On the Co-Design of Components and Racing Strategies in Formula 1

Conference Paper

Author(s):
Neumann, Marc-Philippe; Zardini, Gioele; Cerofolini, Alberto; Onder, Christopher H.

Publication date:
2024

Permanent link:
https://doi.org/10.3929/ethz-b-000655029

On the Co-Design of Components and Racing Strategies in Formula 1

Marc-Philippe Neumann[1], Gioele Zardini[2], Alberto Cerofolini[3], Christopher H. Onder[1]

***Abstract*— We present a study focusing on the joint optimization of the sizing of hardware components as well as strategic decisions for a race car in a Formula 1 setting. Our research leverages a monotone theory of co-design, which allows for hardware and software considerations to achieve optimal, synergistic performance improvements. We aim to identify the Pareto optimal curves that illustrate the optimal balance between conflicting objectives, such as lap time, energy allocation, and component choice, within the tight constraints imposed by the regulations. The results of the study demonstrate the versatility of our framework by showing optimal component sizing on two structurally different track layouts on a single lap. Moreover, by increasing the amount of laps under consideration, we show the ability of our tool to consider strategic energy allocation decisions.**

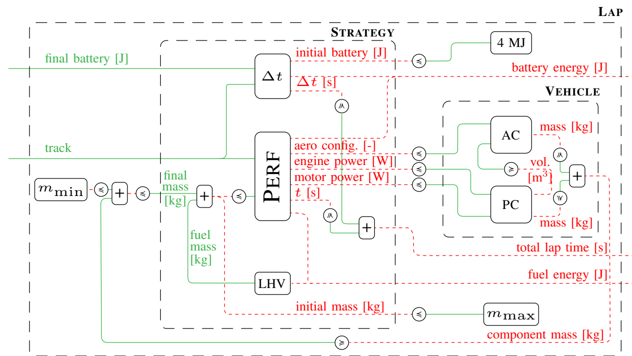

Fig. 1: Single lap co-design diagram for a F1 race car on a given track, reaching a final battery target. The component design is indicated by *vehicle*, and the energy allocations represent the racing *strategy*. The resources to minimize are allocated battery and fuel energies, and the lap time.

I. Introduction

In the past decades, Formula 1 (F1) has been the most prestigious racing category in Europe, gaining more and more importance all over the world. Adding to the euphoria of the regular spectators, the increased design complexity of the high-performance hybrid electric race cars in this competition makes this sport extremely interesting also from an engineering point of view. In particular, stringent constraints on weight, energy availability and deployment, as well as strategically crucial aspects such as the mandatory use of at least two different tire compounds during one race, amplify the existing physical limitations [1], [2]. While the primary focus for the drivers is achieving maximum speed at every instant of the race, the strategists involved in the competition require a more macroscopic view. Clearly, some limitations and constraints are of concurrent nature and a lap time optimal solution might result suboptimal to the problem of minimizing the time for the entire race. For instance, depleting the battery during one lap affects the subsequent lap, leading to a potentially overall worse performance. Purely strategic assessments during a race have already been studied in literature [3]. However, the influence of physical alterations has not been shown. Moreover, although the races during a season appear to be independent of each other, there are couplings which are often neglected. Considerations such as component wear, which during one single race can be neglected, gain crucial importance on a seasonal level. In particular, each team has a maximum amount of units per season (e.g., three internal combustion engines (ICEs)) they can equip their cars with. Failure to comply with such regulations results in grid penalties for subsequent races, compromising the championship position.

This study leverages a recently developed monotone theory of co-design to offer a tool able to assess the influence of component sizing on the lap time optimal strategy. Such a co-design diagram is shown in Figure 1. Additionally, we increase the number of laps to broaden the system boundaries and show the versatility of the framework to jointly solve strategic and sizing optimization problems.

A. Related work

When examining the design of modern F1 cars, the strategic methodologies that come into play align with those explored in the field of hybrid electric vehicles. However, our co-design purposes require a closer analysis of the literature focusing on the joint optimization of hardware and strategic choices. In particular, we focus on two research streams.

The first focuses on the lap time optimal energy management of hybrid electric race cars. In this context, key contributions include high-level lap time energy management via convex optimization [4], and the extended efforts to an entire race [3]. The inclusion of a performance envelope [5], relating aerodynamic specifications to cornering velocity, is particularly important in our investigations. In [6] the authors investigated nonlinear low-level optimization techniques, incorporating a detailed ICE model. The influence of the vehicle mass on lap time has been investigated, providing valuable insights [7]–[9]. Finally, [10] explores the transition from optimization to control, offering practical perspectives on implementing control strategies in future applications.

[1]Institute for Dynamic Systems and Control, ETH Zürich, Zürich, Switzerland, {mneumann,onder}@ethz.ch
[2]Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, USA, gzardini@mit.edu
[3]Power Unit Performance and Control Group, Ferrari S.p.A., 41053 Maranello, Italy, Alberto.Cerofolini@ferrari.com

In a parallel stream, we explore the benefits of joint optimization. An intriguing application for lap time and optimal sizing, specifically for the battery, has been addressed by Riva [11] and Radrizzani [12], presenting a comprehensive perspective on energy management in hybrid electric race cars. The monotone theory of co-design represents a paradigm shift in the optimization of hardware and software design choices, where both components call for joint optimization. This approach is particularly evident in the context of entire Autonomous Mobility on Demand (AMoD) systems, as highlighted in [13], [14]. Notably, the theory extends its application to autonomous systems. In [15]–[17], the authors delved into the co-design intricacies necessary for optimizing the longitudinal and lateral control performance of an autonomous robot. Lastly, in [18] a convex approach was investigated to optimize concurrent quantities.

Overall, the studied literature has three fundamental fallacies. First, it does not consider the effect of physical changes of the vehicle on lap and race performance metrics. Second, research in this area typically resorts to complex problem formulations and solution algorithms, which do not scale well both computationally, and intellectually. Finally, the multi-objective nature of the problem is often neglected.

B. Statement of Contribution

This work is motivated by a gap in the ability to jointly optimize hardware elements as well as strategic choices in the context of circuit racing in F1. We present a framework which leverages a monotone theory of co-design to formulate and solve optimization problems in the world of hybrid electric F1 race car design. First, we show how to model race-critic aspects of the vehicle from a co-design point of view. For instance, the aerodynamic implementation is described by drag and downforce coefficients [19], while we exploit the g-g-diagram to link them to track performance [20]. Second, we showcase the strategic component of the architecture. Starting from a nonlinear optimal control problem (OCP), we are able to link energy allocations to achieved lap times, which represent the main performance metric in this sport. Finally, to showcase the properties of the framework, we present two case studies. In the first, we analyze the choice of aerodynamic setups for two different racetracks. The second case study shows the ability of our framework to include strategic aspects as soon as multiple laps are considered.

C. Paper Structure

In Section II we introduce a monotone theory of co-design, which we applied to the case of the race car. In Section III, we introduce each design problem and its modeling assumptions. Section IV showcases the performance of our framework by means of two case studies where we first analyze the varying optimal sizing of the aerodynamic configuration depending on track choice and further introduce multiple consecutive laps to highlight the strategy optimization of our framework. Finally, in Section V we draw the conclusions, comment on the relevant insights gained by means of our framework, and give an outlook on future research.

II. Monotone Co-Design Theory

We present the main concepts related to a monotone theory of co-design, presented in [21], and more extensively in [22], [23].

We assume the reader to be familiar with basic concepts from order theory. A possible reference is [24]. This design theory is based on the atomic notion of monotone design problem with implementations (MDPI).

Definition 1: (MDPI). Given the partially ordered sets (posets) $\mathcal{F}$, $\mathcal{R}$, representing functionalities and resources, respectively, we define a MDPI d as a tuple $\langle \mathcal{I}_d, \mathrm{prov}, \mathrm{req}\rangle$, where $\mathcal{I}_d$ is the set of implementations, and prov, req are functions mapping $\mathcal{I}_d$ to $\mathcal{F}$ and $\mathcal{R}$, respectively:

$$\mathcal{F} \xleftarrow{\mathrm{prov}} \mathcal{I}_d \xrightarrow{\mathrm{req}} \mathcal{R}$$

To each MDPI we associate a monotone map $\bar{d}$ given by

$$\begin{aligned}\bar{d} &: \mathcal{F}^{\mathrm{op}} \times \mathcal{R} \to \mathcal{P}(\mathcal{I}_d) \\ &\langle f^*, r\rangle \mapsto \{i \in \mathcal{I}_d : (\mathrm{prov}(i) \succeq_{\mathcal{F}} f) \wedge (\mathrm{reqs}(i) \preceq_{\mathcal{R}} r)\},\end{aligned}$$

where $(\cdot)^{\mathrm{op}}$ reverses the order of a poset. A MDPI is represented in diagrammatic form as in Figure 2. The expression $\bar{d}(f^*, r)$ returns the set of implementations $S \subseteq \mathcal{I}_d$ for which functionalities f are feasible with resources r.

Remark 1 (Monotonicity): Consider a MDPI for which we know $\bar{d}(f^*, r) = S$. If we have $f' \preceq_{\mathcal{F}} f$, then $\bar{d}(f'^*, r) = S' \supseteq S$ (i.e., lowering the desired functionalities will not increase the required resources). Conversely, if we have $r' \succeq_{\mathcal{R}} r$, then $\bar{d}(f^*, r') = S'' \supseteq S$ (i.e., increasing the available resources cannot decrease the provided functionalities).

Remark 2 (Populating the MDPIs): In practical cases, one can populate the feasibility relations of MDPIs with analytic relations, numerical analysis of closed-form relations, and simulations.

Definition 2: Given a MDPI d, we define monotone maps $h_d : \mathcal{F} \to \mathcal{AR}$, mapping a functionality to the minimum antichain of resources providing it, and $h'_d : \mathcal{R} \to \mathcal{AF}$, mapping a resource to the maximum antichain of functionalities provided by it.

Solving MDPI requires finding such maps. For this, one can rely on Kleene's fixed point theorem, as detailed in [21].

Individual MDPIs can be composed in many ways to form a co-design problem (i.e., a multigraph of co-design problems), allowing one to decompose large problems into smaller subproblems, and to interconnect them. Series composition describes the case in which the functionality of a MDPI becomes the resource of another MDPI. For instance, the maximum power provided by an electric motor is required by the lap time computation. The relation "$\preceq$" appearing in Figure 3a represents a co-design constraint: The resource one component requires has to be at most the one provided by another component. Parallel composition corresponds to processes happening together. Finally, loop composition describes feedback. Composition operations preserve monotonicity and thus all related algorithmic properties [21]. Populating these models gives rise to a family of multi-objective optimization problems, which is computationally efficient (complexity is only linear in the number of options

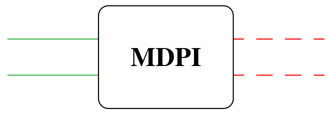

Fig. 2: A MDPI is a monotone relation between partially ordered sets of functionalities (in green) and resources (in red).

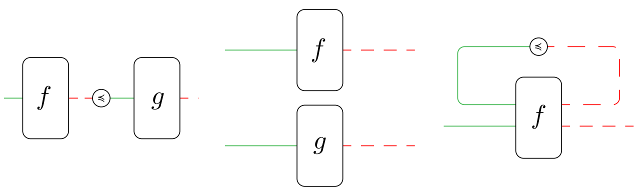

(a) Series. (b) Parallel. (c) Loop.

Fig. 3: MDPIs can be composed in various ways. Composition operations include series, parallel, and loop. Notably, the composition of MDPIs results in a MDPI (closure).

for each design problem), intellectually tractable, and easily manipulable [23].

III. Design Problems

We present co-design models for various components of the system. For each MDPI, we describe functionalities, resources, and the relationship between them. We will then show how to interconnect the components into the complex system under investigation. Given the interconnected relations, in Section IV we will show how to practically populate them, and how to solve the arising optimization problems.

A. Strategy

The strategy MDPI entails a performance component, where strategic decisions are linked to a lap time, and a battery sensitivity component, taking care of suboptimal energy deployment strategies due to physical constraints.

1) Performance Metric: The MDPI describing the performance of the system lies at the core of the strategy, as it relates the assigned energy budgets to a nominal lap time. In particular, its functionalities are the total mass m, and the track choice, which directly defines the curvature profile γ related to the investigated lap. This quantity is an indicator of the deflection of the track. The higher it is, the tighter are the curves, and the more advantageous it is to have high cornering velocities. The resource space consists of assigned battery and fuel budgets $\mathrm{d}E_\mathrm{b}$ and $\mathrm{d}E_\mathrm{f}$, respectively, the aerodynamic configuration, the maximum engine power $P_{\mathrm{e,max}}$, the maximum motor generator unit - kinetic (MGU-K) power $P_{\mathrm{k,max}}$, and the nominal lap time t_{nom}. Note that the aerodynamic configuration consists of a set of incomparable parameters $\{c_0, c_1, c_2, c_3, d_{1,\mathrm{top}}, d_{2,\mathrm{top}}, d_{3,\mathrm{top}}, d_{1,\mathrm{bot}}, d_{2,\mathrm{bot}}, d_{3,\mathrm{bot}}\}$. The first parameter is related to the aerodynamic longitudinal drag, while the others define the maximal acceleration forces in longitudinal and lateral directions due to the tire-road contact [3]. To relate functionalities and resources for this MDPI, we propose the OCP defined in Problem 1.

Problem 1: We define the state vector

$$\vec{x} = \begin{bmatrix} v \\ E_\mathrm{b} \\ E_\mathrm{f} \end{bmatrix}, \tag{1}$$

where v is the vehicle speed, E_b is the battery energy, and E_f is the fuel energy, and the input vector

$$\vec{u} = \begin{bmatrix} P_\mathrm{e} \\ P_\mathrm{k} \\ P_\mathrm{brk} \end{bmatrix}, \tag{2}$$

where P_e is the engine power, P_k is the MGU-K power, and P_brk is the brake power. The nominal lap time is the lowest cost to the following OCP:

$$\min_{\vec{u}} \; t_{\mathrm{nom}} = \int_0^S \frac{\mathrm{d}s}{v(s)}, \tag{3}$$

subject to the dynamic state equations:

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}s} v &= \frac{1}{m} \cdot \frac{P_\mathrm{e} + P_\mathrm{k} - P_\mathrm{brk}}{v^2} - c_0 \cdot v, \\ \frac{\mathrm{d}}{\mathrm{d}s} E_\mathrm{b} &= -\frac{P_\mathrm{k}}{v}, \\ \frac{\mathrm{d}}{\mathrm{d}s} E_\mathrm{f} &= \frac{P_\mathrm{e}}{v \cdot \eta_\mathrm{e}}, \end{aligned} \tag{4}$$

where η_e is a constant efficiency relating engine to fuel power. We consider the following input constraints:

$$\begin{aligned} 0 \leq P_\mathrm{e} &\leq P_{\mathrm{e,max}}, \\ -P_{\mathrm{k,max}} \leq P_\mathrm{k} &\leq P_{\mathrm{k,max}}, \\ 0 \leq P_\mathrm{brk} &\leq \infty, \end{aligned} \tag{5}$$

and the energy related terminal contraints:

$$\begin{aligned} E_\mathrm{b}(S) &= E_\mathrm{b}(0) + \mathrm{d}E_\mathrm{b}, \\ E_\mathrm{f}(S) &= \mathrm{d}E_\mathrm{f}. \end{aligned} \tag{6}$$

Finally, we include also the maximum acceleration constraints in longitudinal and lateral directions given by the track-dependent characteristics:

$$\begin{aligned} \left(\frac{\frac{P_\mathrm{e} + P_\mathrm{k} - P_\mathrm{brk}}{v}}{F_{\mathrm{long,max,bot}}} \right)^2 + \left(\frac{m \cdot v^2 \cdot \gamma}{F_{\mathrm{lat,max}}} \right)^2 &\leq 1, \\ \left(\frac{\frac{P_\mathrm{e} + P_\mathrm{k} - P_\mathrm{brk}}{v}}{F_{\mathrm{long,max,top}}} \right)^2 + \left(\frac{m \cdot v^2 \cdot \gamma}{F_{\mathrm{lat,max}}} \right)^2 &\leq 1, \end{aligned} \tag{7}$$

where

$$\begin{aligned} F_{\mathrm{lat,max}} &= (c_1 \cdot v^2 + c_2 \cdot v + c_3) \cdot m, \\ F_{\mathrm{long,max,top}} &= (d_{1,\mathrm{top}} \cdot v^2 + d_{2,\mathrm{top}} \cdot v + d_{3,\mathrm{top}}) \cdot m, \\ F_{\mathrm{long,max,bot}} &= (d_{1,\mathrm{bot}} \cdot v^2 + d_{2,\mathrm{bot}} \cdot v + d_{3,\mathrm{bot}}) \cdot m, \end{aligned} \tag{8}$$

are the maximum lateral and longitudinal forces allowed on a track.

The monotonicity of the relation given by the OCP can be checked analytically and empirically.

2) Battery Sensitivity: During an optimal lap, the state of charge (SoC) fluctuates between charging and discharging intervals. Therefore, the lap time does not only depend on the allocated energy budgets, but also on the initial battery SoC [3]. We describe the phenomenon of increased lap time as a MDPI, providing the final battery SoC and the chosen track as functionalities, as the suboptimality differs depending on the length and characteristics of the circuit. In turn, the resources include the initial battery SoC, the allocated battery energy, and the lost lap time. To capture the relationship between these quantities, we propose the following affine relations:

$$\begin{aligned} E_{\mathrm{b,end}} &= E_{\mathrm{b},0} + \mathrm{d}E_{\mathrm{b}}, \\ \Delta t &= \max\left(0, r \cdot \mathrm{d}E_{\mathrm{b}} + s \cdot E_{\mathrm{b},0} + q\right), \end{aligned} \tag{9}$$

where r, s, q are track-dependent parameters.

B. Vehicle

The vehicle MDPI includes all elements that are subject to sizing, i.e., the aerodynamic and the propulsive configurations.

1) Aerodynamic Configuration: The aerodynamic configuration (AC) is a crucial component of a racing car. While the teams invest millions of USD to decrease longitudinal drag, its correlation to downforce, and therefore to increased cornering velocities, results in a concurrent behavior of these two quantities [25], [26]. Moreover, additional aerodynamic components might negatively affect the overall mass of the shell. The complexity is further increased due to the influence of the aerodynamic setup on the propulsive configuration (PC) choice in Section III-B.2. In particular, a more slender configuration, which could appear more advantageous aerodynamically, provides a smaller available volume for the propulsive packaging, potentially compromising available power. As introduced in Section III-A.1, the functionality of this MDPI consists of a catalog containing a set of incomparable parameters that are provided to the performance MDPI. Specifically, such parameters quantify the drag coefficient and translate the downforce generated by the airflow to the parameters needed in the performance envelope model introduced in Equation (7). The only resource of this MDPI is the mass of the setup.

2) Propulsive Configuration: The PC MDPI provides engine and MGU-K powers as functionalities. Generally, a higher maximal power appears to be the trivial choice when designing the race car. However, the available volume provided by the AC needs to be considered, as it might exclude some implementations. Furthermore, the resulting mass needs to be considered, as a heavier car will perform worse in the strategic component of the framework. In this study, we decouple the aerodynamic influence on the propulsive component (i.e., the available volume) and only consider the effect of the mass on lap time.

IV. Results

This section showcases the application of our framework on real-world scenarios in the domain of F1. After an introductory section on the design of experiments, two case studies are analyzed.

Conf. #	**Engine Power** [kW]	**Motor Power** [kW]	**Mass** [kg]
1	540.0	120.0	300.0
2	550.0	120.0	310.0

TABLE I: Catalog populating the PC MDPI.

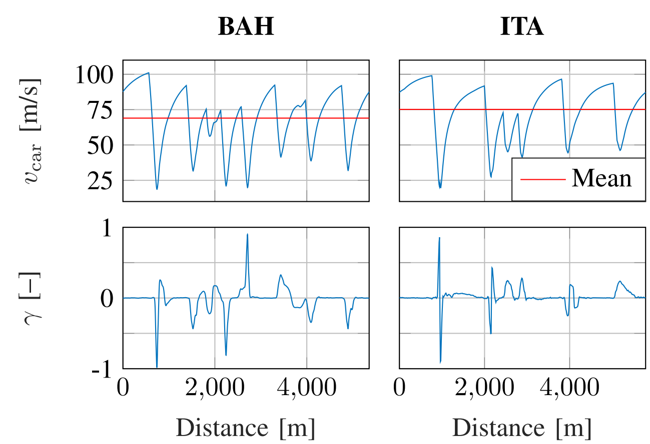

Fig. 4: Velocity and normalized curvature trajectories for BAH and ITA. In both velocity plots we include the mean velocity in red, to show the slightly higher value for the Italian circuit. Moreover, the curvature clearly shows that the circuit of Bahrain offers more cornering situations.

A. Design of Experiments

We populated the AC MDPI with two different setups. First, we consider the *low drag configuration (LDC)*, which should emulate a quite flat rear wing, which is beneficial to longitudinal drag but comes at the expense of downforce, and therefore of cornering velocity. Second, we study the *high drag configuration (HDC)*, which emulates a higher angle of attack at the rear wing, leading to higher longitudinal drag and downforce. To populate the PC MDPI we leveraged a heuristic catalog, where an increase in engine power comes with an increase in mass, as can be seen in Table I. The performance MDPI introduced in Section III-A is populated with optimization results allowing for fuel energy budgets reaching from $94\,\%$ to $102\,\%$ every $2\,\%$, while the battery energy budgets reach from $-0.2\,\mathrm{MJ}$ (discharging) to $0.1\,\mathrm{MJ}$ (charging) with a discretization of $0.1\,\mathrm{MJ}$.

B. Hardware Implementations on Different Race Tracks

In the first case study, we show how the co-design framework answers track-driven optimization queries and adapts to track characteristics. In particular, we optimize for the lap time, energy allocation, and the vehicle's optimal component sizing for one lap, given two different circuits and the same final battery status $4.0\,\mathrm{MJ}$. In Figure 4 we show the main characteristics of the two chosen track layouts: On the left the Bahrain International Circuit (BAH) and on the right the Autodromo Nazionale Monza (ITA). The structural difference between the two is shown by the 8% higher mean velocity in ITA, which suggests that this track is generally faster. On the contrary, BAH exhibits higher fluctuations in the curvature γ, suggesting that high cornering velocities are more beneficial for optimal lap times. Figure 5 shows

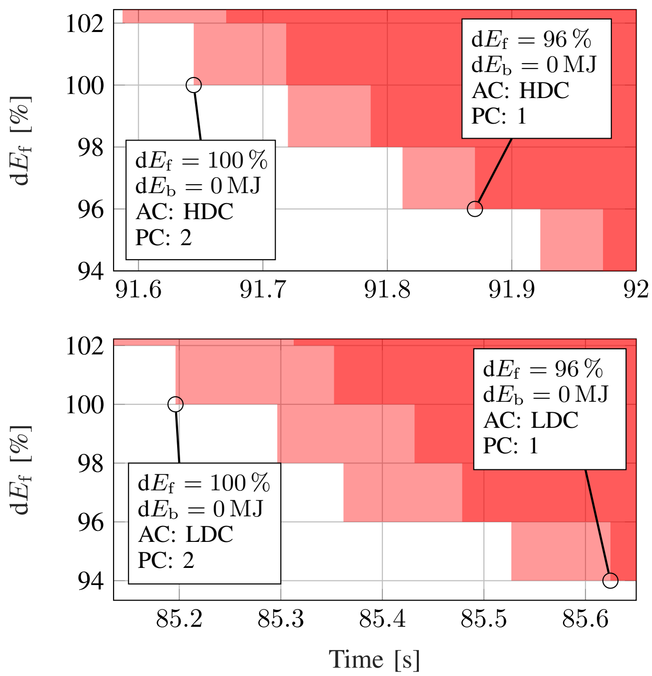

Fig. 5: Fuel energy allocated during one lap versus the lap time for BAH (top plot) and for ITA (bottom plot). The battery energy is equal for all scenarios. The shaded areas represent the feasible solution space, as it is always possible to drive slower for a given energy budget. The different shading is due to two different engine configurations. Implicitly, this figure contains information about the aerodynamic configuration, which for BAH is the HDC and for ITA is the LDC.

the Pareto fronts of optimal solutions in terms of the fuel energy allocated in that lap against the obtained lap time for BAH and ITA. As expected, a higher available fuel energy leads to improved lap times, despite a slightly higher mass to carry around the circuit. Given the requested full final battery, all the points on the front assign a charge-sustaining strategy, meaning that we start with a full battery and assign $\mathrm{d}E_{\mathrm{b}} = 0\,\mathrm{MJ}$. The other feasible solutions featuring a slightly discharged battery at the beginning and assigning a charging strategy clearly result in higher lap times. Furthermore, from each point on the front we can extract related implementation details. Finally, in agreement with the curvature trajectory, on BAH, the optimal AC is the HDC. This differs from ITA, where the LDC is preferred.

C. Energy Allocation vs. Lap Time

In the second case study we increase the number of strategy MDPIs of our framework. Figure 6 shows the general structure for n concatenated laps. In this scenario, we look at the case $n = 3$. In the following, we show the ability of our tool to allocate energy budgets once the number of laps increases. The physical implementation that is included in the vehicle MDPI remains available only once, as the design choices cannot be altered during one race. The functionalities we requested in this case study are a final battery of $4.0\,\mathrm{MJ}$ (meaning that over the three laps, the vehicle needs to race in charge-sustaining mode) on the circuit of BAH. Table II shows the set of feasible battery energy allocation policies that we obtain from the first requested functionality. Note that all policies that include charging the battery prior to at least one discharging lap (i.e., Policy 2) are suboptimal a priori, as they imply an initially not fully charged battery status. In Figure 7 we show the optimal relationship between total fuel allocation versus time to complete three laps. The different shadings arise from different battery energy allocation policies. We see that it is beneficial to first maximally discharge the battery while the vehicle is heaviest, and charge it in the two subsequent laps (Policy 5). Each point on the Pareto front is also characterized by a fuel allocation policy. Apart from the highest and lowest possible fuel allocations (i.e., three laps with $102\,\%$ or $94\,\%$ fuel budget), the policies in between contain various permutations that also affect lap time. For example, the optimal allocation policy for a total fuel consumption of $95.3\,\%$ highlighted on the right, is $\{96\,\%,\ 96\,\%,\ 94\,\%\}$. We explain this trend by looking at the vehicle's mass. With a fuel allocation policy of $\{96\,\%,\ 94\,\%,\ 96\,\%\}$, the mass that needs to be carried around the circuit is decreased slower than in the optimal policy. In the sense of this case study, this physically sensible sensitivity validates our framework.

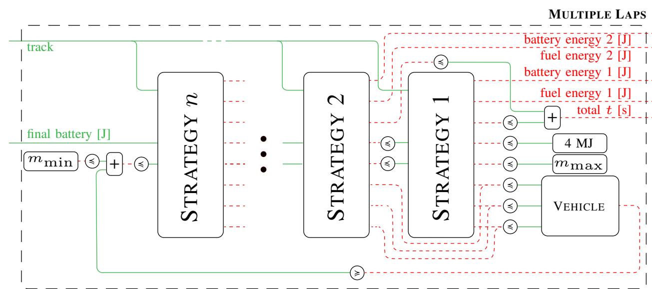

Fig. 6: Co-design diagram for the component design and racing strategy of a racing vehicle on n concatenated laps on a given track while reaching a final battery target. The resources to minimize are allocated battery and fuel energies, and the total time.

Policy #	$\mathrm{d}E_{\mathrm{b},1}$ [MJ]	$\mathrm{d}E_{\mathrm{b},2}$[MJ]	$\mathrm{d}E_{\mathrm{b},3}$[MJ]
Policy 1	-0.1	0.1	0
Policy 2	0.1	-0.2	0.1
Policy 3	0	-0.1	0.1
Policy 4	-0.1	0	0.1
Policy 5	-0.2	0.1	0.1

TABLE II: Set of feasible battery allocation policies.

V. Conclusions

Our work has marked a significant step forward in the exploration of joint optimization of component sizing and strategic decisions within the context of energy allocation in the domain of F1. Through our preliminary study, we have successfully applied a monotone theory of co-design to competitive racing and were able to replicate well-known energy sensitivities, validating our approach. Furthermore,

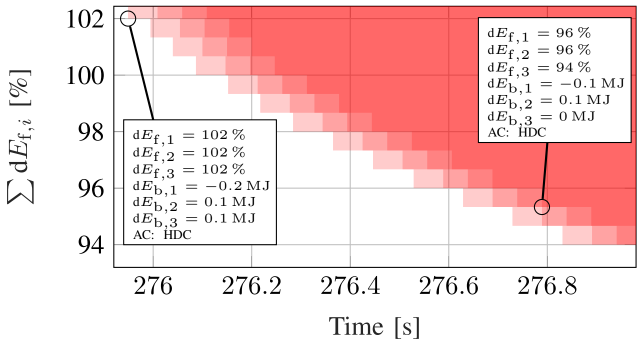

Fig. 7: Total fuel energy allocated during three laps versus the lap time for BAH. The fuel energy allocation policy of each edge on the Pareto front is not necessarily equal. The shaded areas represent the feasible solution space, as there always exists a suboptimal fuel energy policy or implementation that results in a higher time to complete the three laps. The different shading is due to different battery energy allocation policies according to Table II. Implicitly, this figure contains information about the aerodynamic configuration, which for BAH is the HDC.

we demonstrated the framework's versatility in adapting to its environment and making optimal components' design choices considering various track layouts. In particular, we considered the two real-world case studies of the Bahrain International Circuit and the Autodromo Nazionale Monza. Finally, by augmenting the design analysis with additional laps, we were able to compare different energy allocation policies, laying the foundations for future analyses on the interplay between strategy decisions and component sizing in a race. The presented framework poses computational advantages with respect to state of the art methodologies, and is more intellectually tractable and manipulable.

From a sizing perspective, the next phase of our work will focus on a more detailed interconnection between the physical components, and the inclusion of additional components which need to be designed. As for the strategic aspect of our work, expanding our analysis to encompass more laps, potentially extending to an entire race duration, will provide a more comprehensive understanding of the long-term implications of our proposed framework. Finally, the incorporation of other resources (i.e., optimization objectives), such as wear considerations or monetary aspects, will add depth to our model, offering a broader view that could become relevant when optimizing an entire season.

Acknowledgments

We would like to thank Ferrari S.p.A. for the support.

References

[1] FIA, "2024 Formula one sporting regulations," Geneva, Switzerland, Tech. Rep., 2023.

[2] ——, "2024 Formula one technical regulations," Geneva, Switzerland, Tech. Rep., 2023.

[3] P. Duhr, D. Buccheri, C. Balerna, A. Cerofolini, and C. H. Onder, "Minimum-race-time energy allocation strategies for the hybrid-electric formula 1 power unit," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 6, pp. 7035–7050, 2023.

[4] S. Ebbesen, M. Salazar, P. Elbert, C. Bussi, and C. H. Onder, "Time-optimal control strategies for a hybrid electric race car," *IEEE Trans. Control Syst. Technol.*, vol. 26, no. 1, pp. 233–247, 2018.

[5] P. Duhr, A. Sandeep, A. Cerofolini, and C. H. Onder, "Convex performance envelope for minimum lap time energy management of race cars," *IEEE Trans. Veh. Technol.*, vol. 71, no. 8, pp. 8280–8295, 2022.

[6] C. Balerna, M.-P. Neumann, N. Robuschi, P. Duhr, A. Cerofolini, V. Ravaglioli, and C. Onder, "Time-optimal low-level control and gearshift strategies for the formula 1 hybrid electric powertrain," *Energies*, vol. 14, p. 171, 2021.

[7] J. Bekker and W. Lotz, "Planning formula one race strategies using discrete-event simulation," *Journal of the Operational Research Society*, vol. 60, no. 7, pp. 952–961, 2009.

[8] A. Heilmeier, M. Graf, and M. Lienkamp, "A race simulation for strategy decisions in circuit motorsports," in *Proc. 2018 IEEE ITSC*, 2018, pp. 2986–2993.

[9] McLaren-Racing-Limited, "Formula one race strategy," Royal Academy of Engineering, Tech. Rep., 15/02/2022. [Online]. Available: https://www.raeng.org.uk/publications/other/14-car-racing

[10] M.-P. Neumann, G. Fieni, C. Balerna, P. Duhr, A. Cerofolini, and C. H. Onder, "Low-level online control of the formula 1 power unit with feedforward cylinder deactivation," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 7, pp. 8382–8397, 2023.

[11] G. Riva, S. Radrizzani, and G. Panzani, "Battery model impact on time-optimal co-design for electric racing cars: review and application," 2023.

[12] S. Radrizzani, G. Riva, G. Panzani, M. Corno, and S. M. Savaresi, "Optimal sizing and analysis of hybrid battery packs for electric racing cars," *IEEE Transactions on Transportation Electrification*, vol. 9, no. 4, pp. 5182–5193, 2023.

[13] G. Zardini, N. Lanzetti, M. Salazar, A. Censi, E. Frazzoli, and M. Pavone, "On the co-design of av-enabled mobility systems," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, 2020, pp. 1–8.

[14] G. Zardini, N. Lanzetti, A. Censi, E. Frazzoli, and M. Pavone, "Co-design to enable user-friendly tools to assess the impact of future mobility solutions," *IEEE Transactions on Network Science and Engineering*, vol. 10, no. 2, pp. 827–844, 2023.

[15] G. Zardini, D. Milojevic, A. Censi, and E. Frazzoli, "Co-design of embodied intelligence: A structured approach," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 7536–7543.

[16] G. Zardini, A. Censi, and E. Frazzoli, "Co-design of autonomous systems: From hardware selection to control synthesis," in *2021 European Control Conference (ECC)*, 2021, pp. 682–689.

[17] G. Zardini, Z. Suter, A. Censi, and E. Frazzoli, "Task-driven modular co-design of vehicle control systems," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, pp. 2196–2203.

[18] M. Clemente, M. Salazar, and T. Hofman, "Concurrent powertrain design for a family of electric vehicles," *IFAC-PapersOnLine*, vol. 55, no. 24, pp. 366–372, 2022, 10th IFAC Symposium on Advances in Automotive Control AAC 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896322023436

[19] J. Katz, "Aerodynamics in motorsports," *Proceedings of the Institution of Mechanical Engineers, Part P: Journal of Sports Engineering and Technology*, vol. 235, no. 4, pp. 324–338, 2021. [Online]. Available: https://doi.org/10.1177/1754337119893226

[20] R. S. Rice, "Measuring car-driver interaction with the g-g diagram," *SAE Technical Paper 730018*, 1973.

[21] A. Censi, "A mathematical theory of co-design," *arXiv preprint arXiv:1512.08055*, 2015.

[22] A. Censi, J. Lorand, and G. Zardini, *Applied Compositional Thinking for Engineering*, 2024, work-in-progress book. [Online]. Available: https://bit.ly/3H6pwMo

[23] G. Zardini, "Co-design of complex systems: From autonomy to future mobility systems," Doctoral Thesis, ETH Zurich, Zurich, 2023.

[24] B. A. Davey and H. A. Priestley, "Introduction to lattices and order," *Cambridge university press*, 2002.

[25] S. C. Kim and S. Y. Han, "Effect of steady airflow field on drag and downforce," vol. 17, pp. 205–211, 2016.

[26] D. S. Nath, P. C. Pujari, A. Jain, and V. Rastogi, "Drag reduction by application of aerodynamic devices in a race car," vol. 3, 2021.