Carbon-Aware Energy-Efficient Aerial Mobility Control via Large Language Models

Hyojun Ahn, Tae Hoon Lee, and Joongheon Kim
Department of Electrical and Computer Engineering, Korea University, Seoul, Republic of Korea
E-mails: {hyojun,taehoon822,joongheon}@korea.ac.kr

***Abstract*—The rapid expansion of Unmanned Aerial Vehicle (UAV) operations has raised critical concerns about their environmental impact in achieving global net-zero emissions targets. Traditional Aerial Mobility Control (AMC) systems prioritize operational metrics while overlooking carbon footprints from flight operations, battery charging, and computational overhead. This paper presents a comprehensive review of carbon-aware AMC systems that integrate sustainability considerations into UAV decision-making. We systematically analyze three primary paradigms: classical optimization-based methods, Reinforcement Learning (RL) techniques, and emerging Large Language Model (LLM)-guided systems. Through comparative analysis, we identify that LLM-driven approaches offer unprecedented flexibility through natural language reasoning, enabling sophisticated temporal carbon optimization and intuitive fleet coordination. However, significant challenges remain including computational efficiency, response time variability, and safety verification for stochastic language generation. Our analysis reveals key research gaps including the meta-challenge of LLM energy consumption potentially offsetting UAV carbon savings and the need for robust verification mechanisms. This review provides a roadmap for future research in sustainable aerial mobility systems that balance operational efficiency with environmental responsibility.**

***Index Terms*—Unmanned Aerial Vehicles (UAVs), Aerial Mobility Control (AMC), Large Language Models (LLMs), Carbon-aware control**

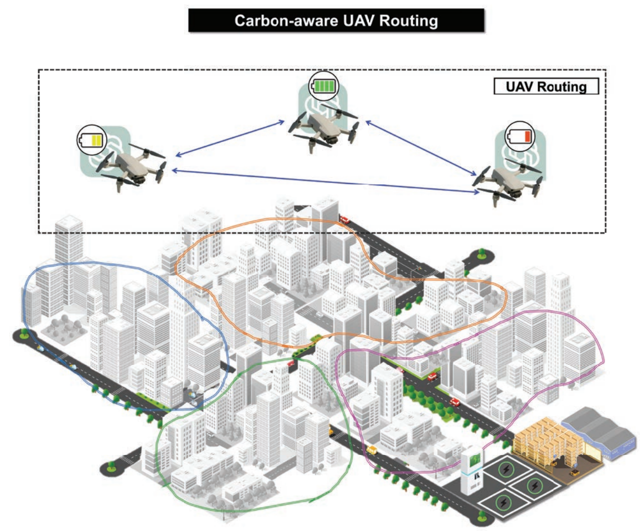

Fig. 1: Carbon-aware UAV routing system coordinating multiple UAVs in urban environments with LLM-guided decision making for sustainable aerial mobility control.

I. Introduction

The rapid proliferation of Unmanned Aerial Vehicles (UAVs) across diverse application domains has transformed modern logistics, surveillance, and urban service delivery [1], [2]. From package delivery systems operated by major e-commerce platforms to emergency response networks and smart city infrastructure, UAVs have become indispensable components of contemporary mobility ecosystems [3]. However, as UAV operations scale to support hundreds of thousands of daily flights, their cumulative environmental impact has become a pressing concern for achieving global net-zero emissions targets [4].

Traditional Aerial Mobility Control (AMC) systems prioritize operational metrics such as flight time minimization, collision avoidance, and service reliability, often overlooking the carbon footprint of UAV operations [5]. This oversight becomes particularly problematic when considering the complete lifecycle emissions of UAV systems, including energy consumption during flight operations, battery charging from carbon-intensive power grids, and the computational overhead of control algorithms running on energy-consuming infrastructure [6], [7]. The emergence of carbon-aware computing paradigms in cloud systems, data centers, and edge networks has demonstrated the feasibility of integrating sustainability metrics into real-time decision-making processes [8]. Building upon these developments, researchers have begun exploring how similar principles can be applied to aerial mobility systems, leading to the development of carbon-aware AMC frameworks that explicitly consider environmental impact alongside traditional performance objectives.

This paper provides the first comprehensive examination of carbon-aware AMC research, offering a systematic taxonomy of existing approaches, identifying key technical challenges, and outlining future research directions. Our contributions include: (1) a comprehensive classification of carbon-aware AMC methodologies, (2) analysis of carbon modeling approaches and their accuracy trade-offs, (3) evaluation of integration

This work was supported by the IITP(Institute of Information & Communications Technology Planning & Evaluation)-ITRC(Information Technology Research Center) grant funded by the Korea government(Ministry of Science and ICT) (IITP-2025-RS-2024-00436887) and this work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT)(RS-2025-00561377) *(Corresponding authors: In-Sop Cho and Joongheon Kim).*

challenges between sustainability and performance objectives, and (4) identification of open research problems and future opportunities.

II. Related Work

A. Classical Control Paradigms

Traditional AMC evolved from optimization-based approaches rooted in vehicle routing and path planning problems, primarily focusing on minimizing flight time, fuel consumption, and collision avoidance [9]. Early approaches adapted Traveling Salesman Problem (TSP) and Vehicle Routing Problem (VRP) formulations to three-dimensional airspace, incorporating constraints such as no-fly zones, altitude restrictions, and payload limitations. These methods typically employed Mixed-Integer Linear Programming (MILP), genetic algorithms, and particle swarm optimization to solve multi-UAV coordination problems [10].

The primary advantage of optimization-based approaches lies in their ability to provide globally optimal solutions under well-defined constraints and their mathematical rigor in handling complex operational requirements. Recent multi-objective optimization frameworks simultaneously consider multiple performance metrics, leading to Pareto-optimal solution sets that reveal trade-offs between competing objectives [11]. However, these methods suffer from significant computational scalability issues as fleet sizes increase, often requiring exponential time complexity that makes real-time application challenging.

Reinforcement Learning (RL) emerged to address dynamic environments and uncertain conditions that traditional optimization methods struggle to handle [12]. Early single-agent approaches used Deep Q-Networks (DQN) and policy gradient methods, evolving to Multi-Agent Reinforcement Learning (MARL) systems like Multi-Agent Deep Deterministic Policy Gradient (MADDPG) for fleet coordination [13]. While RL excels in adaptive learning and handling complex interactions, it faces sample efficiency challenges, safety concerns during exploration, and black-box interpretability issues for regulatory approval [14].

B. LLM-Based Control and Carbon-Aware Computing

Large Language Models (LLMs) recently entered AMC through natural language processing capabilities for high-level mission planning and decision-making [15]. These systems encode operational constraints, environmental conditions, and mission objectives into structured text prompts, generating human-readable flight plans that can be translated into executable control commands [16]. LLMs demonstrate particular strength in handling heterogeneous constraints and providing explainable decisions, but face computational overhead, response time inconsistency, and safety verification challenges [17]. Carbon-aware computing originated in cloud computing and data center management, where researchers developed techniques to minimize the carbon footprint of computational workloads by leveraging temporal and spatial variations in electricity grid carbon intensity. Key techniques include carbon intensity forecasting, workload migration strategies, and demand response that shifts energy consumption to periods of high renewable generation. This paradigm has expanded to edge computing, mobile systems, and IoT networks, inspiring adaptation to transportation networks including AMC systems.

III. LLM-Driven Carbon-Aware AMC

The integration of LLMs into carbon-aware AMC systems represents a fundamental paradigm shift toward natural language-driven decision making that can simultaneously optimize operational efficiency and environmental sustainability [18]. This approach leverages the inherent reasoning capabilities of LLMs to process diverse information sources and generate comprehensive flight strategies that explicitly consider carbon emissions alongside traditional performance metrics.

The core system architecture employs a three-layer design: environmental data aggregation collecting real-time carbon intensity, weather, and airspace information; natural language interface converting structured data into LLM-processable prompts; and execution layer translating generated plans into control commands [19]. Unlike traditional optimization requiring explicit mathematical formulations, LLMs process textual descriptions of operational requirements, regulatory constraints, and environmental considerations within a unified framework, enabling seamless adaptation to new constraints without system redesign. LLM-driven systems excel at temporal carbon optimization through natural language interpretation of grid carbon intensity forecasts. The system understands linguistic concepts such as "low carbon periods during midday solar generation" and "high carbon intensity during evening peak demand," enabling sophisticated charging schedule optimization that balances immediate operational needs against longer-term carbon objectives [20]. This linguistic reasoning provides transparency often lacking in black-box optimization algorithms while enabling direct incorporation of regulatory and policy constraints through textual descriptions. The temporal optimization process incorporates reasoning about uncertainty and risk management, enabling robust decisions under incomplete information about future carbon intensity and operational requirements. Fleet-level coordination through LLMs introduces natural language-based multi-agent communication, facilitating intuitive coordination strategies that consider collective carbon footprints while ensuring individual mission success. The system reasons about resource allocation, charging assignments, and route coordination using natural language logic that mirrors human decision-making processes [21]. This capability enables incorporation of high-level strategic objectives that may be difficult to encode in traditional optimization formulations, such as "prioritize carbon reduction during peak demand periods while maintaining emergency response capability [22]."

However, significant challenges remain in LLM-based carbon-aware AMC. Response time variability due to fluctuating LLM processing complexity threatens real-time performance guarantees essential for safety-critical operations. The stochastic nature of language generation raises reliability concerns, as identical inputs may produce different outputs, potentially compromising operational consistency [23]. Safety verification for LLM-generated plans presents unprecedented challenges, as traditional formal verification methods are not directly applicable to natural language reasoning processes. Additionally, the computational overhead of LLM processing creates a meta-challenge where the energy consumption of language model inference may offset carbon savings from optimized UAV operations, motivating research into energy-efficient architectures and hybrid systems that balance reasoning capability with computational efficiency.

IV. Conclusion

This paper has provided a comprehensive examination of carbon-aware AMC systems, tracing the evolution from traditional performance-oriented approaches to sustainability-integrated frameworks. Through systematic analysis of optimization-based methods, RL techniques, and emerging LLM-guided systems, we have identified the unique advantages and limitations of each paradigm in addressing the dual challenges of operational efficiency and environmental sustainability. The emergence of LLM-driven carbon-aware AMC represents a significant paradigm shift that offers unprecedented flexibility in handling diverse constraints and objectives through natural language reasoning. These systems demonstrate particular strength in integrating heterogeneous data sources, providing explainable decision-making processes, and adapting to evolving regulatory requirements without fundamental system redesign. However, challenges remain in ensuring computational efficiency, maintaining real-time performance guarantees, and providing robust safety verification for stochastic language generation outputs.

The field of carbon-aware AMC is at a critical juncture where technological capabilities are rapidly advancing while regulatory frameworks and industry standards are still evolving. The successful deployment of these systems will require continued collaboration between researchers, industry practitioners, and regulatory bodies to establish appropriate safety standards, performance benchmarks, and environmental accounting frameworks. As UAV operations continue to scale globally and environmental regulations become more stringent, the development of effective carbon-aware AMC systems will be essential for ensuring that the benefits of aerial mobility can be realized without compromising global climate objectives.

References

[1] Y. Y. Chan, K. K. Ng, C. Lee, L.-T. Hsu, and K. L. Keung, “Wind dynamic and energy-efficiency path planning for unmanned aerial vehicles in the lower-level airspace and urban air mobility context,” *Sustainable Energy Technologies and Assessments*, vol. 57, p. 103202, Apr. 2023.

[2] D. S. Lakew, A.-T. Tran, N.-N. Dao, and S. Cho, “Intelligent self-optimization for task offloading in leo-mec-assisted energy-harvesting-uav systems,” *IEEE Transactions on Network Science and Engineering*, vol. 11, no. 6, pp. 5135–5148, Nov. 2024.

[3] H. Ahn, S. Oh, G. S. Kim, S. Jung, S. Park, and J. Kim, “Hallucination-aware generative pretrained transformer for cooperative aerial mobility control,” in *Proc. IEEE Global Communications Conference (GLOBECOM)*, Taipei, Taiwan, Dec. 2025.

[4] M. S. Aslanpour, A. N. Toosi, M. A. Cheema, and M. B. Chhetri, “FAASSHOUSE: sustainable serverless edge computing through energy-aware resource scheduling,” *IEEE Transactions on Services Computing*, vol. 17, no. 4, pp. 1533–1547, Jul. 2024.

[5] S. Park, C. Park, and J. Kim, “Learning-based cooperative mobility control for autonomous drone-delivery,” *IEEE Transactions on Vehicular Technology*, vol. 73, no. 4, pp. 4870–4885, Apr. 2024.

[6] S. Jung, C. Park, M. Levorato, J.-H. Kim, and J. Kim, “Two-stage self-adaptive task outsourcing decision making for edge-assisted multi-uav networks,” *IEEE Transactions on Vehicular Technology*, vol. 72, no. 11, pp. 14 889–14 905, Nov. 2023.

[7] S. Jung, H. Baek, and J. Kim, “Neural myerson auction for truthful and distributed mobile charging in uav-assisted digital-twin networks,” *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 11, pp. 7326–7338, Nov. 2024.

[8] Y. Jiang, R. B. Roy, B. Li, and D. Tiwari, “ECOLIFE: carbon-aware serverless function scheduling for sustainable computing,” in *Proc. International Conference for High Performance Computing, Networking, Storage and Analysis*, Nov. 2024, pp. 1–15.

[9] J. Dou, H. Zhang, Y. Luo, and G. Sun, “Scheduling drone and mobile charger via hybrid-action deep reinforcement learning,” *IEEE Transactions on Mobile Computing*, vol. 24, no. 8, pp. 7788–7804, Aug. 2025.

[10] B. Du, J. Chen, D. Sun, S. G. Manyam, and D. W. Casbeer, “UAV trajectory planning with probabilistic geo-fence via iterative chance-constrained optimization,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 6, pp. 5859–5870, Jun. 2022.

[11] J. Xie and J. Chen, “Multiregional coverage path planning for multiple energy constrained UAVs,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 10, pp. 17 366–17 381, Oct. 2022.

[12] M. Fan, Y. Wu, T. Liao, Z. Cao, H. Guo, G. Sartoretti, and G. Wu, “Deep reinforcement learning for uav routing in the presence of multiple charging stations,” *IEEE Transactions on Vehicular Technology*, vol. 72, no. 5, pp. 5732–5746, May 2023.

[13] X. Dai, Z. Xiao, H. Jiang, and J. C. S. Lui, “UAV-assisted task offloading in vehicular edge computing networks,” *IEEE Transactions on Mobile Computing*, vol. 23, no. 4, pp. 2520–2534, Apr. 2024.

[14] J. Goldverg, H. Jamil, E. Rodrigues, and T. Kosar, “Toward carbon-aware data transfers,” *IEEE Internet Computing*, vol. 29, no. 2, pp. 19–26, Mar. 2025.

[15] Q. Zhou, J. Wu, M. Zhu, Y. Zhou, F. Xiao, and Y. Zhang, “LLM-QL: A LLM-enhanced Q-learning approach for scheduling multiple parallel drones,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 37, no. 9, pp. 5393–5406, Sep. 2025.

[16] X. Cao, G. Nan, H. Guo, H. Mu, L. Wang, Y. Lin, Q. Zhou, J. Li, B. Qin, Q. Cui, X. Tao, H. Fang, H. Du, and T. Q. S. Quek, “Exploring LLM-based multi-agent situation awareness for zero-trust space-air-ground integrated network,” *IEEE Journal on Selected Areas in Communications*, vol. 43, no. 6, pp. 2230–2247, Jun. 2025.

[17] T. Yang, P. Feng, Q. Guo, J. Zhang, X. Zhang, J. Ning, X. Wang, and Z. Mao, “AutoHMA-LLM: efficient task coordination and execution in heterogeneous multi-agent systems using hybrid large language models,” *IEEE Transactions on Cognitive Communications and Networking*, vol. 11, no. 2, pp. 987–998, Apr. 2025.

[18] Z. Lan, L. Liu, B. Fan, Y. Lv, Y. Ren, and Z. Cui, “Traj-LLM: A new exploration for empowering trajectory prediction with pre-trained large language models,” *IEEE Transactions on Intelligent Vehicles*, vol. 10, no. 2, pp. 794–807, Feb. 2025.

[19] H. Liu, Y. Zhu, K. Kato, A. Tsukahara, I. Kondo, T. Aoyama, and Y. Hasegawa, “Enhancing the LLM-based robot manipulation through human-robot collaboration,” *IEEE Robotics and Automation Letters*, vol. 9, no. 8, pp. 6904–6911, Aug. 2024.

[20] M. Kahng, I. Tenney, M. Pushkarna, M. X. Liu, J. Wexler, E. Reif, K. Kallarackal, M. Chang, M. Terry, and L. Dixon, “LLM Comparator: interactive analysis of side-by-side evaluation of large language models,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 31, no. 1, pp. 503–513, Jan. 2025.

[21] T. Sheokand, G. Jain, A. Bahga, and V. K. Madisetti, “Enhancing LLM reasoning capabilities through brokered multi-expert reflection,” *IEEE Access*, vol. 13, pp. 67 993–68 019, Apr. 2025.

[22] L. Wang, H. Zhao, D. Wang, F. Dong, T. Feng, and R. Xiong, “Low-carbon economic dispatch of integrated energy system considering expanding carbon emission flows,” *IEEE Access*, vol. 12, pp. 104 755–104 769, Jul. 2024.

[23] Z. Yao, Z. Tang, W. Yang, and W. Jia, “Enhancing LLM QoS through cloud-edge collaboration: A diffusion-based multi-agent reinforcement learning approach,” *IEEE Transactions on Services Computing*, vol. 18, no. 3, pp. 1412–1427, May 2025.