizes the energy consumption based on the feedback of the battery energy capacity, torque output and actual velocity compared to the target velocity. The implementation of the controlled was realized using the MathWorks MPC Designer Toolbox™.

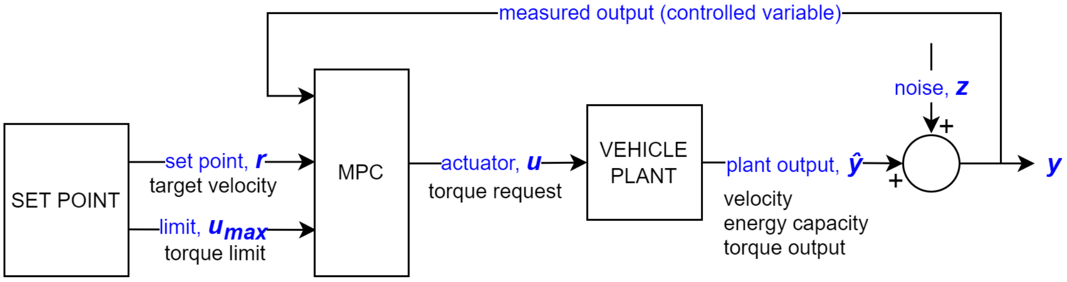

Figure 4: MPC block diagram. Energy consumption is optimized based on the feedback of the battery energy capacity, torque output and vehicle velocity.

3.1 Horizons

The MPC controller calculates an optimal solution to a quadratic control problem. The minimal cost is computed over a receding prediction horizon using a defined plant model at a repeating time step. Two horizons are therefore implemented – visualized previously in Figure 1: *prediction* and *control* horizons. In general, longer horizons increase controller stability but come at higher computational cost because of the increased number of optimized solutions it computes. If control horizons can be lowered without affecting the controller's response, this is typically taken as indication that the prediction and control horizon times are adequate [6]. In this work, a *prediction horizon* of ten time steps and a *control horizon* of two time steps were used.

The relation between the cost function and the number of horizons used by the MPC is dictated by how close the evaluation of the cost function is to the design constraints – in this work: energy consumption and target velocity. A longer prediction horizon is needed to find the optimal solution if the evaluation of the cost function is close to those constraints. On the other hand, if the cost function is not reaching those limits, it is possible to reduce the number of the prediction and control horizons without major effects on the optimal solution, which is advantageous to reduce computational costs. In order to investigate instabilities caused by reducing the predictive and control horizons, these values were reduced to five and one respectively, further discussed in the results section.

3.2 Set Point and Control Signal Weighting

The MPC uses five setpoints: torque request $[N.m]$; rate of change of the torque request $[N.m/\Delta t]$; maneuver target velocity $[km/h]$; target energy consumption $[kW]$;and finally the torque output $[N.m]$ to meet the energy consumption and target velocity constrains. Each set point is weighted differently in the MPC cost function. Multiple weight schemes were tested for this implementation using the MPC Designer Toolbox™. In general, the driver's torque request carried the less weight than the velocity term. This is partially because the units used in the torque request are generally larger than those seen in the velocities. The controller attempts to meet the torque request until the cost weighed against that of the other set points grows too high. The energy term carried greater weight because the deviation from the set point occurs over smaller absolute numbers when compared to torque and velocity. This term reduced torque limits as energy use exceeded targets. Similarly, as speeds rises over the predetermined set point, the cost increases, and the controller reduces torque. The set point weighting parameters used in this work are described in Table 2.

Table 2: Set point weighting parameters

Variable	Weight[1]
Manipulated Variable – Torque request $[N.m]$	0
Manipulated Variable rate $[N.m/\Delta t]$	0.1
Target velocity $[km/h]$	10
Energy consumption $[kW]$	13
Torque output $[N.m]$	3

[1]dimensionless

4 Results and Discussion

The purpose of the MPC controller presented in this paper is to optimize torque delivery during acceleration events. Two scenarios were investigated: 1. straight line wide-open throttle acceleration event over 150 meters, and 2. cornering maneuver of continuous positive torque request over 150 meters – no braking or coasting.

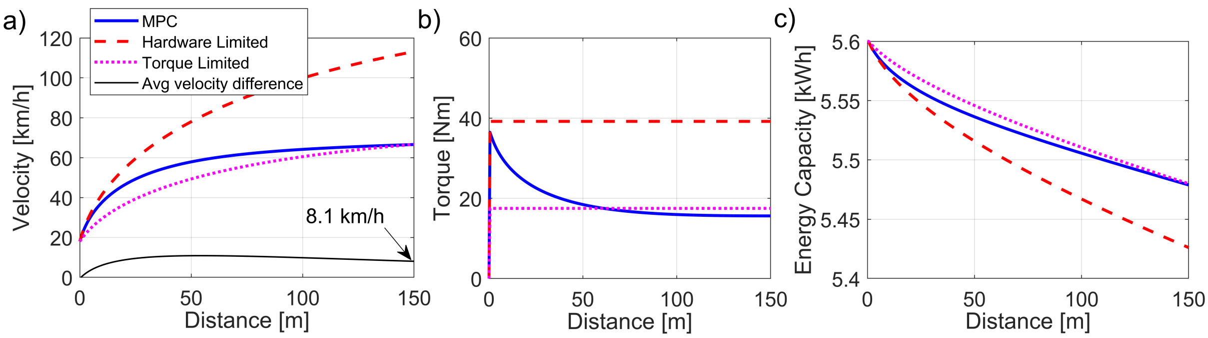

Figure 5: Constant torque request scenario. a) speed profile, b) torque request c) energy remaining in the battery

In the first scenario, the vehicle accelerates from a low speed corner onto a straight line, the length of the straight represents the longest continuous straight a Formula Student car is likely to encounter in an endurance event. In Figure 5, the MPC controller receives an arbitrated constant torque request of $39.8\,Nm$, which is the maximum torque output of the motors (hardware limit). The system response is then compared between a torque-limited approach with an equivalent energy usage.

As expected, the MPC torque setpoint decreased with distance to optimize the energy consumption as shown in Figure 5 b). This is explained because the most efficient use of the energy in an acceleration maneuver occurs at low speeds, as thrust force is resisted by the least amount of drag force, which increases with the square of the velocity ($\frac{1}{2}\rho C_d A V^2$ part of Equation 1) . Compared to the torque-limited approach, an average velocity improvement of $8.1\,km/h$ was observed using the MPC strategy, resulting in a time improvement of 1.31 seconds ($11.4\,\%$) for the same energy consumed.

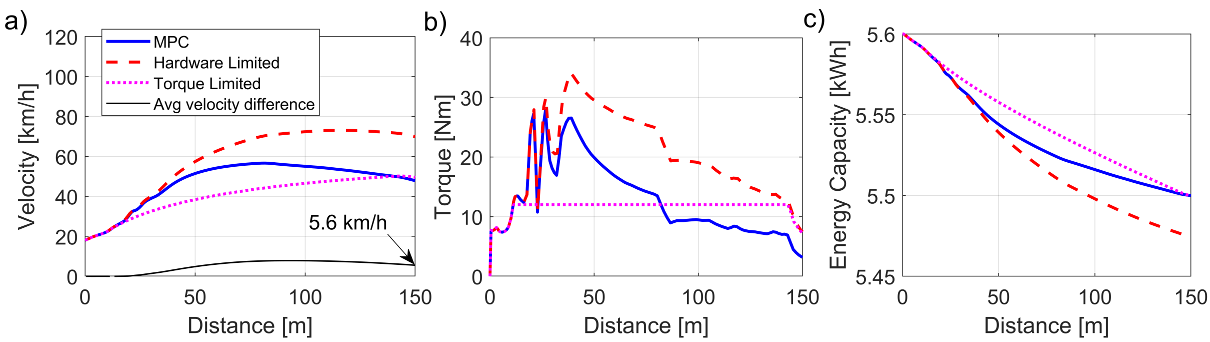

Figure 6: Formula Student maneuver torque request scenario. a) speed profile, b) torque request c) energy remaining in the battery

In the second scenario, the maneuver input is the same torque request used in the vehicle model verification (Figure 2). Figure 6 compares the MPC approach with the torque-limited approach. The MPC results in an average velocity of $42.5\,km/h$, $5.6\,km/h$ faster than the equivalent energy used by the torque-limited system. This yields a 1.99 seconds, or 13.5% lap time improvement over the course of 150 meters.

In order to investigate instabilities caused by reducing the predictive horizon was reduced to five and the control horizon reduced to a single time step. The results are shown in Figure 7, while the controller was less able restrict energy use, no instability or boundary violations in the controller were observed, indicating that with this controller and plant, reducing control and prediction horizons are a viable method of lowering computational cost if shorter time steps are desired.

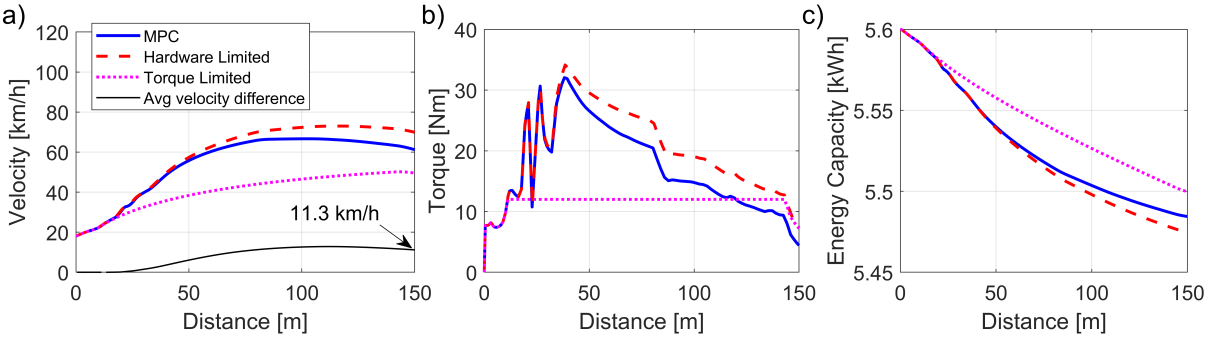

Figure 7: Reduced horizons scenario. a) speed profile, b) torque request c) energy remaining in the battery

5 Conclusions

The MPC torque limiter meets the preliminary requirements of the system – it provides significant improvements in laptime and energy use, while balancing the immediate torque demand and associated energy use against the overall state of energy in the system. Computational cost is well within the limits of the rapid prototyping controller. By using the hard constraints allowed by MPC methods, torque safety is maintained. Therefore, Model Predictive Control methods show significant promise for energy-limited race cars, especially Formula Student-class electric vehicles – which are required by the competition rules to have amateur drivers, this controller is especially helpful to such drivers, since it decreases cockpit workload during acceleration events, eliminating the need for drivers to pick lift and coast points in order to save energy.

Many control architectures are possible to solve a given set of control requirements. While this controller housed the energy management function of the controller within the torque controller, future explorations of this method should consider a low frequency energy management hypervisor to evaluate the energy state of the system and estimated energy use per lap using a plant model generated from a laptime simulation as suggested by [2]. This offloads the computational cost of energy management to a very low frequency controller, freeing up computational headroom to increase the frequency of the torque converter. The actuators of the hypervisor can include the weighting, setpoints, and limitations of the torque controller. The torque controller can then be expanded to include thermal winding costs of the torque request as well as accumulator thermal costs.

This paper outlines the process in which an MPC based control algorithm can be implemented in electric motorsport powertrain optimization, and provides a starting point for additional control algorithm development in Formula Student.

Acknowledgments

The authors wish to thank Oxford Brookes Racing team members, alumni, faculty, and sponsors for their support and guidance, particularly Gordana Collier. The authors would also like to thank Williams Advanced Engineering for their enthusiastic support and feedback.

References

[1] D. J. Limebeer and A. V. Rao, “Faster, higher, and greener: Vehicular optimal control,” *IEEE Control Systems Magazine*, 2015.

[2] M. Salazar, P. Elbert, S. Ebbesen, C. Bussi, and C. H. Onder, “Time-optimal control policy for a hybrid electric race car,” *IEEE Transactions on Control Systems Technology*, 2017.

[3] M. Salazar, P. Duhr, C. Balerna, L. Arzilli, and C. H. Onder, “Minimum lap time control of hybrid electric race cars in qualifying scenarios,” *IEEE Transactions on Vehicular Technology*, 2019.

[4] S. Ebbesen, M. Salazar, P. Elbert, C. Bussi, and C. H. Onder, “Time-optimal control strategies for a hybrid electric race car,” *IEEE Transactions on control systems technology*, 2017.

[5] T. D. Gillespie, *Fundamentals of vehicle dynamics*, 2021.

[6] A. Bemporad, M. Morari, and N. L. Ricker, "Model predictive control toolbox user's guide," *The mathworks*, 2010.

[7] N. Khaled and B. Pattel, *Practical design and application of model predictive control: MPC for MATLAB® and Simulink® Users*. Butterworth-Heinemann, 2018.

[8] A. Liniger, A. Domahidi, and M. Morari, "Optimization-based autonomous racing of 1: 43 scale rc cars," *Optimal Control Applications and Methods*, 2015.

[9] J. L. Vázquez, M. Brühlmeier, A. Liniger, A. Rupenyan, and J. Lygeros, "Optimization-based hierarchical motion planning for autonomous racing," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 2397–2403.

[10] S. L. Brunton and J. N. Kutz, *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press, 2022.

[11] M. Salazar, C. Balerna, P. Elbert, F. P. Grando, and C. H. Onder, "Real-time control algorithms for a hybrid electric race car using a two-level model predictive control scheme," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 12, pp. 10 911–10 922, 2017.

[12] D. J. Limebeer and M. Massaro, *Dynamics and optimal control of road vehicles*. Oxford University Press, 2018.

[13] J. Lee and H.-J. Chang, "Explicit model predictive control for linear time-variant systems with application to double-lane-change maneuver," *Plos one*, vol. 13, no. 12, p. e0208071, 2018.

[14] L. Ljung, "System identification toolbox: User's guide," 2021.

Authors Biography

Tom Maull received his BSc in Mechanical Engineering at Lehigh University and his MSc at Oxford Brookes University. He served as the Vehicle Controls Lead for Oxford Brookes Racing in the 2021 season, and currently works in the Advanced Battery Concepts team at Williams Advanced Engineering

Adriano Schommer received his Mechanical Engineering BSc at Federal University of Santa Maria, MSc at Federal University of Santa Catarina and is currently a PhD candidate at Oxford Brookes University. He is the Oxford Brookes Project Lead and his main area of research is high power Lithium-Ion batteries state estimators for motorsport applications.