HYPERSOLVER

HYPERSOLVER Presentation

20th November 2025

Toulouse, ENAC premises and on-line

SUPPORTED BY sesar JOINT UNDERTAKING

Co-funded by the European Union

Agenda

- Concept Principles – Gilles Gawinowski
- AI Concept – Matteo Cavaletti
- Technical Platform – Jean-Louis Raoul
- AI Performance Results – Matteo Cavaletti
- DCB – Yutong Chen
- Results & Feedback – Victor Tuvesson

Concept Principles

Gilles Gawinowski, EUROCONTROL

Towards Integrated ATFM & ATC Management

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

Key Points:

- SESAR Exploratory Research Project
- Partners : NEOMETSYS, ENAC, EUROCONTROL, LFV, NTU, Univ. Warwick
- Goal: Reinvent air traffic management by integrating ATFM (Air Traffic Flow Management) and ATC (Air Traffic Control) in a single framework
- Powered by Advanced AI and Deep Reinforcement Learning
- Aligned with SESAR’s vision for digital transformation and Human-AI Teaming

The Problem

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

Current Issues:

- ATFM and ATC handled separately : conflicting strategies, capacity loss
- Airspace saturation and traffic growth
- Sector-based approach limitations:
  - Increased coordination
  - Reduced anticipation horizon
- Rising ATFM regulations and ATCO workload

Core Principles of HYPERSOLVER

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

- **Holistic ATFM + ATC management** in one framework
- **New ATM architecture**: Replace current layers with **Bunching**, **Proximity**, **Conflict**
- Shift from **sector-based** to **flight-centric**
- **Single Controller Operation** (**SCO**): one controller per set of flights
- **Control-by-exception** & real-time synchronization
- Leverage **FMS/EPP** capabilities for precise 4D trajectories

Core Principles of HYPERSOLVER

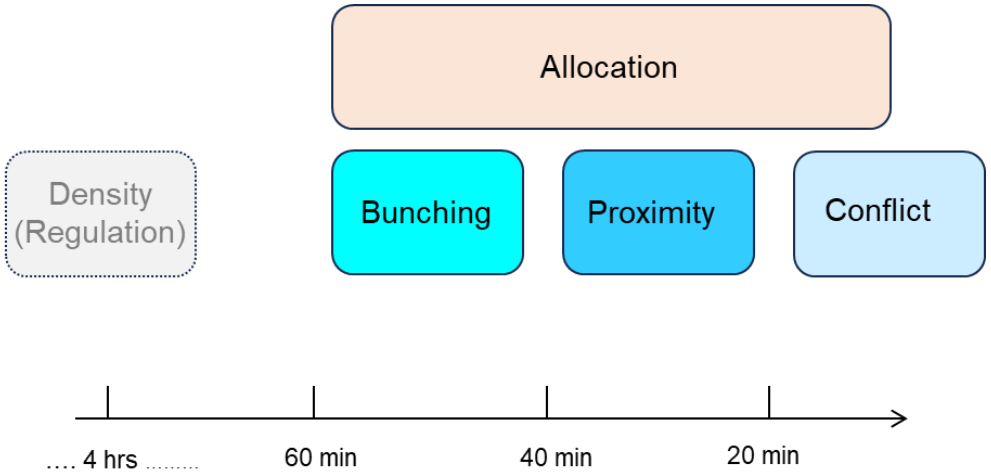

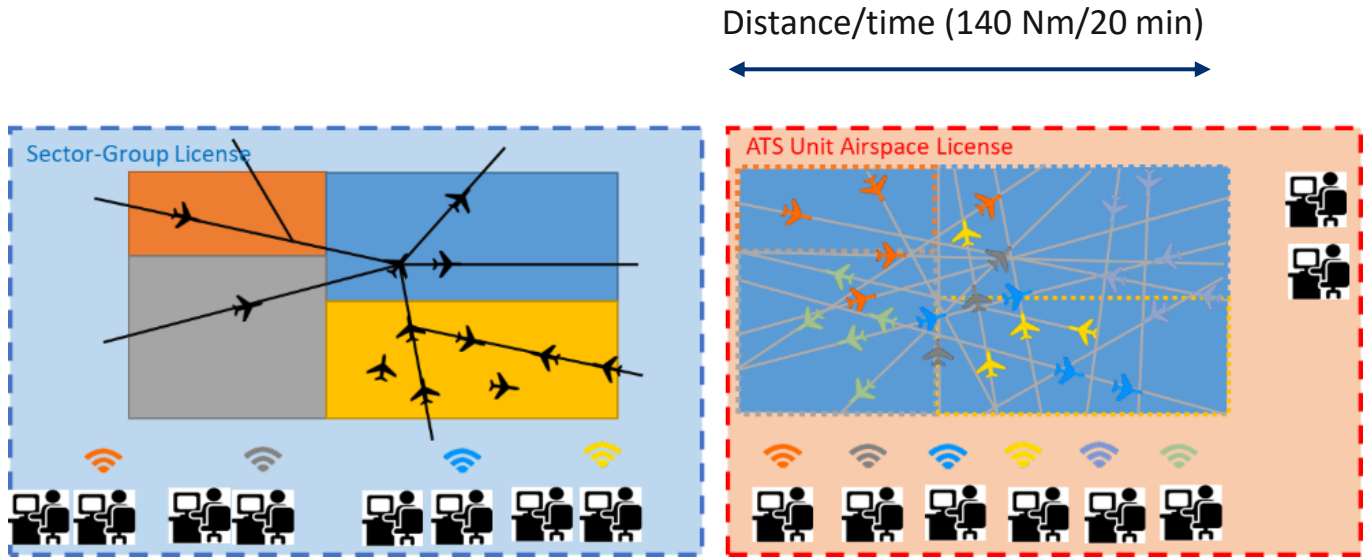

Key Innovations

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

- AI & Deep Reinforcement Learning to anticipate and resolve:
  - Bunching (density)
  - Proximity (conflict prevention)
  - Conflict (tactical resolution)
- Human-AI Teaming: New Supervisor Role
- Collaborative Playbook to adjust goals and AI autonomy
- ATS Unit: large sectorless airspace, dynamic flight allocation to controllers

Key Innovations

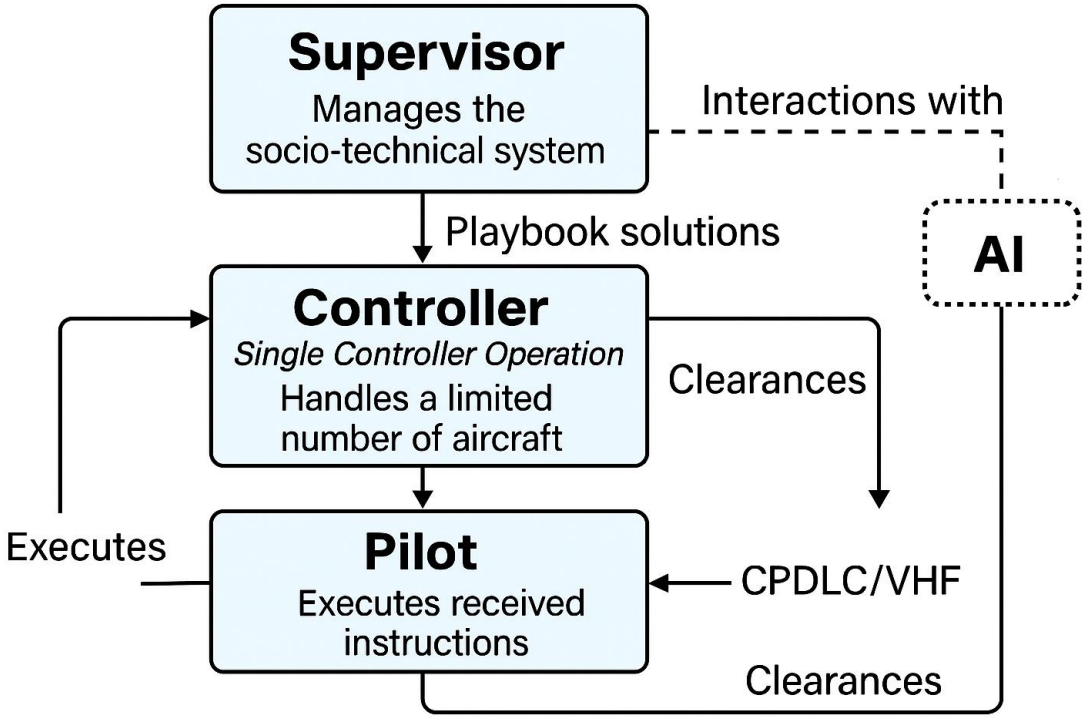

Supervisor: manages the socio-technical system and sends Playbook solutions.

Controller: operates in Single Controller Operation, handles a limited number of aircraft.

Pilot: executes instructions received via CPDLC/VHF.

AI: interacts with both Supervisor and Controller for automation and decision support.

AI concept

Matteo Cavalletti, EUROCONTROL

Fundamental AI concepts for HYPERSOLVER

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

Reinforcement Learning

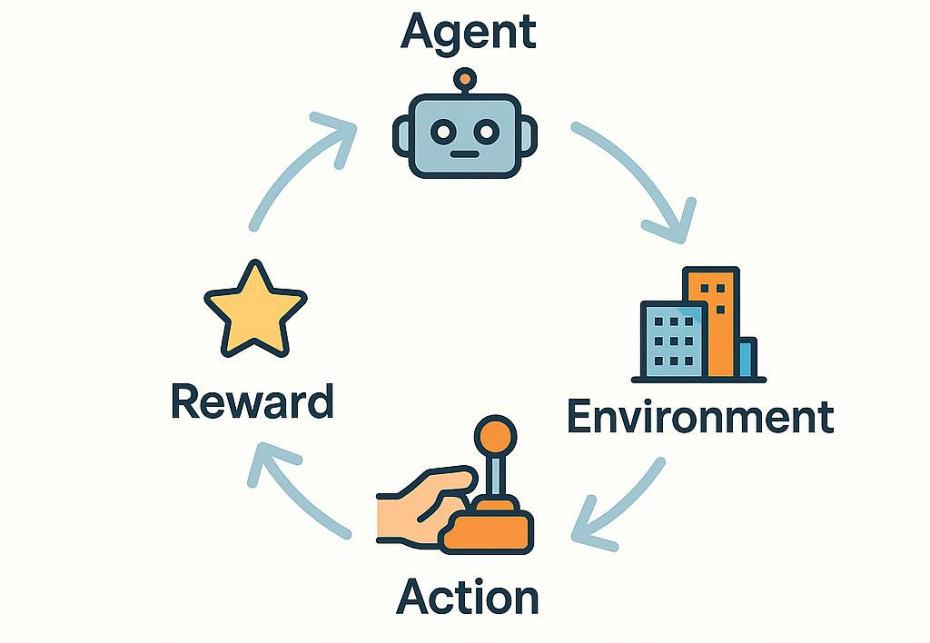

Graph Neural Neworks

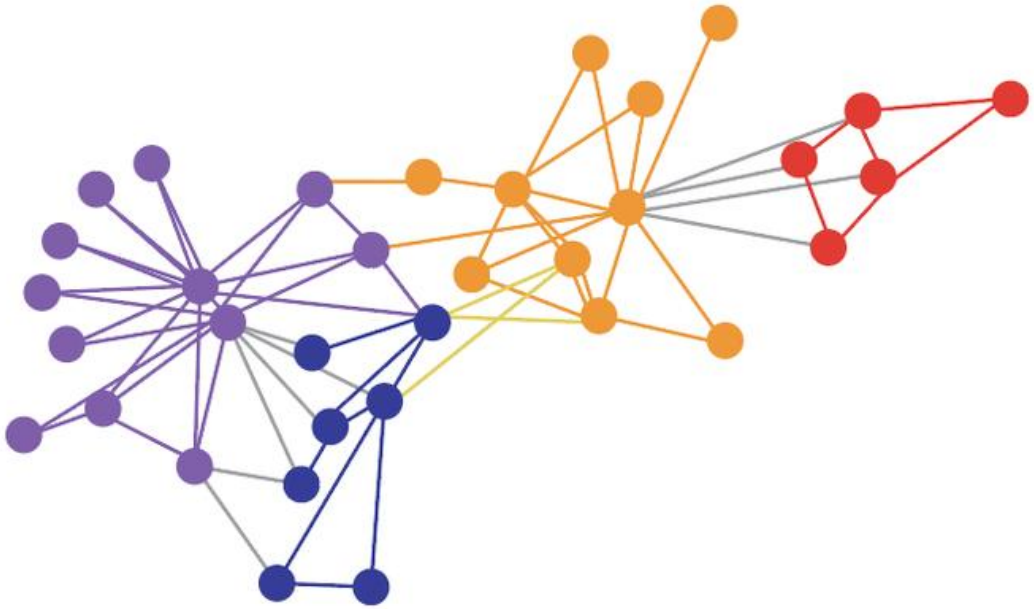

Reinforcement Learning

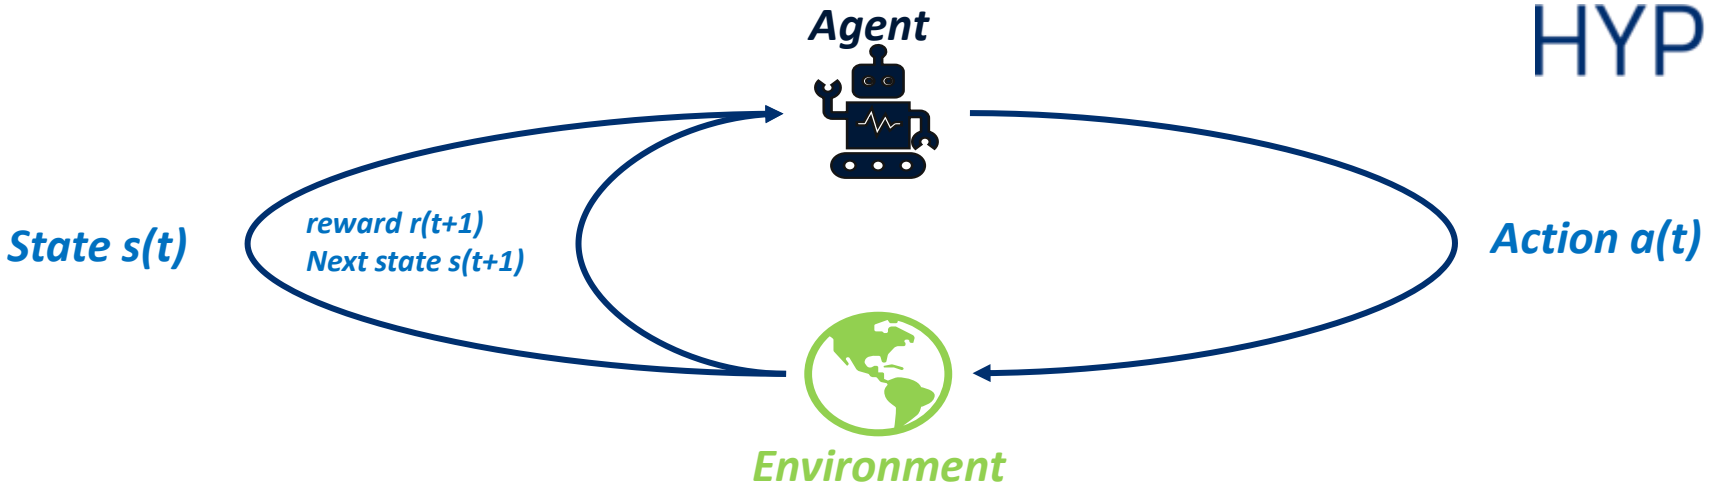

Fundamental Components:

- **Agent:** *An entity capable of performing actions.*
- **Environment:** *The setting in which the agent operates.*
- **Objective:** *The task the agent aims to accomplish.*
- **Reward:** *A numeric signal providing feedback on the agent's actions(the higher the more the action helped to achieve the objective).*
- **Policy:** *The strategy the agent uses to determine its actions.*

HYPERSOLVER:

- **Agent:** *ATCO(mimic system)*
- **Environment:** *controlled sector*
- **Objective:** *avoid conflicts, general optimization*
- **Reward:** *weighted average of conflicts, fuel consumed etc..*
- **Policy:** *The strategy that the ATCO-like system uses to manage the traffic*

*The Goal of **Reinforcement learning(RL)** is to learn the **policy** the **maximizes** the **reward** overtime using the experience gained interacting with the Environment.*

Graph Neural Neworks

*In HYPERSOLVER aircrafts in the space are defined as nodes of a **fully connected graph** where each airplane is a node:*

- *The **nodes features** are the features(e.g. Destination, altitude..) of the airplane*
- *The **edges features** are proximity measures between 2 airplanes*
- *Every node has an edge with every other node(**fully connected graph**)*

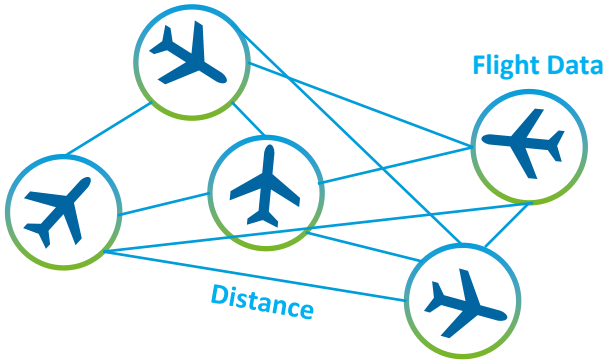

***Graph Neural Networks** are the types of neural networks specifically designed to handle graph structures and relationships.*

Modelling the airspace as a graph:

- *The **Action** is slected per **node**.. the **number** of possible **aircraft** to control is **theoretically infinite***
- *The policy is **spatially agnostic** because it is modeled just considering relative distances between flights*
- *The **Action** is the **next way point** where the aircraft will go before **pointing back to the destination**(close loop)*

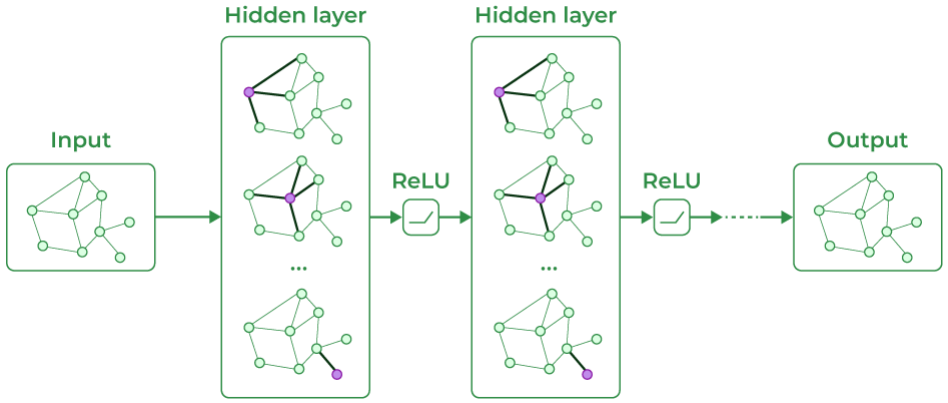

*Combining **RL and GNN** is possible to train a **flexible policy** for a varying number of flights and scenarios*

Technical Platform

Jean-Louis Raoul, Neometsys

Simulation Platform

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

ATSU Geometry

Flight Physics

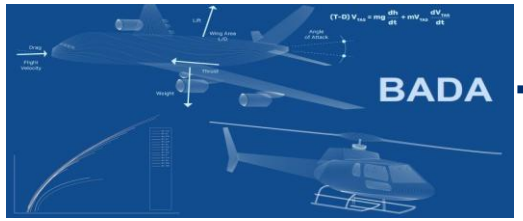

Real past traffic

HyperSolver Simulator

Conflict, Proximity and Bunching actions

Realistic traffic evolution

Reinforcement Learning Agents

ATCO, SUP

ADMIN

POST_OPS

The Simulator choice

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

No ATFM/ATC simulator available

3 options were considered:

- Airsim (Microsoft)
- BlueSky (TU Delft)
- Develop a simulator from scratch

Leave of Microsoft from the consortium

- Difficult to use Airsim

Dissuaded to use BlueSky

- It’s a “monster” / not bug free / the learning curve is slow

Advantages of an in-house simulator

- a light piece of software, only implementing what is strictly necessary
- fast and interactive (as much as we want & can)
- 0 learning curve

Difficulties encountered

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

Data extraction

- Historical flight routes and ATSU geometry

Geometrical calculations

- Is the waypoint in the ATSU?
- Source of software complexity and performance issues

Minimum HMI for the validation

- Specifications and development

Building the validation scenarios

- Required specific adaptations

Integration and platform tuning

- Flight allocation algorithm
- Remote access for LFV
- Platform configuration and tuning

Post-operation analysis

Performances

- Target: perform the training in approximately 1 day

Assessing the performances

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

WARWICK THE UNIVERSITY OF WARWICK

Computational complexity

Times for the current simulator (includes speed and altitude changes)

# Flights	10,000 interactions	1 million interactions	5 million interactions*
2	9 mins	14 hrs 45 mins	*3 days 1 hr 44 mins*
5	22 mins	1 day 15 hrs 11 mins	*8 days 3 hrs 53 mins*
10	45 mins	*3 days 6 hrs 57 mins*	*16 days 10 hrs 44 mins*
50	4 hrs 6 mins	*17 days 2 hrs 5 mins*	*85 days 10 hrs 27 mins*
100	9 hrs 43 mins	*39 days 7 hrs 3 mins*	*196 days 11 hrs 36 mins*
250*	*20 hrs 30 mins*	*85 days 10 hrs 27 mins*	1 year 62 days 4 hrs

* indicates all values in these rows/columns are predicted.

Italics are predicted times

- Times do not include RL training time (negligible) and evaluation of policy (usually 10 episodes every 1000 interactions)
- Simulator is too slow to train for more than 2 flights
- Linear relationship between number of flights and time taken for interacting with the simulator (likely quadratic see plot)

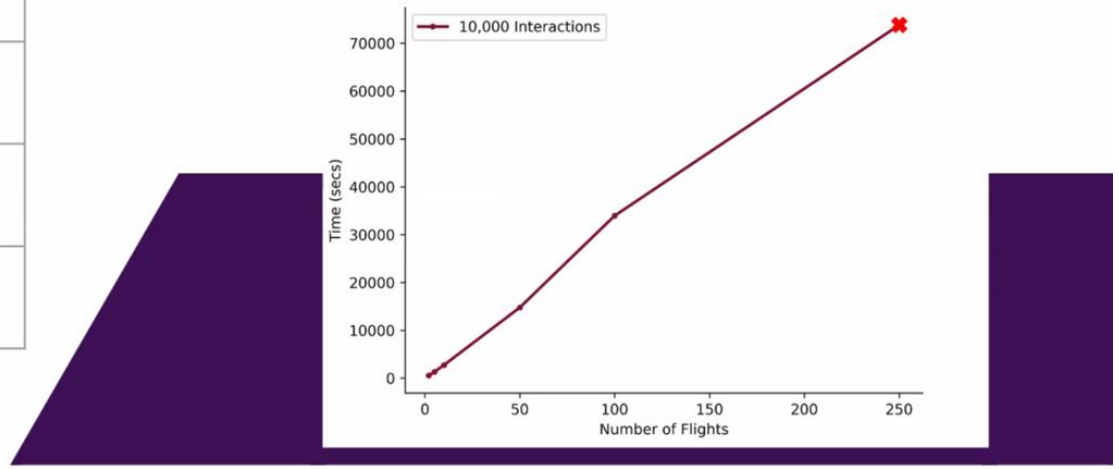

+ 25/09/2024

Optimising the simulator

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

Delay delivery

Software optimisation

- Cache geometrical operations
- Optimize the geometry operations and algorithm
- Simplify the observation space
- Optimize conflict determination
- Parallelize CasADi integrations
- Vectorize all operations

Simplifying the problem

- Limiting the action space (not using speed or altitude change)
- Simplifying the ATSUs geometry (divided by 2 the computation time)
- No ARES
- Reduce the number of flights (not below 80 flights)
- Convexify ATSU
- Simplify trajectories

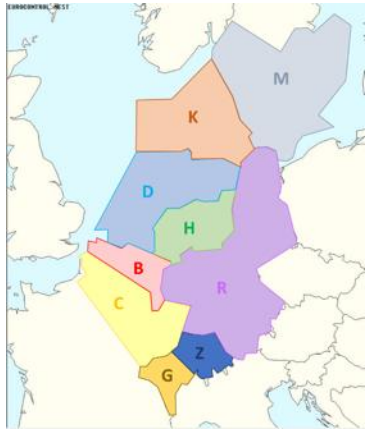

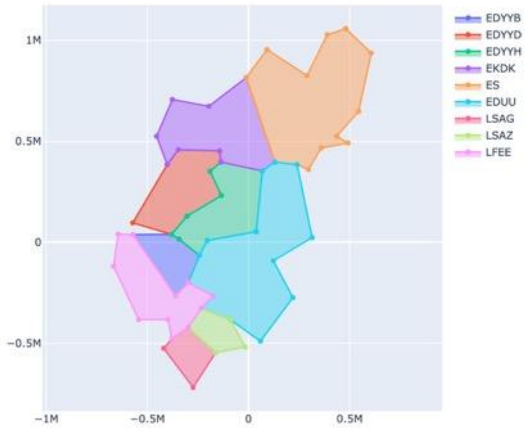

Adapt the approach with the training

- Kinematic flag added to get rid of dynamics and speed up by an extra 2x factor the simulator
- Allow change of trajectory every 30 s (without to wait for the next modified waypoint)
- Consider other RL algorithms
- Start simple and then complexify (training with 10 flights)

Consider another platform

Processing time

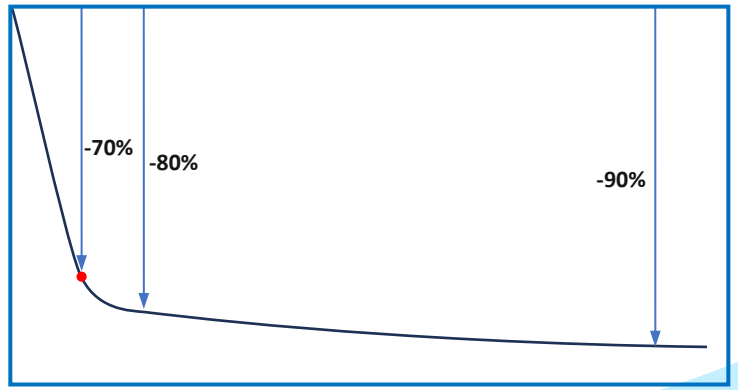

A prototype of fast time/run time simulation

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

- Hypersolver is an air traffic simulator that faithfully reproduces aircraft movements
- Hypersolver is designed to host/train Artificial Intelligence software
- Hypersolver emulates the different layers (ATC, ATFCM) and can apply short, medium and long-term conflict resolution measures
- Hypersolver is lightweight, realistic, flexible and high-performance

AI Performance Results

Matteo Cavalletti, EUROCONTROL

Training Results(1/2)

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

- *The model was directly trained on the exercise scenarios and airspace*
- ***50 flights**.*
- *The **conflict threshold** was set to **10 NM**(double the operational value)*
- *ATCO could modify model decision during the exercise*
- *To simulate this behavior, **model decisions were randomly changed with a 2% probability to mimic ATCO interventions***

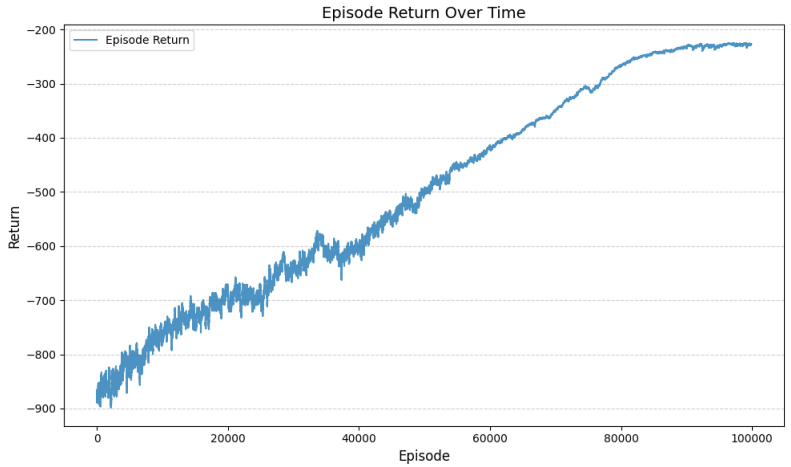

The episode reward reached a plateau and the entropy decreased overtime:

- ***Entropy** measures **uncerainity** in the model, A small entropy means a model more sure of its actions.*
- *An episode **reward** that increaes and reaches a **plateau** means that the model is not improving anymore and it's possible to end the training.*

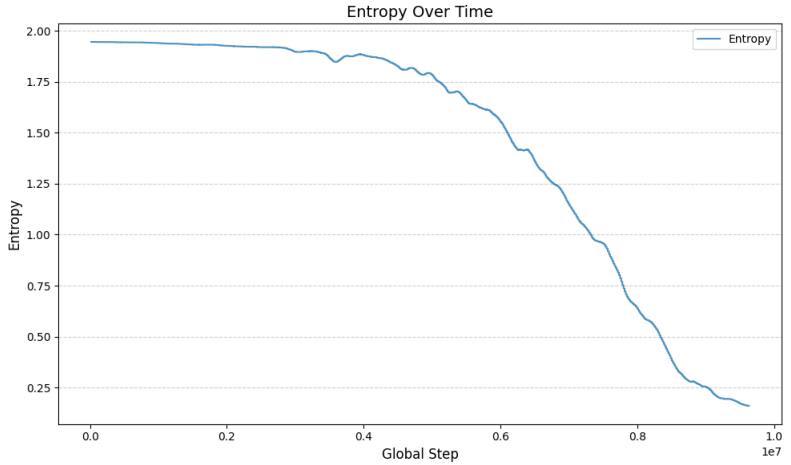

Training Results(2/2)

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

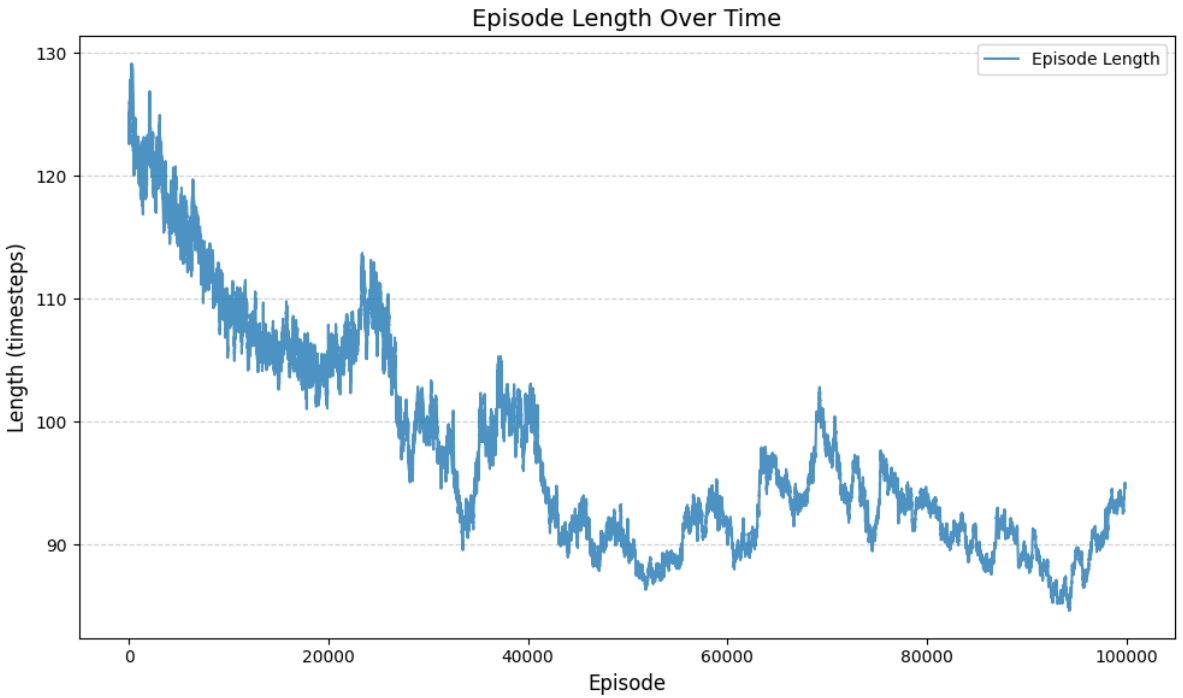

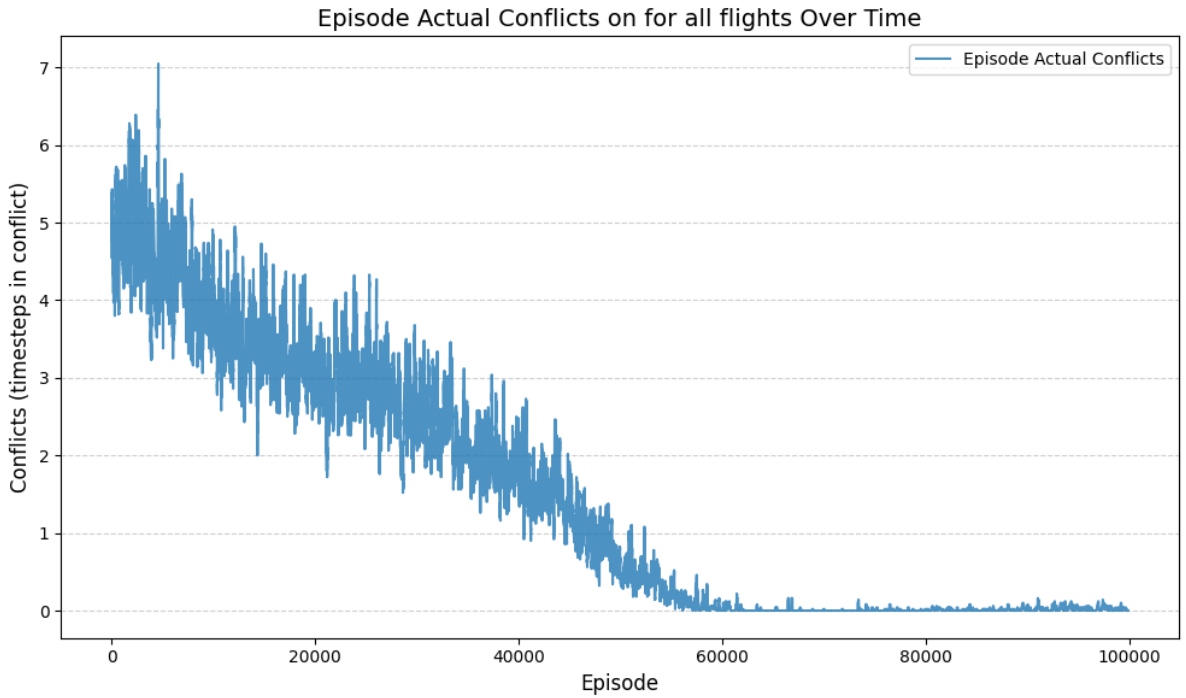

- ***Conflicts*** *over time during episodes decreased* ***to 0 on average***
- ***Average epsiode length*** *also decreased overtime from 130 steps to 90, meaning the algorithm was able to solve the scenarios in less time*

AI Performance in Real Traffic

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

Strengths

Operational feasibility demonstrated in low to medium traffic:

- Manageable workload, acceptable Situational Awareness.
- AI resolutions generally considered safe and acknowledged by ATCOs.

Limitations

Not suitable for high traffic:

- Situation Awarness significantly reduced, cognitive overload.
- AI stops proposing resolutions or suggests late/incoherent actions.
- Low AI transparency → limited trust.

Key Recommendations

- Improve AI: Reliability, relevance of resolutions, robustness in high traffic.
- Integrate Explainable AI: Explanation mechanisms and decision visualization

DCB

Yutong Chen, NTU

Data sources & processing flows

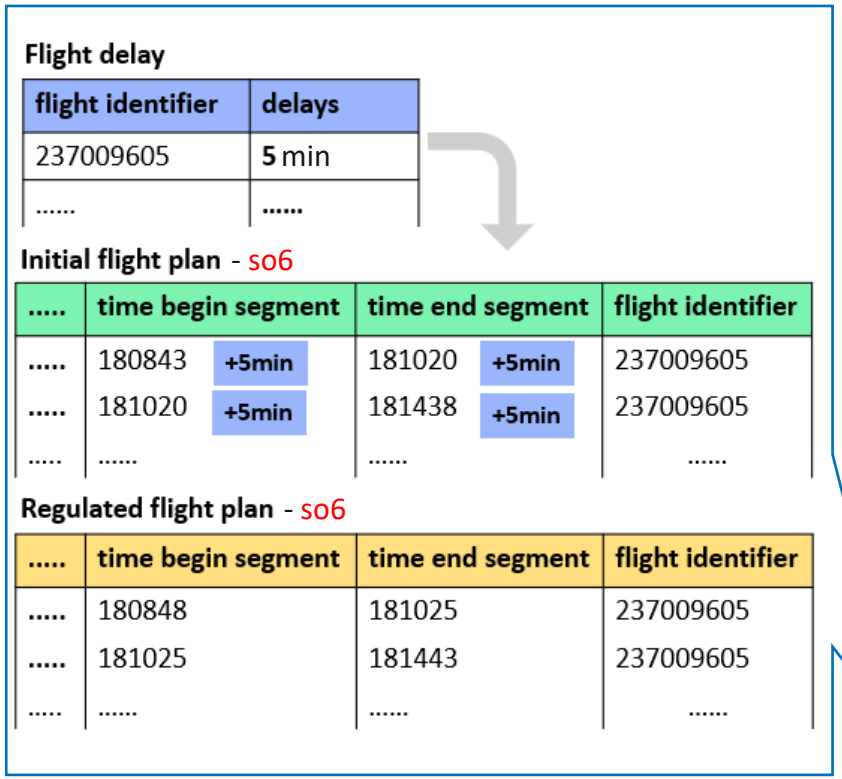

Flight delay

flight identifier	delays
237009605	**5** min
.....	.....

Initial flight plan - so6

.....	time begin segment	time end segment	flight identifier
.....	180843 +5min	181020 +5min	237009605
.....	181020 +5min	181438 +5min	237009605
.....	.....	.....	.....

Regulated flight plan - so6

.....	time begin segment	time end segment	flight identifier
.....	180848	181025	237009605
.....	181025	181443	237009605
.....	.....	.....	.....

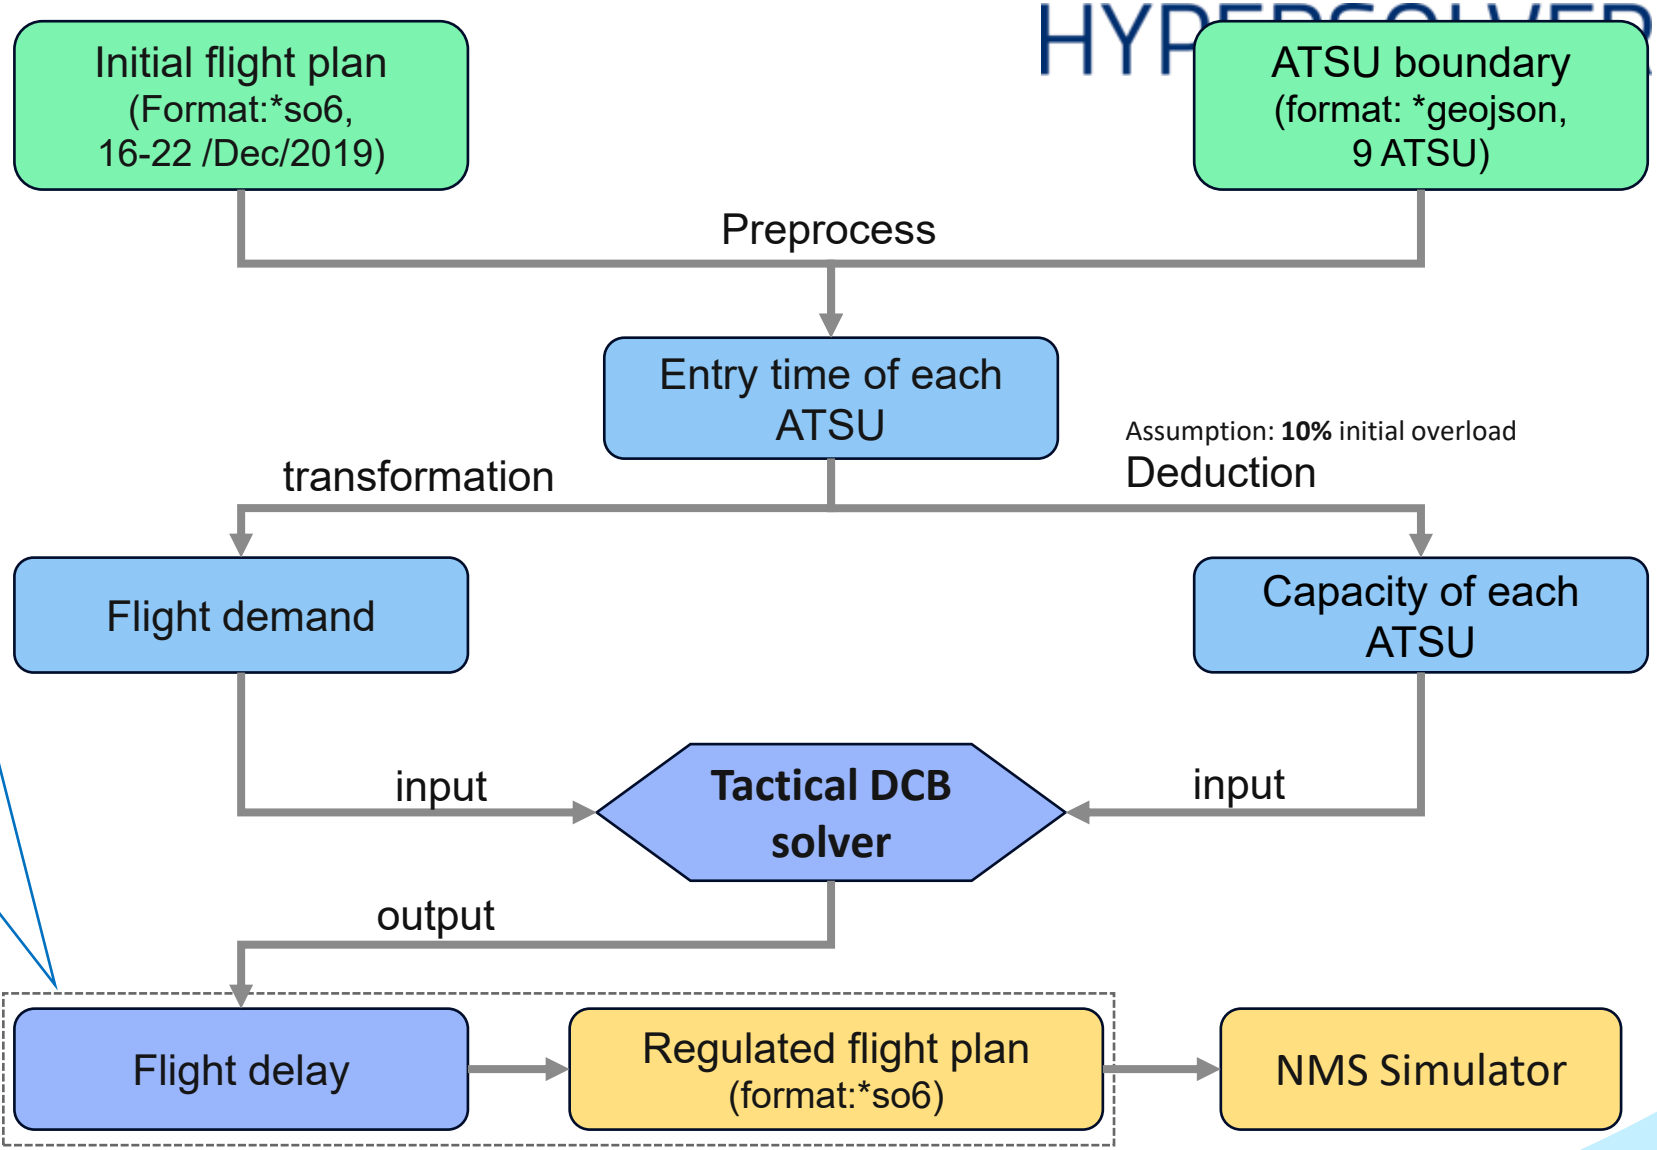

Dataflow of the developed tactical DCB solver

Hotspot elimination

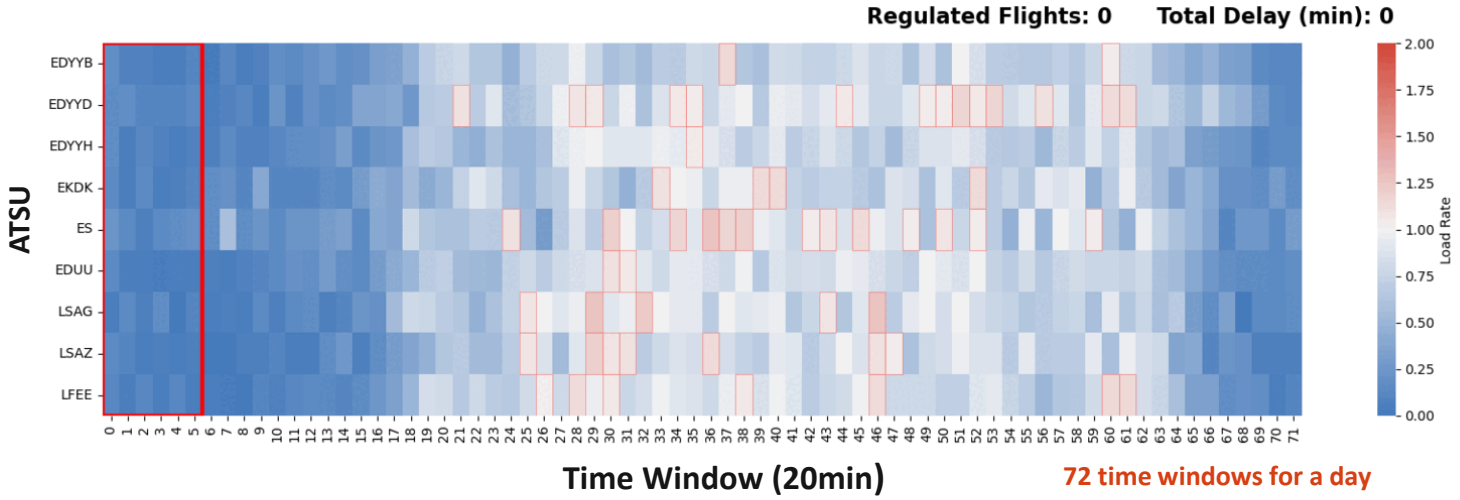

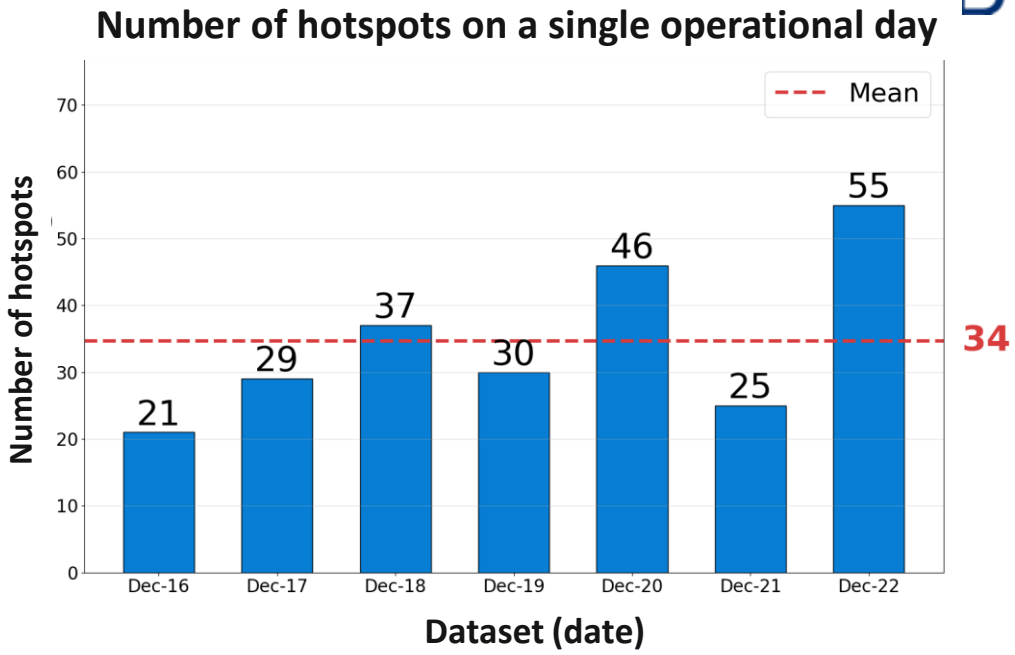

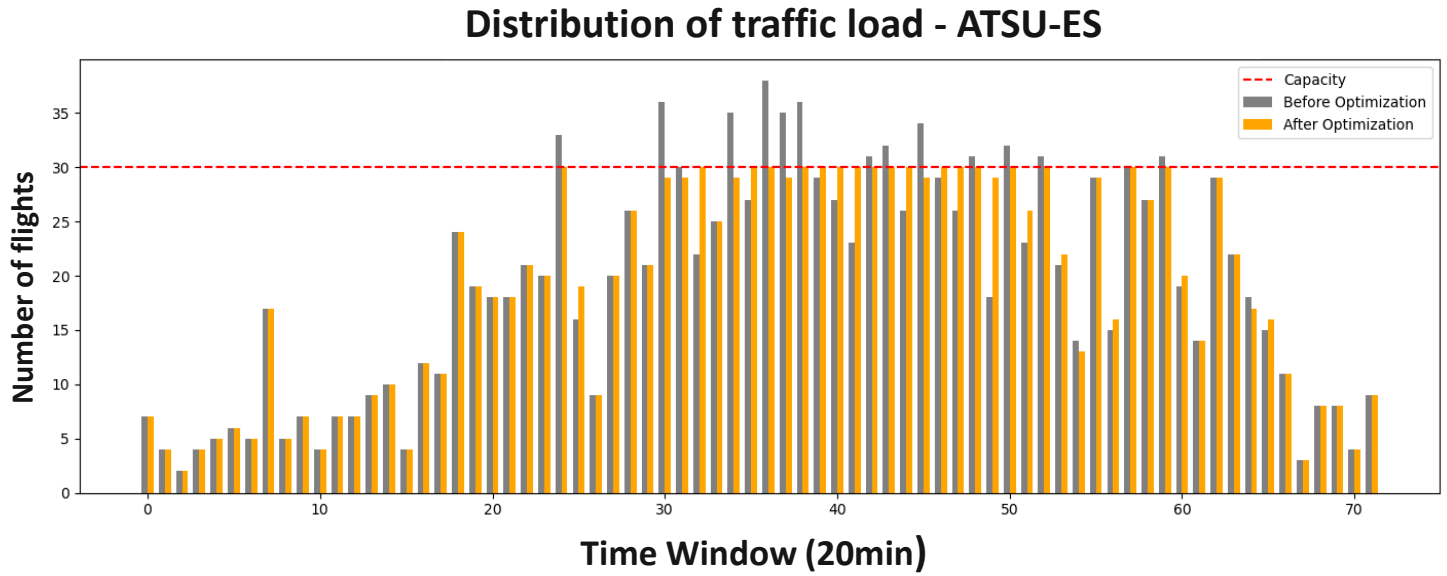

All the hotspots of each regulated time window can be eliminated

Delay indicator analysis

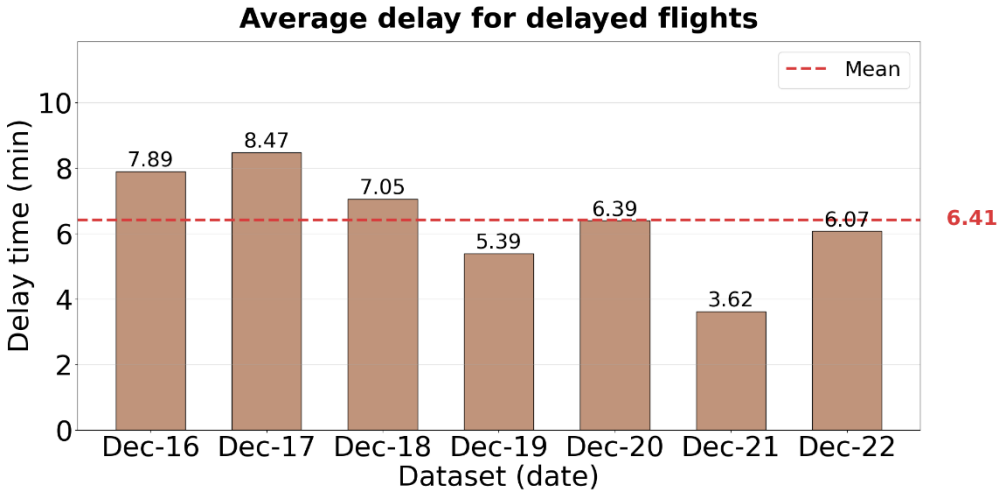

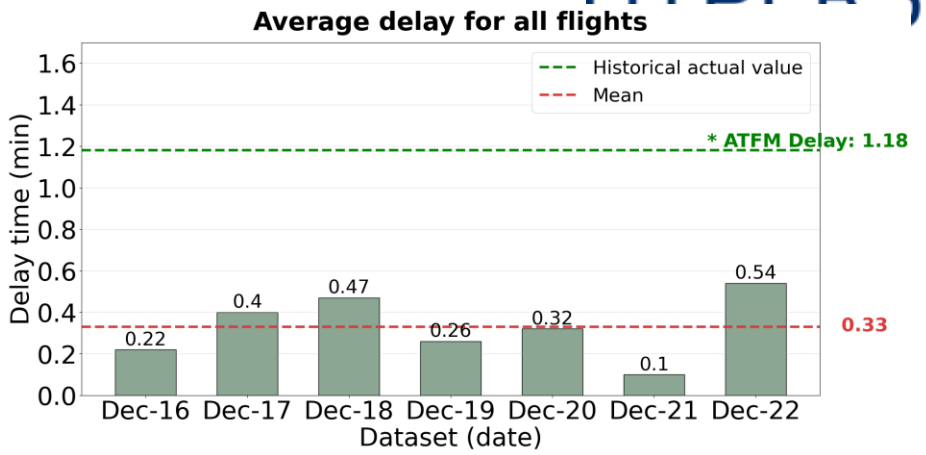

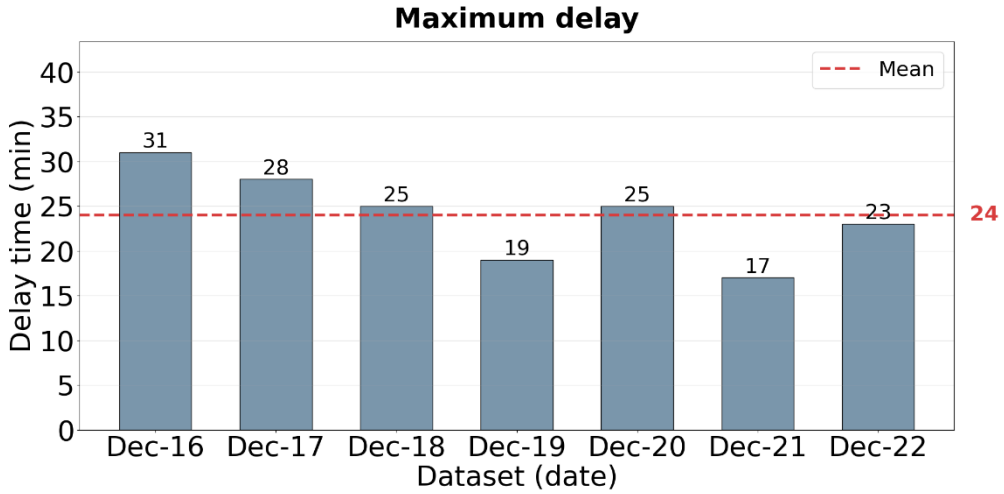

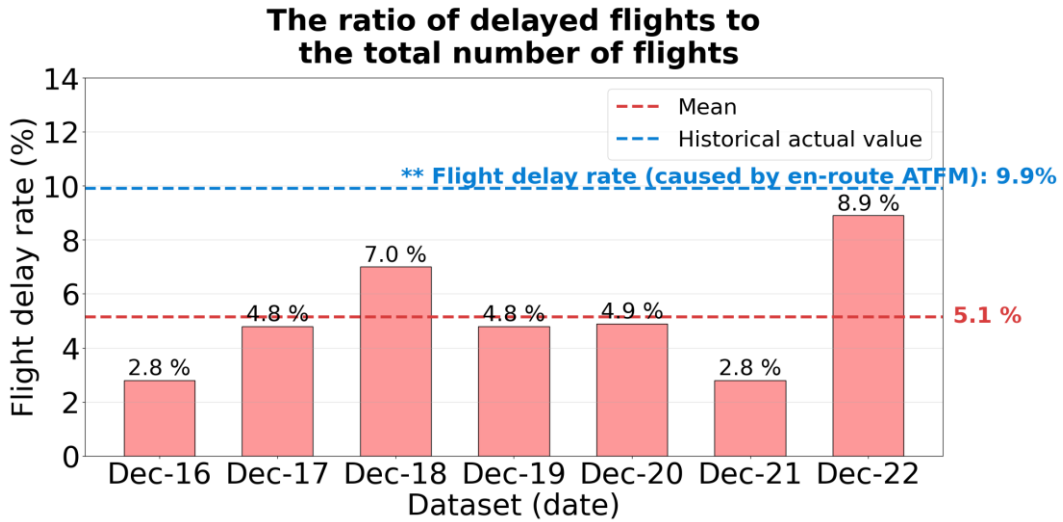

The delay performance is acceptable in practice.

** Monthly Network Operations Report Analysis – December 2019.*

*** Performance Review Report: An Assessment of Air Traffic Management in Europe during the Calendar Year 2019*

Computational performance

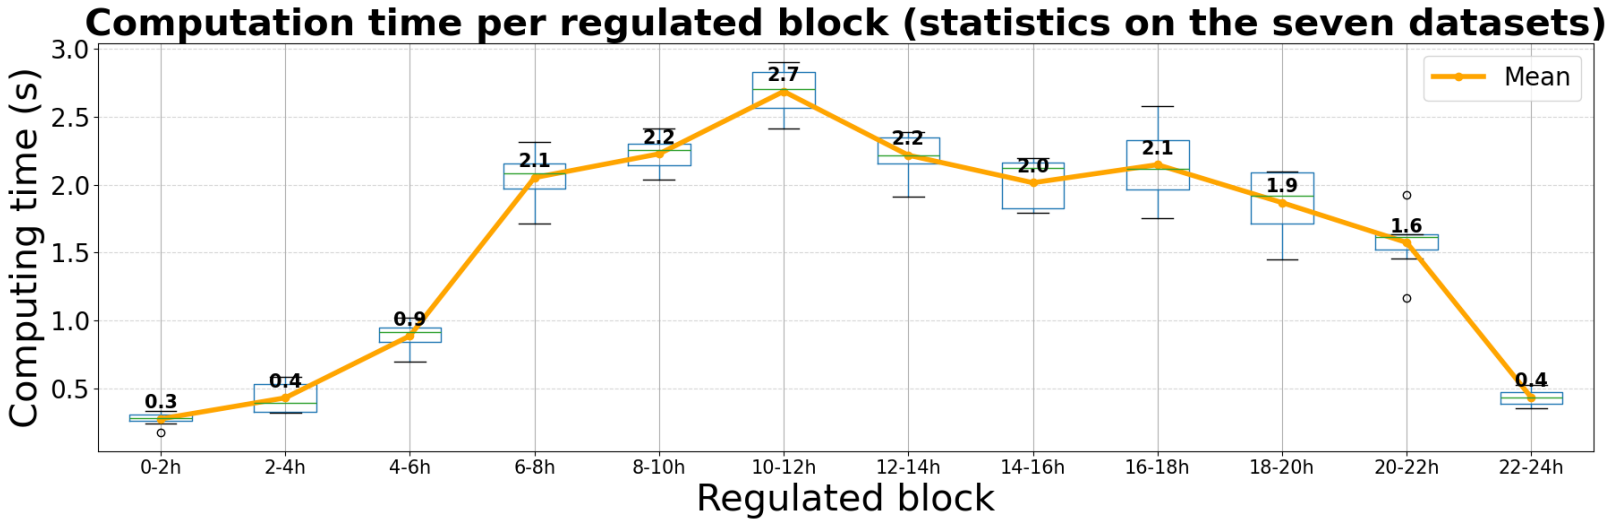

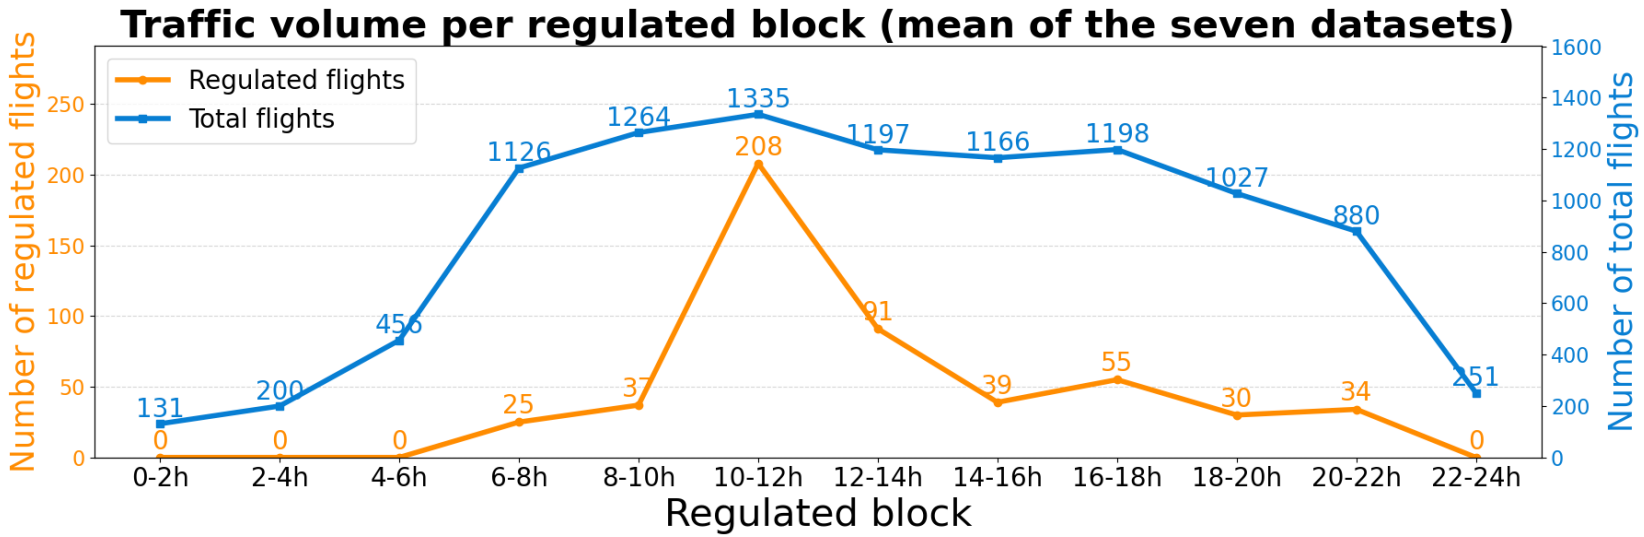

- No more than 2.9 seconds for a solving slot
- The solving time is primarily influenced by the numbers of total and regulated flights.

The computational performance meets the HyperSolver concept requirements.

Results : ATCOs feedback

Victor Tuvesson, LFV

Validation Exercise - Intro

Introduction, scope and limitations

Validation Scope:

- Assessed the HYPERSOLVER concept and support tool across pre-tactical ATFM to tactical ATC phases.
- Evaluated Human-AI Teaming Playbook in Single Controller Operation (SCO) and Flight Centric ATC (FCA) contexts.

Exercise Overview

- Real Time Simulation involving professional ATCOs.
- Included three Controller Working Positions (CWPs) and one Supervisor Working Position (SWP).
- Explored various operational configurations with/without HYPERSOLVER assistance.

Operational Environment

- Airspace: Parts of Germany, Belgium, the Netherlands, Denmark, Sweden, and France.
- En-route environment, from FL250 and upwards.
- Limited to lateral conflict resolution only.

Technical environment

- The simulator included a prototype developed to provide AI capabilities and a CWP interface.
- The integration of this prototype enabled the HYPERSOLVER to interact with traffic, receive real-time data and interface with human operator and decision making process.

Validation Exercise - Intro

Introduction, scope and limitations

Scenario and configuration

- The exercise was performed in a Flight Centric ATC (FCA) setup and Single Controller Operations (SCO) environment together with the Human-AI Teaming aspect.
- Three traffic densities were used (Low, Medium, High)
- Three operational configurations (OPS/Manual, AI/Follow AI, Hybrid) were performed for each traffic density
- Performed 9 Runs. Additional Runs were performed to assess FCA, Allocation Model and Supervisor position separately. In total 16 Runs were performed during the validation exercise.

The following operational configurations were used:

- An ‘’OPS/Manual’’ configuration was performed in low, medium and high traffic density, where ATCOs controlled the traffic manually themselves, without any HYPERSOLVER AI proposals. This was considered the reference configuration/scenario.
- An ‘’AI/Follow AI’’ configuration was performed in low, medium and high traffic density, where the HYPERSOLVER AI was solving the conflicts and the ATCOs were observing the resolutions.
- A ‘’Hybrid’’ configuration was performed in low, medium and high traffic density, where the HYPERSOLVER AI was solving the conflicts, but the ATCOs were able to intervene to modify the AI solutions.

Validation Exercise

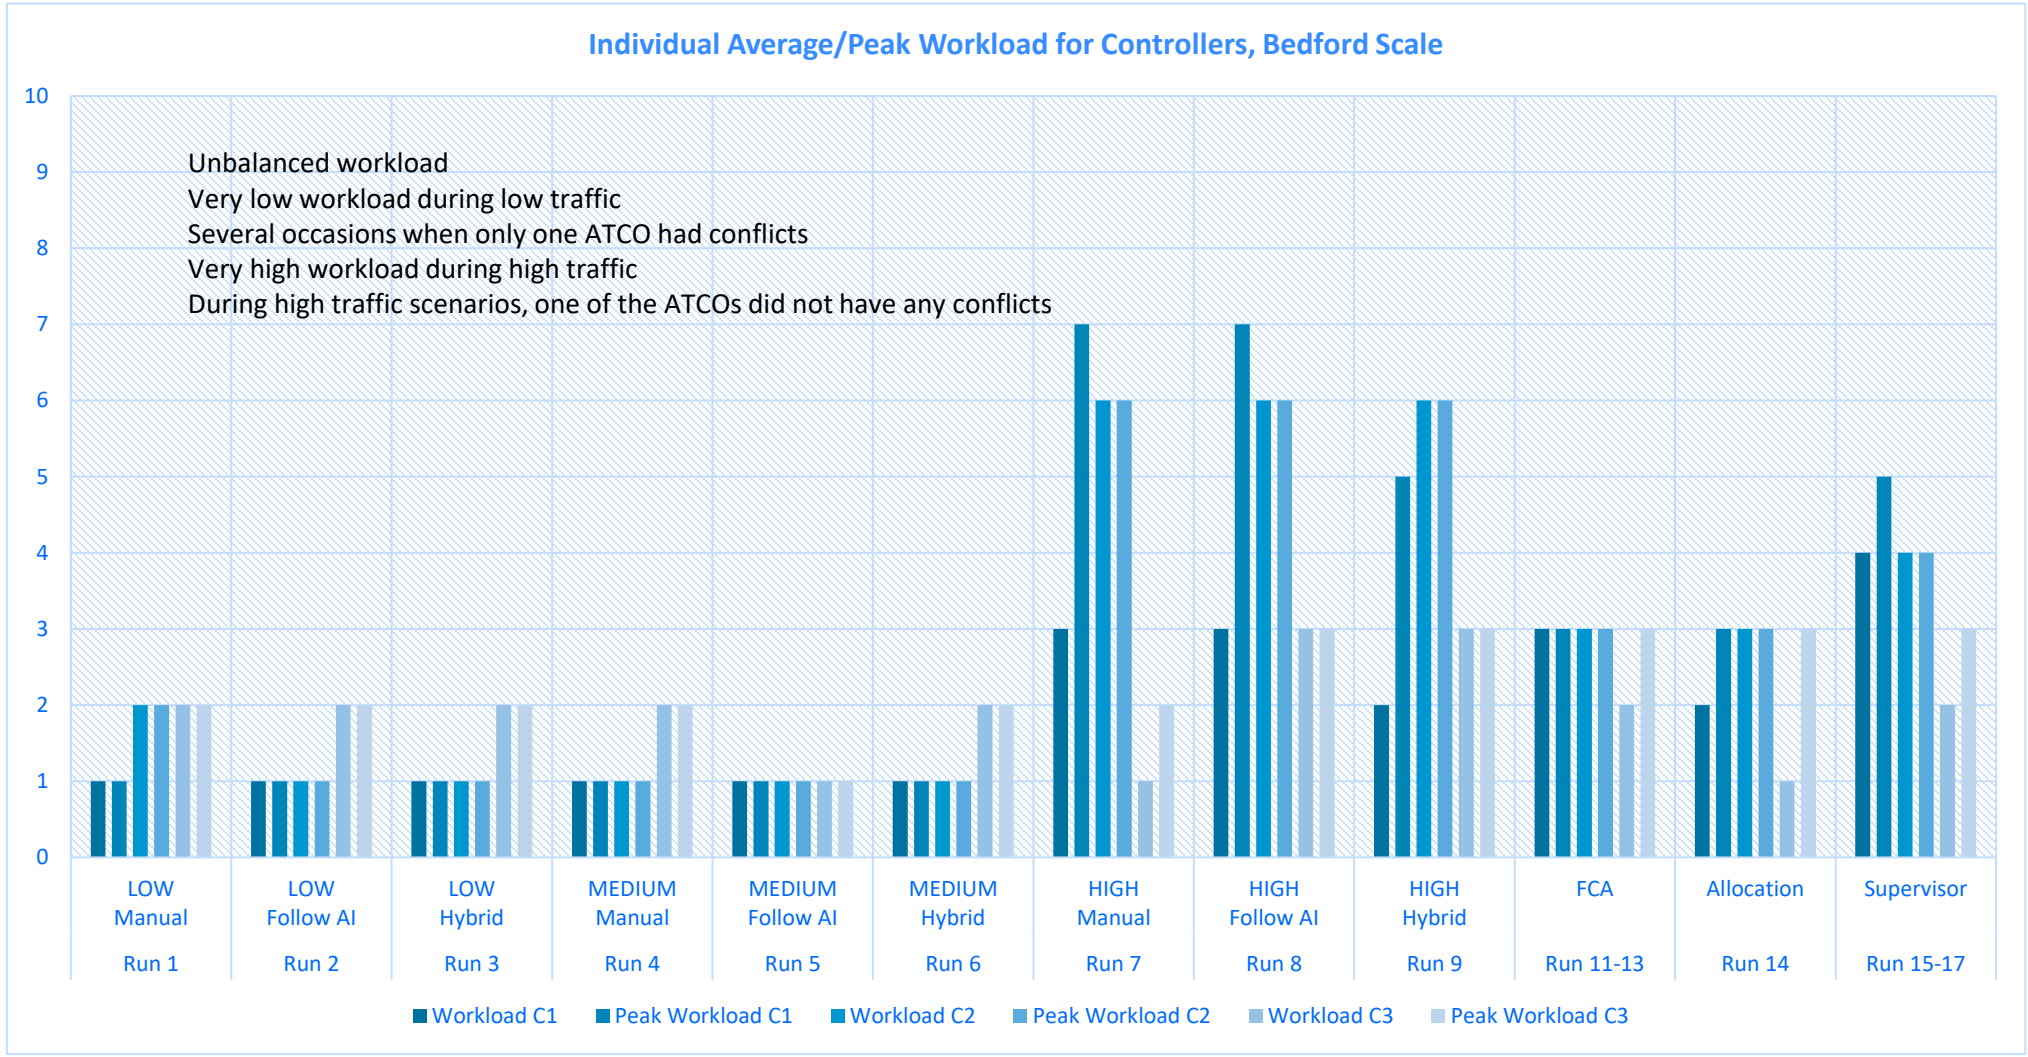

Validation Exercise

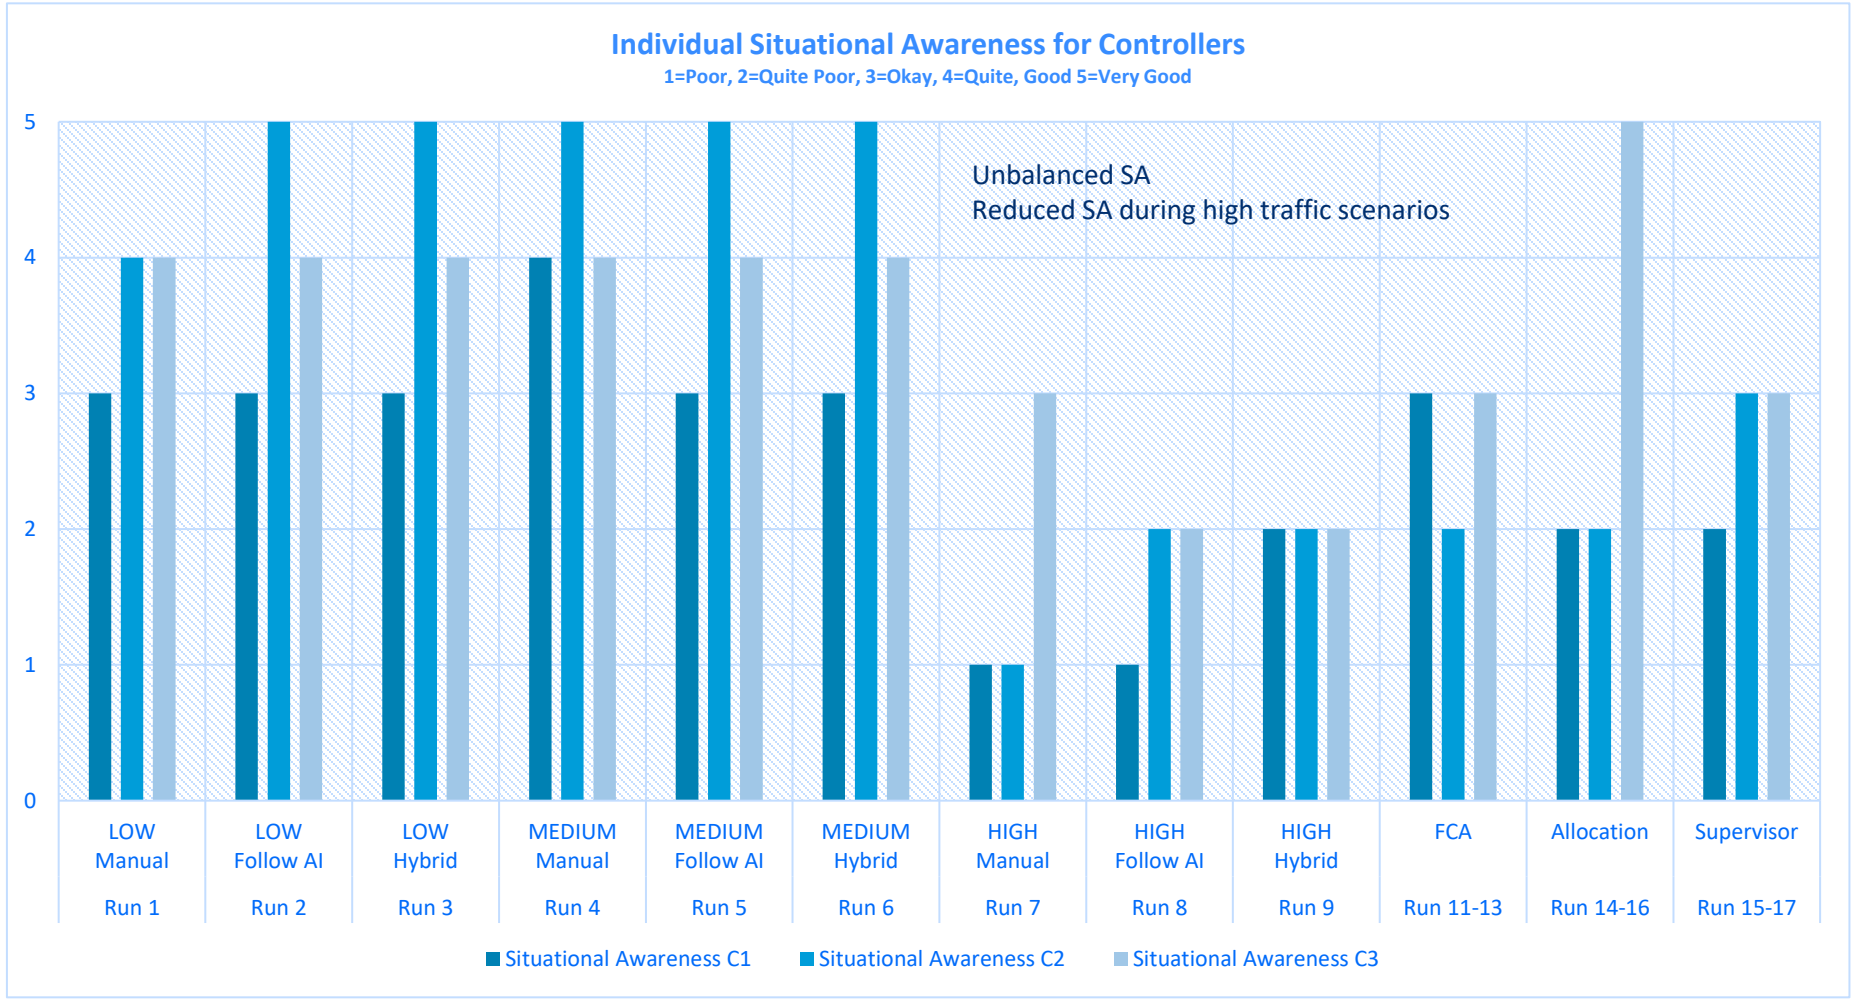

Validation Results Intro

Conclusions and recommendations

- **Hypersolver concept and AI showed some potential for use under simple traffic conditions.**
- **Feedback indicated that most Hypersolver AI proposals were nevertheless considered safe.**
- **Hypersolver AI always included "closed loop" clearances.**

- In several scenarios with high traffic, the Hypersolver AI gave very few suggestions or stopped acting, which forced the ATCO to revert to manual conflict resolution.
- A number of Hypersolver AI proposals were observed as irrelevant or excessive, for example large lateral deviations near sector boundaries or repeated suggestions for the same aircraft when the conflict was already resolved.

Validation Results Intro

Conclusions and recommendations

- SCO and FCA: Both concepts were rated as neutral to positive in terms of operational acceptability, but not yet suitable for high-density traffic due to reduced situational awareness. However, the controllers expressed interest in trying the HS concept in a sector based operations.

- Human-AI teaming: Showed positive signs, but the HYPERSOLVER system provided limited support for team coordination. Situational awareness could be improved with a more advanced HMI and communication tools. The feedback from the ATCOs also include positive experience when allowed to work in hybrid with the HYPERSOLVER AI in order to engage in Human-AI teaming.

Thank You!

SUPPORTED BY sesar JOINT UNDERTAKING

HYPERSOLVER

HYPERSOLVER

NeoMetSys

ENAC

EUROCONTROL

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

LFV AIR NAVIGATION SERVICES OF SWEDEN

WARWICK THE UNIVERSITY OF WARWICK