Traffic Network-Aware Energy Management for FCEVs: Integrating Trip-Specific Control and Long-Run Optimality

Kyunghwan Choi
CCS Graduate School of Mobility
KAIST
Daejeon, Republic of Korea
kh.choi@kaist.ac.kr

***Abstract*—Energy management for fuel cell electric vehicles (FCEVs) is a challenging trajectory optimization problem. Conventional studies primarily focus on trip-specific optimal control, where the power distribution is optimized based on a predicted finite-horizon driving profile. However, these methods often suffer from a limited look-ahead horizon and fail to guarantee long-run optimality within the stochastic traffic network where the vehicle operates. This study proposes a novel framework that integrates finite-horizon optimal control with traffic network-aware long-run average costs. We formulate the problem by embedding the long-run optimality, derived from network-level transition probabilities, into the terminal cost of the trip-specific optimization. This approach enables an adaptive target State of Charge (SOC) that aligns with global network efficiency while satisfying immediate driving constraints. Simulation results in a virtual traffic network demonstrate that the proposed integrated strategy consistently outperforms traditional trip-specific methods, achieving a maximum performance improvement of 11%. These findings highlight the necessity of network-level statistical awareness for maximizing the long-term energy efficiency of electrified mobility.**

***Index Terms*—Energy management strategy, Fuel cell electric vehicles (FCEVs), Traffic network-aware, Trip-specific optimal control, Long-run optimality, Integrated strategy.**

I. INTRODUCTION

The energy management strategy (EMS) for multiple power sources is a fundamental challenge in maximizing the fuel efficiency of fuel cell electric vehicles (FCEVs). Theoretically, the optimal energy management problem can be solved via dynamic programming (DP), but its noncausal nature requires a priori knowledge of the entire future power demand, making real-time implementation difficult. To bridge this gap, predictive energy management (PEM) has emerged, leveraging intelligent transportation systems (ITS) to anticipate upcoming driving conditions [1].

Traditionally, PEM has relied on trip-specific strategies that optimize energy distribution based on a predicted driving profile for a single, intended journey. These methods typically estimate future velocity or power demand using stochastic models like Markov chains [2] or machine learning techniques [3]. While effective for short-term adjustments, these trip-specific approaches are often criticized for their "short-sightedness," as their optimization is confined to a limited prediction horizon [4] and fails to account for the long-term energy implications beyond the immediate trip [5].

To provide more structured guidance, a more advanced methodology known as SOC node planning (or reference SOC planning) has been widely adopted [6], [7]. This approach utilizes navigation data to identify a predicted path and pre-calculates an optimal SOC trajectory for specific spatial nodes—such as intersections or hills—along that route. By assigning fixed SOC targets to these nodes, the vehicle attempts to balance energy consumption according to the anticipated terrain and traffic.

However, even with precise navigation, SOC node planning faces a critical limitation: it is inherently "path-locked." Most existing methods pre-determine a final SOC target value based on a single, fixed destination or a specific predicted trip. Such a rigid target is often decoupled from the global optimality of the broader traffic network. In a stochastic driving environment where a vehicle may encounter unexpected route changes or continuous operations across multiple journeys, a fixed SOC target derived only from the current trip can be counterproductive. It may drive the battery state into a region that is efficient for the immediate destination but highly suboptimal for the unpredicted, subsequent links within the network.

Recent studies have attempted to alleviate the short-sightedness of finite-horizon predictive control by incorporating cost-to-go functions as terminal costs within model predictive control frameworks [8]. In such approaches, the long-term information is typically derived from the remaining portion of a predefined route, where dynamic programming is used to compute a distance-based cost-to-go under deterministic trip assumptions. Although effective for a single planned mission, these methods remain route-dependent and assume that the future driving sequence is known or fixed. Consequently, their optimality is still confined to a specific trip and does not explicitly account for continuous operation within a stochastic traffic network.

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (RS2025-00554087).

These limitations motivate a broader perspective in which the vehicle is viewed as operating persistently within a traffic network rather than along a single predetermined route. In such a setting, the terminal cost should reflect not the remainder of a specific trip, but the long-run value of the SOC state within the entire network. This shift from route-dependent optimality to network-level optimality under uncertainty provides the foundation for the present study.

To resolve the conflict between rigid node-based planning and global network optimality, this paper proposes a traffic network-aware energy management framework that integrates trip-specific optimal control with long-run network statistics. The main contributions of this study are summarized as follows:

- **Integration of Local Control and Global Optimality**: This study presents a new perspective on EMS design by embedding the long-run average cost, derived from the stochastic transition probabilities of a traffic network, into the terminal cost of a finite-horizon optimal control problem. This formulation mathematically bridges the gap between trip-specific efficiency and network-level optimality.
- **Realization of Adaptive Target SOC**: It is revealed that the statistical "value" of each node in a traffic network can be used to determine an adaptive target SOC. Unlike conventional SOC node planning that relies on fixed setpoints, the proposed approach allows the vehicle to autonomously adjust its energy state to remain optimal for both current and unpredicted future links.
- **Validation of Network-Aware Efficiency**: A real-time integrated strategy is proposed and validated through simulations in a stochastic virtual traffic network. The results reveal that incorporating network-level statistical awareness provides a significant performance buffer against limited prediction horizons, achieving a maximum efficiency improvement of 11% over traditional trip-specific methods.

The remainder of this paper is organized as follows. Section II presents the conventional trip-specific strategy as preliminaries for this study. Section III introduces the proposed integrated energy management strategy, while Section IV describes the numerical solvers for both conventional and proposed strategies. Validation results are reported in Section V, and finally, Section VI concludes the paper with an outlook on future work.

II. Preliminaries

A. Trip-Specific Strategy

A trip is represented by a sequence of power demands $\{w_t\}_{t=0}^{T-1}$, where the subscript t denotes the time step from the start ($t = 0$) to the end ($t = T$) of the journey.

For a given trip, the trip-specific energy management problem for a FCEV is formulated as follows:

$$\min_{u_0,\ldots,u_{T-1}} \sum_{t=0}^{T-1} \dot{m}_{fc}(u_t) \tag{1a}$$

subject to

$$x_{t+1} = x_t + f(x_t, u_t, w_t), \tag{1b}$$

$$\underline{u} \leq u_t \leq \overline{u}, \tag{1c}$$

$$\underline{x} \leq x_t \leq \overline{x}, \tag{1d}$$

$$x_T = x^*, \tag{1e}$$

where u_t is the control input (i.e., fuel-cell power); $\dot{m}_{fc}$ is the fuel consumption rate; x_t is the state (i.e., battery state of charge (SOC)); $f(\cdot)$ is the SOC dynamics model whose expression is given in [9]; $\underline{u}$ and $\overline{u}$ are the control limits; $\underline{x}$ and $\overline{x}$ are the state constraints; and x^* is the fixed target SOC at the end of the trip.

As a trajectory optimization problem, solving (1) requires the entire future power demand trajectory $\{w_t\}_{t=0}^{T-1}$ in advance. While a typical trip duration spans tens of minutes, the reliable prediction horizon in practical applications is often limited to tens of seconds. Consequently, existing PEM strategies based on such short-term horizons fail to achieve global optimality.

Furthermore, even with a perfect long-term prediction, the hard constraint in (1e)—which forces the state x_T to a fixed target x^*—can be suboptimal. This rigidity prevents the vehicle from adaptively preserving or depleting energy in response to the stochastic nature of the subsequent, unpredicted traffic network.

III. Proposed Energy Management Strategy

This section introduces a hierarchical framework that bridges local trip-specific control and global network optimality. We first present the long-run average strategy derived from the traffic network statistics, followed by the integrated strategy that embeds this long-run cost into a finite-horizon optimization.

A. Long-Run Average Strategy

To account for the infinite-horizon operation within a stochastic traffic network, the energy management problem is formulated as a discounted cost minimization:

$$\min_{\pi} J_\pi(x_0) = \lim_{T\to\infty} \mathbb{E}\left[\sum_{t=0}^{T-1} \gamma^t \dot{m}_{fc}(\pi(x_t))\right] \tag{2a}$$

subject to

$$x_{t+1} = x_t + f(x_t, \pi(x_t), w_t), \tag{2b}$$

$$\underline{u} \leq \pi(x_t) \leq \overline{u}, \tag{2c}$$

$$\underline{x} \leq x_t \leq \overline{x}, \tag{2d}$$

where π represents the control policy optimized for the traffic network and $\gamma \in (0, 1)$ is the discount factor. Notably, the final SOC constraint (1e) is omitted in this formulation, as the infinite-horizon perspective naturally evaluates the "value" of each SOC state x in terms of long-term fuel economy rather than aiming for a specific terminal setpoint.

The resulting cost-to-go, $J_\pi(x)$, represents the expected fuel consumption from state x under the optimal network policy.

B. Integrated Strategy: Trip-Specific Control with Long-Run Awareness

Leveraging the long-run cost $J_\pi(x)$ obtained from (2), we propose an integrated energy management strategy. This strategy optimizes the immediate trip while remaining cognizant of future network-level efficiency by incorporating $J_\pi(x)$ as a terminal cost:

$$\min_{u_0,\ldots,u_{T-1}} \sum_{t=0}^{T-1} \dot{m}_{fc}(u_t) + J_\pi(x_T) \tag{3a}$$

subject to

$$x_{t+1} = x_t + f(x_t, u_t, w_t), \tag{3b}$$

$$\underline{u} \le u_t \le \overline{u}, \tag{3c}$$

$$\underline{x} \le x_t \le \overline{x}. \tag{3d}$$

In this formulation, the hard constraint $x_T = x^*$ is replaced by the terminal cost $J_\pi(x_T)$.

Unlike the conventional trip-specific problem (1), the terminal SOC x_T is adaptively optimized; it is determined such that the sum of the immediate fuel consumption and the long-term expected cost is minimized. Consequently, the vehicle can "invest" SOC during the current trip if the long-run network value at x_T justifies the expenditure, thereby ensuring global optimality across continuous operations.

IV. SOLVERS

Problems (1)–(3) are trajectory optimization problems defined over multiple time steps. In particular, Problem (2) is formulated over an infinite horizon within a stochastic traffic network, which renders direct solution intractable.

To improve computational tractability while preserving optimality, we adopt a link-level reformulation based on the approach in [10]. This reformulation transforms the time-domain control problem into an equivalent link-wise optimization problem using Pontryagin's Minimum Principle (PMP). Under the convexity of the fuel consumption model and the lumped energy representation within each link, the reformulated problem is theoretically equivalent to the original time-level formulation.

A. Link-Level Representation

A road link is defined as a continuous road segment connecting two nodes at which significant speed transitions occur (e.g., intersections, highway ramps, bridges, or tunnels). The link index k denotes the link step, where each link contains multiple time steps t.

Let q_k denote the link at step k, and x_k denote the SOC state when exiting link q_k. The trajectory information within link q_k is aggregated into four lumped parameters:

- $E^+(q_k)$: positive energy demand,
- $E^-(q_k)$: negative energy demand,
- $\Delta t(q_k)$: total duration,
- $\Delta t^+(q_k)$: positive power duration.

Applying PMP to the original time-level optimal control problem yields that the optimal fuel-cell power within each link can be characterized by a constant costate λ_k. Under standard convexity assumptions on the fuel consumption function, the resulting link-level cost becomes quadratic in λ_k. Accordingly, the SOC dynamics can be reformulated as

$$x_k = x_{k-1} + g(\lambda_k, E^+(q_k), E^-(q_k), \Delta t(q_k), \Delta t^+(q_k)), \tag{4}$$

where $g(\cdot)$ denotes the lumped SOC update model derived from the original time-domain dynamics.

B. Trip-Specific Strategy

Based on the link-level representation, Problem (1) is reformulated as

$$\min_{\lambda_1,\ldots,\lambda_N} \sum_{k=1}^{N} \Delta t^+(q_k)\lambda_k^2 \tag{5a}$$

$$\text{s.t. } x_k = x_{k-1} + g(\lambda_k, \cdots), \tag{5b}$$

$$\underline{\lambda} \le \lambda_k \le \overline{\lambda}, \tag{5c}$$

$$\underline{x} \le x_k \le \overline{x}, \tag{5d}$$

$$x_N = x^*. \tag{5e}$$

For a fixed terminal SOC x^*, Problem (5) is a convex quadratic program (QP) in $\{\lambda_k\}_{k=1}^N$. It is solved using MATLAB's `quadprog` function. The computational complexity scales linearly with the number of links N.

C. Long-Run Average Strategy

In the link-level formulation, the infinite-horizon problem (2) becomes a constrained Markov Decision Process (MDP) defined over the augmented state space (x, q).

The reformulated problem is

$$\min_{\pi} \quad J_\pi(x_0, q_1) = \lim_{N\to\infty} \mathbb{E}\left[\sum_{k=1}^{N} \gamma^{k-1} \Delta t^+(q_k)\left(\pi(x_{k-1}, q_k)\right)^2\right] \tag{6a}$$

$$\text{s.t.} \quad P(q_{k+1} = j \,|\, q_k = i) = p_{ij}, \tag{6b}$$

$$x_k = x_{k-1} + g(\pi(x_{k-1}, q_k), \cdots), \tag{6c}$$

$$\underline{\lambda} \le \pi(x_{k-1}, q_k) \le \overline{\lambda}, \tag{6d}$$

$$\underline{x} \le x_k \le \overline{x}. \tag{6e}$$

The state space is discretized over the SOC grid and finite link set. The policy $\pi(x, q)$ and optimal cost-to-go $J(x, q)$ are obtained using value iteration (VI). Convergence is declared when the maximum Bellman residual falls below a predefined tolerance.

D. Integrated Trip-Specific and Long-Run Strategy

The integrated problem (3) is reformulated as

$$\min_{\lambda_1,\ldots,\lambda_N,x^*} \quad \sum_{k=1}^{N} \Delta t^+(q_k)\lambda_k^2 + \mathbb{E}_{q_{N+1}|q_N}\left[J_\pi(x^*, q_{N+1})\right] \tag{7a}$$

$$\text{s.t. } x_k = x_{k-1} + g(\lambda_k, \cdots), \tag{7b}$$

$$\underline{\lambda} \le \lambda_k \le \overline{\lambda}, \tag{7c}$$

$$\underline{x} \le x_k \le \overline{x}, \tag{7d}$$

$$x_N = x^*. \tag{7e}$$

Compared to Problem (5), the terminal SOC x^* is treated as an optimization variable. For a fixed x^*, the problem is convex in $\{\lambda_k\}$ and solved as a QP. Therefore, the overall problem reduces to a one-dimensional search over x^*.

Specifically, x^* is discretized over the admissible SOC grid. For each candidate x^*, Problem (5) is solved to obtain the optimal $\{\lambda_k\}$ and corresponding trip-level cost. The expected terminal cost is then evaluated using the long-run cost-to-go J_π. The x^* that minimizes the combined cost is selected as the optimal adaptive terminal SOC.

V. Validation

The proposed integrated strategy was evaluated in a virtual traffic network and compared against the trip-specific strategy and the long-run average strategy.

The virtual traffic network consists of nine nodes and twenty four directed links. The transition probabilities p_{ij} and the link-level parameters are illustrated in Fig. 1(a) and Fig. 1(b), respectively. The link parameters include the positive and negative energy demands, total duration, and positive power duration.

To evaluate statistical performance, 100 independent travel sequences were generated. Each travel starts from node 5 and evolves for $N_{\text{travel}} = 1000$ links according to the transition probabilities p_{ij}. Each long travel is partitioned into multiple finite trips, each consisting of N links. The parameter N represents the prediction horizon of the trip-specific component and is varied as $N = 1, 2, 5, 10, 15$, and 20 to investigate the effect of the finite-horizon length.

The target FCEV model is based on the specifications of the Hyundai NEXO. The SOC bounds were set to $\underline{x} = 0.4$ and $\overline{x} = 0.7$, and the initial SOC was set to $x_0 = 0.55$. For the trip-specific strategy, the terminal constraint was set to $x^* = 0.55$.

For fair comparison, all three strategies are evaluated using the same total cost metric:

$$J_{\text{total}} = \sum_{k=1}^{N_{\text{travel}}} \Delta t^+(q_k)\lambda_k^2. \tag{8}$$

Although each strategy is derived from different optimization formulations, their performance is assessed using this unified cost definition.

The total costs of the three strategies are shown in Fig. 2(a). The trip-specific strategy exhibits strong dependence on the

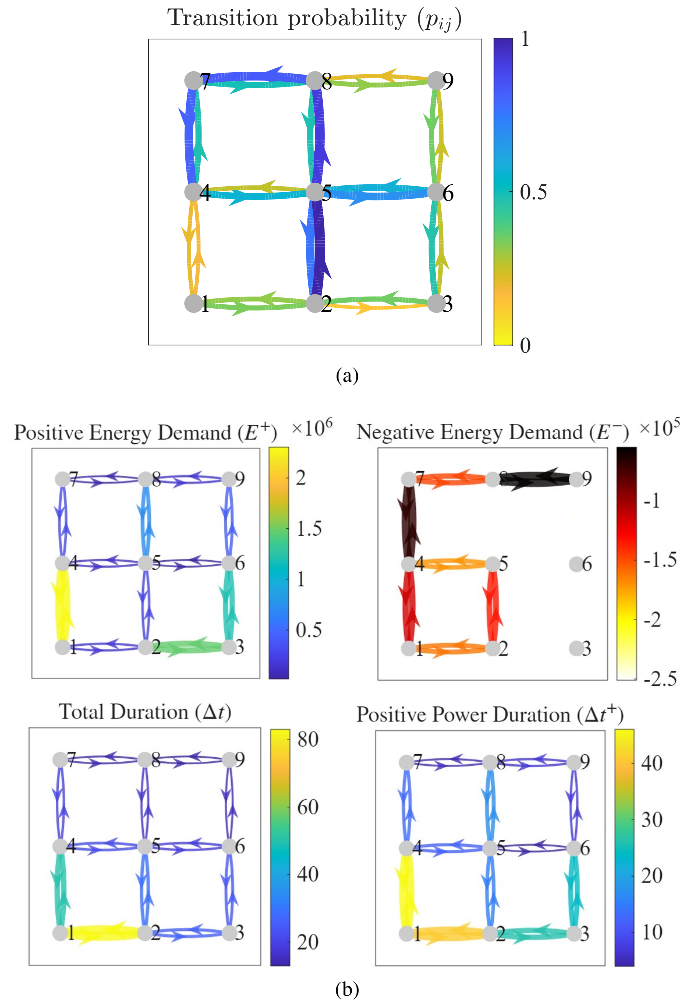

Fig. 1. Virtual traffic network used in the validation: (a) link transition probabilities p_{ij} and (b) link-level parameters including energy demands and travel durations.

prediction horizon N. When N is small, the limited look-ahead capability leads to suboptimal decisions and the highest total cost. As N increases, performance improves due to richer trip information.

In contrast, the proposed integrated strategy consistently outperforms the trip-specific strategy across all tested values of N. The improvement is particularly significant for short horizons, achieving a maximum cost reduction of approximately 11% at $N = 1$, as shown in Fig. 2(b). This result demonstrates that embedding long-run network information into the terminal cost effectively compensates for limited prediction horizons.

The long-run average strategy outperforms the trip-specific strategy when $N = 1, 2, 5$, and 10, indicating that network-level statistical awareness alone can enhance performance when trip information is scarce. However, for larger N, the trip-specific strategy slightly surpasses the long-run average strategy because detailed trip-level information becomes dominant in the optimization.

In practical scenarios, accurately predicting more than 10 future links is technically challenging. Therefore, the long-run

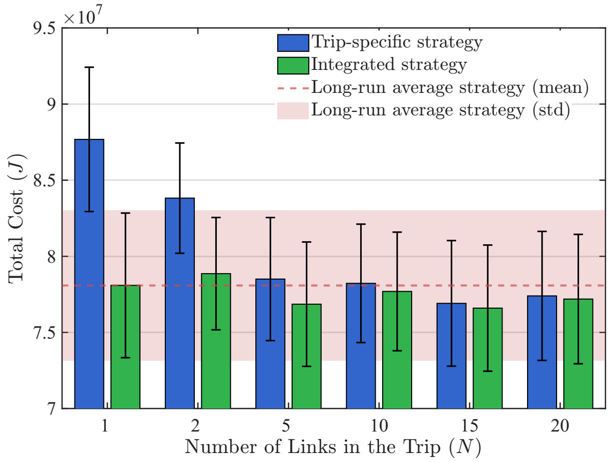

(a)

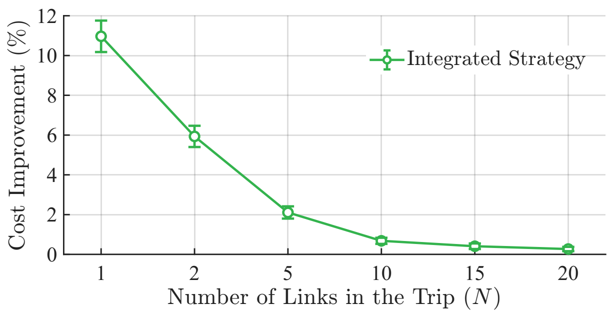

(b)

Fig. 2. Validation results in the virtual traffic network: (a) total cost averaged over 100 independent travels and (b) relative cost improvement of the integrated strategy over the trip-specific strategy as a function of the trip length N.

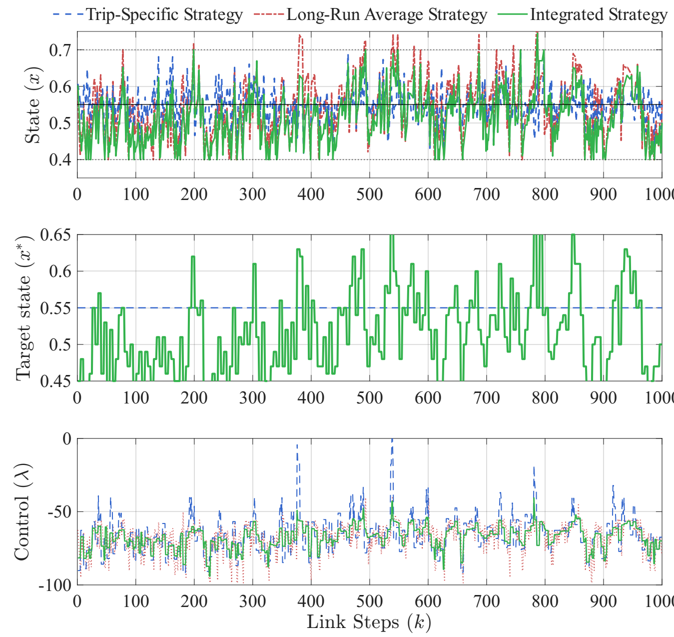

Fig. 3. State and control trajectories of the three strategies for a representative travel scenario with $N = 5$.

perspective provides meaningful robustness against prediction uncertainty, and the integrated strategy further improves performance by combining both perspectives.

Figure 3 presents the state and control trajectories for a representative travel scenario with $N = 5$. The trip-specific strategy tightly regulates the SOC around the fixed terminal target $x^* = 0.55$, resulting in clustered state trajectories. In contrast, both the long-run average strategy and the integrated strategy allow the SOC to utilize a broader admissible range.

Notably, the integrated strategy adaptively adjusts its terminal SOC target at each trip segment based on the long-run cost-to-go. This adaptive behavior enables the vehicle to temporarily "invest" or "reserve" energy depending on the statistical value of the subsequent network links.

This characteristic fundamentally distinguishes the proposed approach from conventional trip-specific strategies and even advanced SOC node planning methods, which typically rely on predetermined SOC references derived solely from trip-level information.

VI. Conclusion

This paper proposed a traffic network-aware energy management strategy for FCEVs by integrating finite-horizon trip-specific optimal control with long-run network-level optimality. Unlike conventional approaches that impose a fixed terminal SOC target, the proposed framework embeds the long-run cost-to-go derived from stochastic traffic transitions into the terminal cost of the finite-horizon problem. This formulation enables adaptive terminal SOC selection, thereby bridging local trip efficiency and global network optimality.

Simulation results in a virtual traffic network demonstrated that the proposed integrated strategy consistently outperforms both the trip-specific and long-run average strategies, with performance gains of up to 11% under short prediction horizons. The results highlight that incorporating network-level statistical awareness significantly improves robustness against limited prediction horizons and enhances long-term energy efficiency in stochastic driving environments. Future work will include validation on realistic traffic network data, robustness analysis under prediction errors in finite-horizon parameters, and extension of the framework to time-varying traffic networks where the long-run average cost must be updated adaptively.

References

[1] J. Guo, H. He, C. Jia, and S. Guo, "The energy management strategies for fuel cell electric vehicles: An overview and future directions," *World Electric Vehicle Journal*, vol. 16, no. 9, p. 542, 2025.

[2] X. Tian, Y. Cai, X. Sun, Z. Zhu, and Y. Xu, "A novel energy management strategy for plug-in hybrid electric buses based on model predictive control and estimation of distribution algorithm," *IEEE/ASME Transactions on Mechatronics*, vol. 27, no. 6, pp. 4350–4361, 2022.

[3] W. Zheng, M. Ma, E. Xu, and Q. Huang, "An energy management strategy for fuel-cell hybrid electric vehicles based on model predictive control with a variable time domain," *Energy*, vol. 312, p. 133544, 2024.

[4] S. D. Cairano, D. Bernardini, A. Bemporad, and I. V. Kolmanovsky, "Stochastic MPC with learning for driver-predictive vehicle control and its application to HEV energy management," *IEEE Transactions on Control Systems Technology*, vol. 22, no. 3, pp. 1018–1031, 2014.

[5] Y. Zhou, H. Li, A. Ravey, and M.-C. Péra, “An integrated predictive energy management for light-duty range-extended plug-in fuel cell electric vehicle,” *Journal of Power Sources*, vol. 451, Art. no. 227780.

[6] D. Chen, Y. Kim, and A. G. Stefanopoulou, “Predictive equivalent consumption minimization strategy with segmented traffic information,” *IEEE Transactions on Vehicular Technology*, vol. 69, no. 12, pp. 14377–14390, 2020.

[7] W. Tang, X. Jiao, and Y. Zhang, ”Hierarchical energy management control for connected hybrid electric vehicles in uncertain traffic scenarios,” *Energy*, vol. 315, p. 134291, 2025.

[8] S. Kofler, Z. P. Du, S. Jakubek, and C. Hametner, ”Predictive energy management strategy for fuel cell vehicles combining long-term and short-term forecasts,” *IEEE Transasctions on Vehicular Technology*, 2024.

[9] K. Choi and W. Kim, ”Real-time predictive energy management strategy for fuel cell-powered unmanned aerial vehicles based on the control-oriented battery model,” *IEEE Control Systems Letters*, vol. 7, pp. 943–948, 2022.

[10] K. Choi, G. Park, and D. Kum, ”An analytical approach to the predictive energy management of connected hevs: What information do we need to guarantee global optimality?,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 9, pp. 12749–12761, 2024.