

Airlift Challenge: A Competition for Optimizing Cargo Delivery

Adis Delanovic*, Carmen Chiu*, and Andre Beckus

Air Force Research Lab, Rome, NY 13441

{adis.delanovic, carmen.chiu, andre.beckus}@us.af.mil

* Authors contributed equally to this work

Abstract

Airlift operations require the timely distribution of various cargo, much of which is time sensitive and valuable. However, these operations have to contend with sudden disruptions from weather and malfunctions, requiring immediate rescheduling. The Airlift Challenge competition seeks possible solutions via a simulator that provides a simplified abstraction of the airlift problem. The simulator uses an OpenAI gym interface that allows participants to create an algorithm for planning agent actions. The algorithm is scored using a remote evaluator against scenarios of ever-increasing difficulty. The second iteration of the competition was underway from November 2023 to April 2024. In this paper, we describe the competition and simulation environment. As a step towards applying generalized planning techniques to the problem, we present a temporal PDDL domain for the Pickup and Delivery Problem, a model which lies at the core of the Airlift Challenge.

Introduction

An airlift operation comprises movement of cargo, supplies and/or personnel often within strict time limits. This is hampered by various factors, ranging from excessive amounts of cargo/personnel to conditions at the airports themselves, e.g., very limited space for taxiing or parking (Owen, 1997). Usually, the air cargo consists of time sensitive and valuable supplies. Even just a small change can result in a large amount of disruption, requiring extensive replanning which takes time and resources. This drives an increased need for automation within the planning process. To answer this need, we present the Airlift Challenge competition¹.

The airlift planning problem is difficult because of the number of constraints and factors. Often, Mixed Integer Linear Programming (MILP) solutions are utilized, but these require recalculation of the entire solution when a disruption occurs (Bertsimas, Chang, Mišić, & Mundru, 2019). However, reinforcement learning has also shown promise in quick replanning. Even if a solver can find an optimal solution to a problem, this solution must be created quickly in order to deal with dynamic elements.

The first iteration of the Airlift Challenge was held with the SPIE conference in 2023 and attracted 20 registered users and two high quality submissions that received 1st and

2nd place respectively (Delanovic, et al., 2023). In the remainder of this paper, we describe the second iteration held in association with ICAPS 2024 and concluded in April 2024. We then present a Pickup and Delivery PDDL (Planning Domain Definition Language) domain as a starting point for a potential submission to future IPCs (International Planning Competitions).

Simulation Environment Details

The goal of the competition’s simulation environment is to provide a simplified model that captures the major factors influencing an airlift operation. This approach is based off the Flatlands competition which was focused on train vehicle routing (Mohanty, et al., 2020). The Flatlands competition focused on “flattening”, or simplifying, a complex environment into a set of the most important parameters. The advantage of a simplified environment is it abstracts away details in the environment, while the design of the starter kit allows many types of solutions to be used with it.

The model consists of an air network graph, where each node corresponds to an airport and each edge denotes the route that agents can traverse. The agents are airplanes that travel preset routes between airports. The agents move the same predetermined distance at every time step based on the speed of their airplane type. Airplanes are limited by fuel, routes available, maximum loading weight and speed. Airports themselves can only process a certain number of airplanes at a time, creating possible bottlenecks. A summary of the simulated objects and associated parameters is shown in Table 1.

¹ See main website <https://airliftchallenge.com/>

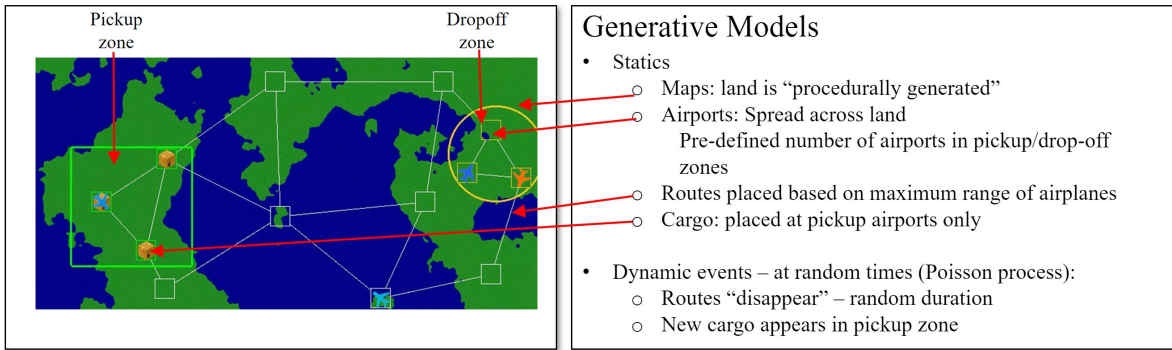


Figure 1: Simulator visualization highlighting aspects of generative models

Object	Parameters
Airport	<ul style="list-style-type: none"> Location Working capacity (# of airplanes that can process/load/unload at same time)
Route	<ul style="list-style-type: none"> Flight Cost (determined based on distance) Flight Time (determined based on distance and flight speed) Start/end airports Availability status (routes may become unavailable at random intervals)
Airplane	<ul style="list-style-type: none"> Cargo weight capacity Time to process Route Map: Contains all routes available to travel Processing priority (determines order in which planes process/load/unload)
Cargo	<ul style="list-style-type: none"> Weight Source/destination airports Soft delivery deadline Hard delivery deadline

Table 1: Model Parameters

At each time-step all agents receive an observation and corresponding reward. The agents then decide what action to take next using the policy provided by the competition participant, and pass it to the environment. This continues until either 1) all the cargo is delivered, or 2) a maximum number of time steps has elapsed. For the competition, each scenario is limited to 5,000 time steps.

An agent can be in one of four states at any given time: waiting, processing, ready-for-takeoff, or moving. An agent is put into a waiting state after landing at an airport. While in the waiting state it is held in a priority queue. Each airport has a finite processing capacity. When an airport has capacity and the airplane is next in the priority queue, it will transition into the processing state. While in this state, cargo loading or unloading actions will take place. After processing cargo, the airplane will then transition to the ready-

for-takeoff state. An airplane will remain in the ready-for-takeoff state until it receives actions specifying its next destination. Then, the airplane will transition to the moving state, restarting the entire process over again.

Plans are implemented by issuing actions to change the state of an airplane. We briefly summarize the actions and their associated preconditions and effects in Table 2. The *Process* action transitions the airplane from the waiting state into the processing state. If processing capacity is reached, the airplane waits in priority queue based on a priority given by the agent policy. Processing must be performed at least once after landing, to represent re-fueling and other maintenance, even if there is no cargo to load/unload. Concurrent with processing, the *Load* and *Unload* actions may be used to transfer cargo. The *Takeoff* action receives a destination and transitions from ready-for-takeoff into a moving state. A valid destination is defined as any neighboring node of the current airport with respect to the route graph. Note that the simulator offers a consolidated action space which allows actions to be issued while in flight to the next airport, with sequencing at the airport handled automatically. More details are provided in the competition documentation².

Action	Preconditions	Effects
<i>Process</i>	There is available capacity to process.	Sends airplane into processing state.
<i>Load</i>	New cargo will not exceed airplane capacity. Airplane is in processing state.	Loads a piece of cargo.
<i>Unload</i>	Airplane has cargo onboard and is in processing state.	Unloads a piece of cargo.
<i>Takeoff</i>	Airplane has finished processing.	Sends airplane into moving state.

Table 2: Key components of action. All actions are undertaken while an airplane is landed at an airport.

² See https://airliftchallenge.com/chapters/ch2_model/main.html and https://airliftchallenge.com/chapters/ch3_interface/main.html for more details.

A malfunction generator will randomly cause disruptions to airplane routes, forcing the planner to reschedule. Malfunction are generated according to a Poisson distribution with the mean number of malfunctions per time step λ predetermined before generating the scenario. Throughout each scenario, cargo will be dynamically generated at random times. In addition, the simulator has the option to place cargo and airplanes in random initial places.

Cargo is placed within a designated area called the “pick up zone”. The pick-up zone contains a subset of airports. The cargo must be picked up and delivered to the “drop off zone” that is generally on the other side of the map. An illustration of these zones is shown in Figure 1. The drop-off and pick-up zones contain a subset of airports within a specific area.

Cargo is assigned a weight, a soft deadline and a hard deadline. The weight of the cargo varies. As shown in Figure 2, cargo is considered to be on-time if it is delivered within the soft deadline. If the cargo is delivered past the soft deadline, it is considered late. If the agents are unable to deliver the cargo within the hard deadline, the cargo is considered missed, and the agent is penalized heavily for it.

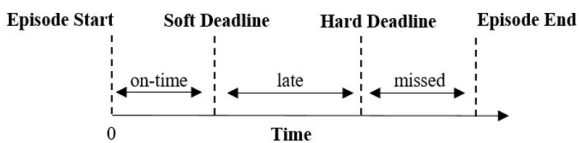


Figure 2: Deadlines for delivery.

Maps are generated using the Perlin noise algorithm, whereby a certain threshold value decides if an area is considered land or water, and airports are placed uniformly at random on the landmasses (Perlin, 1985). The generator also ensures that airports are not placed too close to each other. An example of a scenario can be seen in Figure 1.

The environment supports multiple airplane types, since this is required for successful airlift operations given the variety of airplanes and potential landing areas. Airplane types have speed, maximum travel range and maximum carrying capacity. The edges on the graph network correspond to the maximum range of the airplane type. The disparate graph networks for each airplane type are combined to form a single multi-graph where the edges contain data on which airplane type can traverse it. Certain airplane types may not be able to reach all airports and would have to work together with the others to complete cargo delivery.

For the ICAPS Airlift Challenge we introduced agent priority. Given the finite processing capacity of airports, this addition enables solutions to prioritize agents for processing, regardless of their arrival time. Proper prioritization of agents provides the ability to improve cargo delivery outcomes from missing the hard deadline to late or even on-time.

Baselines

There are three baselines provided. These are the random agent, shortest paths as well as a MILP. The random agent samples the observation and generates a random valid action. The shortest paths agent only traverses airports using the shortest path available. If a route along a path becomes unavailable, the agent will re-route to the next shortest path.

The MILP solution uses constraints as formulated in (Bertsimas, Chang, Mišić, & Mundru, 2019). It utilizes the GNU Linear Programming Kit (GLPK), an open-source solver that can be used for linear programming or MILPs (Makhorin, 2012). This solution does not scale well when additional agents are added to the scenarios and currently only runs with one agent and simple scenarios.

Scoring

The competition evaluator is hosted on Amazon Web Services (AWS). The services do not allow communication with any outside resources. This ensures that all participants receive an even playing field and that there is no bias in scoring the submission. A participant will use the CodaLab platform to make a submission (Pavao, et al., 2022)³. The submission is immediately sent to AWS, where it will start the evaluation process. Results are shown on a leaderboard.

For each phase of the competition, over four hundred scenarios are generated to score the submitted algorithms. Scenarios increase in difficulty. This is accomplished by changing “static” as well as “dynamic” parameters. The static parameters include the quantity of agents, airports and cargo. The dynamic parameters include the rate at which new cargo is generated, as well as how often routes become unavailable. The evaluation scenarios are not released to the public, although a test set is provided.

A submission will run until it either reaches the allotted four-hour time limit, exceeds an overall 30% missed deliveries, or completes all scenarios within the time constraint. Upon either of these occurring, a score is generated and placed on the leaderboard. Figure 3 shows visualizations of the environment for four tests, ranging from the simplest to

³ See <https://codalab.lisn.upsaclay.fr/competitions/16103>

most difficult scenarios. The blue and white colored edges represent different airplane types.

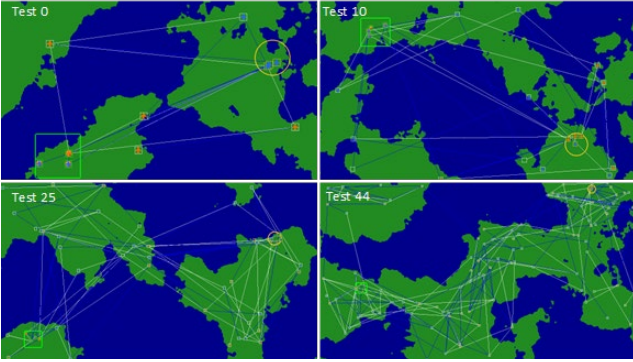


Figure 3: Example scenarios exhibiting difficulty increase.

To generate an overall score for any submission, we first assign an episode score. For each level that the submission completes, a score is assigned based on missed cargo, late cargo and total flight cost. An overall score is calculated by summing the normalized score over all tests and levels.

A starter-kit is provided and hosted on the competition GitHub page with detailed description on what a submission should include⁴. Additionally, instructions are included on generating custom-made scenarios, and we provide a large number of pre-generated scenarios stored in serialized form as pickle (.pkl) files.

Current Competition

The competition hosted on CodaLab began in November 2023, attracting 40 registered users. The competition was divided into Phase 1 and Phase 2. Phase 1 of the competition utilized one airplane type and assisted the team in debugging any environment issues that the participants may encounter. Phase 1 concluded on January 12th, 2024, followed immediately by the start of Phase 2. Unlike Phase 1, Phase 2 introduces two airplane types: a larger aircraft with greater carrying capacity and a smaller one with a reduced carrying capacity.

The competition concluded on April 18th, 2024, and we are happy to announce the following winners of the competition:

- 1st place: Team JLH4E – 1201.01 points
- 2nd place: Team FKIE – 507.70 points
- 3rd place: The Down Underdogs – 372.31 points

Detailed results and approaches are deferred to an extended version of this paper⁵.

⁴ See <https://github.com/airlift-challenge/airlift-starter-kit>

⁵ We plan to post the extended version as a revision at the following link: <https://arxiv.org/abs/2404.17716>

PDDL Domain: Pickup and Delivery with Time Windows

As a first step towards applying AI Planners to the Airlift Planning problem, we have developed a PDDL domain that captures some core deterministic aspects of the airlift problem (Fox & Long, 2003). This model essentially embodies a Pickup and Delivery with Time Windows (PDPTW) problem (Berbeglia, Laport, & Cordeau, 2010). We refer to the new domain as the PDPTW domain.

The basic pickup and delivery problem (without time windows) has been explored through the LOGISTICS domain in IPC going back to the AIPS-98 Planning Competition, see (Long, et al., 2000) and (Seipp, Fawcett, Masataro, Muise, & Blonet, 2022). A domain for the closely related Vehicle Routing Problem has also been proposed in (Cheng & Gao, 2014). However, these domains do not include temporal aspects.

The PDPTW domain incorporates temporal elements in the form of durative actions (movement, loading, an unloading) and time windows via Timed Initial Literals (Haslum, Lipovetzky, Magazzeni, & Muise, 2019). This model can be considered a temporal variant of the TRANSPORT_{∞+} planning task introduced in (Helmert, 2008), which represents a transportation problem with no fuel constraints, no vehicle capacity constraints, and one vehicle type. For simplicity, we initially consider a fully connected transport graph with uniform travel times between all locations. We also leave multiple vehicles and action costs as a future extension.

It could be interesting to use PDDL solvers in order to compare the results between the OpenAI version and the PDDL version, or even to solve sub-components of the overall problem. There is promise in not only using other generalized planners but possibly using Large Language Models (LLMs) for planning tasks using PDDL. See for example (Silver, et al., 2023) and (Valmeekam, Sreedharan, Marquez, Olmo, & Kambhampati, 2013)

In order to more closely match the current environment, probabilistic elements of malfunctioning routes and dynamic cargo generation would have to be added. An ideal candidate for expressing this model may be RDDDL (Relational Dynamic Influence Diagram Language)⁶. RDDDL could also model rewards which are an important element of the OpenAI Gym interface for techniques such as reinforcement learning.

The PDPTW domain is listed in the appendix and is also available under the Airlift Challenge Github account⁷.

⁶ See https://users.cecs.anu.edu.au/~ssanner/IPPC_2011/RDDL.pdf

⁷ See <https://github.com/airlift-challenge/PDDL-Domain>

Future Work

Both iterations of the Airlift Challenge have brought in several high-quality submissions. Despite this success, a number of future formats are possible which may allow better or different solutions as well. One of the criticisms we had in the past from participants interested in reinforcement learning was that the randomized environments create entirely new air route maps for each episode, which required reinforcement learning algorithms to work in a very general manner. In fact, this was not as indicative of real-world situations where airports would stay in the same areas and the air network would remain relatively unchanged.

The PDDL domain may serve as an initial step towards a domain submission for a future IPC. It could allow for general planners as possible solutions for this challenge.

References

- Berbeglia, G., Laport, G., & Cordeau, J.-F. (2010). Dynamic pickup and delivery problems. *European Journal of Operational Research*, 202(1), 8-15.
- Bertsimas, D., Chang, A., Mišić, V. V., & Mundru, N. (2019). The Airlift Planning Problem. *Transportation Science*, 773-795. Retrieved November 30, 2021
- Cheng, W., & Gao, Y. (2014). Using PDDL to Solve Vehicle Routing Problems. *Intelligent Information Processing VII*, 207-215.
- Delanovic, A., Chiu, C., Kolen, J., Gnanasekaran, A., Surana, A., Srivastava, K., . . . Beckus, A. (2023). Results of the airlift challenge: a multi-agent AI planning competition. *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*. SPIE.
- Fox, M., & Long, D. (2003). PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains. *Journal of Artificial Intelligence Research*, 61-124.
- Haslum, P., Lipovetzky, N., Magazzeni, D., & Muise, C. (2019). *An Introduction to the Planning Domain Definition Language*. Switzerland: Springer Cham.
- Helmert, M. (2008). *Domain Complexity and Heuristic Decomposition*. Berlin: Springer-Verlag Berlin Heidelberg.
- Long, D., Kautz, H., Selman, B., Blonet, B., Geffner, H., Koehler, J., . . . Fox, M. (2000, June 1). The AIPS-98 Planning Competition. *AI Magazine*, 21(2), 13-34.
doi:<https://doi.org/10.1609/aimag.v21i2.1505>
- Makhorin, A. (2012, June 23). *GLPK GNU Linear Programming Kit*. Retrieved from GNU:
<https://www.gnu.org/software/glpk/#introduction>
- Mohanty, S., Nygren, E., Laurent, F., Schneider, M., Scheller, C., Bhattacharya, N., . . . Spigler, G. (2020). Flatland-RL : Multi-Agent Reinforcement Learning on Trains. arXiv:2012.05893.
- Owen, R. C. (1997). The Airlift System: A Primer. *Air Power Journal*.
- Pavao, A., Guyon, I., Letournel, A.-C., Escalante, H., Escalera, S., Thomas, T., & Xu, Z. (2022). *CodaLab Competitions: An open source platform to organize scientific challenges*.
- Perlin, K. (1985). An image synthesizer. (pp. 287-296). ACM SIGGRAPH Computer Graphics.
- Seipp, J., Fawcett, C., Masataro, A., Muise, C., & Blonet, B. (2022). PDDL Generators. Zenodo.
doi:<https://doi.org/10.5281/zenodo.6382173>
- Silver, T., Dan, S., Srinivas, K., Tenenbaum, J. B., Kaelbling, L. P., & Katz, M. (2023). Generalized Planning in PDDL Domains with Large Language Models. *arXiv*.
- Valmeekam, K., Sreedharan, S., Marquez, M., Olmo, A., & Kambhampati, S. (2013). On the Planning Abilities of Large Language Models: A critical Investigation with a Proposed Benchmark. arXiv:2302.06706.

Acknowledgments

We would like to thank ICAPS for allowing us to host the Airlift Challenge competition using their venue, and facilitating outreach to the wider planning community.

We extend our gratitude to the Flatland challenge for its influence on our competition and code base.

We also extend our appreciation and thank you to all the registered users and participants of the challenge, and most importantly, to the winners. Their contributions and bug reports played the most significant role in the success of this competition.

Appendix: PDDL Domain listing

The PDPTW domain incorporates temporal elements from the PDDL 2.1 specification (Fox & Long, 2003). The temporal PDPTW domain is defined as follows:

```
;; Simple Pickup and Delivery with Time Windows (PDPTW) Domain

(define (domain PDPTW)
  (:requirements :strips :typing :durative-actions :timed-initial-literals)
  (:types package
    vehicle - physobj
    place
    physobj - object)
  (:predicates
    (objat ?obj - physobj ?loc - place)
    (in ?pkg - package ?veh - vehicle)
    (available ?pkg - package))

  (:durative-action LOAD
    :parameters (?pkg - package ?vehicle - vehicle ?loc - place)
    :duration (= ?duration 1)
    :condition (and (at start (objat ?pkg ?loc))
      (over all (and (objat ?vehicle ?loc)
        (available ?pkg))))))
    :effect (and (at start (not (objat ?pkg ?loc)))
      (at end (in ?pkg ?vehicle))))

  (:durative-action UNLOAD
    :parameters (?pkg - package ?vehicle - vehicle ?loc - place)
    :duration (= ?duration 1)
    :condition (and (at start (in ?pkg ?vehicle))
      (over all (and (objat ?vehicle ?loc)
        (available ?pkg))))))
    :effect (and (at start (not (in ?pkg ?vehicle)))
      (at end (objat ?pkg ?loc))))

  (:durative-action MOVE
    :parameters (?vehicle - vehicle ?loc-from - place ?loc-to - place)
    :duration (= ?duration 2)
    :condition (and (at start (objat ?vehicle ?loc-from))
      (at end (objat ?vehicle ?loc-to))))
    :effect (and (at start (not (objat ?vehicle ?loc-from)))
      (at end (objat ?vehicle ?loc-to))))
  )
)
```

An example problem is as follows (this problem contains three locations, three packages, and two vehicles):

```
(define (problem PDPTW-3-3-2)
  (:domain PDPTW)
  (:objects
    veh1 veh2 - vehicle
    loc1 loc2 loc3 - place
    pkg1 pkg2 pkg3 - package)

  (:init
    ; Vehicle locations
    (objat veh1 loc1)
    (objat veh2 loc1)
    ; Package pickup locations
    (objat pkg1 loc1)
    (objat pkg2 loc2)
    (objat pkg3 loc3)
    ; Package times available for pickup
    (at 10 (available pkg1))
    (at 1 (available pkg2))
    (at 2 (available pkg3))
    ; Package delivery deadlines
    (at 15 (not (available pkg1)))
    (at 10 (not (available pkg2)))
    (at 7 (not (available pkg3)))
  )

  ; Package delivery locations
  (:goal (and (objat pkg1 loc2)
              (objat pkg2 loc3)
              (objat pkg3 loc1)
            )
  )
)
```