

Multi UAVs Preflight Planning in a Shared and Dynamic Airspace

Amath Sow
Linköping University
Linköping, Sweden
amath.sow@liu.se

Mauricio Rodriguez
Universidade Estadual de Campinas
Campinas, Brazil
m272321@dac.unicamp.br

Fabiola M. C. de Oliveira
Universidade Federal do ABC
São Paulo, Brazil
fabiola.oliveira@ufabc.edu.br

Mariusz Wzorek
Linköping University
Linköping, Sweden
mariusz.wzorek@liu.se

Daniel de Leng
Linköping University
Linköping, Sweden
daniel.de.leng@liu.se

Mattias Tiger
Linköping University
Linköping, Sweden
mattias.tiger@liu.se

Fredrik Heintz
Linköping University
Linköping, Sweden
fredrik.heintz@liu.se

Christian Rothenberg
Universidade Estadual de Campinas
Campinas, Brazil
chesteve@unicamp.br

ABSTRACT

Preflight planning for large-scale Unmanned Aerial Vehicle (UAV) fleets in dynamic, shared airspace presents significant challenges, including temporal No-Fly Zones (NFZs), heterogeneous vehicle profiles, and strict delivery deadlines. While Multi-Agent Path Finding (MAPF) provides a formal framework, existing methods often lack the scalability and flexibility required for real-world Unmanned Traffic Management (UTM). We propose DTAPP-IICR: a Delivery-Time Aware Prioritized Planning method with Incremental and Iterative Conflict Resolution. Our framework first generates an initial solution by prioritizing missions based on urgency. Secondly, it computes roundtrip trajectories using SFIPP-ST, a novel 4D single-agent planner (Safe Flight Interval Path Planning with Soft and Temporal Constraints). SFIPP-ST handles heterogeneous UAVs, strictly enforces temporal NFZs, and models inter-agent conflicts as soft constraints. Subsequently, an iterative Large Neighborhood Search, guided by a geometric conflict graph, efficiently resolves any residual conflicts. A completeness-preserving directional pruning technique further accelerates the 3D search. On benchmarks with temporal NFZs, DTAPP-IICR achieves near-100% success with fleets of up to 1,000 UAVs and gains up to 50% runtime reduction from pruning, outperforming batch Enhanced Conflict-Based Search in the UTM context. Scaling successfully in realistic city-scale operations where other priority-based methods fail even at moderate deployments, DTAPP-IICR is positioned as a practical and scalable solution for preflight planning in dense, dynamic urban airspace.

KEYWORDS

Unmanned Aircraft System Traffic Management; PreFlight Planning in dynamic environment; Multi-Agent Path Finding; Heterogeneous Agents with kinemodynamic constraints; UAV; Drones; Route Planning; No-Fly Zones; UTM



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/KPWJ5508>

ACM Reference Format:

Amath Sow, Mauricio Rodriguez, Fabiola M. C. de Oliveira, Mariusz Wzorek, Daniel de Leng, Mattias Tiger, Fredrik Heintz, and Christian Rothenberg. 2026. Multi UAVs Preflight Planning in a Shared and Dynamic Airspace. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/KPWJ5508>

1 INTRODUCTION

The rapid growth of Unmanned Aerial Vehicle (UAV) operations in low-altitude airspace has significantly increased airspace density and complexity. UAVs are now deployed across a wide range of civilian and governmental applications, including package delivery, search and rescue, and surveillance [7, 12, 13, 20]. Global adoption is accelerating, with millions of UAVs expected to operate concurrently across the globe in the coming decades [1, 5, 23, 31]. This complexity highlights the need for Unmanned Traffic Management (UTM) systems to ensure conflict-free navigation in dense airspace.

In the context of UTM, UAVs are allowed to operate below 120 meters (400 feet), as stipulated by the European Union Aviation Safety Agency (EASA) [2]. The U-Space framework [3] is specifically designed to ensure the safe integration of UAVs into the airspace, particularly in urban and complex environments. To achieve this objective in a shared airspace, an effective *preflight planner* is essential. In this study, we focus on a service delivery scenario in which quadcopter UAVs transport goods from centralized hubs to target locations such as homes or hospitals. UAV Service Providers (Hubs) periodically submit operation requests to the UTM system. The central controller processes these requests and approves flights based on compliance with operational requirements and airspace availability. Accepted flights are then authorized and coordinated within the UTM system to ensure safe and efficient operations.

Within this controlled urban environment, we model the problem as a Multi-Agent Path Finding (MAPF) variant in a 4D space (3D + time). Missions are prioritized by urgency and must strictly respect temporal No-Fly Zones (NFZs). Unlike standard point-to-point pathfinding, these operations require executing a complete roundtrip (outbound, wait, and return). This multi-stage structure is similar to Multi-Agent Pickup and Delivery (MAPD) [22], which

generalizes MAPF to handle sequential tasks. However, our formulation differs by treating the entire roundtrip as a single atomic mission that must be planned holistically under strict 4D regulatory constraints, rather than as a sequence of tasks. Furthermore, existing work in the UTM context is largely limited to traditional batch methods, such as Conflict-Based Search (CBS) and Enhanced CBS (ECBS) [15], which lack the scalability and flexibility required for such dense and dynamic environments.

To address these challenges, we propose a new two-tier planning framework tailored to dynamic and high-density UAV operations. Our contributions are described as follows:

- (1) **A Novel 4D Single-Agent Planner:** We propose Safe Flight Interval Path Planning with Soft and Temporal Constraints (SFIPP-ST), which natively supports heterogeneous UAV profiles (varying start times, speeds, and sizes), strictly enforces temporal NFZs, and models inter-agent conflicts using soft, penalizing constraints;
- (2) **A Scalable Prioritized Planning Framework:** We introduce the Delivery-Time Aware Prioritized Planning method with Incremental and Iterative Conflict Resolution (DTAPP-IICR), an urgency-aware preflight planner for heterogeneous fleets. This framework employs an iterative conflict-resolution approach based on Large Neighborhood Search (LNS), enhancing scalability in dense airspace and replanning the most critical agents guided by a geometrical heuristic;
- (3) **Efficiency via Directional Pruning:** To accelerate the path search in a 3D grid with 26-connected voxels, we develop a completeness-preserving directional pruning strategy that reduces the branching factor by prioritizing neighbors whose direction aligns with the goal vector, with a fallback mechanism that guarantees search completeness.

Supported by its theoretical guarantees and validated through experimental evaluation, our approach offers a robust, scalable, and collision-free planning framework for high-density UAV deployments in shared, dynamic urban airspace. The implementation is made reproducible and available as an open-source project ¹.

2 RELATED WORK

MAPF is a core problem in multi-agent systems, modeling collision-free coordinated navigation of agents in shared environments [30]. Research in this area includes optimal and bounded-suboptimal solvers. Optimal solvers such as CBS [28] and Improved CBS (ICBS) [8] guarantee to find the least cost solution but often struggle to scale in dense environments, while bounded-suboptimal variants like ECBS [6] and Explicit Estimation CBS (EECBS) [19] relax optimality to improve scalability. Prioritized Planning (PP) [11] improves scalability by sequentially planning agents according to pre-defined orders, while Priority-Based Search (PBS) [21] dynamically explores priority orderings. While these approaches are formulated for general graph problems, they do not fully address UTM-specific challenges.

Anytime approaches, particularly those based on LNS [16–18, 24, 25], have been shown to scale well by iteratively destroying and repairing portions of candidate solutions. LNS2 [18], in particular,

has demonstrated significant success in grid benchmarks by utilizing Safe Interval Path Planning with Soft constraints (SIPPS) [18] to resolve conflicts. These ideas inspire our framework. However, standard LNS2 typically rely on discrete-time vertex/edge conflicts, homogeneous agents, and synchronous start times. Our framework, DTAPP-IICR, adapts the LNS2 architecture but fundamentally re-designs the conflict model and repair operators. Unlike the discrete operator, our approach identifies conflicts using continuous-time geometric envelopes to handle heterogeneous agent sizes and speeds. Furthermore, while LNS2 utilizes SIPPS with reservation tables, our single-agent planner, SFIPP-ST, constructs safe intervals directly from dynamic NFZ constraints and geometric conflict checking, as interval reuse is invalid with heterogeneous velocities.

In UTM-specific contexts, only a few studies have directly addressed preflight multi-UAV planning. Batch extensions of CBS and ECBS [15] support heterogeneous agents but assume symmetric roundtrips and cannot accommodate temporal NFZs. More recently, Rolling Continuous-time CBS (RCCBS) [10] extended MAPF to continuous time by decomposing tasks into windowed batches and resolving conflicts through geometric overlap detection between agents. While RCCBS enables dynamic task arrivals and continuous-space planning, it remains constrained to semi-urban environments with static airspace conditions. Other studies have considered heterogeneous edges and roundtrips [4], or combined takeoff scheduling with speed adjustment [14], to improve realism, but these approaches still rely on static maps. In contrast, our framework integrates urgency-aware prioritization, temporal NFZ handling, and LNS-style conflict repair to provide scalable, regulation-compliant preflight planning for urban UAV operations. A comparison of the algorithms is presented in Table 1.

3 PROBLEM DEFINITION

We consider a set of UAVs $\mathcal{U} = \{U_1, \dots, U_k\}$ executing delivery operations. The environment is discretized as an undirected 26-connected cubic grid represented as a graph $G = (V, E)$; vertices are unit cubic voxels (V) and edges (E) connect all 6 face-, 12 edge-, and 8 corner-adjacent voxels, allowing diagonal motion.

Definition 3.1 (UAV model). Each UAV U_i is described by

- a hub (start and return) vertex $s_i \in V$ and a delivery vertex $g_i \in V$;
- a departure time $t_i^{\text{init}} > 0$;
- a constant cruise speed $v_{u_i} > 0$;
- a physical UAV radius $r_{u_i} > 0$ with spherical UAV profile.

Definition 3.2 (Roundtrip path). Each mission is a *roundtrip* composed of an outbound path, an on-site waiting period (e.g., package handover), and a return path. A path for U_i is a sequence of space-time tuples:

$$p_i = [(v_{i,0}, t_{i,0}), \dots, (v_{i,\tau}, t_{i,\tau}), \dots, (v_{i,\ell_i}, t_{i,\ell_i})], \quad (1)$$

where $v_{i,0} = v_{i,\ell_i} = s_i$; ℓ_i is the path length and t_{i,ℓ_i} is the arrival time of UAV U_i . There exists a unique index τ that marks the *first* time the UAV reaches its delivery vertex, i.e., $v_{i,\tau} = g_i$. For some $w \geq 0$, the consecutive tuples $v_{i,\tau}, \dots, v_{i,\tau+w}$ are identical, representing a w -timestep waiting period at g_i . The trajectory between consecutive distinct tuples is traversed at the constant UAV speed v_{u_i} , forming the outbound and return segments.

¹URL: <https://github.com/amathsow/4DPlanning>

Table 1: Comparison of MAPF methods based on key features relevant to urban air mobility: Heterogeneous Agents (HA), Conflict Types, Roundtrip, Temporal NFZs, and Urban Environment (Urban Env).

| Method | HA | Conflict Types | Roundtrip | Temporal NFZs | Urban Env |
|--------------------------|----|------------------|-----------------------------------|---------------|-------------------------------|
| Batch CBS/ECBS | ✓ | Geometric | Symmetric outbound/return | × | Semi-urban (Japan study) |
| PP/PBS | × | Vertex/Edge | × | × | Grid map |
| MAPF-LNS | × | Vertex/Edge | × | × | Grid map |
| DTAPP-IICR (Ours) | ✓ | Geometric | Asymmetric (realistic UTM) | ✓ | Urban (dense airspace) |




Figure 1: Types of conflicts between two agents of different sizes and speeds. a) Pursuit: agents move closely towards the same direction; b) Head-on: agents move into opposite directions; c) Intersection: agents cross paths [15].

Definition 3.3 (Temporal NFZs). The airspace may contain a finite set \mathcal{F} of NFZs. Each NFZ is characterized by the pair $(R_f, [t_f^{\text{start}}, t_f^{\text{end}}])$ with $R_f \subseteq V$ being a spatially connected subset of voxels representing the restricted region and $[t_f^{\text{start}}, t_f^{\text{end}}] \subset \mathbb{R}_{\geq 0}$ is its active time interval. Once activated at t_f^{start} , an NFZ remains in force until t_f^{end} , after which it is removed from the operational airspace. We explicitly *exempt* a UAV’s hub and delivery vertices from NFZ restrictions, reflecting regulatory practice that allows authorized take-off and landing inside controlled areas.

Solution validity constraints:

- NFZ compliance. Let $pos_i(t)$ denote the position of UAV U_i at time t . For any time t ,

$$pos_i(t) \notin F(t) := \bigcup_{f \in \mathcal{F}} \{v \in R_f \mid t \in [t_f^{\text{start}}, t_f^{\text{end}}]\} \quad (2)$$

unless $pos_i(t) \in \{s_i, g_i\}$.

- Collision avoidance. In this paper, conflicts are classified as pursuit, head-on, or intersection [15] as shown in Fig. 1. Additionally, to ensure physical collision avoidance between UAVs, we introduce a safety buffer $\gamma > 0$. Thus, for any two UAVs U_i and U_j , the following must always hold:

$$\forall t, \quad \text{dist}(pos_i(t), pos_j(t)) > r_{u_i} + r_{u_j} + \gamma. \quad (3)$$

Objective. Given the set \mathcal{U} of UAVs, a *solution* is a set of collision-free, NFZ-respecting roundtrip paths $P = \{p_i \mid U_i \in \mathcal{U}\}$. We minimize the flowtime (total flight time at constant speed):

$$\min_P \sum_{U_i \in \mathcal{U}} T_i = \min_P \sum_{U_i \in \mathcal{U}} (t_{i,t_i} - t_{i,0}). \quad (4)$$

4 DTAPP FRAMEWORK

We now introduce the *DTAPP* (Delivery-Time Aware Prioritized Planning) framework for scalable, conflict-free multi-UAV preflight planning in dynamic airspace. DTAPP adopts a prioritized planning paradigm, where UAVs are ordered by delivery urgency (start time t^{init}). To support this, we develop *SFIPP-ST*, a 4D single-agent planner that computes NFZ-compliant, collision-free roundtrip trajectories while respecting heterogeneous UAV dynamics. The architecture of the system is described in Fig. 2.

4.1 Single Agent Planning with Soft and Temporal Constraints

We extend the path planning framework from PMDO (Pathfinding with Mixed Dynamic Obstacles) as introduced in [18], by integrating additional constraints such as *temporal NFZs* and UAV profiling (*different start times, speeds, and sizes*).

Formally, the planner must compute a roundtrip path p for a given UAV while avoiding three categories of restrictions: **hard obstacles** O_h , which are static obstacles in the environment that must be strictly avoided; **temporal NFZs** \mathcal{F} , which impose regulatory constraints that must be strictly respected, as formally defined by the NFZ compliance condition in Equation (2); and **soft constraints** O_s , which represent other UAV trajectories. In this case, SFIPP-ST penalizes potential conflicts with O_s but does not strictly forbid them, allowing iterative conflict resolution to refine the solution.

4.2 SFIPP-ST

SFIPP-ST is a single-agent planner for 4D trajectory planning in UTM systems. It extends SIPPS [18] by incorporating temporal NFZs, heterogeneous UAV profiles, and continuous geometric conflict modeling, making it suitable for UTM.

Definition 4.1 (Safe Flight Interval (SFI)). A SFI for a vertex $v \in V$ is defined as a contiguous time interval $[t_{\text{start}}, t_{\text{end}}]$ during which a UAV can safely traverse or wait at v without: (1) Colliding with any hard obstacle O_h . (2) Violating any active temporal NFZ \mathcal{F} . We define the indicator function for NFZ exclusion as:

$$\mathbb{1}_{\mathcal{F}}(v, t) = \begin{cases} 1 & \text{if } \exists (R_f, [t_f^{\text{start}}, t_f^{\text{end}}]) \in \mathcal{F}, v \in R_f, t \in [t_f^{\text{start}}, t_f^{\text{end}}] \\ 0 & \text{otherwise.} \end{cases}$$

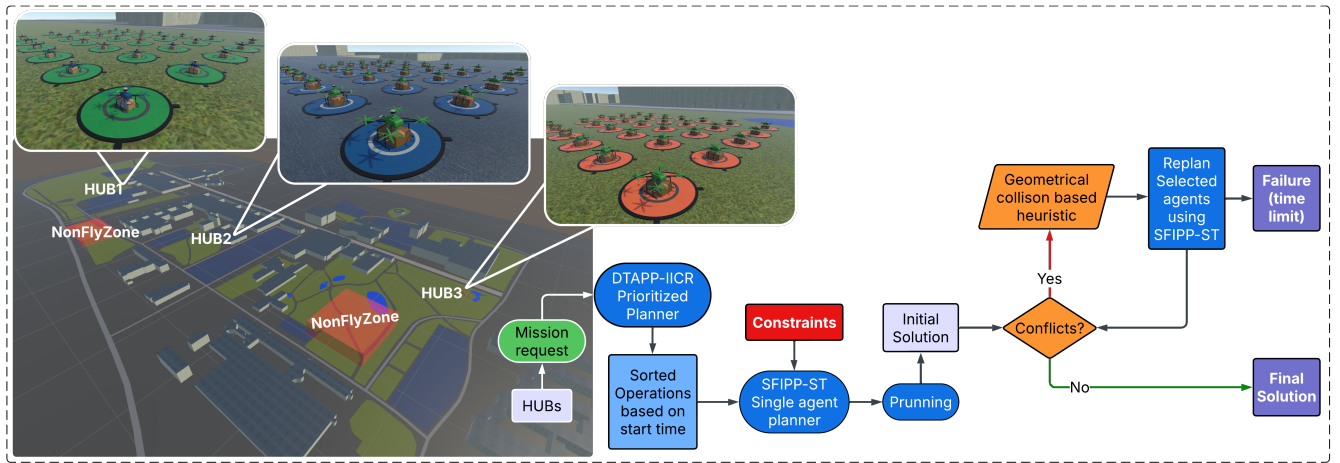


Figure 2: Overview of Preflight Planning in UTM. We have three hubs that send periodic operation requests to the central controller. We use DTAPP-IICR as a prioritized planner, where operations are sorted based on delivery urgency, t^{init} . We use SFIPP-ST as the single-agent planner, which receives the UAV profile $(t^{\text{init}}, v_{u_i}, r_{u_i})$, precomputed hard and regulatory constraints, and computes a 4D roundtrip trajectory from the hub to the delivery location and back. We add pruning during neighbor expansion to reduce the search. The initial solution might have soft conflicts, and we iteratively replan the conflicting agents until the solution is free of conflict or the time limit is reached.

The set of SFIs at vertex v is given by:

$$T(v) = \{ [t_{\text{start}}, t_{\text{end}}] \subseteq \mathbb{R}_{\geq 0} \mid \forall t \in [t_{\text{start}}, t_{\text{end}}], v \notin O_h \text{ and } \mathbb{1}_{\mathcal{F}}(v, t) = 0 \}. \quad (5)$$

That is, SFIs are all maximal contiguous time intervals where the vertex v is free from both hard obstacles and active NFZs.

Definition 4.2 (Search Node Definition). SFIPP-ST represents each search node as a tuple:

$$n = (v, I, i, g(n), h(n), c(n), r_{u_i}, v_{u_i}),$$

where: $v \in V$: the current vertex location; $I = [t_{\text{start}}, t_{\text{end}}]$: the associated SFI free from O_h and \mathcal{F} ; i : the identifier of I in the precomputed SFI table $T(v)$; $g(n)$: the earliest arrival time at node n ; $h(n)$: the heuristic cost-to-go; $c(n)$: the cumulative conflict cost, representing the estimated number of geometric collisions with other UAV paths O_s ; r_{u_i} and v_{u_i} are respectively the radii and the speed of UAV U_i . The overall cost function for node n is defined as: $f(n) = g(n) + h(n)$.

Main Algorithm: Algorithm 1 starts by constructing the SFI table T for all voxels V in the 3D grid, where each $T(v)$ lists contiguous time intervals free from hard obstacles (O_h) and active NFZs (\mathcal{F}), reducing the 4D search to time-constrained intervals (Line 1). The earliest safe interval I_0 at the start vertex s after t_i^{init} is queried; if none exists, FAILURE is returned (Lines 2–4). Otherwise, the departure time is set to $t_s = \max(t_i^{\text{init}}, I_0.\text{start})$, and the root node n_{root} is created and inserted into the OPEN list with CLOSED set initialized empty (Lines 5–7). OPEN nodes are prioritized by ascending soft collision cost c , with ties broken by total cost f , minimizing conflicts before travel time.

At each iteration, a node n is popped from OPEN (Line 9). If n is the goal g , the trajectory is reconstructed via BACKTRACKPATH(n)

(Lines 10–11). Otherwise, n is expanded using a pruning strategy aligned with the goal (Line 12). For each promising neighbor v' and interval $I' \in T[v']$, the travel time δ and earliest feasible arrival t' are computed; intervals where $t' \notin I'$ are discarded (Lines 13–17). Geometric conflicts with soft obstacles O_s are evaluated along $[n.v \rightarrow v']$ using radius r_{u_i} to obtain cumulative conflict cost c' , and the heuristic h' is estimated as the minimum travel time to g at speed v_{u_i} (Lines 18–19). A new node n' is then inserted into OPEN (Lines 20–21).

In addition to movement, the algorithm also considers *waiting* at the current vertex $n.v$, since this action may be required when future safe intervals exist. If a subsequent safe interval I_{wait} is available after $n.I$ (Line 22), the wait time is set to $t_{\text{wait}} = I_{\text{wait}}.\text{start}$ (Line 24), and then we compute geometric conflicts for $[n.v \rightarrow n.v]$ over $[g(n), t_{\text{wait}}]$ to get c_{wait} (Line 25). We estimate h_{wait} , create a wait node n_{wait} , and insert it into OPEN (Lines 26–27). The algorithm terminates with a conflict-minimizing path or FAILURE if OPEN is exhausted (Line 28).

Directional Pruning for Neighbor Expansion: In a 3D grid with 26 connected voxels, a brute-force expansion of all neighbors at each step is computationally expensive. To accelerate the search, we introduce a directional pruning heuristic that reduces the branching factor. The core idea is to prioritize expanding neighbors that represent clear, monotonic progress toward the goal. The heuristic works by comparing the normalized direction vector from the current location to a neighbor with the direction vector to the goal. This vector, which we call the *direction sector*, is computed as:

$$S = (\text{sign}(\Delta x), \text{sign}(\Delta y), \text{sign}(\Delta z)) \in \{-1, 0, 1\}^3$$

where $(\Delta x, \Delta y, \Delta z)$ is the displacement vector. For example, $S = (1, 0, -1)$ implies a move along positive x , no change in y , and negative z .

Algorithm 1: SFIPP-ST

Input: Graph $G = (V, E)$, start s , goal g , UAV profile $(t_i^{\text{init}}, v_{u_i}, r_{u_i})$, hard obstacles O_h , temporal NFZs \mathcal{F} , soft obstacles O_s

Output: Path p or FAILURE

```

1  $T \leftarrow \text{BuildSFITable}(V, O_h, \mathcal{F}, r_{u_i}, v_{u_i});$ 
2  $I_0 \leftarrow \text{FirstSafeInterval}(T[s], t_i^{\text{init}});$ 
3 if  $I_0 = \emptyset$  then
4   return FAILURE;
5  $t_s \leftarrow \max(t_i^{\text{init}}, I_0.\text{start});$ 
6  $n_{\text{root}} \leftarrow \text{CreateStartNode}(s, I_0, t_s, r_{u_i}, v_{u_i});$ 
7  $OPEN \leftarrow \{n_{\text{root}}\}; CLOSED \leftarrow \emptyset;$ 
  // Order by lowest  $c(n)$ ; tie-break by  $f(n)$ 
8 while  $OPEN \neq \emptyset$  do
9    $n \leftarrow OPEN.\text{pop}();$ 
10  if  $n.v = g$  then
11    return BacktrackPath( $n$ );
12  foreach  $v' \in \text{PrunedNeighbors}(n.v, g)$  do
13    foreach  $I' \in T[v']$  do
14       $\delta \leftarrow \text{TravelTime}(n.v, v', v_{u_i});$ 
15       $t' \leftarrow \max(g(n) + \delta, I'.\text{start});$ 
16      if  $t' \notin I'$  then
17        continue;
18       $c' \leftarrow c(n) + \text{GeoConflict}([n.v \rightarrow$ 
19         $v'], [g(n), t'], r_{u_i}, O_s);$ 
20       $h' \leftarrow \text{Heuristic}(v', g, v_{u_i});$ 
21       $n' \leftarrow (v', g = t', h = h', c = c', I');$ 
22       $OPEN \leftarrow OPEN \cup \{n'\};$ 
23  if NextWaitIntervalExists( $n.v, n.I, T$ ) then
24     $I_{\text{wait}} \leftarrow \text{NextWaitInterval}(n.v, n.I, T);$ 
25     $t_{\text{wait}} \leftarrow I_{\text{wait}}.\text{start};$ 
26     $c_{\text{wait}} \leftarrow c(n) + \text{GeoConflict}([n.v \rightarrow$ 
27     $n.v], [g(n), t_{\text{wait}}], r_{u_i}, O_s);$ 
28     $h_{\text{wait}} \leftarrow \text{Heuristic}(n.v, g, v_{u_i});$ 
29     $n_{\text{wait}} \leftarrow (n.v, g = t_{\text{wait}}, h = h_{\text{wait}}, c = c_{\text{wait}}, I_{\text{wait}});$ 
30     $OPEN \leftarrow OPEN \cup \{n_{\text{wait}}\};$ 
31 return FAILURE;

```

A neighbor n is considered promising if its direction sector, S_{nb} , aligns with the goal's direction sector, S_{goal} , under conditions that favor direct paths achieved through 1) XY-Plane Alignment: The neighbor lies in the same general direction as the goal on the 2D plane. This includes: *Diagonal moves* that match the goal's quadrant, for instance, if the goal is north-east, ($S_{\text{goal}.x} = 1, S_{\text{goal}.y} = 1$), any neighbor in that quadrant is promising; and *Cardinal moves* along an axis if the goal also lies primarily on that axis (e.g., moving straight north if the goal is directly north); 2) Z-Axis Alignment: The neighbor involves a direct vertical movement toward the goal's altitude ($S_{nb.z} = S_{\text{goal}.z} \neq 0$). This rule prioritizes changing flight levels, a common requirement in UAV navigation.

Algorithm 2: PRUNEDNEIGHBORS($V.loc, g$)

Input: Set of neighbors V , current location loc , goal location g

Output: Pruned neighbor set $\mathcal{N}(loc)$

```

1  $S_{\text{goal}} \leftarrow \text{DirectionSector}(loc, g);$ 
2  $\mathcal{N}(loc) \leftarrow \{loc\};$ 
  // Always include wait action
3 foreach  $n \in V$  do
4   if  $n = g$  then
5      $\mathcal{N}(loc) \leftarrow \mathcal{N}(loc) \cup \{n\};$ 
6   else
7      $S_{nb} \leftarrow \text{DirectionSector}(loc, n);$ 
8     // Check alignment conditions
9     if  $(S_{nb}.x = S_{\text{goal}.x} \wedge S_{nb}.y = S_{\text{goal}.y})$  or
10       $(S_{nb}.x = S_{\text{goal}.x} \wedge S_{nb}.x \neq 0 \wedge S_{nb}.y = 0)$  or
11       $(S_{nb}.y = S_{\text{goal}.y} \wedge S_{nb}.y \neq 0 \wedge S_{nb}.x = 0)$  or
12       $(S_{nb}.z = S_{\text{goal}.z} \wedge S_{nb}.z \neq 0)$  then
13         $\mathcal{N}(loc) \leftarrow \mathcal{N}(loc) \cup \{n\};$ 
14 if  $|\mathcal{N}(loc)| = 1$  then
15    $\mathcal{N}(loc) \leftarrow V \cup \{loc\};$ 
  // Fallback to all neighbors
16 return  $\mathcal{N}(loc);$ 

```

To guarantee completeness, our approach includes two important exceptions: any adjacent neighbor that is the goal is always included, and the "wait" action is always an option. If these pruning rules result in no valid moves, the algorithm reverts to expanding the full 26-neighbor set, preserving global completeness. The full logic is detailed in Algorithm 2.

4.3 Theoretical Properties of SFIPP-ST

We now establish two theoretical properties of SFIPP-ST planner: *completeness* and *optimality*. The proof is omitted in the main paper as it follows the proof of SIPP [26] and SIPPS [18].

THEOREM 4.3 (COMPLETENESS). *SFIPP-ST is complete; it returns a valid path if one exists and FAILURE otherwise.*

THEOREM 4.4 (OPTIMALITY). *If a conflict-free path exists, SFIPP-ST returns the shortest such path with zero soft collisions.*

4.4 DTAPP-IICR

We now introduce *DTAPP-IICR*, which combines sequential round-trip planning with iterative conflict resolution. *DTAPP-IICR*, presented in Algorithm 3, begins by sorting UAVs based on t_i^{init} , ensuring urgent deliveries are planned first. Ties among UAVs with equal start times are broken randomly to avoid bias and deadlocks (Lines 1–3). Each UAV U_i computes its roundtrip path using *PLAN-ROUNDTrip*, treating all previously planned paths P as dynamic obstacles (Line 6). The resulting paths are stored in the solution set P (Line 7), and any conflicts between UAV trajectories are recorded in the collision graph (Line 8). The iterative conflict resolution phase (Lines 9–19) systematically eliminates conflicts until a globally feasible solution is achieved or the time limit is exceeded. At each

Algorithm 3: DTAPP-IICR

Input: UAV set $U = \{U_1, \dots, U_n\}$, 3D grid \mathcal{G} , temporal NFZs $F(t)$, time limit T_{\max}

Output: Conflict-free paths P or FAILURE

```

1  $U' \leftarrow \text{sortAsc}(U, t^{\text{init}})$ ;
2 foreach group  $G \subseteq U'$  with equal  $t^{\text{init}}$  do // break tie
3    $\lfloor$  shuffle  $G$  within  $U'$ ;
4  $P \leftarrow \emptyset$ ;
5 foreach  $U_i \in U'$  do
6    $p_i \leftarrow \text{PLANROUNDTrip}(U_i, \mathcal{G}, F(\cdot), \emptyset)$ ;
7    $P \leftarrow P \cup \{p_i\}$ ;
8  $\text{num\_conflicts} \leftarrow \text{COUNTCONFLICTS}(P)$ ;
9 while  $\text{num\_conflicts} > 0$  and  $t < T_{\max}$  do
10   $S_c \leftarrow \text{CONFLICTHEURISTIC}(P, \text{num\_conflicts})$ ;
11   $P_{\text{old}} \leftarrow \{p_i : U_i \in S_c\}$ ;
12   $P \leftarrow P \setminus P_{\text{old}}$ ;
13  foreach  $U_i \in S_c$  do
14     $p_i \leftarrow \text{PLANROUNDTrip}(U_i, \mathcal{G}, F(\cdot), P)$ ;
15    if  $p_i$  valid then
16       $P \leftarrow P \cup \{p_i\}$ 
17    else
18       $\lfloor$   $P \leftarrow P \cup P_{\text{old}}$ ; break
19   $\text{num\_conflicts} \leftarrow \text{COUNTCONFLICTS}(P)$ ;
20 if  $\text{num\_conflicts} = 0$  then
21   return  $P$ 
22 else
23   return FAILURE
24 Function  $\text{PLANROUNDTrip}(U_i, \mathcal{G}, F(\cdot), P)$ :
25    $p_{\text{out}} \leftarrow \text{SFIPP-ST}(s_i \rightarrow g_i, \mathcal{G}, F(\cdot), P)$ ;
26    $p_{\text{wait}} \leftarrow \text{GENERATEWAIT}(g_i, w_i)$ ;
27    $p_{\text{ret}} \leftarrow \text{SFIPP-ST}(g_i \rightarrow s_i, \mathcal{G}, F(\cdot), P \cup \{p_{\text{out}}, p_{\text{wait}}\})$ ;
28   return  $\text{COMBINE}(p_{\text{out}}, p_{\text{wait}}, p_{\text{ret}})$ ;

```

iteration, the algorithm first identifies the set S_c of collided UAVs in the solution set P (Line 10). Then we remove all paths $P_{\text{old}} \subset P$ belonging to the selected UAVs in S (Lines 11–12). Each UAV U_i in this prioritized subset is then replanned sequentially using the PLANROUNDTrip function (Line 14), treating all remaining paths in P (from non-selected UAVs) as fixed spatio-temporal obstacles. This ensures that newly computed paths avoid conflicts with UAVs not currently being replanned. If a UAV successfully finds a valid conflict-free path, it is immediately added back to the solution set P , making it a constraint for subsequent UAVs in the current iteration (Lines 15–16). However, if any UAV fails to find a feasible path during replanning, the algorithm performs a rollback operation: all previously backed-up paths P_{old} are restored to P , and the current iteration terminates (Line 18). We update the number of collision counts (Line 19). If there are no more collisions, we return the final path P , or we return failure otherwise (Lines 20–23).

4.5 Geometric Conflict-Based Heuristic

We adapt the neighborhood selection strategy from LNS2 [18] to the continuous 4D domain. Instead of discrete graph collisions, we maintain a *geometric collision graph* $G_c = (V_c, E_c)$, where each vertex denotes a UAV and an edge $(U_i, U_j) \in E_c$ exists if the pair violates the minimum separation $\text{dist}(\text{pos}_i(t), \text{pos}_j(t)) \leq r_i + r_j + \gamma$ at any time t .

Given a target neighborhood size N , we construct a replanning set \mathcal{A}_s as follows. First, sample a vertex $v \in V_c$ with $\text{deg}(v) > 0$ and extract the connected component $G_c^0 = (V_c^0, E_c^0)$ containing v . Then, following the logic of [18], we apply one of two expansion cases:

Case 1 (Small conflict region, $|V_c^0| \leq N$): Initialize $\mathcal{A}_s \leftarrow V_c^0$ and expand it via spatial random walks to include nearby agents that may restrict repair. Until $|\mathcal{A}_s| = N$ or no new agents can be found: 1) Sample $U_a \in \mathcal{A}_s$ uniformly and a random time index t along its current path; 2) From the corresponding 3D location, perform a random walk over the 26-connected neighborhood (including wait); and 3) If the walk reaches a location where the required safety buffer overlaps with another UAV U_b , add U_b to \mathcal{A}_s .

Case 2 (Large conflict region, $|V_c^0| > N$): Perform a random walk on the *collision graph* G_c^0 itself, collecting distinct vertices until $|\mathcal{A}_s| = N$. This biases selection toward densely connected sub-graphs, prioritizing UAVs engaged in tightly coupled conflicts.

5 EXPERIMENTAL EVALUATION

To evaluate the effectiveness, efficiency, and robustness of our proposed methods, we conduct experiments in two settings: a Monte Carlo Simulation (MCS) environment and a realistic city-scale map.

In the MCS setup, the environment is a $100 \times 100 \times 10$ 3D grid with 1 m^3 voxels providing high-precision collision detection and path planning. The lower four levels (0–3) contain static obstacles (e.g., buildings, trees), while the upper six levels (4–9) provide open airspace, reflecting typical urban layering. UAV profiles are randomly sampled: start times in [1 s, 1000 s], radii in [0.5 m, 2.0 m], and speeds in [1.0, 5.0 m/s]. NFZs are activated over intervals randomly drawn from [100s, 500s], capturing the time-varying nature of restricted airspace. Across trials, we vary the number of UAV operations, obstacle densities, and NFZs to test scalability and robustness. The duration of package handover is fixed at $w = 10$ and the neighborhood size used by DTAPP-IICR for iterative conflict resolution is fixed at $N = 10$.

For external validation, we also evaluate our methods on a realistic city-scale map (Fig. 2) with designated UAV hubs, delivery destinations (e.g., homes, hospitals), and regulatory constraints. This case study demonstrates the applicability of our approach to practical urban air mobility scenarios.

As evaluation metrics, we measure: (1) *Success rate*: percentage of instances solved within the time limit, and (2) *Average runtime*: mean planning time, using the time limit for unsolved instances with their standard deviations. All experiments were run on an Intel Core i9-10900X CPU with 128GB RAM.

Experiment 1: SFIPP-ST and Pruning Evaluation. In this experiment, we evaluate the effectiveness of the SFIPP-ST algorithm and the proposed pruning optimization. We compare Priority Planning (PP) with a modified version of SFIPP-ST (where soft obstacles

Table 2: Comparison of solver performance across varying agent counts (N). PP and DTAPP-IICR denote versions with pruning. (W/O) indicates no pruning.

| Solver | N | Success Rate (%) | Runtime (s) |
|------------------|-----|------------------|---------------------|
| PP (W/O) | 10 | 100 | 6.05 ± 2.04 |
| PP | | | 3.06 ± 1.09 |
| DTAPP-IICR (W/O) | | | 2.31 ± 0.82 |
| DTAPP-IICR | | | 1.31 ± 0.57 |
| PP (W/O) | 50 | 100 | 115.81 ± 11.64 |
| PP | | | 54.53 ± 5.90 |
| DTAPP-IICR (W/O) | | | 10.97 ± 1.32 |
| DTAPP-IICR | | | 7.23 ± 1.06 |
| PP (W/O) | 100 | 100 | 300.0 ± 0.00 |
| PP | | | 223.22 ± 16.14 |
| DTAPP-IICR (W/O) | | | 25.91 ± 2.15 |
| DTAPP-IICR | | | 13.09 ± 1.26 |

are treated as hard obstacles) against DTAPP-IICR. Both solvers are tested with and without pruning across scenarios with 10, 50, and 100 agents in environments with 5% obstacle density. For each configuration, we generate 20 instances and set a fixed time limit of 300s (5 min). Results are summarized in Table 2. DTAPP-IICR demonstrates superior scalability, maintaining a 100% success rate across all agent counts, while PP without pruning fails completely at 100 agents. DTAPP-IICR achieves up to 15 times faster runtimes than PP, highlighting the fundamental advantage of soft collision handling over SFIPP-ST’s hard obstacle treatment. Our direction-based pruning provides significant speedup for both algorithms, with DTAPP-IICR achieving up to 34-50% faster performance and PP with pruning remains viable at scale.

Experiment 2: Scalability Evaluation. This experiment evaluates the performance of DTAPP-IICR as the number of operations increases. We consider our previous environment, where the number of UAVs increases up to 1000. We implemented batch CBS and ECBS from [15] since the original code is unavailable and compared them against our DTAPP-IICR and PP. To get a fair comparison, we consider only the outbound path without temporal NFZs. The return path will then be symmetric to the outbound. The time limit is fixed at 480s (8 min). Figure 3 presents the results. While DTAPP-IICR maintains a 100% success rate even with 1000 UAVs, batch CBS fails to scale beyond small instances due to its exponential complexity. Notably, PP also suffers from poor scalability, failing beyond 250 agents. This overly conservative treatment results in over-constrained planning, making feasible solutions unlikely in dense scenarios. ECBS remains more robust, sustaining higher success rates than CBS/PP thanks to its bounded-suboptimal, batch-search strategy. In contrast, DTAPP-IICR leverages soft constraints with iterative conflict resolution, allowing temporary conflicted paths that are later resolved, resulting in significantly higher success rates and reduced computation times. These results highlight DTAPP-IICR’s effectiveness for large-scale UAV traffic scenarios where flexibility and scalability are critical.




Figure 3: Scalability of DTAPP-IICR compared to PP, batch CBS/ECBS under varying obstacle densities. a) Using 5% of obstacles. b) Using 10% of obstacles.




Figure 4: Success Rate and Average Runtime for PP, PP (W/O), DTAPP-IICR, and DTAPP-IICR (W/O). (a) 100 agents with 5% obstacles, with a time limit of 500 seconds. (b) 1000 agents with 5% obstacles, with a time limit of 900 seconds.

Experiment 3: Dynamic Constraint Handling. We evaluate robustness under dynamic airspace conditions by introducing temporal NFZs that are active only for specific time intervals. For each environment (5% obstacles), we fix the obstacle layout and agent scenarios while varying the number of NFZs in 1, 2, 4. Each NFZ is activated over a random interval within the mission horizon. All results are averaged over 20 independent instances for every configuration. W/O denotes planner without pruning. The results show that DTAPP-IICR performs the best when the number of NFZs increases compared to other planners (Fig. 4 a). The runtime grows with more activated NFZs, and pruning provides substantial speedups for DTAPP-IICR up to about 50% faster than DTAPP-IICR (W/O) (Fig. 4 b).

Experiment 4: Realistic Urban Use Case Simulation. To demonstrate applicability to real-world conditions, we simulate city-scale UAV delivery scenarios using an urban environment in UNetyEmu simulator [9, 27]. The environment dimensions are 400 × 400 × 25 with three distributed delivery hubs positioned to create realistic traffic patterns with both intra-hub and inter-hub missions. We evaluate three configurations with {0, 2, 4} temporal NFZs, where each NFZ consists of 500-2000 voxels positioned in the upper airspace ($z \geq 12$) to avoid ground obstacles. NFZ activation intervals are randomly distributed between 0-1800 seconds (30 min) with durations of 200-600 seconds (3-10 min), simulating dynamic airspace restrictions. Agent operations span 0-900 seconds (15 min)

Table 3: Success rate (%) and runtime (s, mean \pm std) for different algorithms across NFZ settings and agent counts.

| NFZ | Agents | DTAPP-IICR | | DTAPP-IICR (W/O) | | PP | | PP (W/O) | |
|-----|--------|--------------|-------------------------------------|------------------|--------------------|---------|--------------------|----------|------------------|
| | | Success | Runtime | Success | Runtime | Success | Runtime | Success | Runtime |
| 0 | 10 | 100.0 | 14.5 \pm 8.7 | 100.0 | 28.8 \pm 17.1 | 90.0 | >44.2 \pm 14.7 | 90.0 | >67.7 \pm 27.2 |
| | 50 | 100.0 | 91.7 \pm 14.7 | 100.0 | 191.1 \pm 34.4 | 50.0 | >1142.8 \pm 69.1 | 0.0 | 1200.0 \pm 0.0 |
| | 100 | 100.0 | 192.8 \pm 27.2 | 100.0 | 412.3 \pm 63.5 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 |
| | 200 | 100.0 | 401.2 \pm 34.3 | 100.0 | 839.5 \pm 85.9 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 |
| | 500 | 100.0 | 980.9 \pm 67.7 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 |
| 2 | 10 | 100.0 | 19.0 \pm 7.9 | 100.0 | 42.4 \pm 17.8 | 95.0 | 54.2 \pm 15.4 | 95.0 | 95.9 \pm 31.7 |
| | 50 | 100.0 | 117.6 \pm 22.4 | 100.0 | 258.4 \pm 57.9 | 15.0 | 1180.5 \pm 40.4 | 0.0 | 1200.0 \pm 0.0 |
| | 100 | 100.0 | 253.8 \pm 33.2 | 100.0 | 557.9 \pm 79.9 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 |
| | 200 | 100.0 | 509.7 \pm 43.1 | 95.0 | >1065.6 \pm 84.3 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 |
| | 500 | 89.0 | 1076.4 \pm 51.7 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 |
| 4 | 10 | 100.0 | 19.7 \pm 10.1 | 100.0 | 53.9 \pm 24.8 | 95.0 | 54.4 \pm 18.6 | 95.0 | 99.0 \pm 38.1 |
| | 50 | 100.0 | 116.9 \pm 22.2 | 100.0 | 313.2 \pm 55.6 | 5.0 | 1182.6 \pm 52.9 | 0.0 | 1200.0 \pm 0.0 |
| | 100 | 100.0 | 235.8 \pm 37.1 | 100.0 | 620.2 \pm 89.7 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 |
| | 200 | 100.0 | 490.4 \pm 35.9 | 40.0 | >1163.1 \pm 54.1 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 |
| | 500 | 30.0 | 1195.4 \pm 13.7 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 | 0.0 | 1200.0 \pm 0.0 |

with heterogeneous properties including speeds of 1.5-3.0 m/s and radius of 1.0-2.0 m. For each configuration, we generate 20 scenarios and test four algorithm variants: PP, PP (W/O), DTAPP-IICR, and DTAPP-IICR (W/O) across varying agent counts (10, 50, 100, 200, 500). The time limit is fixed at 1200s (20 min). The results are presented in Table 3. Without NFZ, DTAPP-IICR maintains 100% success rates across all configurations while DTAPP-IICR (W/O) fails to complete 500-agent operations. In contrast, PP algorithms show significant scalability limitations, achieving only 50% success at 50 agents with no NFZs and failing completely at large scale, while PP (W/O) fails at 50+ agents across all NFZ configurations. As the number of NFZs increases, the runtime of DTAPP-IICR also increases, reflecting the additional computational effort required for safe navigation around restricted zones. For instance, with 4 NFZs, DTAPP-IICR achieves a 30% success rate at 500 agents, indicating that the algorithm remains functional but requires more time. Increasing the allowed runtime, which is feasible in UTM scenarios, would likely improve success rates further. Overall, the pruning optimization provides substantial improvements, with DTAPP-IICR consistently outperforming DTAPP-IICR (W/O) by 2-3 times in runtime. These results indicate that DTAPP-IICR’s soft constraint handling and iterative conflict resolution approach provides better scalability for large-scale and dynamic urban environments compared to traditional priority-based methods with hard obstacle treatment, which become impractical beyond small-scale deployments.

6 CONCLUSION AND FUTURE WORK

This work addressed the critical challenge of scalable preflight planning for heterogeneous UAV fleets operating in shared, dynamic airspace with temporal NFZs. We introduced DTAPP-IICR, a novel framework that integrates urgency-aware prioritization, iterative conflict resolution, and efficient 4D path planning. Our key contributions include SFIPP-ST, a single-agent planner that handles

temporal NFZs and hard obstacles while supporting heterogeneous UAV profiles (varying speeds, sizes, and start times) and modeling inter-agent conflicts through continuous geometric cost functions. Furthermore, our iterative conflict resolution mechanism leverages LNS guided by geometric conflict graphs to systematically resolve residual conflicts, while directional pruning accelerates search in 3D grids by up to 50% without sacrificing completeness.

Experimental validation across Monte Carlo simulations and realistic urban scenarios demonstrates that DTAPP-IICR’s superior scalability and robustness compared to the previous batch CBS/ECBS method in the UTM context. These results establish DTAPP-IICR as a foundational solution for safe, efficient UAV integration in dense urban airspace under regulatory constraints.

While our current method prioritizes finding conflict-free solutions, future work will explore advanced neighborhood selection strategies inspired by LNS [17, 29, 33]. This includes designing a set of heuristics tailored for 4D environments and an adaptive learning mechanism, such as a multi-armed bandit (MAB), to dynamically choose the most effective heuristic during the search [24, 25]. This will allow the planner to not only ensure safety but also optimize for operational efficiency. Another promising direction would be to integrate SDP/HMAPP [32, 34] techniques that decompose the environment into smaller sub-regions, enabling scalable planning across very large fleets.

ACKNOWLEDGMENTS

This work is supported in part by Sweden’s Innovation Agency Vinnova (Projects 2022-02671 DyMuDRoP and 2024-01322 AHA-IMPULS), the Excellence Center at Linköping-Lund in Information Technology (ELLIIT), Ericsson Telecomunicações Ltda., and the São Paulo Research Foundation (FAPESP), grant 2021/00199-8, CPE SMARTNESS.

REFERENCES

- [1] FAA Aerospace Forecast Fiscal Years 2024–2044. 2024. *Emerging Aviation Entrants: Unmanned Aircraft System and Advanced Air Mobility*. FAA Aerospace Forecast, Washington, DC. <https://www.faa.gov/dataresearch/aviation/aerospaceforecasts/uas-and-aam.pdf>
- [2] European Union Aviation Safety Agency. 2019. *Commission Implementing Regulation (EU) 2019/947 of 24 May 2019 on the rules and procedures for the operation of unmanned aircraft*. Official Journal of the European Union, European Union. http://data.europa.eu/eli/reg_impl/2019/947/oj
- [3] European Union Aviation Safety Agency. 2021. *Commission Implementing Regulation (EU) 2021/664 of 22 April 2021 on a regulatory framework for the U-space*. Official Journal of the European Union, European Union. http://data.europa.eu/eli/reg_impl/2021/664/oj
- [4] Bing Ai, Jiuchuan Jiang, Shoushui Yu, and Yichuan Jiang. 2021. Multi-Agent Path Finding with heterogeneous edges and roundtrips. *Knowledge-Based Systems* 234 (2021), 107554. <https://doi.org/10.1016/j.knsys.2021.107554>
- [5] Federal Administration Aviation. 2023. *FAA Aerospace Forecast Fiscal Years 2023–2043*. Federal Administration Aviation, Washington, DC. <https://www.faa.gov/dataresearch/aviation/aerospaceforecasts/aa-aerospace-forecast-fy-2023-2043>
- [6] Max Barer, Guni Sharon, Roni Stern, and Ariel Felner. 2014. Suboptimal Variants of the Conflict-Based Search Algorithm for the Multi-Agent Pathfinding Problem. In *Proceedings of the Seventh Annual Symposium on Combinatorial Search (SoCS)*. SoCS, NLD, 961–962. <https://doi.org/10.1609/socs.v5i1.18315>
- [7] Igor Bisio, Chiara Garibotto, Halar Haleem, Fabio Lavagetto, and Andrea Sciarone. 2024. RF/WiFi-based UAV surveillance systems: A systematic literature review. *Internet of Things* 26 (2024), 101201. <https://doi.org/10.1016/j.iot.2024.101201>
- [8] Eli Boyarski, Ariel Felner, Roni Stern, Guni Sharon, David Tolpin, Oded Betzalel, and Eyal Shimony. 2015. ICBS: Improved Conflict-Based Search Algorithm for Multi-Agent Pathfinding. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*. IJCAI, South America, 740–746. <https://www.ijcai.org/Proceedings/15/Papers/110.pdf>
- [9] Mauricio Rodriguez Cesen, Ariel Góes de Castro, Ibini Santana, Ramon Fontes, Fabricio R. Cesen, and Christian Esteve Rothenberg. 2025. UNetyEmu: Unity-based simulator for aerial and non-aerial vehicles with integrated network emulation. In *Anais Estendidos do XLIII Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (Natal/RN)*. SBC, Porto Alegre, RS, Brasil, 100–111. https://doi.org/10.5753/sbrc_estendido.2025.7122
- [10] Ming Chen, Ning He, and Chen Hong. 2024. Continuous Multi-Agent Path Finding for Drone Delivery. In *Proceedings of the 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (PRCV)*. Springer, Singapore, 129–141. https://doi.org/10.1007/978-981-97-8795-1_9
- [11] M. Erdmann and T. Lozano-Perez. 1986. On multiple moving objects. In *Proceedings. 1986 IEEE International Conference on Robotics and Automation*, Vol. 3. IEEE, San Francisco, CA, 1419–1424. <https://doi.org/10.1109/ROBOT.1986.1087401>
- [12] Hossein Eskandaripour and Enkhsaikhan Boldsaikhan. 2023. Last-Mile Drone Delivery: Past, Present, and Future. *Drones* 2, 7 (2023), 77. <https://doi.org/10.3390/drones7020077>
- [13] Ridha Guebsi, Sonia Mami, and Karem Chokmani. 2024. Drones in Precision Agriculture: A Comprehensive Review of Applications, Technologies, and Challenges. *Drones* 8, 11 (2024), 686. <https://doi.org/10.3390/drones8110686>
- [14] Florence Ho, Artur Alves Gonçalves, Bastien Rigault, Ruben Geraldes, Alexandre Chicharo, Marc Cavazza, and Helmut Prendinger. 2021. Multi-Agent Path Finding in Unmanned Aircraft System Traffic Management With Scheduling and Speed Variation. *IEEE Intelligent Transportation Systems Magazine* 14 (2021), 2–15.
- [15] Florence Ho, Ana Salta, Ruben Geraldes, Artur Goncalves, Marc Cavazza, and Helmut Prendinger. 2019. Multi-Agent Path Finding for UAV Traffic Management. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 131–139. <https://doi.org/10.5555/3306127.3331684>
- [16] Taoan Huang, Jiaoyang Li, Sven Koenig, and Bistra Dilkina. 2022. Anytime Multi-Agent Path Finding via Machine Learning-Guided Large Neighborhood Search. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 9 (Jun. 2022), 9368–9376. <https://doi.org/10.1609/aaai.v36i9.21168>
- [17] Jiaoyang Li, Zhe Chen, Daniel Harabor, Peter J. Stuckey, and Sven Koenig. 2021. Anytime Multi-Agent Path Finding via Large Neighborhood Search. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI)*. Association for the Advancement of Artificial Intelligence (AAAI), Montreal, Canada, 4127–4135. <https://doi.org/10.24963/ijcai.2021/568>
- [18] Jiaoyang Li, Zhe Chen, Daniel Harabor, Peter J. Stuckey, and Sven Koenig. 2022. MAPF-LNS2: Fast Repairing for Multi-Agent Path Finding via Large Neighborhood Search. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 9 (Jun. 2022), 10256–10265. <https://doi.org/10.1609/aaai.v36i9.21266>
- [19] Jiaoyang Li, Wheeler Ruml, and Sven Koenig. 2021. EECBS: A Bounded-Suboptimal Search for Multi-Agent Path Finding. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 14 (May 2021), 12353–12362. <https://doi.org/10.1609/aaai.v35i14.17466>
- [20] Mingyang Lyu, Yibo Zhao, Chao Huang, and Hailong Huang. 2023. Unmanned Aerial Vehicles for Search and Rescue: A Survey. *Remote Sensing* 15, 13 (2023), 3266. <https://doi.org/10.3390/rs15133266>
- [21] Hang Ma, Daniel Harabor, Peter J. Stuckey, Jiaoyang Li, and Sven Koenig. 2019. Searching with Consistent Prioritization for Multi-Agent Path Finding. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (Jul. 2019), 7643–7650. <https://doi.org/10.1609/aaai.v33i01.33017643>
- [22] Hang Ma, Jiaoyang Li, T.K. Satish Kumar, and Sven Koenig. 2017. Lifelong Multi-Agent Path Finding for Online Pickup and Delivery Tasks. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (São Paulo, Brazil) (AAMAS '17)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 837–845.
- [23] Civil Aviation Administration of China. 2023. Statistical Bulletin of Civil Aviation Industry Development in 2023. *China Civil Aviation Annual Report* 1, 1 (2023), 12. https://www.caac.gov.cn/English/Research/Reports/Statistical/202412/t20241211_226085.html
- [24] Thomy Phan, Taoan Huang, Bistra Dilkina, and Sven Koenig. 2024. Adaptive Anytime Multi-Agent Path Finding Using Bandit-Based Large Neighborhood Search. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 16 (Mar 2024), 17514–17522. <https://doi.org/10.1609/aaai.v38i16.29701>
- [25] Thomy Phan, Benran Zhang, Shao-Hung Chan, and Sven Koenig. 2025. Anytime Multi-Agent Path Finding with an Adaptive Delay-Based Heuristic. *Proceedings of the AAAI Conference on Artificial Intelligence* 39, 22 (Apr. 2025), 23286–23294. <https://doi.org/10.1609/aaai.v39i22.34495>
- [26] Mike Phillips and Maxim Likhachev. 2011. SIPP: Safe interval path planning for dynamic environments. In *2011 IEEE International Conference on Robotics and Automation*. IEEE, Shanghai, China, 5628–5635. <https://doi.org/10.1109/ICRA.2011.5980306>
- [27] Mauricio Rodriguez, Ariel Goes de Castro, Ramon Fontes, Fabricio Rodriguez, and Christian Rothenberg. 2025. An Integrated Framework for Network Emulation and Multi-vehicle Algorithm Testing. In *Proceedings of the ACM SIGCOMM 2025 Posters and Demos (Coimbra, Portugal) (ACM SIGCOMM Posters and Demos '25)*. Association for Computing Machinery, New York, NY, USA, 167–169. <https://doi.org/10.1145/3744969.3748436>
- [28] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R. Sturtevant. 2015. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence* 219 (2015), 40–66. <https://doi.org/10.1016/j.artint.2014.11.006>
- [29] Paul Shaw. 1998. Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems. In *Principles and Practice of Constraint Programming – CP98*, Michael Maher and Jean-Francois Puget (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 417–431. https://doi.org/10.1007/3-540-49481-2_30
- [30] Roni Stern. 2019. *Multi-Agent Path Finding – An Overview*. Springer International Publishing, Cham, 96–115. https://doi.org/10.1007/978-3-030-33274-7_6
- [31] Single European Sky ATM Research 3 Joint Undertaking. 2017. *European drones outlook study – Unlocking the value for Europe*. Publications Office of the European Union, European Union. <https://doi.org/10.2829/085259>
- [32] Christopher Wilt and Adi Botea. 2014. Spatially Distributed Multiagent Path Planning. *Proceedings of the International Conference on Automated Planning and Scheduling* 24 (05 2014), 332–340. <https://doi.org/10.1609/icaps.v24i1.13618>
- [33] Setyo Tri Windras Mara, Rachmadi Norcahyo, Panca Jodiawan, Luluk Lusiantoro, and Achmad Pratama Rifai. 2022. A survey of adaptive large neighborhood search algorithms and applications. *Computers & Operations Research* 146 (2022), 105903. <https://doi.org/10.1016/j.cor.2022.105903>
- [34] Han Zhang, Mingze Yao, Ziang Liu, Jiaoyang Li, Lucas Terr, Shao-Hung Chan, T. K. Satish Kumar, and Sven Koenig. 2021. A Hierarchical Approach to Multi-Agent Path Finding. *Proceedings of the International Symposium on Combinatorial Search* 12, 1 (Jul 2021), 209–211. <https://doi.org/10.1609/socs.v12i1.18586>