26th International Congress on Modelling and Simulation, Adelaide, SA, Australia, 30 Nov–4 Dec 2025
mssanz.org.au/modsim2025

Journey costing based scheduling of uncrewed aerial vehicles with energy constraints and wind effects

Gregory D Sherman [a], William LH Bai [b], Benjamin C Campbell [a],

[a] *Defence Science and Technology, Edinburgh, South Australia,* [b] *University of Adelaide (Vacation scholar),*
Email: gregory.sherman@defence.gov.au

Abstract: .

Efficient scheduling of uncrewed aerial and ground vehicles (UAVs/UGVs) is a critical challenge in autonomous operations, particularly under constraints related to energy consumption, return-to-base requirements, and environmental interference. This paper presents a modular scheduler–transitioner system designed to optimise multi-agent operations with a cost matrix that integrates spatial positioning, impedance effects (e.g., wind), and agent-specific parameters such as velocity, battery level, and energy loss rates.

Schedules are generated separately for each agent type using a modified Travelling Salesman Problem approach with Two-Opt heuristics. All agents operate in a round-robin fashion, with compulsory returns to a single base where recharging is handled by a rapid battery swap system. The impedance model simulates environmental effects, such as wind fields, via a spatially varying matrix that directly influences cost calculations during both scheduling and transitions.

The transitioner module matches agents between sequential scheduling blocks and applies a penalty-based function to determine whether repositioning occurs directly or via the base. A “breath” parameter is introduced to allow controlled slack in transition timing, enhancing robustness under variable conditions. The use of this terminology is inspired by biological systems and the ability to undertake additional effort in such models.

The specific case modelled in this work involves known task positions that represent *loitering zones*, where agents undertake monitoring or surveillance activities. The scheduler allocates these positions into cyclic task blocks and shuffles available agents across them to satisfy energy and battery constraints. Vehicles are required to return to base for energy replenishment and may be repositioned between blocks to maintain persistent task coverage. The system is designed to respond flexibly to real-time or planned changes, including environmental disturbances such as *wind, rain, or boggy terrain*, which alter journey times and feasibility. Such responsiveness is vital in systems focused on *trusted autonomy*, especially in applications like *military surveillance (see Figure 1)*, *border patrol*, or *logistics supply chains* in remote or contested terrain. With recent advances in environmental modelling (e.g., fast-updating wind fields) and predictive methods in AI/ML, journey time estimation can increasingly be refined using either simulation or real-world sensor inputs. Our system is structured to leverage such data, enabling more robust and energy-aware scheduling processes.

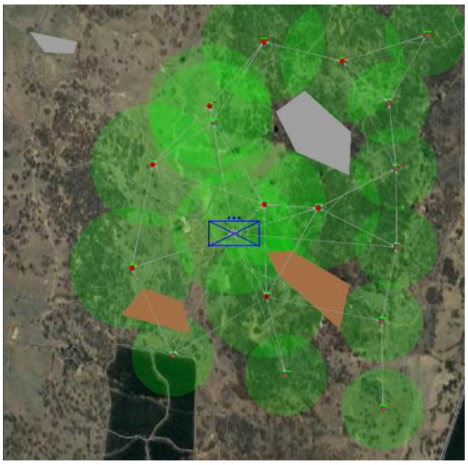

Figure 1: An image from a simulation for UAV/UGV dispersal in a setting with overlapping sensing regions for surveillance

This work contributes a flexible and extensible framework for persistent, autonomous scheduling, with the capacity to incorporate complex environmental feedback and real-world operational requirements.

Keywords: *Trusted autonomy, scheduling, environmental simulation, battery-constrained routing.*

1. INTRODUCTION

The scheduling of uncrewed aerial and ground vehicles (UAVs/UGVs) for persistent task coverage is a critical capability in autonomous systems deployed for monitoring, surveillance, and logistical support. In this paper, we focus on UAV operations, using accumulative wind resistance as an exemplar environmental factor influencing journey time and feasibility. This case illustrates how environmental dynamics can be integrated into the scheduling process to support more robust and energy-aware planning across agent types. We rely on our previously published work (Elfeky et al., 2024), to provide the UAV's loitering surveillance positions.

Effective task allocation must consider operational constraints including energy limitations, recharging logistics, and environmental variability. We consider a base protection task where multiple UAVs (30 in the first experiment), with homogeneous, sensing, communication, movement and energy consumption characteristics, are being utilised as loitering sensors to surveil an area in order to give early warning of the approach of an adversary. Co-ordinating the movements of many UAVs, to ensure battery recharging occurs in a timely fashion, without creating gaps in surveillance or such that drones have insufficient charge to return to base, becomes especially difficult in fluctuating environmental conditions that effect flight times and energy consumption - often too difficult for a single human operator to be effective.

We present a modular scheduling framework that assigns agents to known loitering task locations and dynamically transitions them between blocks while respecting battery constraints and environmental interference. The scheduler incorporates a cost matrix built from distances and impedance values derived from mock wind simulations. A transitioner module handles inter-block repositioning with a penalty function and "breath" buffer for timing flexibility, usually represented as proportion of battery charge set aside for this purpose.

2. BACKGROUND

The core of our approach builds on the Travelling Salesman Problem (TSP), with routing initially formulated via the classic 2-opt heuristic (Croes, 1958) and adapted through block-based scheduling akin to split vehicle routing problems (Dorling et al., 2017). Our focus on trusted autonomous systems aligns with the growing literature on UAV/UGV persistence under energy constraints and rotation-based scheduling. Cassandras et al. (2013) offer foundational insights into loitering UAVs and round-robin rotation logic. Energy-aware planning frameworks, such as that proposed by Mathew et al. (2015), directly inform our return-to-base constraint and scheduling cycle. Environmental influence is modelled through wind impedance, extending ideas from Zhou et al. (2020) and predictive journey estimation from Wang et al. (2021). Our agent transitioner draws on the classical Hungarian assignment method (Kuhn, 1955) and the Santa Claus problem's fairness formulation (Bansal & Sviridenko, 2006). The agent-task matching during transitions is modelled as a transportation problem, solved using minimum-cost bipartite assignment methods (Kuhn, 1955; Burkard et al., 2009), rooted in early formulations of optimal transport (Kantorovich, 1942). Multi-objective reassignment and penalty-based feasibility logic are grounded in the work of Liu et al. (2019) which we redesign and ignore the communication constraints they introduce in their paper.

We aim to build on these approaches by incorporating tuneable redundancy into our energy aware planning framework. By adding additional energy requirements to the cost estimate of each task we build in resilience to changes in environmental conditions and unplanned repositioning. We refer to this additional energy cost as "Breath". Modern machine learning techniques are enabling faster, more accurate and more granular wind models and weather forecasting (Allen et al., 2025). Data-driven techniques enable better estimations for energy consumption and journey times (Góra et al., 2022). In this work we seek to utilise these techniques to balance resilience from expected fluctuations in wind conditions with the need to minimize the number of agents so as to avoid unnecessary re-planning and repositioning (Elfeky et al., 2024) by selecting an optimal breath setting given the predicted wind conditions.

3. MODELLING

There are three main components to the model: the scheduler, the location transitioner, and the wind model. Each simulation instance includes a predefined set of elements: a list of desired loitering locations, a group of agents (modelled here as UAVs), a central base for energy replenishment and spare storage, and a simplified 2D terrain environment in which all operations occur.

3.1 Scheduler

The scheduling component operates on the predefined groups of agents, each representing a distinct type with homogeneous movement and energy characteristics. All agents within a type share the same velocity, energy consumption rate, and recharge behaviour. They are assumed to move across flat, obstacle-free terrain via straight-line paths and return to a single central base for recharge using a fast battery-swap mechanism, with any delay absorbed into the modelled costs.

For each agent group, the scheduler partitions target positions into feasible loops (blocks), ensuring that each position within a loop can be assigned an agent and that each assigned agent can move to their target position and sequentially return to base without any agent exceeding a maximum operational cost threshold. This cost includes:

- Motion-related energy loss, scaled by a movement coefficient (α_1),
- Stationary energy losses, scaled by a residual loss coefficient (α_2),
- Recharge requirements and delay penalties,
- A spatial cost buffer ("breath"), ensuring loops remain well within safe maximum bounds.

We parameterise breath as a reserve fraction $b \in [0.1]$ of an agent's total loop energy/time budget E and allow amount bE for each loop is held in reserve for contingencies; only the non-loitering motion term α_1 is wind-modulated. Stationary losses during loitering or at base (scaled by α_2) are independent of wind in this model.

The core cost calculations are derived from energy-distance relationships, with all travel assumed to occur at a fixed airspeed and energy rate. An impedance matrix modifies these travel costs by embedding spatial difficulty, notably wind-induced travel resistance across agent movement paths.

Intra-loop ordering is optimised using a local search strategy inspired by the Travelling Salesman Problem (TSP), further reducing motion costs. The breath parameter enables sensitivity testing of how conservative loop allocations impact agent efficiency.

3.2 Location Transitioner

Agent transitions are triggered when a UAV must relocate from its current position to a new task location. This may occur due to task completion, reassignment, environmental changes, or system-level scheduling updates—not solely at the end of a predefined schedule block. The transitioner handles this repositioning by computing a cost-minimised mapping between current agent and required location positions.

Transition costs incorporate energy used for direct travel, detour via base, recharge delay, and penalty-weighted feasibility. Assignments are made using bipartite graph matching methods: either the Hungarian algorithm (for total cost minimisation) or a bottleneck strategy (minimising worst-case cost; Bansal & Sviridenko, 2006). All costs are modulated by an impedance matrix (e.g., wind), which reflects environmental resistance. Motion is assumed to be linear, and recharging is considered instantaneous via battery swap.

Each possible transition (agent position → new task) is evaluated as either a direct trip using residual battery charge or a two-leg journey via the base for recharging. If the agent's remaining energy is insufficient or is close to the new schedule's requirement for its new assigned location for direct travel, a penalty function inflates the transition cost to reflect energy risk. This cost in the assignment process may divert the agent, if necessary, to base en route or reassign it to another location. In this implementation, we use a capped inverse-square weighted sum, where the penalty increases sharply as the residual charge approaches the required threshold (set by new schedule). This shaping ensures marginal transitions are strongly disincentivised. The form of the penalty function affects assignment quality and feasibility. A portion of the reserve bE may be temporarily borrowed during a transition; after recharge at base, the full reserve is restored

3.3 Wind Model for resistance and journey costings (impedance relation)

To simulate environmental effects on UAV deployment, wind was modelled using a grid-based approach in which wind speed and direction were assumed constant within each mapped grid cell and over time. These modelled wind fields consisted of a 10-by-10 grid of wind vectors that were generated about some mean wind speed and added variance. The UAV's instantaneous groundspeed vector was defined as the sum of its constant-magnitude airspeed vector and the local wind vector. We assume UAV airspeed is regulated to a fixed magnitude (typical for energy modelling), so ground speed varies with wind via the vector sum, and all motion occurs in a two-dimensional plane. To keep scheduling tractable, wind is treated as time-invariant; a time-varying wind can be handled via horizon re-planning or periodic impedance updates

To calculate the travel time between any two points, let C be the set of grids which some portion of the UAV's path must cross to travel between those two points. The total travel time t_{total} is therefore given by the sum of δt_i for $i \in C$ constrained to align with the straight-line path between two points, with the direction of the airspeed vector adjusted accordingly. Here, δt_i is the time taken for the UAV to traverse the portion of the path that passes through the ith grid in C. From this, and assuming a one way journey between the two points, δt_i could be calculated as follows:

$$[\delta t_i]_{one\,way} = \frac{\delta s_i}{\hat{\boldsymbol{e}}_{\boldsymbol{GS}} \cdot \overrightarrow{\boldsymbol{v_w}} + \sqrt{(\hat{\boldsymbol{e}}_{\boldsymbol{GS}} \cdot \overrightarrow{\boldsymbol{v_w}})^2 - |\overrightarrow{\boldsymbol{v_w}}|^2 + v_{AS}^2}},$$

Whereby δs_i is the length of the portion of the path passing through the ith grid point, $\hat{e}_{GS} = \overrightarrow{v_{GS}} / |\overrightarrow{v_{GS}}|$, for $\overrightarrow{v_{GS}}$ the ground speed velocity vector, $\overrightarrow{v_w}$ the wind velocity vector and v_{AS} the magnitude of the UAV's airspeed. Return journey times were derived assuming symmetric paths and constant wind, yielding:

$$[\delta t]_{return} = \frac{2\delta s}{v_{AS}} \frac{v_{AS}\sqrt{(\hat{\boldsymbol{e}}_{\boldsymbol{GS}} \cdot \overrightarrow{\boldsymbol{v_w}})^2 - |\overrightarrow{\boldsymbol{v_w}}|^2 + v_{AS}^2}}{v_{AS}^2 - |\overrightarrow{\boldsymbol{v_w}}|^2}.$$

This formulation demonstrates that any nonzero wind increases the round-trip journey time: tailwind benefits in one direction are outweighed by headwind penalties on return. These directional wind effects are incorporated into journey cost estimates via an impedance matrix, which adjusts both scheduling and transition costs accordingly.

The impedance matrix encodes the spatial and directional cost effects of wind by adjusting the nominal Euclidean distance between node pairs. For each transition between a source and a target location, the effective cost is scaled based on the wind vector projected along the direction of travel. Specifically, the impedance multiplier is derived from the computed journey time under wind influence relative to the time required in still-air conditions. This multiplier is applied element-wise to the direct distance matrix, resulting in a wind-adjusted cost matrix used. We calculate both one way and two-way travel impedance matrices for scheduling and transitioning respectively.

3.4 Operational considerations and physical assumptions.

To maintain tractability, the model abstracts away certain physical complexities such as vehicle mass, acceleration, turning dynamics, and nonlinear energy effects. Travel costs are based on Euclidean distances modified by an impedance matrix, which captures environmental resistance—such as wind—without explicitly modelling vehicle physics. These effects can be incorporated into future impedance formulations if required.

Operationally, it is assumed that each loop has sufficient agents stationed at the base for immediate deployment, with spare capacity included to support resilience and adaptation. This over-provisioning enables rescheduling under changing environmental conditions without requiring repositioning, provided that loitering assignments remain feasible. However, if updated journey costs violate operational constraints, a full transition is triggered, reallocating agents to new positions. The number of required spares is implicitly linked to the breath parameter: greater breath values introduce slack, reducing the need for transitions but increasing agent count, while tighter allocations raise efficiency at the expense of flexibility.

4. ILLUSTRATIVE RUNS AND RESULTS

We now perform experiments under varying simulated wind fields, with agents stationed in loitering positions plus spares located at the base or on those in a replacement route to and from those places. Each UAV agent is set simulated charge values for transition purposes appropriate to its location.

4.1 Illustrative runs

In the first experiment, we illustrate how the journey-costing scheduler and location transitioner is applied under two environmental conditions. Specifically, we simulate two wind field scenarios—one with negligible wind and another with high wind.

In our first case, there are 20 task locations and a nominal fleet of 30 UAVs available for scheduling, replenishment, and sparing. If more than 30 UAVs are needed to maintain a viable schedule, alternative strategies are explored. Figure 2 shows the breath versus sparing profile of the scheduler under different impedance conditions whereby wind has increased in the second (right) plot, this causes left shift of the curve and the need for more spares as we increase our breath. The elbow point in red is where the marginal increase in spares per 0.05 breath step exceeds 1 (i.e., diminishing-returns onset), and is used to indicate a breath level to choose, Under negligible wind this occurs at b ≈ 0.2; noting the stronger wind shifts the elbow left. We therefore use b ≈ 0.2 the elbow under negligible wind in Figure. 2 for our further experimentation.

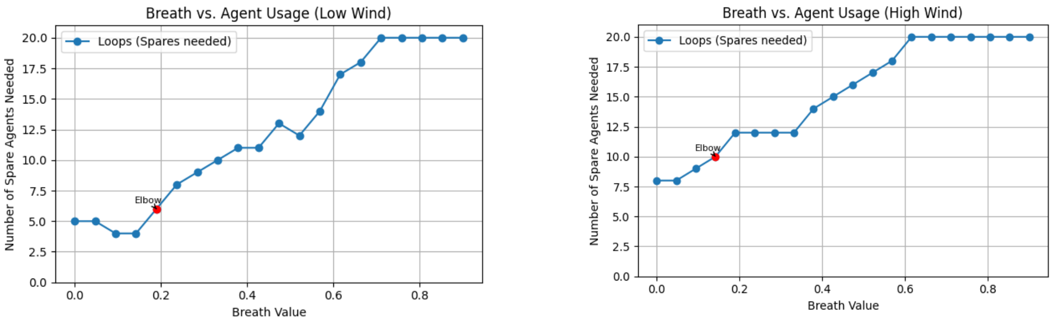

Figure 2: Breath profiles under different environmental conditions (increasing impedance changes)

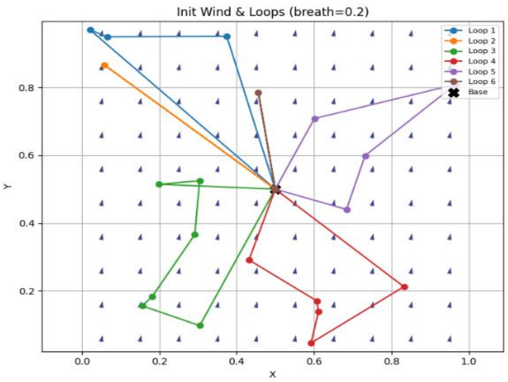

(a) negligible wind on initial 20 locations (6 loops)

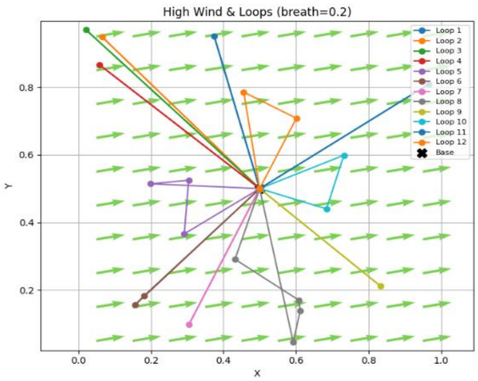

(b) high wind on initial 20 locations (11 loops)

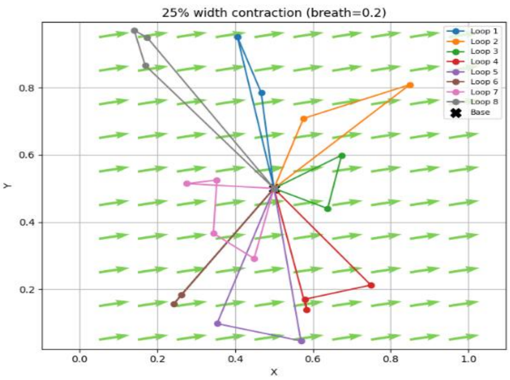

(c) high wind horizontal contraction by 25% (8 loops)

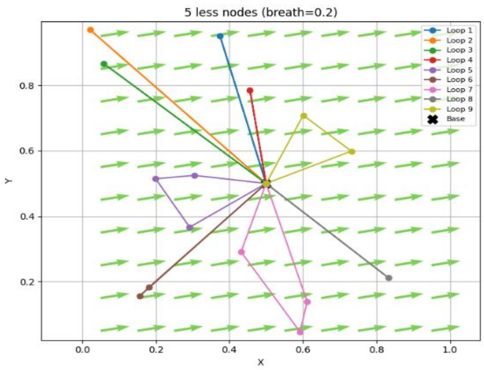

(d) high wind with 5 locations removed (9 loops)

Figure 3: Scheduling loops over wind fields from left to right no wind, rescheduling (more loops) and location transitioning examples

Figure 3 presents the resulting schedules as loop plots under the different wind scenarios, along with two adjustment strategies for the high-wind case. The subfigures represent: (a) Negligible wind: the scheduler generates loops from an initial state with all UAVs located at the base. This scenario requires 6 spare UAVs, (b) High wind: using the same locations as in (a), the strong wind results in an infeasible schedule unless 11 additional UAVs (beyond the available 30) are introduced, (c) High wind with contracted locations: the location transitioner modifies the task layout by contracting loitering positions by 25% horizontally (closely in line with the "high" wind direction). This spatial adjustment reduces the required spare UAVs to 8, (d) High wind with task reduction: 5 locations are removed entirely, which brings the spare requirement down to 9.

4.2 Findings from experimentation of wind conditions

To evaluate the impact of wind on scheduling, we conducted further simulation trials across a range of wind ratios and breath values. For each configuration, 30 trials were run—comprising six randomly generated sensor layouts and five wind field conditions—using fixed seeds for reproducibility. Each trial involved 100 randomly placed locations within a radius of 0.5 units from the base, with UAVs operating at an airspeed of 1.0 unit and a battery capacity allowing 10 units of travel. Wind fields were directed consistently from left to right.

Figure 4 shows the mean number of loops generated per configuration, each loop requiring a spare for replenishment. If a configuration was infeasible (i.e., one or more locations could not be reached), data collection was halted at that wind ratio for the corresponding breath value. This truncation occurs earliest for higher breath values, since a larger reserve bE leaves less time available for planned loops, making schedules infeasible sooner as wind increases.

The results demonstrate that low breath values (0.0–0.2) lead to minimal loop counts under calm conditions and degrade as wind increases. Conversely, higher breath values (0.6–0.8) may allow more capacity in reserve to abrupt changes but result in fragmented schedules and higher spare agent costs even when wind is mild. As the wind ratio approaches 1.0, all configurations show asymptotic growth in loop count, highlighting the theoretical infeasibility of long-range scheduling under extreme wind conditions.

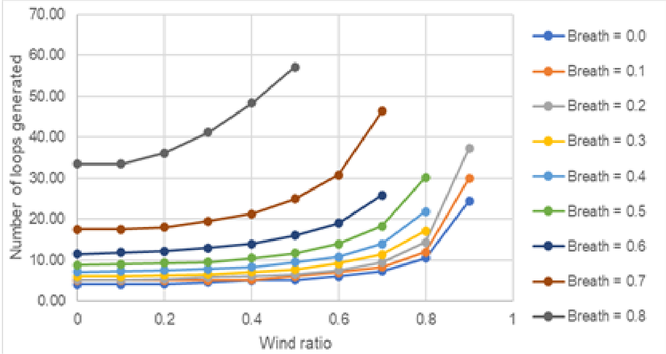

Figure 4: Varying breath values, wind ratios (to max air speed) and loops required for scheduler to operate

4.3 Other Findings from Wind Analysis and Location Transitioning

Wind begins to affect scheduling once it exceeds ~20% of airspeed; block counts rise > 5% beyond this threshold and > 10% once wind surpasses ~28%. Below these levels, wind has little influence. Variability in wind magnitude increases loop counts, whereas direction variability has little effect. With six equal-mean-distance layouts (circle, triangle, square, imbalanced circle, two ellipses), shape matters little until wind > 50% of airspeed. Return efficiency improves by minimising travel parallel to the wind rather than spreading perpendicular to it. A value $b \approx 0.2$ yields compact schedules in calm–moderate wind. Higher breath adds a transition buffer and, in several high-wind trials (e.g., 0.1→0.3), restored feasibility at the cost of more spares; however, in loop-limited regimes the reduced plan energy $(1 - b)E$ means larger b can precipitate earlier infeasibility. Thus, breath is a tunable reserve. Under high wind, spatial contraction preserved persistent loitering better than removing locations.

For transition assignments (as shown in Figure 5 using the Hungarian algorithm "Transitioner"), the importance of residual charge versus travel distance becomes significant, primarily when energy costs are high or breath allocations are constrained. In most typical configurations, distance dominates, especially when higher breath values allow for wider energy margins. Even when spare agents are available, the breath parameter remains essential to prevent UAV agents losing all charge and becoming unable to return for recharge under dynamic conditions. Nonetheless, when conditions are favourable, optimising breath usage can extend the reach and density of loitering positions.

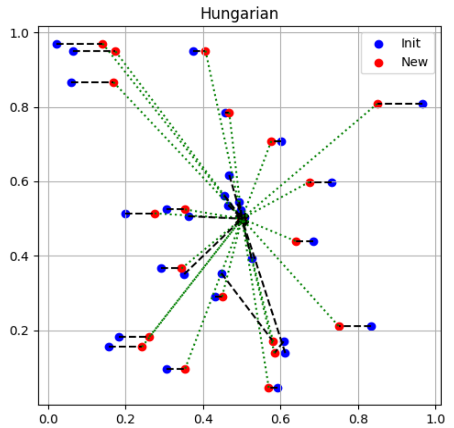

Figure 5: The movement of UAV agents with varying charge levels for a horizontal contraction of 25% as specified by location Transtioner using the Hungarian alogrithm for assignment

5. DISCUSSION AND RECOMMENDATIONS

The scheduling and transition framework presented in this work demonstrates a flexible, modular approach for managing UAV operations under varying environmental and energy constraints. Its plug-and-play design supports wind-aware impedance modelling, energy-aware scheduling, and cost-sensitive agent reassignment. While this implementation utilised a capped inverse-square penalty function, alternative forms such as exponential or piecewise-linear functions may offer better alignment with mission-specific tolerances. The breath parameter plays a role in bridging scheduling and transition phases, acting as a dynamic buffer that enhances resilience to environmental fluctuations. Beyond persistent loitering, this framework also opens the door to investigating non-persistent or periodic task satisfaction, where agents service locations in a staggered or duty-cycled manner based on revisit intervals, rather than constant presence.

This system can be extended in several directions: supporting heterogeneous agent types through customised schedulers, refining penalty functions to suit different energy profiles, and introducing look ahead capabilities based on microclimate forecasts. Further improvements include sequencing transitions using multi-objective criteria or reinforcement learning, and leveraging surrogate models—such as neural networks trained on energy or transition cost data—to accelerate runtime evaluation. Additionally, the model could be adapted for UGV scheduling, incorporating track-based or graph-constrained movement, where travel is restricted to roads or predefined paths, and affected by terrain-based impedance and weather effects such as rain, dust, or terrain degradation. Future work may also explore terrain-aware planning, richer vehicle dynamics (e.g., acceleration and speed variability), and layered multi-constraint impedance models to further enhance operational realism and robustness.

6. CONCLUSION

This work presents a modular scheduling and transition framework for UAV operations under energy and environmental constraints, using wind-aware journey costing and flexible agent reassignment. Through simulated experiments, we demonstrate that modest breath values support efficient and resilient schedules, while high wind conditions necessitate spatial adaptation or task reduction. The transitioner effectively reallocates agents using cost-sensitive matching, with breath acting as a tuneable reserve. Our findings underscore the value of integrating environmental modelling into operational planning and highlight the system's extensibility to UGVs, heterogeneous fleets, and terrain-aware routing in real-world autonomous deployments.

REFERENCES

Allen, A., Markou, S., Tebbutt, W. *et al.*, 2025. End-to-end data-driven weather prediction. *Nature*, 641, pp.1172–1179.

Bansal, N. and Sviridenko, M., 2006. The Santa Claus problem. *In Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pp.31–40.

Cassandras, C.G., Lin, X., Ding, Y. and Pajic, M., 2013. Optimal persistent monitoring with loitering drones. *In Proceedings of the 52nd IEEE Conference on Decision and Control*, pp.7181–7186.

Croes, G.A., 1958. A method for solving traveling-salesman problems. *Operations Research*, 6(6), pp.791–812.

Dorling, K., Heinrichs, J., Messier, G.G. and Magierowski, S., 2017. Vehicle routing problems for drone delivery. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 47(1), pp.70–85.

Elfeky, E., Sherman, G., Elsayed, S., Shovon, M.H.I., Lodge, R., Campbell, B., Essam, D. and Sarker, R., 2024. *Differential Evolution Algorithm for Battlefield Surveillance Sensor Placement*. In: *Proceedings of the 2024 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, pp.1–9

Góra, K., Smyczyński, P., Kujawiński, M. and Granosik, G., 2022. Machine learning in creating energy consumption model for UAV. *Energies*, 15(18), p.6810.

Kantorovich, L.V., 1942. On the translocation of masses. Dokl. Akad. Nauk USSR, 37(7-8), pp.199–201.

Koopmans, T.C. and Beckmann, M., 1957. Assignment problems and the location of economic activities. *Econometrica*, 25(1), pp.53–76.

Kuhn, H.W., 1955. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1–2), pp.83–97.

Liu, C., Xu, X. and Wang, Y., 2019. Multi-objective task assignment and path planning for heterogeneous UAVs with sensing and communication constraints. *Sensors*, 19(3), p.582.

Mathew, N., Smith, S.L. and Waslander, S.L., 2015. Planning paths for package delivery using multiple drones with energy constraints. *IEEE Transactions on Automation Science and Engineering*, 12(4), pp.1296–1307..

Wang, H., Chen, J., Xu, S. and Du, Q., 2021. Environment-aware UAV path planning using wind field prediction. *IEEE Transactions on Vehicular Technology*, 70(3), pp.2445–2456.

Zhou, B., Ma, Y., Li, W. and Han, Z., 2020. A wind-aware path planning method for UAVs in complex environments. *Aerospace Science and Technology*, 99, p.105771.

