Hierarchical Co-Design for Multi-Race Strategy Optimization in Formula 1

Working Paper

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

Publication date:
2025-10-29

Permanent link:
https://doi.org/https://doi.org/10.3929/ethz-c-000784888

Hierarchical Co-Design for Multi-Race Strategy Optimization in Formula 1

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

Abstract— **This work presents a hierarchical optimization framework for season-level decision-making in Formula 1, combining monotone co-design with sequential optimization. By integrating component design, energy deployment, and long-term degradation within a unified structure, the approach captures both local (lap and race) trade-offs and global (seasonal) constraints. The co-design layer generates track-dependent Pareto-optimal mappings between performance and wear, which are then composed into a finite-horizon optimization problem governing component usage and replacement decisions. Applied to a hybrid-electric Formula 1 power unit, the framework determines optimal battery sizing, deployment, and replacement timing across an entire season. Results show that exploiting regulatory constraints strategically, accepting local penalties to enable global gains, can increase cumulative championship points. Furthermore, while race order does not alter the attainable total reward, it modifies the optimal control policy, illustrating the temporal coupling inherent to multi-stage decision problems. Beyond motorsport, the proposed formulation provides a general template for long-horizon resource management and co-design of complex engineering systems.**

I. Introduction

Formula 1 (F1) is one of the most advanced and competitive forms of open-wheel circuit racing, renowned for its technological innovation, global reach, and competitive intensity since the championship began in 1950. Today, the series spans 24 races across five continents, featuring 10 teams, each represented by two cars driven by the most skilled drivers in motorsport. The diversity of the calendar plays a central role in the sport's strategic and engineering challenges. Track layouts vary significantly, from high-speed circuits like the Italian Grand Prix (GP) in Monza with long straights and sweeping corners, to narrow and twisty street tracks such as the Monaco GP, where agility and mechanical grip are crucial. As a result, both vehicle design and race strategy must be adapted to the unique requirements of each venue, while complying with the overarching racing regulations [1], [2].

Modern F1 race cars are propelled by a hybrid electric power unit, composed of a 1.6 L V6 internal combustion engine (ICE), two electric machines, and a high-voltage battery. While the interplay between these components determines lap-wise performance, this work focuses on the battery, which is both a performance-critical and regulation-constrained element across the season. Ultimately, each driver competes to accumulate championship points awarded according to their finishing position, as shown in Table I.

[1]Institute for Dynamic Systems and Control, ETH Zürich, Zürich, Switzerland, mneumann@ethz.ch

[2]Power Unit Performance and Control Group, Ferrari S.p.A., 41053 Maranello, Italy, Alberto.Cerofolini@ferrari.com

[3]Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, USA, gzardini@mit.edu

Pos.	**P1**	**P2**	**P3**	**P4**	**P5**	**P6**	**P7**	**P8**	**P9**	**P10**
Points	25	18	15	12	10	8	6	4	2	1

TABLE I: Rewarded points depending on the finishing race position. From P11 onwards, no points are awarded.

The goal is to maximize the seasonal point tally, balancing immediate race performance with long-term strategic considerations, especially in a context where technical and regulatory constraints evolve across the season.

Beyond the regulations imposed by the Fédération Internationale de l'Automobile (FIA) that constrain energy flows and power unit operation, additional constraints span across race and season levels. On a seasonal level, only two (potentially different) battery units are permitted. Should one of them be replaced, a penalty of 10 starting positions is applied for the first infraction, and 5 positions for each subsequent replacement. This intertwining of tactical and strategic constraints calls for a framework that can jointly optimize performance and resource allocation.

Co-design methods are traditionally applied to optimize both physical configuration and operational strategy, as shown in [3] for the lap and race optimization. However, when modeling an entire racing season as a sequence of races, the computational effort of classical co-design approaches becomes prohibitive. Therefore, we propose dynamic programming (DP) as a compelling alternative: by exploiting problem structure and the principle of optimality, DP allows for tractable and globally consistent decision-making across temporally coupled stages, making it especially suitable for seasonal-level planning [4].

A. Related work

Research in this area can be divided into two streams: optimization of hybrid-electric vehicle design and control, and the application of multi-objective methods for strategic decision-making in motorsports. The former focuses on component sizing and energy management trade-offs, while the latter explores race-level optimization problems involving stochastic events and tactical choices.

In the context of hybrid vehicle design, [5] presents a joint optimization of battery sizing and energy management for a city bus, while [6] extends such co-optimization to include motor sizing, minimizing operational cost. Works such as [7] and [8] apply similar methods to optimize lap times by jointly considering battery size and control. In motorsports applications, [3] and [9] introduce a co-design formulation that integrates hybrid-electric F1 car design with race strategy, leveraging a monotone optimization framework. DP has also been widely applied to race strategy optimization: [10] develops a deterministic model for pit-stop and tire decisions, while [11] extends this approach to a stochastic setting, accounting for random events such as safety cars and weather variations. Beyond race-level optimization, [12] addresses race order optimization from an environmental perspective.

Overall, the studied literature has two fundamental fallacies. First, optimal component design is typically analyzed only from a lap or race-level perspective. Second, the optimal deployment of components from a season-wide point of view, capturing degradation, replacement, and long-term strategy, has not been investigated.

B. Statement of Contribution

Addressing a gap in literature on seasonal design and component employment optimization, this work extends the prior lap- and race-focused contribution [3] by merging monotone co-design with dynamic programming. This combination removes the computational limitations of each method, enabling tractable optimization across laps, races, and an entire F1 season. By means of a fictitious F1 season we show the optimal employment strategy under various scenarios. Moreover, we show the change in strategy if the seasonal calendar changes.

C. Paper Structure

In Section II, we present our methodology, introducing monotone co-design theory for lap and race optimization and DP for seasonal optimization. Section III details the design problems and assumptions underlying the race car and season models. Section IV demonstrates the framework through two case studies: optimal component deployment over a full F1 season and the impact of race order on strategy. Finally, Section V summarizes the findings and outlines future research directions.

II. Methodology

Our framework combines two well-established methods to solve the problem of component design optimization from a seasonal perspective. As DP is extensively covered in the literature, we do not elaborate on its theory here and refer the reader to [4], [13] for details. In this section we introduce the foundational principles underlying a monotone theory of co-design [14]–[16]. Our discussion assumes familiarity with fundamental notions from order theory. For a comprehensive introduction, see [17].

A. Monotone Co-Design Theory

The co-design framework is built around a core construct referred to as the monotone design problem with implementations (MDPI).

Definition 1. (MDPI) Let $\mathcal{F}$ and $\mathcal{R}$ be partially ordered sets (posets) representing functionalities and resources, respectively. A MDPI d is characterized by a tuple $\langle \mathcal{I}_d, \text{prov}, \text{req} \rangle$, where $\mathcal{I}_d$ denotes the set of implementations, prov: $\mathcal{I}_d \xrightarrow{\text{prov}} \mathcal{F}$ assigns to each implementation the functionality it delivers, and req: $\mathcal{I}_d \xrightarrow{\text{req}} \mathcal{R}$ indicates the resources it consumes. This structure induces a monotone map $\bar{d}$:

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

where $\mathcal{F}^{\text{op}}$ is the poset $\mathcal{F}$ with reversed ordering. The set $\bar{d}(f^*, r)$ contains all implementations that deliver at least functionality f while requiring no more than resource r.

Remark 1 (Monotonicity)**.** The monotone nature of $\bar{d}$ guarantees that:

- If $f' \preceq_{\mathcal{F}} f$, then $\bar{d}(f', r) \supseteq \bar{d}(f, r)$, i.e., relaxing the functional requirements broadens the set of feasible implementations.
- If $r' \succeq_{\mathcal{R}} r$, then $\bar{d}(f, r') \supseteq \bar{d}(f, r)$, i.e., increasing available resources can only expand feasibility.

Remark 2 (Populating an MDPI)**.** In practical applications, the feasibility relation $\bar{d}$ can be specified using analytical models, numerically evaluated functions, or in a data-driven fashion (e.g., via simulations, optimizations, experiments), depending on the complexity and nature of the system under study.

Definition 2. For a given MDPI d, one can define two monotone maps:

- $h_d : \mathcal{F} \to \mathcal{A}(\mathcal{R})$, associating to each functionality the minimal antichain of resources needed to realize it.
- $h'_d : \mathcal{R} \to \mathcal{A}(\mathcal{F})$, associating to each resource budget the maximal antichain of achievable functionalities.

These maps capture the Pareto-optimal trade-offs inherent to the design problem.

Solving a design problem amounts to computing these monotone maps, which can be done via iterative methods based on fixed-point theory, such as Kleene's fixed point theorem as described in. Individual MDPIs can be systematically connected to model complex systems, forming structured networks of design problems, often referred to as co-design graphs. This modular composition allows large-scale problems to be split into manageable subproblems.

- **Series composition** links problems where the functionality of one serves as the resource for another (e.g., the power output of a motor acting as input for a lap time computation).
- **Parallel composition** models independent subsystems operating concurrently.
- **Feedback (loop) composition** introduces interdependencies between functionalities and resources, modeling closed-loop behavior.

Those composition rules retain the underlying monotonicity, preserving tractability and enabling efficient computation. When these models are instantiated with real system data, they naturally lead to multi-objective optimization problems. These problems maintain low computational complexity, linear in the number of implementation options per component, while being conceptually clear and structurally modular.

III. Modeling

In this section we introduce the model for an F1 season in our framework. After briefly recalling the adapted co-design formulation of [3] in Section III-A, we present the data-driven position dependency model in Section III-B, and finally connect these components to the seasonal optimization via DP in Section III-C.

A. Lap- and Race-Level Co-Design

The work of [3] models the F1 car at two interconnected abstraction levels: the *lap level*, which captures instantaneous physical trade-offs between component configuration and trade-off, and the *race level*, which aggregates these dynamics over multiple laps to derive track-specific performance mappings. Both levels are expressed in the formalism of co-design, providing feasibility relations between functionalities and resources.

a) Lap level: At the lap scale, the objective is to determine the optimal component configuration that minimizes lap time t_{lap} while constraining battery wear w_b. Each feasible implementation $i \in \mathcal{I}_{\text{lap}}$ yields a tuple

$$\left(t_{\text{lap}}^{(i)},\, w_b^{(i)}\right) = h_{\text{lap}}(i,\, \theta_{\text{track}}),$$

where θ_{track} denotes circuit-dependent parameters such as curvature distribution and elevation profile. The underlying design problem involves the following submodules:

- **Optimal control problem (OCP)**: An optimal control problem linking component choices and control trajectories to lap time, subject to track-dependent dynamics and constraints.
- **Aerodynamic configuration (AC)**: Defines the aerodynamic setup, including drag and cooling parameters, with two configurations considered: low drag configuration (LDC) and high drag configuration (HDC).
- **Engine cooling**: Constrains sustainable engine power. Larger radiator apertures improve thermal dissipation at the cost of higher drag and mass.
- **Battery health**: A degradation model capturing the coupling between deployed electrical energy, temperature, and cumulative wear.

Each of these modules is represented as an individual MDPI and composed via the operators of the co-design theory, yielding a tractable composite mapping from aerodynamic and powertrain parameters to lap performance. To ensure compatibility with the seasonal optimization, the battery type is treated as a fixed parameter, not a design variable, resulting in independent Pareto-optimal solutions $\mathcal{P}_{\text{lap}}^{(b)}$ for each available battery b.

b) Race level: The race-level problem extends the lap-scale formulation by integrating temporal and strategic decisions over the full race horizon. Given a race length L_{race}, the expected total time and cumulative wear are obtained by aggregating lap-level outcomes:

$$(t_{\text{race}}^{(i)},\, \Delta w_b^{(i)}) = h_{\text{race}}^{(b)}(i, \bar{w}_b^{\text{in}}, p_{\text{strat}}, \theta_{\text{track}}),$$

where $\bar{w}_b^{\text{in}}$ is the initial battery age, p_{strat} indexes feasible strategy implementations (e.g., energy deployment and aerodynamic setup), and $h_{\text{race}}^{(b)}$ denotes the race-level co-design map for battery b. The race model incorporates the following strategic layers:

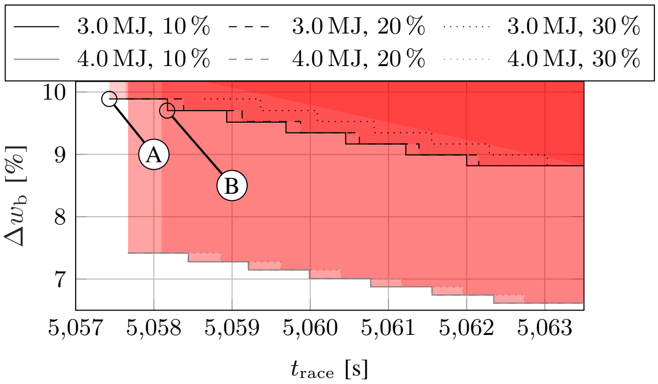

Fig. 1: Race results for every choice of battery and initial age, on the Circuit Paul Ricard (FRA). The red areas define the feasible space of the respective choice. The nodes on the Pareto fronts shown by means of solid, dashed and dotted lines depict the respective optimal implementations. The highlighted nodes are explained in the text.

- **Energy allocation**: Distributes available energy between fuel and electric deployment per lap, balancing lap time against energy consumption and vehicle mass.
- **Pit stops**: Captures mandatory tire-change rules and time losses, using a catalog of precomputed optimal stop strategies derived via dynamic programming.

c) Race solution: For each track and battery configuration, the mapping $h_{\text{race}}^{(b)}$ produces a Pareto front relating race time, cumulative wear, and feasible component and strategy implementations. An example for the Circuit Paul Ricard (FRA) is shown in Fig. 1. Specifically, for each battery size and initial age (10–30%), the model yields a discrete set of nondominated solutions $\{(t_{\text{race}}, \Delta w_b, i, p_{\text{strat}})\}$ forming the track-dependent Pareto frontier. Two points on such a frontier may exhibit nearly identical race times but differ in degradation due to variations in maximum deployable electric power, which arise from the battery health model within the corresponding MDPI. Each Pareto point thus represents an optimal combination of aerodynamic configuration, energy strategy, and thermal setup. Notably, solutions that appear dominated at the single-race level (e.g., higher wear for equal race time) are retained, as they may become optimal when aggregated in the season-level dynamic program.

B. Position Estimation Model

Unlike previous analyses, which evaluate performance solely in terms of race time, we consider the expected championship points Λ_{cs} as the primary performance metric. Since Λ_{cs} depends on the final finishing position Pos_{end} rather than raw time, we introduce a stochastic mapping that links predicted race times to expected points while accounting for the influence of the starting grid position Pos_0. The model decomposes Pos_{end} into a deterministic component driven by race time and a probabilistic correction capturing position uncertainty.

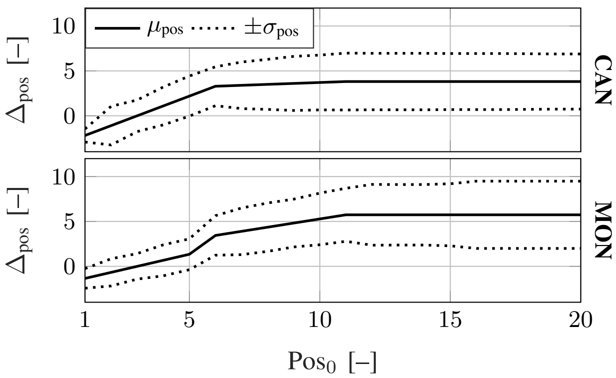

Fig. 2: Expected values and standard deviations of the position penalty, as a function of the initial position Pos_0, for CAN and MON. The latter's expected penalty is higher, aligning with the difficulty of overtaking on that track.

a) Deterministic race-time component: Let t_{race} denote the race time obtained from the co-design mapping $h^{(b)}_{\text{race}}(i, \bar{w}^{\text{in}}_b, p_{\text{strat}}, \theta_{\text{track}})$. The deterministic position function

$$\text{Pos}_{\text{end,det}} = f_{\text{gap}}(t_{\text{race}}),$$

maps race time to an expected finishing position using empirically derived time gaps between consecutive ranks. The function $f_{\text{gap}}(\cdot)$ is fitted from historical FIA race data over the past decade and normalized such that the minimum attainable time corresponds to $\text{Pos}_{\text{end,det}} = 1$ (P1, e.g., point Ⓐ in Fig. 1). Notably, this deterministic model neglects the influence of grid position, implying that both a pole-sitter and a driver starting from P10 would achieve P1 given identical t_{race}. Moreover, since the minimum times are associated with new components (lowest wear), winning becomes impossible once aging effects increase t_{race}. To mitigate these biases, we introduce a probabilistic correction term conditioned on Pos_0.

b) Probabilistic starting-position component: The probabilistic correction (also called position penalty) models the empirical dependency between starting and finishing positions as a random variable Δ_{pos}, representing the position offset due to grid starting order, stochastic race events, and track-specific overtaking difficulty. We assume

$$\Delta_{\text{pos}}(\text{Pos}_0) \sim \mathcal{N}\big(\mu_{\text{pos}}(\text{Pos}_0),\, \sigma^2_{\text{pos}}(\text{Pos}_0)\big),$$

where $\mu_{\text{pos}}(\text{Pos}_0)$ and $\sigma_{\text{pos}}(\text{Pos}_0)$ are the corrected mean finishing position and its standard deviation, estimated from race data of the top-performing teams (typically the four strongest constructors). The correction stems from a reference position $\text{Pos}_{\text{offset}}$, defining zero expected penalty, corresponding to a baseline performance for which $\mu_{\text{pos}}(\text{Pos}_{\text{offset}}) = 0$. In our case studies, we will set $\text{Pos}_{\text{offset}} = 3$. The expected penalties and variances for the Circuit Gilles Villeneuve (CAN) and the Circuit de Monaco (MON) are illustrated in Fig. 2. As seen, penalties saturate beyond P12, reflecting the performance dominance of top cars, while for positions better than $\text{Pos}_{\text{offset}}$, the mean offset $\mu_{\text{pos}} < 0$ yields a performance bonus.

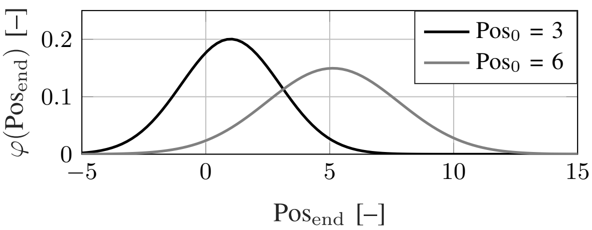

Fig. 3: Probability distribution of the final position, assuming starting positions P3 and P6, and the fastest implementation for FRA.

c) Combined finishing-position distribution: The final finishing position is then modeled as a continuous random variable Pos_{end} with distribution

$$\text{Pos}_{\text{end}} \sim \mathcal{N}\big(f_{\text{gap}}(t_{\text{race}}) + \mu_{\text{pos}}(\text{Pos}_0),\, \sigma^2_{\text{pos}}(\text{Pos}_0)\big).$$

The corresponding probability density $\varphi(\text{Pos}_{\text{end}})$ quantifies the likelihood of finishing in a given position given race time and grid start. As an illustration, Fig. 3 shows $\varphi(\text{Pos}_{\text{end}})$ for FRA under two starting positions (P3 and P6). The distribution for P3 peaks at P1, consistent with the baseline offset, whereas the P6 case shifts toward P5 due to the positive $\mu_{\text{pos}}(6)$.

d) Expected points: Expected championship points follow from the expectation of Λ_{cs} with respect to Pos_{end}:

$$\mathbb{E}[\Lambda_{\text{cs}}] = \int_{-\infty}^{\infty} \Lambda_{\text{FIA}}(\text{pos})\, \varphi(\text{pos})\, \text{dpos},$$

where $\Lambda_{\text{FIA}}(\cdot)$ is the FIA scoring function defined in Table I. Between integer positions, Λ_{FIA} is linearly interpolated to ensure differentiability. Values below P1 may appear probabilistically due to Gaussian tails but are clipped to the maximum of 25 points. The scalar $\mathbb{E}[\Lambda_{\text{cs}}]$ constitutes the expected reward associated with each feasible implementation in the seasonal dynamic program.

C. Dynamic Programming

Once the expected points $\mathbb{E}[\Lambda_{\text{cs}}]$ are assigned to each feasible implementation and starting position, the season-long optimization is formulated as a finite-horizon Markov decision process (MDP). We solve it using Bellman's dynamic programming [18], extended here to incorporate component wear, replacement penalties, and grid-dependent stochastic rewards under the 2025 FIA regulations.

a) MDP formulation: The MDP is defined by the tuple

$$\mathcal{M} = (\mathcal{X}, \mathcal{U}, \mathcal{D}, f, J),$$

where $\mathcal{X}$ is the state space, $\mathcal{U}$ the control space, $\mathcal{D}$ the disturbance space, f the deterministic transition mapping, and J the expected reward at each stage.

b) State space: Regulations allow to use at most two distinct battery units per season, which may be deployed and replaced strategically. The system state at race k is

$$\vec{x}_k = \begin{bmatrix} w_{\text{b},1,k} & w_{\text{b},2,k} & x_{\text{ex},k} \end{bmatrix}^\top,$$

where $w_{\text{b},j,k} \in [0, 0.3]$ denotes the fractional wear of battery $j \in \{1, 2\}$ (discretized at 0.25 %), and $x_{\text{ex},k} \in \{0, 1\}$

indicates whether any replacement penalty has already been incurred. The state therefore encodes both the physical condition of components and the regulatory context.

c) Control inputs: At each race, the control vector

$$\vec{u}_k = \begin{bmatrix} u_{\mathrm{b},k} & u_{\mathrm{imp},k} & u_{\mathrm{ex},k} \end{bmatrix}^\top$$

comprises the battery choice $u_{\mathrm{b},k} \in \{1, 2\}$ to be employed, the implementation index $u_{\mathrm{imp},k} \in \{1, \ldots, n_{\mathrm{imp}}\}$, referring to a specific Pareto-optimal configuration from $h^{(b)}_{\mathrm{race}}$ (e.g., point Ⓑ in Fig. 1), and the replacement decision $u_{\mathrm{ex},k} \in \{0, 1\}$, where $u_{\mathrm{ex},k} = 1$ triggers the installation of a new unit and a grid penalty according to FIA rules.

d) Disturbances: External conditions affecting performance are grouped in the disturbance vector

$$\vec{d}_k = \begin{bmatrix} d_{\mathrm{r},k} & d_{\mathrm{pos},k} \end{bmatrix}^\top,$$

where $d_{\mathrm{r},k}$ encodes the track-specific parameters $\theta_{\mathrm{track},k}$ and $d_{\mathrm{pos},k}$ denotes the starting position. These variables are exogenous and drawn from historical or simulated distributions.

e) Transition dynamics: Each stage k corresponds to one race weekend. The next-state mapping $f : \mathcal{X} \times \mathcal{U} \times \mathcal{D} \to \mathcal{X}$ is defined as

$$w_{\mathrm{b},u_{\mathrm{b},k},k+1} = \begin{cases} 0, & \text{if } u_{\mathrm{ex},k} = 1, \\ w_{\mathrm{b},u_{\mathrm{b},k},k} + \Delta w_{\mathrm{b}}^{(u_{\mathrm{b},k}, u_{\mathrm{imp},k}, d_{\mathrm{r},k})}, & \text{otherwise}, \end{cases}$$
$$w_{\mathrm{b},j \neq u_{\mathrm{b},k},k+1} = w_{\mathrm{b},j,k},$$
$$x_{\mathrm{ex},k+1} = x_{\mathrm{ex},k} \vee u_{\mathrm{ex},k}.$$

The wear increment Δw_{b} is obtained from the co-design race-level map $h^{(b)}_{\mathrm{race}}$, tabulated for discrete wear levels up to end-of-life (EOL). If the accumulated wear exceeds this threshold, the corresponding battery becomes unavailable for subsequent stages.

f) Stage reward: The reward at stage k is the expected number of championship points earned in that race:

$$J_k(\vec{x}_k, \vec{u}_k, \vec{d}_k) = \mathbb{E}[\Lambda_{\mathrm{cs},k}(\vec{x}_k, \vec{u}_k, \vec{d}_k)],$$

where $\Lambda_{\mathrm{cs},k}$ is computed from the position estimation model in Section III-B. If a replacement occurs ($u_{\mathrm{ex},k} = 1$), the corresponding grid penalty is incorporated by reducing the expected points by $\Delta\Lambda_{\mathrm{pen}}$, determined from the empirical impact of a starting-position drop.

g) Optimization objective: The goal is to maximize the cumulative expected reward over the entire season of N_{races} events:

$$\max_{\pi} \mathbb{E}\left[\sum_{k=1}^{N_{\mathrm{races}}} J_k(\vec{x}_k, \pi(\vec{x}_k), \vec{d}_k)\right],$$

where $\pi : \mathcal{X} \to \mathcal{U}$ is the control policy. The optimal policy satisfies the Bellman recursion:

$$V_k(\vec{x}_k) = \max_{\vec{u}_k \in \mathcal{U}} \Big(J_k(\vec{x}_k, \vec{u}_k, \vec{d}_k) + V_{k+1}\big(f(\vec{x}_k, \vec{u}_k, \vec{d}_k)\big) \Big),$$

with terminal condition $V_{N_{\mathrm{races}}+1}(\vec{x}) = 0$. Backward iteration over the discretized state grid yields a globally optimal sequence of deployment and replacement decisions across the season.

Nr.	Venue	Pos_0	Nr.	Venue	Pos_0
1	Australian GP	3	10	British GP	3
2	Bahrain GP	2	11	German GP	4
3	Chinese GP	3	12	Belgian GP	2
4	Azerbaijan GP	1	13	Italian GP	3
5	Spanish GP	3	14	Singapore GP	2
6	Monaco GP	2	15	Japanese GP	3
7	Canadian GP	2	16	Mexican GP	3
8	French GP	3	17	United States GP	2
9	Austrian GP	3	18	Abu Dhabi GP	3

TABLE II: Season schedule and corresponding starting positions for each race.

	Setup No.		
	Setup 1	**Setup 2**	**Setup 3**
Energy Battery 1 [MJ]	3.0	3.0	4.0
Energy Battery 2 [MJ]	3.0	4.0	4.0

TABLE III: Battery setups evaluated in the case studies. Setups 1 and 3 use identical batteries; Setup 2 combines the two sizes.

IV. RESULTS

This section presents the application of the proposed framework to realistic F1 scenarios. Two case studies are discussed. The first evaluates the effect of different battery configurations and replacement policies on seasonal performance. The second investigates the sensitivity of the optimal policy to race-order permutations, illustrating the temporal coupling inherent to multi-stage decision problems.

A. Design of Experiments

The analysis considers a fictitious 18-race F1 season with venues and nominal starting positions summarized in Table II. Each race corresponds to one stage in the finite-horizon optimization problem. Three battery configurations, summarized in Table III, are evaluated: two with identical batteries (Setups 1 and 3) and one mixed configuration (Setup 2). For each setup, two optimization runs are performed: one prohibiting replacements and one allowing a single penalty-inducing replacement per season.

B. Part One – Optimal Battery Setup

Table IV reports the total expected championship points obtained for each configuration. Without replacement, Setup 1 is infeasible, as two 3.0 MJ batteries reach their EOL before season completion. Among feasible cases, Setup 3 (two 4.0 MJ units) achieves the highest score, while Setup 2, combining one small and one large battery, performs slightly worse due to cumulative degradation effects.

Fig. 4 illustrates the race-by-race strategy for Setups 2 and 3 in the case without replacement. Setup 2 alternates between high-wear/low-time (i.e., implementation 1) and conservative implementations to balance short-term and long-term performance, as reflected in the battery age trajectories of Fig. 5 (top). Battery 1 reaches its EOL approximately five races before season end, forcing exclusive reliance on Battery 2. In contrast, Setup 3 maintains consistently aggressive (high-wear) implementations throughout, leveraging

	Total Expected Points		
	Setup 1	**Setup 2**	**Setup 3**
Replacement Disabled	–	383.6	394.6
Replacement Enabled	387.5	388.4	394.6

TABLE IV: Expected championship points achieved with and without the possibility to perform an additional replacement (incurring a grid penalty) for the various setups.

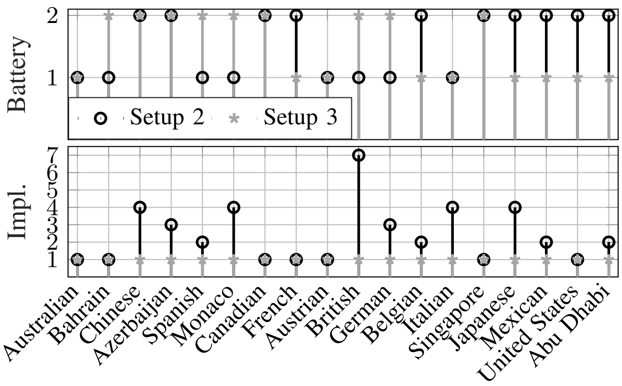

Fig. 4: Race-level strategies for Setups 2 and 3 without replacement. Top: active battery; bottom: selected implementation index. Lower indices correspond to higher wear and faster lap times.

its higher nominal capacity and thus achieving the highest cumulative reward.

When replacements are permitted, Setup 1 becomes feasible but remains suboptimal. Setup 2 gains approximately five championship points, while Setup 3 remains unchanged, as its larger batteries make replacement unnecessary. The optimal policy for Setup 2 (shown in Fig. 6) replaces the smaller battery at the Austrian GP, enabling high-wear strategies for both units at all races. This illustrates the framework's ability to exploit regulatory penalties optimally, sacrificing a short-term grid position for long-term performance gain.

Overall, Setup 3 consistently yields the best seasonal outcome, confirming that higher energy capacity enhances robustness to cumulative degradation. However, allowing a single replacement narrows the performance gap between Setups 2 and 3 to less than seven points, suggesting diminishing returns beyond a certain energy capacity. This trade-off between battery sizing, degradation, and regulatory penalties exemplifies the strength of the proposed hierarchical optimization framework.

C. Part Two – Race Order Permutation

The second case study examines the sensitivity of the optimal policy to the ordering of events in the season. To this end, the sequence in Table II is modified by swapping the French and Austrian Grands Prix. The total expected points change marginally (from 388.4 to 388.3), confirming that the achievable cumulative reward is invariant to race order under deterministic dynamics.

However, the *structure of the optimal policy* changes. In the baseline schedule, the smaller $3.0\,\mathrm{MJ}$ battery is replaced

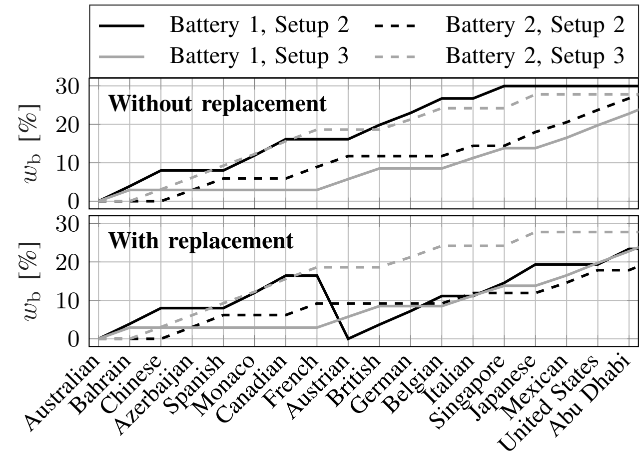

Fig. 5: Battery wear evolution for Setups 2 and 3. Top: replacement disabled; bottom: replacement enabled. In the latter, Setup 2 resets Battery 1 in Austria (Race 9).

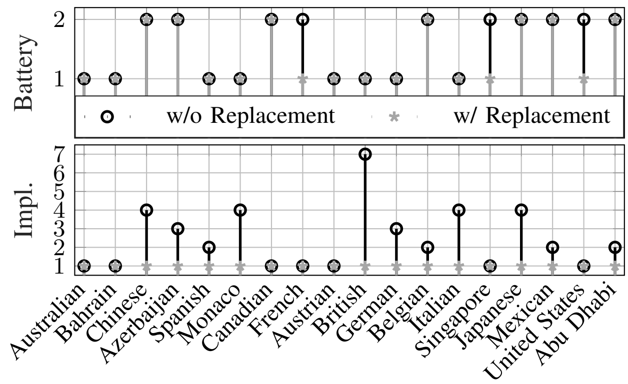

Fig. 6: Optimal deployment strategy for Setup 2 with replacement. The replacement of Battery 1 occurs at the Austrian GP (Race 9), after which high-wear implementations dominate.

at the Austrian GP (Race 9). After swapping the events, the replacement shifts to the British GP (Race 10). This demonstrates that the optimal replacement decision depends not on a specific circuit but on the overall temporal context, battery state, remaining races, and track sequence.

From a control-theoretic perspective, this experiment highlights that the problem exhibits the *principle of optimality*: the current decision depends only on the present state (battery ages and replacement status), but the resulting policy is still sensitive to future race orderings through the value function. Hence, while the achievable total reward remains constant, the optimal policy adapts dynamically to the temporal structure of the season, confirming the suitability of the MDP formulation for such long-horizon coupled problems.

V. Conclusions

This work presented a unified optimization framework that integrates *monotone co-design* and *dynamic programming* to solve long-horizon decision problems in the presence of coupled physical and regulatory constraints. By exploiting the structural properties of both methods, the proposed approach achieves a decomposition that preserves global consistency while maintaining tractable computational complexity. In particular, monotone co-design provides compact representations of subsystem trade-offs, which can be seamlessly embedded into a dynamic program governing season-level decisions.

Two case studies demonstrated the capability of the framework to capture realistic interactions between component design, energy deployment, and regulatory penalties. The first study highlighted how optimal battery replacement and usage strategies emerge naturally from the dynamic interplay between short-term performance and long-term degradation. The second study analyzed the effect of race-order permutations, showing that while total attainable championship points remain invariant, the optimal policy adapts locally, illustrating the temporal coupling intrinsic to multi-stage optimal control problems.

Beyond the application context, the results offer three key insights:

1) Structured co-design enables modular modeling of complex systems, where heterogeneous subsystems interact through formally defined resource–functionality relations.
2) Coupling such representations with dynamic programming extends classical control methods to decision spaces combining discrete design and continuous degradation states.
3) The resulting policy is interpretable, computationally tractable, and robust to variations in system configuration and scheduling constraints.

Future research will focus on increasing the fidelity and generality of the framework. First, incorporating stochastic components such as weather, safety-car events, or random failures would transform the formulation into a stochastic dynamic program, allowing explicit reasoning about uncertainty. Second, the integration of higher-resolution position models could better capture overtaking dynamics and qualifying performance. Finally, the general structure of the proposed approach, linking design optimization and sequential decision-making, can be extended to broader engineering domains, from fleet-level energy management to autonomous system operations under regulatory constraints.

ACKNOWLEDGMENTS

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

REFERENCES

[1] FIA, "2025 Formula one sporting regulations," Geneva, Switzerland, Tech. Rep., 2025.

[2] ——, "2025 Formula one technical regulations," Geneva, Switzerland, Tech. Rep., 2025.

[3] M.-P. Neumann, G. Fieni, F. Furia, A. Cerofolini, V. Ravaglioli, C. H. Onder, and G. Zardini, "Strategic co-design in Formula 1: Balancing physical configuration and race tactics," *ETH Zurich, ETH Library preprint*, 2025.

[4] R. Bellman, "Dynamic programming," *Science*, vol. 153, pp. 34 – 37, 1957.

[5] N. Murgovski, L. Johannesson, J. Sjöberg, and B. Egardt, "Component sizing of a plug-in hybrid electric powertrain via convex optimization," *Mechatronics*, vol. 22, no. 1, pp. 106–120, 2012.

[6] 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

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

[8] 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.

[9] M.-P. Neumann, G. Zardini, A. Cerofolini, and C. H. Onder, "On the co-design of components and racing strategies in Formula 1," in *2024 IEEE Intelligent Vehicles Symposium (IV)*, 2024, pp. 2876–2881.

[10] O. F. C. Heine and C. Thraves, "On the optimization of pit-stop strategies via dynamic programming."

[11] F. Aguad and C. Thraves, "Optimizing pit stop strategies in formula 1 with dynamic programming and game theory," *European Journal of Operational Research*, vol. 319, no. 3, pp. 908–919, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0377221724005484

[12] A. Vajirkar and T. Lee, "Improving formula 1 sustainability by ordering the f1 season calendar to optimize for carbon emissions."

[13] D. Bertsekas, *Dynamic programming and optimal control: Volume I*. Athena scientific, 2012, vol. 4.

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

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

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

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

[18] O. Sundstrom and L. Guzzella, "A generic dynamic programming matlab function," in *2009 IEEE Control Applications, (CCA) & Intelligent Control, (ISIC)*, 2009, pp. 1625–1630.