Developing a complexity-based flight allocator for a shared human-automation air traffic control environment

Thesis Report

Thijs Verkade

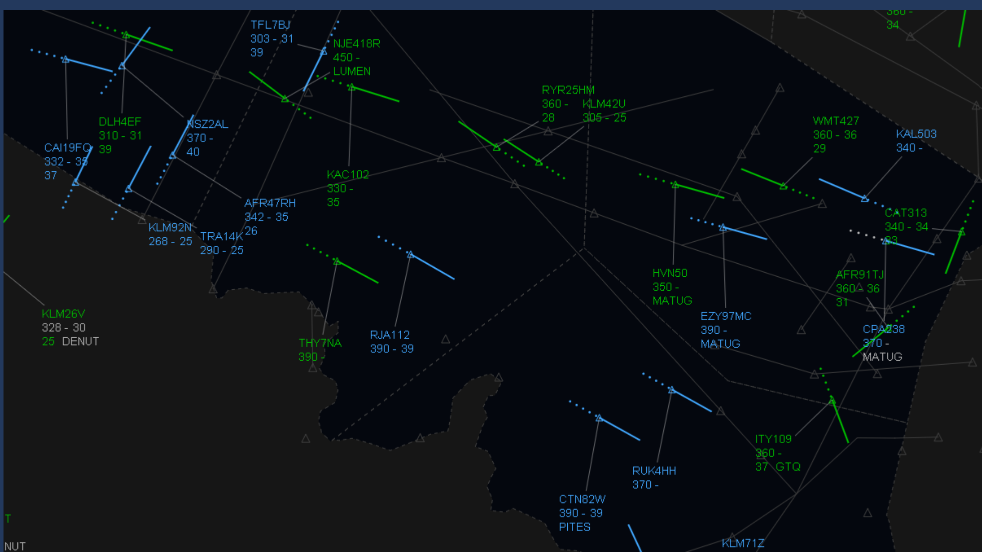

TU Delft

Contents

List of Figures

Part I

Scientific Article

Developing a complexity-based flight allocator for a shared human-automation air traffic control environment

Thijs Verkade
Supervisors: Dr. Ir. Clark Borst, Prof. Dr. Ir. Max Mulder and Dr. Ir. Rene van Paassen

Faculty of Aerospace Engineering, Delft University of Technology

ABSTRACT

The continued growth of air traffic demand is placing increasing pressure on current air traffic control (ATC) systems, prompting the need for alternative ATC strategies. A promising approach is a shared ATC environment between a human controller and an automated controller, where basic, low-complexity traffic is delegated to automation while complex traffic remains under human control. This concept requires a reliable method for predicting the operational complexity of individual flights. This research presents the design and evaluation of a complexity-based flight allocation algorithm for an en-route shared human–automation ATC environment. The allocator classifies incoming aircraft based primarily on the predicted number of interactions along their trajectories, using a flight-filtering mechanism derived from existing models. Additional allocation metrics include the expected number of interactions between human and automation-directed flights and a minimum number of flights controlled by each controller. The allocator was evaluated using offline simulations with real traffic data, followed by a human-in-the-loop experiment with two professional air traffic controllers. Results show that the allocator can consistently assign more complex flights to the human controller while maintaining a balanced workload distribution. The human-in-the-loop experiment saw substantial manual re-allocation and revealed low trust in both the allocator and the automation, indicating the need for further refinement and closer integration with automation capabilities.

I. INTRODUCTION

Global increases in air traffic demand are putting strain on current air traffic control systems, which are currently reaching capacity limits [1]. With increasing reports of controller shortages [2] in Europe and a significant portion of delayed flights being attributed to Air Traffic Control (ATC) capacity [3], a reform of Air Traffic Management (ATM) operations is necessary to meet ever increasing air traffic demands.

In Europe, the Single European Sky ATM Research Project (SESAR), a public-private partnership of relevant parties in the ATC sector, aims to develop and deploy solutions to create a fully scalable ATM system that can handle the growing air traffic while remaining safe and secure. One of the goals of SESAR is to improve the productivity of Air Navigation Services (ANS) through increased levels of automation support [4].

Automation in the ATC environment means that certain tasks that are currently being performed by air traffic controllers (ATCOs) will be performed by an automated system. De Rooij et al. [5] discuss two strategies for allocation of tasks to automation, function-based automation and constraint-based automation. In function-based automation, serialized interactions occur whereby the operator needs to monitor the computer and/or approve/reject solutions proposed by the computer [5]. This is often inefficient as the operator may need to spend analyzing the proposed solution, similar to supervising an ATCO trainee, automation is then degraded to a 'student' that requires close monitoring [5]. Several vigilance studies indicated that humans have difficulty maintaining focus and attention in such situations [6]. In some cases, monitoring tasks had negative consequences on the situational awareness and decision time of the humans involved [7].

In constraint-based automation, higher levels of automation are attained in a constrained environment and the environment is gradually expanded [5]. This leads to a 'parallel system', whereby both the human and automation system can take initiative and execute actions. A parallel ATC system approach suggests that humans should be kept in the loop by establishing teamwork between the two agents. By increasing the exchange of information between the two agents to establish a common frame of reference, the overlap of their teamwork is increased [5]. However, care must be taken to not make the overlap too large, as this may contribute significantly to workload for establishing and maintaining teamwork between the two agents.

To avoid the negative consequences of increased automation support on the ATCO, Maastricht Upper Area Control Centre (MUAC), an Air Navigation Service Provider (ANSP) responsible for ATC in the upper flight levels of the BENELUX area as well as north-west Germany, has proposed an alternative strategy for higher levels of automation in the ATC environment. MUAC's automation strategy is to allocate 'basic' traffic to an automated system, while giving more complex 'non-basic' traffic to a human ATCO. In this context, basic traffic generally refers to flights that have little interactions and are relatively simple to manage in terms of workload. Non-basic flights have more interactions and require greater attention thus producing a larger workload for the responsible controller. Dividing the traffic in a sector in this way ensures that the human is still kept engaged with complex tasks, while the automation can relieve workload by completing simpler tasks. Additionally, complex tasks for the automation may be difficult to reliably manage and solve and may lead to more automation failures. For this allocation strategy to work, a clear definition of the complexity of a flight is required, to be able to distinguish between a basic and non-basic flight.

Several studies have been carried out to determine factors that contribute to the complexity of an individual flight [8] [9]. In an experiment carried out by de Rooij et al. [9] ATCOs from MUAC were asked to determine which background flights influenced the complexity assessment of a selected flight for the controller. While the experiment was carried out in a static situation, and the model performance was weak, it was concluded that altitude overlap was by far the most important feature affecting perceived complexity, which is closely related to the number of background flights [9].

In a separate experiment, de Rooij et al. examined how ATCOs allocated flights in a shared human-automation en-route airspace [5]. In the experiment, the initial allocation was based on geographical or structural properties such as the required flight level change of an aircraft or based on the sector where the flight first entered the airspace. The ATCO had full control over which flights to give to automation, and were able to switch the allocation at any point. Findings from this study showed that traffic in close proximity of the aircraft as well as traffic along the route of the aircraft were among some of the most important factors used in the allocation strategy of the ATCOs. Further recommendations of the study were to investigate an allocation strategy based on the number of spatiotemporal interactions an aircraft has with the surrounding traffic [5].

In this paper, the complexity of a flight is described based on the number of interactions it will encounter along its planned trajectory. This paper proposes various metrics to be used in an allocation algorithm that will allocate incoming flights in a sector to a Human ATCO (H-ATCO) or a digitally automated controller (D-ATCO). These metrics will be used to dynamically determine the responsible controller for the flights. Previous studies into allocation algorithms used static allocation schemes based on the incoming flight characteristics [5] rather than the nature of the interactions of the incoming aircraft with the traffic in the sector.

The design of the allocator is then evaluated in a simulator using live data of real air traffic to determine the robustness of the allocator as well as the impact of design parameters on the allocation distribution and the potential consequences of the allocation. Furthermore, an experiment with H-ATCOs is carried out to determine the effects of design parameters on the experienced workload of the controller as well as evaluating their ability to safely and efficiently manage the traffic in the sector.

The paper is structured as follows. Section II gives some background information on current flight filtering models used to determine potential interactions between flights, as well as information on current research of automated ATM systems. Information is also provided on some allocation strategies that have been proposed. Section III provides a description of the design metrics used in the allocator model, as well as a description of how the allocator operates. The simulation tests and results are described and discussed in section IV. The human-in-the-loop experiment and results is described and discussed in sections V and VI. The paper ends with a discussion of the results and recommendations for future research in section VII followed by a conclusion in section VIII.

II. BACKGROUND

In this section, relevant topics to the flight allocator design are discussed. First, research regarding the flight filtering model concept is explained, which forms the basis for the flight allocation model used in this paper. Research into the automation of basic traffic is discussed, and the findings of studies carried out by MUAC are summarized. Allocation strategies are investigated to gain an understanding of methods that already been tested and what remains to be investigated.

A. Flight Filtering

A study carried out by Kumbhar et al. [10] looked at determining the relevance of background flights for a particular flight of interest (FOI) in order to produce a flight-centric filtering mechanism . The filter was created by mapping out the decisions carried out by an ATCO during enroute traffic. The flowchart for a typical ATCO's conflict detection strategy can be found in Figure 1.

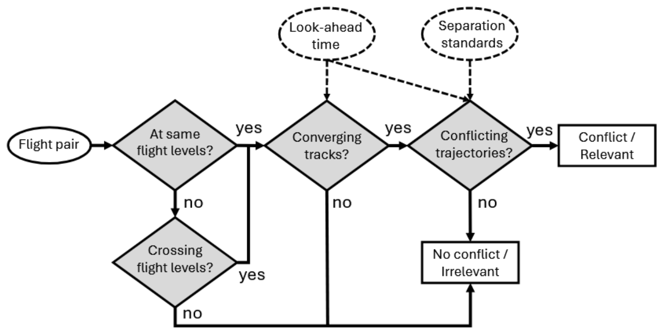

Figure 1: ATCO conflict detection flowchart (Dashed-lines are dependent on the ATCO) [10].

A flight filtering mechanism was then modeled using three key criteria illustrated in Figure 2. The flight filtering model checks for altitude, spatial and temporal overlap for a flight pair. Should all three conditions be satisfied for the flight pair, then the flight is deemed relevant for a FOI, otherwise the flight is deemed irrelevant [10]. A flight filtering model was then developed which makes use of both state-based extrapolation, whereby flight profiles are projected forward based on their current positions, headings and speeds, and intent-based extrapolation where the position of a flight is extrapolated using the planned altitude and speed profiles. These are used together to determine all relevant flights for a FOI.

Following the development of the filter, Kumbhar performed an experiment with professional ATCOs to evaluate the validity of the flight filtering model. In this experiment, ATCOs were presented with static scenarios in which the ATCO must indicate all relevant flights for a highlighted FOI. The flight filtering model was found to perform well in correctly predicting relevant flights from all background flights, with a success rate of 90.46% [10]. Recommendations for further research included using the flight filtering mechanism to determine the number of relevant flights for incoming flights. This information could then be used as the basis for determining whether a flight can be considered basic or non-basic. This paper focuses on incorporating the flight filtering mechanism developed by Kumbhar as the basis for the allocator algorithm.

B. Automating Basic Traffic

MUAC has been researching and developing an automated flight handling system named the ATC Real Ground-breaking Operational System (ARGOS). One of the goals of the ARGOS system is to be able to fully automate control of the least complex tactical traffic scenarios, meaning that the ARGOS system is capable of identifying, isolating and handling basic traffic [11]. The goal of ARGOS is to make use of Controller Pilot Data Link Communications (CPDLC) capabilities on the ground and on board the aircraft to relay information to aircraft. In the case that an aircraft does not have access to CPDLC, ARGOS can still be used as a decision support tool for ATCOs [11].

Currently, MUAC has decided to implement ARGOS in three different stages based on the SESAR levels of automation taxonomy [12]. In the first level, ARGOS acts as a decision support tool, providing conflict resolution overviews and optimal trajectory proposals. In the second level, ARGOS additionally provides CPDLC flight handling commands upon ATCO approval. In the third level, ARGOS is a fully automated mode of operation, handling basic traffic without ATCO supervision. At the first two levels, the ATCO is responsible for controlling and managing the traffic while in the the third level, ARGOS is solely responsible for the control and management of the basic traffic. However, an option remains for ARGOS to indicate that human supervision is required for situations that ARGOS is unable to handle.

In order for such a system to work, the ARGOS system needs to be able to identify basic traffic before the traffic has entered the sector. To do this, a prediction of the complexity of the incoming flight is required, after which the flight will be allocated to either the ATCO or the ARGOS system, depending on the calculated complexity. In a study done by de Rooij et al. an experiment was conducted where an airspace sector was being managed by a human controller and an automated controller, both responsible for managing different flights within the same sector [5]. In this experiment, different flight allocation schemes were used to initially assign incoming flights. The ATCOs were able to re-allocate flights at any time.

Conclusions of this study were that the ATCOs strongly preferred to delegate flights to automation under the conditions that flights were on direct routes and free of conflicts. The study found that the majority of participants welcomed the idea of an automated allocator, noting that the automation should be able to handle and perform all tasks required for the flights assigned to it to prevent the automated flights from requiring human supervision. The allocation schemes used in this study were based on simple pragmatic schemes which the ATCOs generally did not follow [5]. Recommendations for further research suggested exploring the complexity of individual flights in order to classify aircraft as either basic or complex. This should then be used as the basis for an allocation scheme, better aligning with the allocation preferences indicated by ATCOs as well as the vision of ARGOS. This paper explores the concept of using complexity as a metric to allocate aircraft.

C. Allocation Strategies

An alternative approach to ATC currently being researched is Flight-Centric ATC (FCA) in which ATCOs control flights for their entire flight duration. In such an approach, sector boundaries are dissolved [13]. In order for FCA to succeed, a suitable allocation strategy is required to equally distribute workload among controllers [13]. Research into FCA has led to the development of several approaches for allocation strategies to determine the appropriate ATCO for a particular flight [13][14][15]. In general, a distinction can be made in terms of the degree of automation and in terms of the flexibility of the allocation [13].

Finck et al. mention three different levels of allocation: automatic, semi-manual and manual allocation. In automatic allocation, all assignments are made autonomously following pre-defined rules to determine the appropriate allocation. Advantages of automatic allocation are that no additional personnel are required to manage the automatic allocation system, with the disadvantage that the automatic allocation is unable to adapt to unique situations or the current workload experienced by the controller [13]. A further distinction is made between static and dynamic allocation, whereby static allocation allocates aircraft according to predefined rules such as allocation based on flight direction. Dynamic allocation allocates based on predicted complexity calculations of each aircraft.

Several of the allocation strategies have already been studied further, including a dynamic automatic allocation using a workload model [16] and a static automatic allocation considering flows, clusters and conflicts [13]. Dynamic allocation through the use of a comprehensive workload model was found to be well suited to achieve an even workload distribution among controllers [13]. Such a strategy was found to be more well-suited for smaller airspaces as predicted workloads become increasingly inaccurate with increasing distances.

The static automatic allocation strategy made use of a multi-layer allocation strategy, first checking whether an aircraft belongs to a cluster, then checking the entry and exit points of a traffic flow. Finally, the aircraft is checked for potential conflicts with other aircraft. If any of the conditions are satisfied along the way, then the aircraft is assigned to the control responsible for the respective cluster,traffic flow, or the other aircraft involved in conflict. If the three criteria do not lead to an assignment, then the aircraft is given to the controller with the lowest aircraft count. The study showed that allocation strategies based on flows and clustering could have a positive effect on the controller workload. The combined strategy did not show the expected effects, though it was noted that the sample size was relatively small and further testing needed to be done to draw any conclusions. It was also noted that rearranging the hierarchy could lead to different results and would also be a conceivable approach [13].

While some allocation strategies have been explored by researchers, further investigation is required to draw strong conclusions on their feasibility and applicability to ATC. Particularly in the case of automating basic traffic, allocation strategies remain uninvestigated.

III. FLIGHT ALLOCATOR METRICS

The complexity of an aircraft is a combination of several different factors. These factors are to be used as metrics in the flight allocator to determine whether an aircraft can be classified as basic or non-basic, and subsequently assign the aircraft to the appropriate controller. While there is no general consensus on the contribution of each complexity metric to the total complexity of the flight, general findings suggest that some characteristics are more important than others. From [9] [5] [8], the most important metric for the complexity of a flight is the number of interactions along the trajectory of the flight.

A. Number of Relevant Flights

In order to determine the number of relevant flights, the flight filtering mechanism designed by Kumbhar [10] is used to calculate the number of relevant flights for a FOI. Kumbhar's flight filtering mechanism calculates the number of relevant flights based on state-based and intent-based extrapolation.
The flight filtering mechanism designed by Kumbhar has been further expanded to calculate the number of relevant flights based on a direct route (DCT) to the sector exit point (XCOP) of the aircraft. Providing DCT commands to aircraft are within the capabilities of the automation software used in this paper. The flight filtering mechanism therefore produces two integer indicators of the number of interactions expected for a FOI. A diagram showing the relevant geometries in determining the number of relevant flights can be seen in Figure 2.

In Figure 2, the state-based, intent-based and DCT-based trajectories are all used to determine the number of relevant flights. State-based extrapolation is used to project flights forward based on their current positions, headings and speeds. This is represented by the dotted lines in Figure 2. The

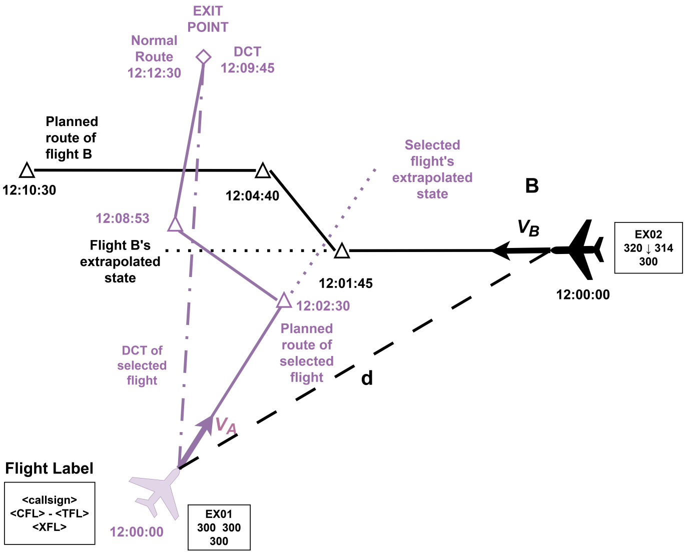

Figure 2: The geometry of state-, intent- and DCT-based predictions. The purple flight and lines correspond to the FOI selected by the ATCO. (d represents the distance between the two aircraft).

extrapolated states are used to predict the closest point of approach between flights while ignoring intended flight plans.

Intent-based projections are made by extrapolating flight positions along their intended trajectory. The intended trajectories are then used to assess spatial and temporal overlaps. Line segments between waypoints are considered for spatial overlap and the estimated arrival time to the waypoints for the temporal overlaps. If a spatial overlap and temporal overlap are present for a flight pair, then the CPA is calculated by predicting the state of both flights at the start of the critical route segment.

The DCT-based projection is also made in the same manner as the intent-based projection. However, instead of determining the line segments between intended waypoints, the intended trajectory is simply crafted using the trajectory from the flights current position to the exit point of the flight. As shown in Figure 2, the DCT projection is only done for the FOI, and is not considered for other flights.

Following the calculation of the state-,intent- and DCT-based projections, the state and intent based projections are used to determine the number of relevant flights for the FOI. The state and DCT based projections are also used to determine the number of relevant flights for the FOI. With these two methods both producing a number of relevant flights, the lowest of the two is then used as an indicator of the number of interactions for an aircraft. This would allow for an aircraft with a high number of interactions based on its state and flight plan, to still potentially be assigned to automation, given that the DCT has a lower number of interactions such that the aircraft could be considered basic along this trajectory.

The parameters for determining whether or not a flight is relevant can be found in Table 1. If there is any vertical overlap between the FOI and another flight including cleared (current), target and sector exit flight levels of both flights, the parameters in Table 1 are used to determine relevant flights. If the calculated parameters for a flight pair are all less than the thresholds specified in Table 1, then the flight pair is added to the relevant flight list. If any of the parameters are above the threshold, the flight pair is deemed irrelevant. The same process is done for the intent/DCT filter parameter. Combining the results of both gives a final list of relevant flights.

Table 1: State and Intent/DCT based flight filtering thresholds.

State filter parameters	
Distance	160 NM
Prediction look-ahead time	600 s
Distance at CPA	12 NM
Time to CPA	11 min
Intent/DCT filter parameter	
Crossing distance	12 NM

B. Number of Mixed Conflicts

While the number of interactions acts as an indication of the initial complexity of the FOI, other factors not covered by the flight filter also affect the complexity of the FOI. One of the factors influencing the complexity of an FOI is the number of interactions with aircraft being controlled by different controllers. While this can occur in the current ATC environment, such as when military aircraft travel through civil airspace, such a situation would become significantly more common in alternative ATC approaches such as FCA and human-automation combined control in a sector. Studies have shown that ATCOs tend to pay less attention to flights that are not under their control and update their mental model less frequently. This can potentially lead to slower conflict detection and resolution [17]. In the case of mixed conflicts, occurring when one flight in the conflict pair is being controlled by the H-ATCO and the other flight is being controlled by the D-ATCO, the lack of attention of the awareness of the H-ATCO could have negative consequences on the resolution of these conflicts [18].

Potential solutions exist that can help alleviate the chance of H-ATCOs being unprepared for mixed conflicts, such as automatically flagging mixed conflicts in a timely manner so that the H-ATCO can update their mental model and be prepared to solve the conflict. In order to reduce the workload of the H-ATCO further, one of the complexity factors used in the allocator is the number of mixed interactions present for a FOI. A large number of mixed interactions could unintentionally add more workload to the H-ATCO, who is then responsible for updating their mental model and solving all potential mixed conflicts. This factor could significantly impact the perceived complexity of an aircraft for a controller, and should be accounted for when deciding which controller should control the incoming aircraft.

C. Minimum Controller Workload

While the first two metrics deal with estimating the complexity of individual flights, the final metric deals with ensuring acceptable minima are met for both the workload of the human controller, as well as ensuring a minimum contribution of the automation software to the sector. Workloads outside of appropriate ranges can deteriorate the ability of the H-ATCO to carry out their tasks, either through stress and overwhelming them in the cases of too much workload, and boredom and lack of focus in cases of too little workload. The third metric introduces a minimum required percentage of flights assigned to each controller. In the case of the H-ATCO, a minimum would ideally provide enough workload to keep them focused and engaged. Incorporating a minimum for the D-ATCO is also preferred, as the D-ATCO must contribute substantially to managing the traffic in a sector such that the overall workload of the H-ATCO is reduced by a significant degree. This reduction in workload of the ATCO would allow for the growth in air traffic demand to be managed by ATC.

D. Flight Allocator Logic

A flowchart of the allocator logic used to derive which controller to allocate an incoming aircraft to can be found in Figure 3. One minute before an aircraft enters the controlled sector, a series of computations are done to determine the appropriate controller for the incoming aircraft. Firstly, the number of relevant flights using both the flight plan as well as the DCT to the XCOP. The lowest number of relevant flights of the two trajectories is then used. If both the original route and the DCT provide the same number of relevant flights, the original route is used. The resultant number is then used to initially describe an aircraft as either basic or non-basic, depending on the set threshold value for the number of relevant flights.

Following an initial allocation, the total distribution of flights already assigned within the sector is checked to determine whether both controllers are controlling a minimum percentage of flights within the sector. If either controller is not meeting the minimum percentage of flights controlled as described by the set threshold, the allocator assigns the aircraft to the controller not meeting the minimum requirement. Should both controllers not be meeting their minimum threshold, priority is given based on the complexity status of the aircraft, with non-basic aircraft then being assigned to the H-ATCO while basic aircraft are assigned to the D-ATCO.

If both controllers meet the minimum requirement for the percentage of flights under their control within the sector, then the next computation is to determine the number of relevant flights being controlled by a specific controller. For initially basic aircraft, this means calculating the percentage of relevant flights controlled by the H-ATCO. Conversely, for initially non-basic aircraft, the percentage of relevant flights being controlled by the D-ATCO is calculated. This metric is used to identify the potential number of mixed interactions for the FOI. If the calculated number of mixed interactions is above the set threshold, then the flight is assigned opposite to the initial allocation in order to reduce the number of mixed interactions.

If the number of mixed interactions is below the threshold such that it is deemed acceptable, then the allocator assigns the flight according to the initial allocation based on the number of relevant flights. Basic aircraft will then be assigned to the D-ATCO while non-basic aircraft will be assigned to the H-ATCO.

In the flowchart, each rectangular shaped step represents a computation step whereby the calculated value of the relevant parameter is compared to the threshold value of that parameter. Based on this comparison, the allocator makes a decision to assign the aircraft or continue with the next step in the computation process. After an aircraft has been allocated, the system recomputes the allocation every ten seconds to see if any changes arise. An allocation by the system only ever occurs once but if the recomputation provides a different allocation outcome, this will be displayed as a suggestion next to the callsign of the aircraft. This can be used by the H-ATCO to identify possible suitable candidates for re-allocation should they want to increase/decrease their workload.

An initial experiment was conducted with several different threshold values to investigate the consequences of these thresholds on the performance of the allocator. The various values used for the thresholds can be found in Table 2. The following section describes how these threshold values were varied in the experiment as well as explaining the results of the experiment.

Table 2: Various thresholds used in flight allocator testing (Bold values indicate default used in testing).

Allocator Metric	Threshold Values		
Minimum number of flights to be considered complex	3	**5**	7
Minimum percentage of flights for each controller	15	**25**	35
Maximum percentage of potential mixed interactions allowed	50	60	**70**

The threshold values were not selected to optimize allocator performance, but rather to examine the sensitivity of the allocation behavior to variations in each metric threshold. For each allocator metric, three representative threshold levels were defined: a low, middle, and high value. This approach allows the effect of relaxing or tightening each constraint to be observed independently, and provides insight into how the allocation spread changes as the allocator metrics become more or less conservative.

The middle values for the complexity metric and minimum flights per controller, as well as the high value for the mixed interaction threshold correspond to the nominal configuration used for baseline testing. The high variant of the mixed interaction threshold was chosen for the nominal configuration due to the primary focus of the allocator being the number of interactions along the trajectory. By comparing allocator performance across these threshold levels, the robustness of the allocation strategy to parameter selection can be assessed.

IV. OFFLINE SIMULATION EXPERIMENT

Having established the design metrics for the flight allocator, this section discloses the experiment setup of the offline simulation runs as well as the performance of the proposed allocator. The effects of varying the threshold values on certain performance metrics were studied and evaluated. This evaluation assesses whether the aircraft assigned to the H-ATCO are generally more complex than those assigned to the D-ATCO. Further performance metrics include the number of mixed interactions, the robustness and consistency of the allocator to provide a consistent number of flights to the H-ATCO and retain a balanced distribution of flights in the sector. Lastly, the nature of short term conflict alerts (STCA) and loss of separation (LOS) alerts provided by the system between a digitally controlled flight pair or mixed flight pair were examined.

The offline simulation experiment utilized SectorX, a Java-based simulator developed by the TU Delft, mimicking the enroute radar display currently used at MUAC. In this experiment, air traffic movement was simulated based on a traffic sample. The allocator would allocate all incoming flights to either an H-ATCO or D-ATCO depending on the calculations and thresholds of the allocator. In this offline experiment, no H-ATCO was present, so traffic assigned to the human controller simply followed the trajectory indicated in the flight

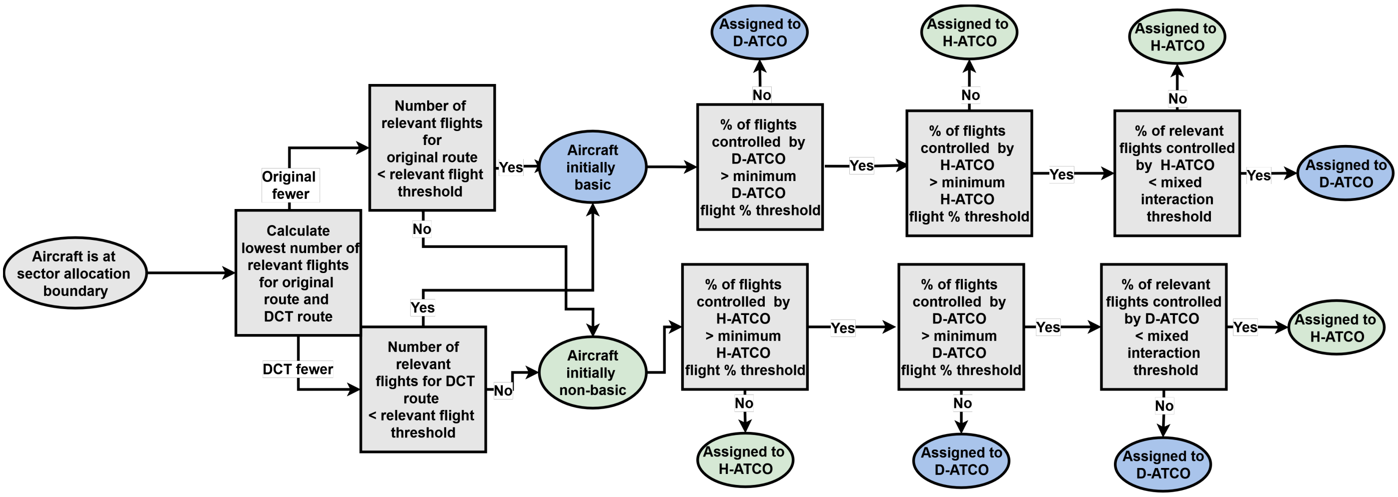

Figure 3: Flight Allocator flowchart logic describing how the allocator assigns an aircraft.

plan, not solving conflicts or climbing/descending to meet exit point requirements. The traffic assigned to the D-ATCO is being controlled by automation. To initially test the threshold parameters of the allocator as well as determine the stability of the allocator, a human controller was not deemed necessary. A follow-up experiment was conducted and is discussed in Section V and VI, whereby two professional controllers were present to manage their assigned traffic.

A. Traffic Scenario

The offline simulations made use of two traffic samples of air traffic above FL245 in the entire Brussels sector group. The traffic was sampled from a relatively busy day in September 2023. Both traffic samples are taken from the same day. The traffic sample includes overflying traffic as well as arrivals and departures to several airports within or close to the sector. The total number of aircraft in the sector over the 25 minute duration of the simulation sample was 70.

B. Automation

During the offline simulation experiment, a digital controller was present in the form of basic automation software. The automation was capable of acting within the experiment scenarios without human involvement, automatically executing actions to ensure safe air traffic. When an aircraft was assigned to automation, the computer was able to perform the following actions:

- Ensure sufficient separation between flights (5NM horizontally, 1000ft vertically).
- Deliver flights at their exit point and transfer level, climbing as early as possible and descending as late as possible.
- Descend arrivals to FL260 to be transferred to lower area control.
- Issue a DCT to flight if available.

Automation solved conflicts in the vertical plane only. Automation was capable of solving mixed conflicts by giving digitally controlled flights vertical commands. The automation mimics ATCO best-practices by keeping aircraft at their cruise level for as long as possible by providing climb commands as early as possible and descend commands as late as possible. Aircraft at cruise level are more predictable and easier to manage for H-ATCOs. Studies show that air traffic with changes in aircraft state were found to be more complex and contributed more to workload for controllers than aircraft flying regular routes. [19] [20].

C. Experimental factors

There were two experimental factors that were varied for the offline experiment:

- Allocator metric threshold values.
- Initial allocation distribution of the sector.

The allocator metrics are varied according to the values in Table 2. The initial allocation distribution indicates the distribution of the aircraft already in the sector when the simulation starts. Three variations were possible, a fully automated sector start, a fully human sector start, or no initial allocation whereby the aircraft were allocated using the allocator system. The variations of both experimental factors can be found in Table 3.

Each variation found in Table 3 was tested twice using two different traffic samples of the traffic scenario described but fifteen minutes apart. This led to 42 simulation runs of different combinations of allocator thresholds and initial sectors. Additionally to this, three variations of the allocator were experimented after the results of the initial simulation runs whereby certain metrics were disabled entirely.

- Number of interactions only.
- No minimum percentage of flights for each controller.
- No maximum percentage of potential mixed interactions allowed.

These were also tested with the three initial allocation distributions of the sector but were only tested with one traffic sample giving a total of 51 simulation runs. The list of experimental variations when disabling certain metrics can be found in Table 4.

Table 3: Variations of allocator threshold values and initial sector allocations tested during the offline simulation.

Number of interactions threshold	Min. flight threshold (%)	Mixed interaction threshold (%)	Initial allocation
Default threshold values			
5	25	70	Automated
5	25	70	Human
5	25	70	None
Interaction threshold varied			
3	25	70	Automated
3	25	70	Human
3	25	70	None
7	25	70	Automated
7	25	70	Human
7	25	70	None
Minimum flight threshold varied			
5	15	70	Automated
5	15	70	Human
5	15	70	None
5	35	70	Automated
5	35	70	Human
5	35	70	None
Mixed interaction threshold varied			
5	25	50	Automated
5	25	50	Human
5	25	50	None
5	25	60	Automated
5	25	60	Human
5	25	60	None

Table 4: Additional allocator configurations evaluated after the initial simulation runs, in which individual allocator metrics were disabled. Each configuration was evaluated under three initial sector allocation conditions.

Number of interactions threshold	Min. flight threshold (%)	Mixed interaction threshold (%)	Initial allocation
Number of interactions only			
5	N/A	N/A	Automated
5	N/A	N/A	Human
5	N/A	N/A	None
No minimum flight threshold			
5	N/A	70	Automated
5	N/A	70	Human
5	N/A	70	None
No mixed interaction threshold			
5	25	N/A	Automated
5	25	N/A	Human
5	25	N/A	None

D. Dependent Measures

The following measures were established to assess the allocator robustness, sector flight distribution, allocator metric weighting, and safety of the sector.

- Sector aircraft distribution: One minute updates of the current aircraft count for each controller.
- Per controller total assignments: The total number of aircraft assigned to each controller.
- Reason for allocation: The relevant step in the Allocator flowchart at which an aircraft was assigned to a controller
- Number of relevant flights at moment of assignment
- Number of Mixed pair STCA/LOS occurrences
- Number of Automation pair STCA/LOS occurrences

E. Control Variables

The following variables were controlled during the simulation experiment in order to limit their influence on the measurements of the dependent variables.

- Sector characteristics: The Brussels sector group was used as the controllable sector for the full duration of the experiment.
- Automation characteristics: The logic used by the automation to direct traffic and solve conflicts.

F. Offline Simulation Results

After the completion of the offline simulation runs for all 51 variations of variables, two variations were identified to be investigated further with a follow-up experiment with human controllers. These two variations were chosen due to the largest discrepancy in calculated complexity of allocated flights to the human controller and automation. The results of these two variations in the offline simulation are discussed below. The changes to the complexity of human controlled flights and automated flights for both allocator models will be discussed, as well as changes in flight allocation distribution.

Lastly, a general discussion of any STCA/LOS occurrences between mixed pairs or automation pairs across all simulation runs gives more insight on the shortcomings of the allocator logic as well as the automation software used in this simulation.

The first scenario is the complexity only variant, in which aircraft are assigned purely based on the number of interactions calculated upon sector entry. The minimum number of flights to be considered complex is then the only metric that is used to allocate flights and the threshold value is the default of 5 flights. Both other allocator metrics are disabled for this setup.

The second scenario is when the allocator has a mixed interaction threshold of 50%, while the other threshold values maintain their default values. In this scenario, aircraft are primarily assigned according to which controller they will have more interactions with. The two allocators are sometimes referred to as allocator A and Allocator B for simplicity.

- **Allocator A:** Only complexity-based allocation. Based on the number of interactions threshold default value.
- **Allocator B:** Reduced mixed interactions allocation using a mixed interaction threshold value of 50% and all other default thresholds.

Figures 4a and 4b breakdown the per-minute distribution of aircraft currently in the sector. Allocator B allocates more aircraft to the human at the beginning of the simulation and consistently maintains this distribution until the later stages of the simulation run where the distribution evens out to nearly equal. Allocator A begins with more aircraft allocated to the automation, but slowly changes to a distribution with more human-controlled aircraft.

These can be used to give a general indication of the allocation distribution throughout the full duration of the experiment, and it was observed that there are no periods in which one controller has very little assignments compared to the other. The larger number of human aircraft, especially in the earlier stages of the simulation run are possibly due to a lack of human controller present in the simulation. Traffic controlled by the automation are given commands to steer them efficiently to the exit point after which they are transferred while the human controlled flights simply follow the flight plan without any further commands. This could lead to human controlled aircraft not taking efficient routes through the sector, inflating the number of human controlled aircraft in the sector compared to automated aircraft.

Figure 4c shows the reason for allocation for all aircraft assigned during the simulation run. For Allocator A, Figure 4c shows that a larger number of aircraft in this traffic sample had less than 5 relevant flights at the moment of allocation, leading to more flights being allocated to the automation overall. For Allocator B, a wider variety of allocation reasons are present due to all allocator metrics being active. The largest reason for allocating an aircraft is reducing the number of mixed interactions, accounting for 57% of assignments across the whole run. 38% of assignments are based on the number of relevant flights and the remaining aircraft are assigned due to the minimum flight distribution requirement.

The impact of the allocation reasons on the complexity of assignments at allocation time can be seen in Figure 5a. Here it can be seen that the complexity of aircraft assigned to human are more complex than those assigned to the automation for Allocator A. This is the expected result as complexity is the only metric responsible for allocation. With Allocator B, it was observed that aircraft assigned to the human are generally less complex than those assigned to the automation. In an effort to reduce the mixed interactions, the complexity of an aircraft is sometimes not considered when making an assignment which in this traffic sample led to easier assignment for the human overall.

Similarly to looking at the complexity of aircraft at allocation moment, the mixed interactions at allocation was also investigated and is plotted in Figure 5b. Interestingly, the number of mixed interactions at allocation is generally higher for the allocator that focuses on reducing mixed interactions than the complexity only allocator. Despite Allocator B allocating 57% of assignments according to this metric, it failed to produce an overall trend in reducing mixed interactions.

Investigating the mixed interactions of aircraft further, it was found that 18 out of the 70 assigned aircraft had a mixed interaction percentage of 50% at the moment of allocation in the experiment with Allocator B. Comparing this to allocator A, only 8 aircraft had a mixed interaction percentage of 50% at the moment of allocation. When the value is 50%, then the eventual allocation to either the digital or human controller has no influence on the number of mixed interactions as a significant number of mixed interactions will occur regardless. This is likely the reason as to why more mixed interactions occurred in the Allocator B setup.

In the experiment with human controllers it will be investigated whether the differences in allocator models are noticeable by the participants. Furthermore, the experiment with controllers will highlight changes to allocator behavior with the introduction of both controllers managing their assigned traffic.

The number of mixed pairs or automation pairs STCA/LOS occurrences gives an indication of the shortcomings of the combination of the allocator and the automation software. In this paper, the allocator is not coupled to the capabilities of the automation software. The allocator therefore independently assigns aircraft to the digital controller without a check to determine whether the digital controller is able to manage the aircraft.

Across the 51 simulation runs that were carried out, 57 instances of an STCA/LOS event occurred giving an average of 1.12 STCA/LOS events per scenario. 17 scenarios were event free, 17 scenarios had one event, 11 scenarios had two events, and 6 scenarios had three events. Of these 57 instances of an STCA/LOS event, only 6 of these were between two automated

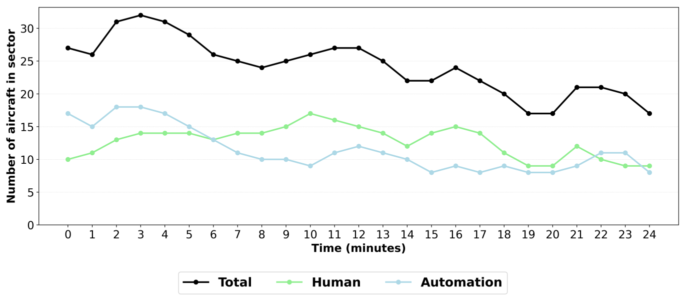

(a) Allocator A sector aircraft distribution per minute.

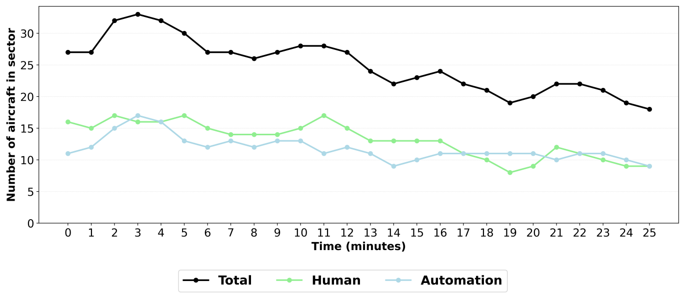

(b) Allocator B sector aircraft distribution per minute.

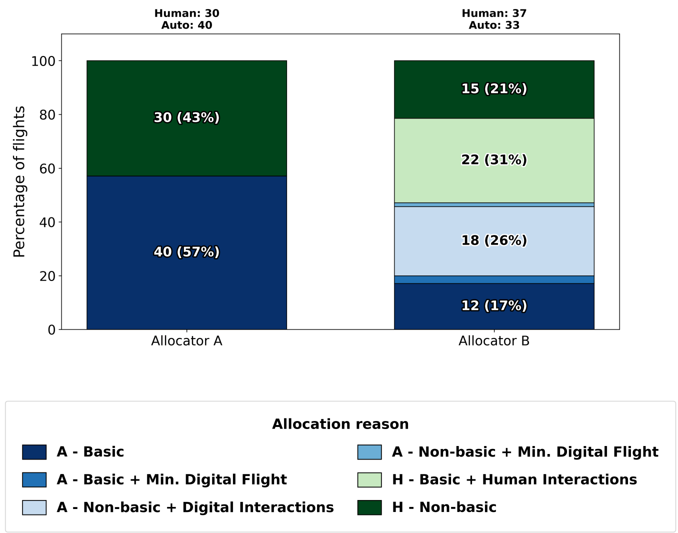

(c) Reasons for allocation for both allocators.

Figure 4: (a) and (b) show the evolution of the aircraft distribution in the sector over time for Allocator A and B respectively. The distribution only includes aircraft currently in the sector. (c) shows the distribution of reasons for allocating a particular aircraft to a controller with both allocators.

aircraft while the rest were mixed pairs. A automation pair event never occurred more than once per scenario.

In the majority of cases, conflicts arose shortly after an aircraft was assigned to the digital controller 60 seconds before the sector boundary. This was true for both mixed pairs as well as the automated pairs. The aircraft would be assumed upon sector entry and would immediately require a maneuver to avoid a conflict with an aircraft already in the sector. Due to the limited solution space of the automation software, due to its lack of ability to solve conflicts in the lateral plane, a STCA/LOS alert would occur.

Due to the relatively late assignment of the aircraft, and the lack of information of the allocator regarding the capabilities of

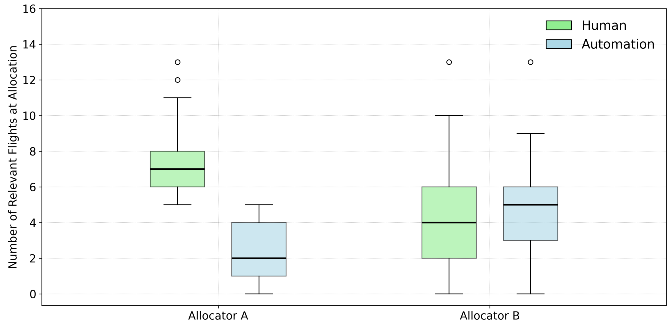

(a) Distribution of number of relevant flights at aircraft allocation moment for both allocators.

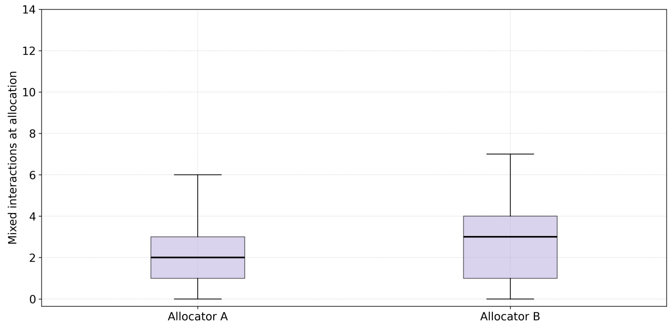

(b) Distribution of number of mixed interactions at allocation moment for both allocators.

Figure 5: (a) shows the distribution of the number of relevant flights at the moment of allocation divided by automated and human aircraft for both allocators. (b) shows the distribution of the number of mixed interactions at the moment of allocation for both allocators.

the digital controller means that some situations arise in which automation is not able to solve a conflict. In a live traffic sample, mixed pair conflicts will likely be spotted and handled by the human controller before STCA/LOS occurs. For automation pair conflicts, the human controller would need to step in to take over the aircraft and solve the conflict between the two. This is less favorable as the human will have to spend focus gaining a mental image of the two aircraft and the surrounding traffic to provide a safe maneuver.

Possible ways to avoid this include assigning and assuming aircrafts at an earlier time, such that the digital controller is able to carry out conflict resolution maneuvers before an STCA/LOS event occurs. A more suitable solution that would generally be beneficial for the allocator is to link the capabilities of the automation to the allocator. If the allocator is aware of situations where an aircraft has an incoming conflict that cannot be resolved by the digital controller, it can take this into account in the allocation process. These aircraft would only be assigned to the human controller, regardless of the allocation algorithm, ensuring that all aircraft that are eventually assigned to the digital controller are also manageable by the digital controller.

V. SIMULATOR FEATURES

Following the offline simulations, a human-in-the-loop experiment was carried out to gain more insight on the operations of the flight allocator in a more realistic scenario. The experiment utilized SectorX, which included several simulator features that the participants were able to use. The features available in the simulation are described in this section.

A. Relevant flights

The flight filtering mechanism is used by the allocation system to determine the number of relevant flights for each individual flight. Additionally, the flight filtering mechanism can also be used by participants as a visual support tool to indicate the relevant flights for a flight of interest. Figure 6 shows how the flight filtering mechanism can be used to fade irrelevant flights for a flight of interest.

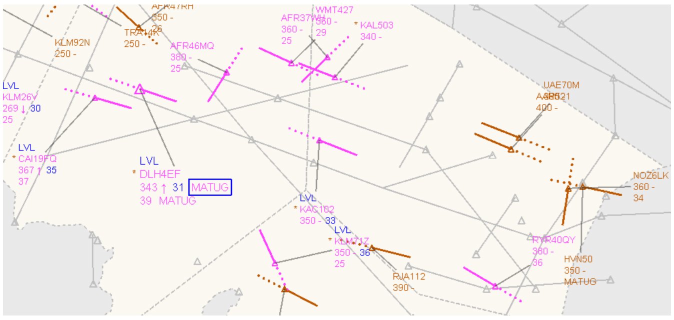

(a) Normal radar screen with no filtering (inverted colors).

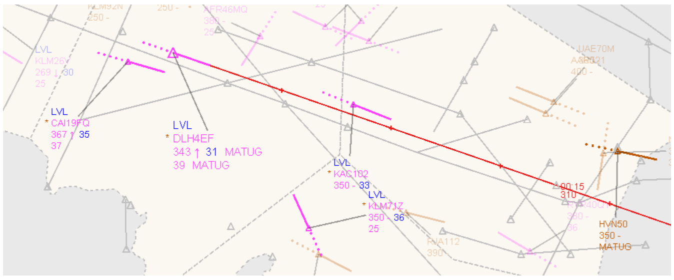

(b) Radar screen with flight filtering for flight DLH4EF (inverted colors).

Figure 6: (a) shows the normal radar screen in SectorX. (b) highlights the flight filtering mechanism whereby irrelevant flights are faded.

The flight filtering mechanism automatically activates upon interaction with the flight label of a flight. Additionally to displaying the relevant flights along its current trajectory, when a participant wants to provide heading or altitude changes, the impact of these commands on the relevant flights can also be previewed.

B. Advice for re-allocation

Once an aircraft has been allocated by the system, it will never be re-allocated unless the participant manually re-allocates the aircraft. However, the allocation system continuously calculates the relevant allocator parameters for each aircraft and provides an allocation advice for aircraft that could be good candidates to give from the participant to the automation or to take from the automation as the participant. This advice would be given in the form of a star to the left of the callsign of the aircraft. An example of this feature in SectorX can be seen in Figure 7.

If the allocator calculates that an aircraft is now better-suited for the opposite controller, the star will appear in the color of the advised controller (blue for advice to give to automation and green for advice to give to human). Figure 7 shows three aircraft,

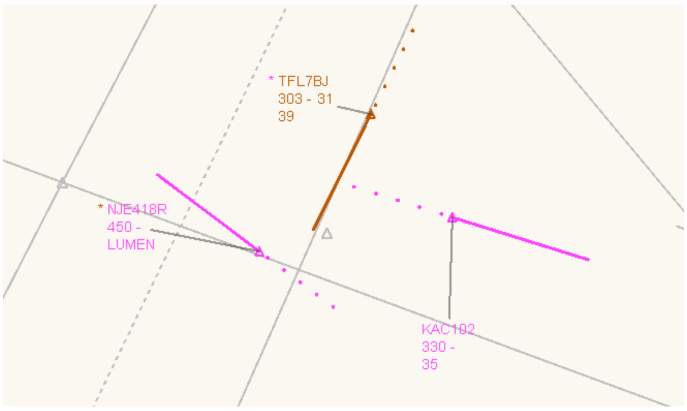

Figure 7: Re-allocation advice in SectorX (inverted colors). Pink aircraft are controlled by the H-ATCO and the auburn aircraft are controlled by the D-ATCO.

in which both forms of advice are present as well as an aircraft for which no re-allocation advice is provided. The participant could choose to use this advice to find suitable candidates for re-allocation but could also ignore the advice given by the system completely.

VI. EXPERIMENT SETUP

To gain more insight on the operations of the flight allocator tool, as well as its usefulness in dividing traffic responsibilities between two controllers, an exploratory experiment was conducted to test the concept in the enroute ATC domain. SectorX was used, a Java-based simulator developed by TU Delft, mimicking the enroute display currently used at MUAC. The aim of the experiment is to gather objective and subjective data to comprehensively assess the effectiveness of the flight allocator in meeting the operational demands of controlling traffic safely and efficiently. More insights are gained on the requirements and considerations needed for creating a live-traffic allocator.

A. Participants and Tasks

Two professional MUAC ATCOs participated. Both participants were working in the Brussels sector group, the same sector group used for the simulation. The trials were scheduled during office days of the ATCOs, when the ATCOs perform tasks other than ATC, or at the end of their ATC shift. During the experiment, the standard MUAC radar screen was mimicked using SectorX, which could be controlled with a computer mouse and keyboard inputs.

Participants performed the full range of enroute ATCO duties of managing the traffic flow safely and efficiently within the Brussels sector group. This involves assuming incoming flights, issuing any clearances necessary to ensure all flights can safely reach their exit point at the correct transfer flight level (TFL) and then transferring control of outgoing flights. Participants were able to manually re-allocate flights at any given time, transferring control from themselves to the digital controller or vice versa.

Participants were not informed about the allocation strategies prior to the experiment. Instead, participants were asked to observe the incoming flight allocations to determine whether they could identify factors that were used in the allocation scheme.

During the experiment, participants had access to MUAC's VERA tool, assisting participants in monitoring a flight pair over an extended period and assessing if current headings result in conflicts. Participants were informed that the automated flights would try to issue maneuvers that do not cause conflicts with human controlled flights. In the event of a mixed conflict, the participant would be responsible for solving the mixed conflict. Alongside the VERA tool, participants also had access to the flight filtering tool designed by Kumbhar [10]. This could be used to fade all flights that are deemed irrelevant to the flight of interest.

Throughout the experiment, participants provided Instantaneous Self Assessment (ISA) scores every 3 minutes. These scores indicated their perceived workload at that specific moment in time while managing the simulated traffic. Additionally, after completing a simulation run, participants were asked to fill out a survey regarding their experience with that simulated run. After both simulation runs were completed, the participant completed an additional survey comparing their experience with the two traffic allocator models.

B. Independent Variables

The experiment aimed to determine the effect of using different allocator models on the safety and efficiency of operations in the sector as well as the effect of the different models on the workload experienced by the human controller. Two allocator models were used, as described in Table 5. Allocator A looks at complexity of incoming flights while Allocator B prioritizes reducing the number of mixed interactions for controllers.

Table 5: Experiment allocation models

Number of interactions threshold	Min. flight threshold (%)	Mixed interaction threshold (%)	Allocator
5	N/A	N/A	Allocator A
5	N/A	50	Allocator B

C. Dependent Measures

The following measures were established to assess the sector safety, efficiency and workload experienced by the participants while controlling the traffic with both allocator models:

- VERA Count: The count of flight pairs for which the VERA tool is activated by the participant to monitor and detect conflicts.
- Total Clearances: The number of clearances issued by participants while simulating air traffic.
- Clearance menu inspections: The total number of times and the duration participants opened the clearance menu to provide clearances.
- Number of manual allocations: The total number of manual allocations carried out by the controller.
- Time between model allocation and manual re-allocation: The average time taken before the participant decides to manual re-allocate an aircraft that has been assigned by the system.
- STCA duration: The total duration during which the STCA tool remained active during conflicts.
- LOS duration: The total duration of a LOS violation.

- Average perceived workload: Workload scores provided by participants throughout the simulation. Compared per participant between the two simulations.
- Subjective Responses: The responses to the survey regarding the trust and reliability of the flight allocator and the digital controller, and the usability of the flight allocator in delegating traffic responsibilities to either controller.
- Open Feedback: The responses to the open feedback questions explaining strategies used by the participants and suggestions to improve the flight allocator model and the automation software.
- Sector aircraft distribution: One minute updates of the current aircraft count for each controller
- Total controller assignments: The total number of aircraft assigned to each controller
- Reason for allocation: The relevant step in the Allocator flowchart at which an aircraft was assigned to a controller.
- Number of relevant flights at moment of assignment.

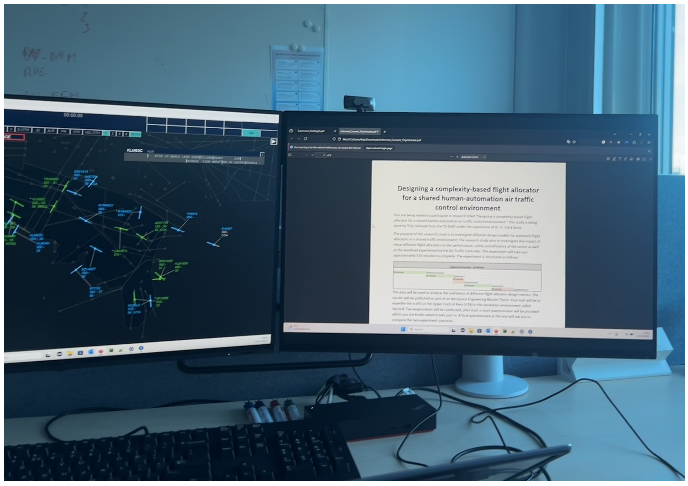

Figure 8: Hardware setup for the experiment

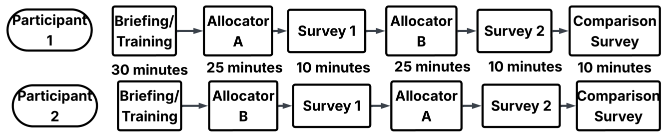

Figure 9: Experimental Procedure

D. Control Variables

The following variables were controlled during the experiment to account for the variability in the enroute ATC domain with the introduction of the automatic flight allocator concept:

- Airspace and traffic sample: As described in Section IV.A.
- Automation characteristics: As described in Section IV.B.
- ISA scale update: Participants utilized the ISA scale to input their perceived workload at a given time, with updates occurring every 180 seconds.
- Secondary Tools: The VERA tool, STCA, and the flight filtering tool were available
- Radar Update: The radar update frequency was set to every 10 seconds.
- Simulation Speed: The simulation speed was real-time.
- Work Environment: The apparatus used for the experiment was kept consistent.

E. Apparatus

The experiment utilized a single computer with a multi-display setup as shown in Figure 8. A single monitor with a 16:9 aspect ratio (1920 x 1080) was used for the experiment. A secondary monitor contained the survey, as well as info about the simulation software that the participant could use if needed. Throughout the experiment, participants issued clearances through data links and no radio communication. The simulation environment was controlled using a standard mouse and keyboard. Participants could issue flight level, heading, and direct-to-exit clearances but could not alter the speed or route of simulated traffic.

F. Procedure

Both participants followed the procedure outlined in Figure 9, starting with a short briefing indicating the purpose of the experiment and the role of the participant. This was followed by a training session in which the participant was exposed to the simulator environment with a greatly reduced traffic sample. During the training, the participant could familiarize themselves with the interface and practice their designated tasks and familiarize themselves with the tools available to them. The automation was also introduced in the training session allowing the participant to familiarize themselves with the automation software. Conflicts were shown to demonstrate how automation would handle the situation.

Following the training, the first experimental run began with the first flight allocator model. The participant was responsible for managing the traffic inside the sector for 25 minutes. After the first experimental run the participant evaluated their experience in a brief survey. Then, the participant continued with the second experimental run with the second allocator model. Concluding the experimental run was an identical survey to the first one, followed by a final survey comparing the two experimental runs. The order in which the flight allocator model was shown to the participants was switched between the two participants. During the full experiment procedure, the participant could ask any questions about the simulator environment and automation logic.

VII. RESULTS

The two participants each carried out the experiment, managing the traffic in the Brussels sector for two different allocator models. The two participants will be distinguished by referring to them as Participant 1 and Participant 2. The two experiment scenarios are also distinguishable by referring to them based on the allocator model. The first is the complexity only allocator model (sometimes referred to as Allocator A), only using the number of interactions to assign incoming aircraft. The second allocator model has a mixed interaction threshold of 50%, i.e. the model that aims to reduce mixed interactions (sometimes referred to as Allocator B).

During Participant 1's trial with the reduced mixed interactions allocator, the simulation crashed approximately 21 minutes into the run. While data was still recorded for those 21 minutes, data for the last 4 minutes of the trial could not be collected. The remaining three experiment runs were carried out to completion. Participant 1 first experimented with the complexity model, followed by the reduced mixed interaction model. Participant 2 had the opposite order, first using the reduced mixed interaction model, followed by the complexity model.

Both participants had vastly different approaches to the experiment. Participant 1 struggled to understand and deal with behaviors from the allocation system and automation that did not align with their own vision and expectations. Participant 1 chose to 'fight' with the system when they discovered flaws and weaknesses in the automation software as well as the allocation algorithms. This led to an overall more intense experience for Participant 1, requiring a large amount of mental work to manage the traffic safely. Participant 2 also saw choices made by the automation and allocation system that they did not agree with or understand. However, instead of trying to fight with the system to enforce their own vision, Participant 2 took a more adaptive approach, accepting most of the choices made by the system and automation and choosing to alter their approach to align with the automated systems in place.

These differences in approach led to two very different experiences for both controllers, which will be examined further in this section. This section will discuss the workload experienced, safety of the sector, manual re-allocations, and survey responses provided by the participants.

A. Workload

The workload is evaluated through subjective ISA score responses provided every 3 minutes by the participant, an overall workload rating scale at the end of the experiment, and by looking at the number of clearances and menu inspections were carried out by the participant. The ISA scores can be found in Figure 10. Participant 1 noted similar scores throughout both experiment trials, with the reduced mixed interaction model having a slightly lower workload score than the complexity model. Participant 1 noted that the experiment was quite challenging due to the high traffic density of the scenario. While it is difficult to tell from the ISA scores alone whether Participant 1 experienced a generally lower workload with the less mixed interactions allocator, the responses to the survey found in Figures 11 and 23 indicate that Participant 1 indeed experienced a lower workload with the reduced mixed interactions allocator.

Participant 2 found the scenario with the complexity only allocator model generally easier to manage, with a more clear distinction in ISA scores across the majority of the experiment run. This is also supported by their responses to the survey found in Figures 11 and 23. Both participants found the workload to be in an acceptable range for the reduced mixed interactions allocator, while Participant 1 experienced a slightly high workload and Participant 2 a slightly low workload in the complexity only allocator.

In general, Participant 2 found both experiment scenarios easier to manage in terms of workload than Participant 1. This is possibly due to the approach mentioned at the beginning of this section, in which Participant 2 adapted to the quirks and shortcomings of the automation system as well as the allocator logic while Participant 1 challenged a lot of the flaws. Both participants experienced a lower workload in their second experiment run, possibly suggesting that the familiarization of the automation system and the simulator environment due to exposure and experience with the simulator reduced the workload for both participants.

Furthermore, the scenarios used in this experiment were identical except for the allocator model present. This could mean that participants gained some familiarity with the scenario and traffic sample, although this was not explicitly mentioned by either participant. From the results of the subjective workload responses, it is unclear whether one allocator would generally result in a lower workload over the other.

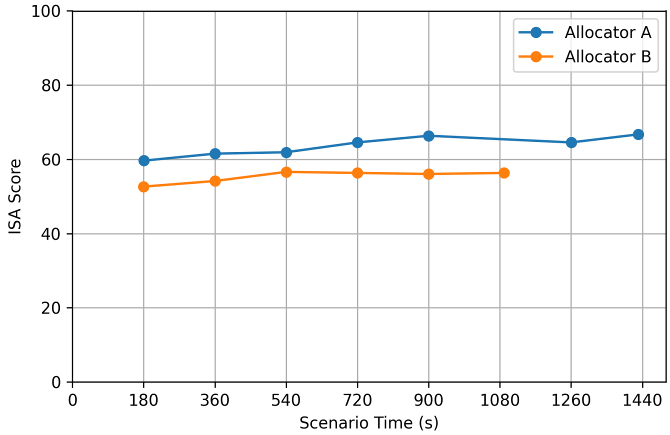

(a) Participant 1

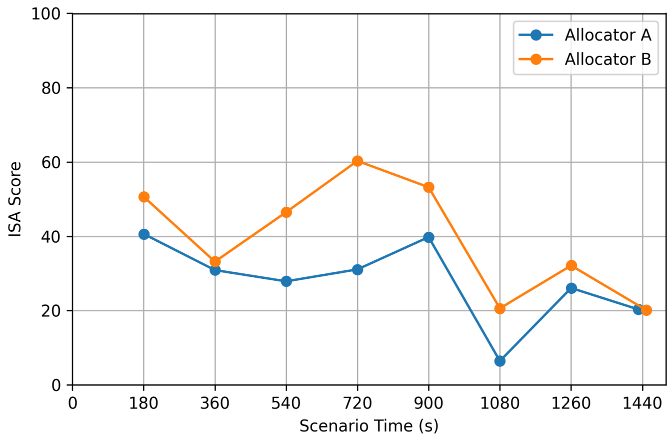

(b) Participant 2

Figure 10: ISA scores of both participants across both experiment scenarios. Participant 1 (a) has a missing value for t=1080s for the complexity allocator and missing values for t=1260s and t=1440s due to simulation crash.

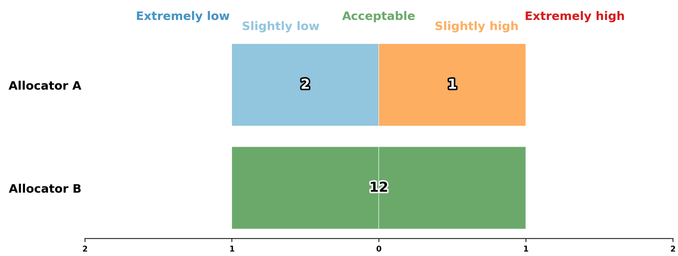

Figure 11: Overall workload rating provided by each participant at the end of the experiment.

Alongside the subjective workload ratings provided by both participants, the number of interactions each participant had with the simulator environment can also provide an indication of their workload. Data collected regarding this included the number of times a flight label was selected, the number of times a preview of a flight command was carried out, the number of flight commands executed, and the number of VERA checks carried out by each participant. The results of this can be found in Table 6.

Looking at these results, it can be seen that the number of interactions with the simulator environment was greatly reduced for Participant 1 in their second scenario run with the reduced mixed interactions model (Allocator B). They executed slightly more flight commands, but the other variables tracked were all significantly lower. Participant 2 had a slightly more consistent number of interactions with the simulator environment. However, their interactions during the second scenario run in which they deal with Allocator A, the complexity only model, they had consistently less interactions with the simulator environment across all variables tracked.

These results show the opposite results of the ISA scores provided by both participants, in which Participant 1 indicated experiencing a slightly lower workload with Allocator B than with Allocator A, though the interaction counts show a significantly reduced number of interactions with Allocator B over Allocator A. The reverse is true for Participant 2, who indicated a more significant difference in ISA scores between Allocator A and B, yet had a smaller change in the number of interactions across both scenarios. This suggests that additional unmeasured factors contributed to the mental load experienced by the participants.

Table 6: Interaction counts with the simulator environment for both participants across both experiment scenarios. (PX: Y, represents participant X with allocator Y)

	P1: A	P1: B	P2: A	P2: B
Flight Label Selection	626	479	430	498
Flight Command Preview	367	266	222	249
Flight Commands	52	60	49	55
VERA Checks	27	6	32	47

B. Safety

The safety of the sector is evaluated through the investigation of any STCAs or LOS occurrences during the course of the experiment. The minimum separation between the two aircraft and the duration of the STCA or LOS occurrence can give an indication of the severity of the situation. Participant 1 had one STCA that turned into a LOS situation during the first experiment run with the complexity-based allocator. Participant 2 had a single STCA in both scenarios, with the STCA during the experiment with the reduced mixed interactions allocator also having LOS.

Having reviewed all three STCAs that arose, one created a dangerous situation. This situation was between a mixed interaction pair, and while a conflict resolution maneuver was possible, the lack of clarity regarding the automation's inability to solve the conflict made it such that the controller did not act on the STCA.

The STCA occurring in Figure 12b shows two converging aircraft with the aircraft controlled by the H-ATCO climbing to FL390 from FL226 while the aircraft controlled by the D-ATCO is descending to FL250 from FL262. The command to descend aircraft KLM24X to FL250 was given by the H-ATCO, shortly before the aircraft entered the sector. Participant 1 assumed the aircraft before the allocator had allocated the aircraft 60 seconds prior to sector entry. Once the aircraft reached the allocation point, the allocator allocated the aircraft to the automation, despite the aircraft already being assumed by Participant 1. This was not an intended feature of the allocator and will be discussed further in the subsection regarding manual re-allocations.

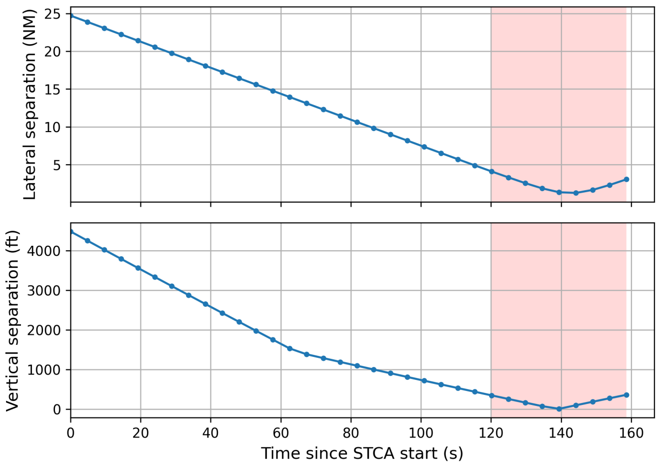

(a) Participant 1, Allocator A: STCA Duration

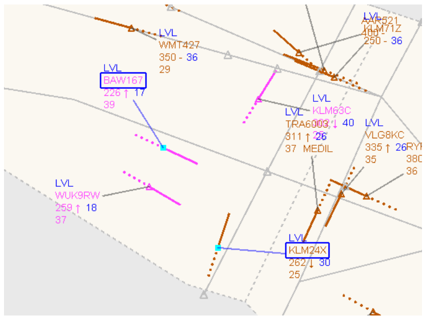

(b) Participant 1, Allocator A: Conflict Geometry (Inverted colors)

Figure 12: STCA duration and conflict geometry for Participant 1 and Allocator A. Shaded red areas indicate the time period for which a LOS has occurred. Pink aircraft are controlled by the H-ATCO and the auburn aircraft are controlled by the D-ATCO.

Once the aircraft was assumed by the automation, the automation also gave a direct to the exit point, altering the trajectory of the aircraft to have the heading currently displayed in Figure 12b. This led to the potential STCA as both aircraft involved were closing in on each other both vertically and laterally. The conflict could have been resolved by leveling off one of the aircraft temporarily until it would be safe to climb/descend again. The conflict was not resolved by either controller, leading to a very dangerous and close encounter, with the LOS being breached for almost a minute.

The closest point of approach (CPA) was less than 1NM at the peak of the conflict and both aircraft were on the same FL. It is unclear why this conflict was not resolved by Participant 1. The human controller should have noticed that the digital controller was not solving the conflict and provided a resolution but did not do so. A possible explanation is that the participant wanted to investigate the response of the automation and was potentially waiting for an indication that the conflict could not be resolved by the automation.

The simulator experiment did not include an indication of when automation was unable to perform conflict resolution tasks. In general, the automated traffic was responsible for avoiding conflicts with human aircraft when providing maneuvers to aircraft, although maneuvers made by the human could result in conflicts that the automation was incapable of solving. This conflict could be a consequence of the ambiguity related to which controller is responsible for solving a mixed conflict.

C. Manual Re-allocations

Participants were able to make use of manual re-allocation at any point during the experiment. Participants could choose to take aircraft from the automation and control it themselves, or to give aircraft under their control to the automation. Participants were also asked to provide a general reason for manual re-allocating flights in the survey following the experiment.

One issue that arose during the experiment is that when a participant chose to assume an aircraft before the system had allocated it, i.e. the aircraft was assumed more than a minute before sector entry, the system would still allocate the aircraft once it had reached the 1 minute boundary. This led to automation 'stealing' certain aircraft from the participant. A stolen aircraft refers to an aircraft that the participant had assumed prior to the allocation moment of the system but was assigned to the automation at the system allocation moment thus resulting in a re-allocation by the system.

Table 7 shows the number of times that Participant 1 assumed an aircraft before the system allocated the aircraft, the number of aircraft that were re-allocated by the participant at some point after the system allocation and the number of aircraft that maintained their system allocation throughout their whole flight path. Table 8 shows the same data for Participant 2.

Participant 1 assumed a significant number of flights before the 1 minute pre-sector boundary at which the system would allocate the aircraft in both experiment scenarios. Participant 1 experienced more difficulties and frustration with the allocation system due to the 'stolen' aircraft. Participant 1 noted that they were not happy with the overrides of their allocations. They particularly did not enjoy that they would be working on their plan for a particular aircraft only to then have it taken over by automation and it carrying out a different plan. This meant that Participant 1 had to adapt their plans often, which may have contributed to the higher workload they experienced.

Table 7: Participant allocation actions for Participant 1 (counts with percentage of total allocated per allocator).

	Allocator A System Allocation		**Allocator B** System Allocation	
	Automation	Human	Automation	Human
Assumed before allocation	7 (19%)	8 (24%)	8 (33%)	17 (42%)
Manual re-allocation	12 (32%)	24 (73%)	4 (17%)	15 (38%)
Not re-allocated	18 (49%)	1 (3%)	12 (50%)	8 (20%)
Total allocated	37	33	24	40

The number of aircraft manually re-allocated by Participant 1 for the complexity-based allocator was 36, roughly 51% of all aircraft in the simulation were re-allocated at some point by Participant 1. In the reduced mixed interactions allocator, 19 aircraft were re-allocated, roughly 30% of all aircraft in the simulation. While the number of aircraft re-allocated can be an indicator of the acceptance of the allocator assignment, it is also important to consider at what point in time the manual re-allocation occurs. System allocations that are quickly overridden by the human indicate a distrust in the allocation made by the system.

A common tactic for the participants was to re-allocate human controlled aircraft when all potential conflicts and maneuvers required had already been carried out. In this way, the automation was simply left with transferring the aircraft as it left the sector. Re-allocations of this nature do not necessarily reflect poorly on the initial allocation, but are used as a tool by controllers to free up mental load for other aircraft.

In order to evaluate the 'survival rate' of an aircraft allocation, a Kaplan-Meier curve has been constructed. The Kaplan-Meier curve indicates the probability that an allocation remains unchanged as a function of time [21]. The survival curve ends at the latest moment between a system allocation and manual re-allocation for that scenario. This is calculated using the data of all aircraft that are either manually re-allocated or not re-allocated at all and excludes aircraft that were assumed before allocation. Equation 1 shows how the probability that the system allocation survives longer than a time t is calculated. The results of this calculation be found in Figure 13.

$$\hat{S}(t) = \prod_{t_i \leq t} \left(1 - \frac{d_i}{n_i}\right) \quad (1)$$

where $t_1 < t_2 < \cdots < t_k$: are the ordered times at which an allocation is overridden, with k the total number of re-allocations,
d_i : is the number of re-allocations occurring at time t_i,
n_i : is the number of aircraft not re-allocated immediately prior to t_i.

Each step in the curve indicates a moment in which an aircraft has received a manual re-allocation. Large steps indicate that multiple aircraft have been manual re-allocated at similar time frames since initial allocation. A time of 120 seconds after system allocation will be used as a metric to determine how many assignments are overridden in the early stages of aircraft entry in the sector. The first minute the aircraft is in the sector will be as the cut-off to determine whether the allocator has made an appropriate allocation according to the participant.

From Figure 13, it can generally be seen that the chance that an allocation is overridden for any aircraft 120 seconds after assignment is roughly 40%. Aircraft assigned to the automation in both experiment scenarios that had their allocation overridden always occurred within the first 80 seconds of the system allocating it. Figure 13 only includes aircraft that were not assumed before allocation of which 40% of the aircraft allocated to the automation with Allocator A did not last more than 80 seconds before being re-allocated. The remaining 60% of aircraft assigned to automation maintained that assignment throughout the experiment.

With Allocator B, all automated aircraft that were manually reallocated were done within the first minute. 25% of aircraft assigned to the automation were overridden within the first minute, while the remaining 75% were never reallocated. While the majority of automated assignments in both scenarios are maintained throughout, a significant number of automated assignment are manually re-allocated by the human quickly, indicating that Participant 1 is often not in agreement with the assignment provided by the system.

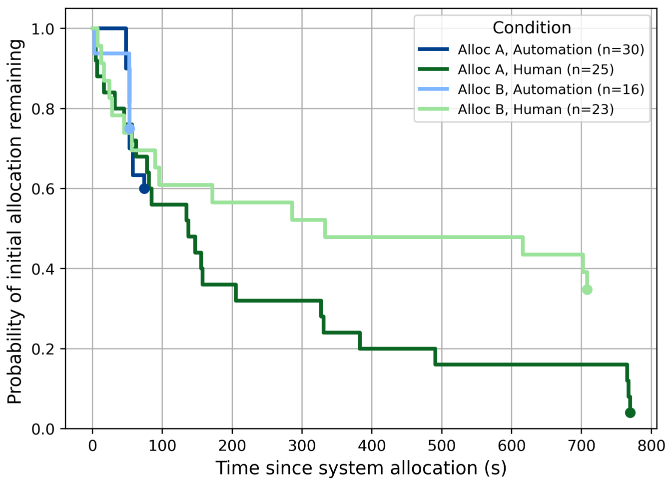

Figure 13: Kaplan-Meier survival curve for Participant 1 experiments. Split by allocator and assignment. The number of aircraft used for the calculation per curve are included in the legend.

Looking at the human assigned aircraft, a similar trend can be seen in the first 120 seconds. The chance that an aircraft is re-allocated within this time frame is 40%. As the aircraft continues to move through the system, the chance of re-allocation increases. With the complexity-based allocator, an aircraft assigned to Participant 1 had an 80% chance of being reassigned by the time 400 seconds had passed since the initial allocation. With Allocator B, this chance was around 50% in the same time frame. This highlights the trend seen in both participants to manually re-allocate their assignments after they had solved all maneuvers. This is consistent with the reasoning provided in the survey in which Participant 1 indicated that they transferred flights to automation after "solving the complex part of their flight path."

Table 8 indicates the manual re-allocation data for the experiments of Participant 2. Participant 2 assumed far less aircraft before the system allocation than Participant 1. Across both experiment scenarios only 8.6% of aircraft were assumed prior to system allocation, with the majority of these occurring during Participant 2's first experiment trial with the reduced mixed interaction allocator. Participant 2 chose to manually re-allocate aircraft more often with the complexity-based allocator. In this experiment scenario, 67% of aircraft are re-allocated Participant 2 at some point during the run. With the reduced mixed interactions allocator,roughly 47% of aircraft are re-allocated by Participant 2 at some point during the experiment run.

Table 8: Participant allocation actions for Participant 2 (counts with percentage of total allocated per allocator).

	Allocator A System Allocation		**Allocator B** System Allocation	
	Automation	Human	Automation	Human
Assumed before allocation	1 (2%)	2 (7%)	5 (18%)	4 (10%)
Manual re-allocation	24 (60%)	23 (77%)	11 (39%)	22 (52%)
Not re-allocated	15 (38%)	5 (17%)	12 (43%)	16 (38%)
Total allocated	40	30	28	42

Figure 14 indicates the probability of a system allocation not being overridden as a function of time for the experiments with Participant 2. For the complexity-based allocator, it can be seen that the probability of an automated aircraft not being re-assigned within the first 120 seconds is less than 50%. If the automated aircraft is not manually assumed by Participant 2 within the first 120 seconds of allocation, then its chance to keep its system allocation stabilizes, dropping to roughly 40% at 400 seconds after allocation time.

For aircraft initially assigned to the human in the complexity-based allocator scenario, the chance of maintaining allocation within the first 120 seconds is roughly 40%. Beyond this, the chance of maintaining allocation stabilizes to just above 30%, slowly dropping to roughly 20% beyond 400 seconds.

In the case of the reduced mixed interactions allocator, it was observed that all automation re-allocations occur within the first 60 seconds, with the chance of an aircraft not being re-assigned is approximately 50% within the first minute. For aircraft assigned to the human, the chance to not be reassigned within the first 120 seconds is approximately 45%, and only slightly drops as the last manual re-allocation of a human controlled aircraft to the automation occurs at around 150 seconds after system allocation.

During Participant 2's first experiment run with the mixed interactions allocator, they chose to maintain all their assignments past 150 seconds of allocation. Participant 2 disagreed with a lot of initial allocations provided by the system, seen in the large number of re-allocations carried out within the first 120 seconds of system allocation. In their experiment with the complexity-based allocator, Participant 2 chose was more comfortable with assigning aircraft to the automation in the later stages.

Participant 2 noted that their re-allocation strategy for the reduced mixed interactions allocator was to "transfer low-workload flights to automation, and assume flights that were relevant for each other (succeeding departures)". For the complexity-based allocator, a slightly different strategy was adopted: "once I had solved the conflicts of certain flights and they became low workload, I would transfer them to the automation (departures reaching cruise level)." The difference in the two strategies is that with Allocator B, Participant 2 would evaluate the workload at the start and transfer flights to automation accordingly while with allocator A, Participant 2 would re-evaluate workload constantly and transfer flights when the complex parts were solved. This can also be seen in the survival curves of the allocation for both experiment scenarios.

D. Allocation Distribution

The allocation distribution statistics that were discussed in the offline simulation experiment chapter are also investigated for the human-in-the-loop experiment. Figure 15 shows the distribution of aircraft assigned to each controller by the system as well as the actual distribution of aircraft in the sector for Participant 1 and both allocators. Participant 1 consistently chose to allocate more aircraft to the digital controller, particularly in the later stages of the simulation run.

Figure 16 shows the assigned and actual allocation distribution per minute across both allocators for Participant 2. Participant 2 also chose to allocate more aircraft to the digital controller with allocator A. During the experiment with Allocator B, Participant 2 kept the distribution very similar to the system allocation distribution until 15 minutes, after which the

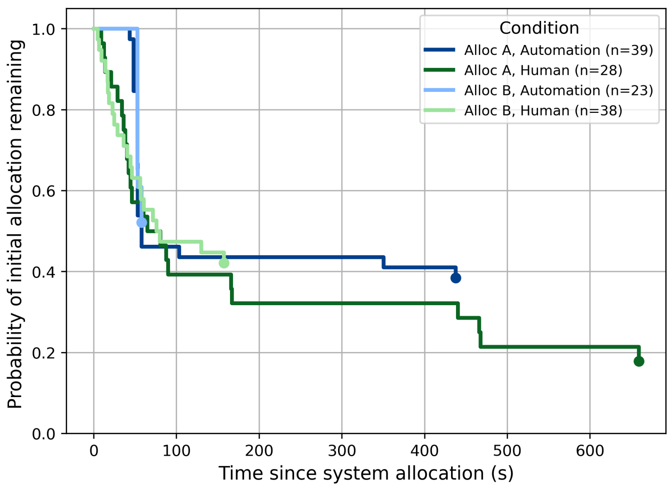

Figure 14: Kaplan-Meier survival curve for Participant 2 experiments. Split by allocator and assignment. The number of aircraft used for the calculation per curve are included in the legend.

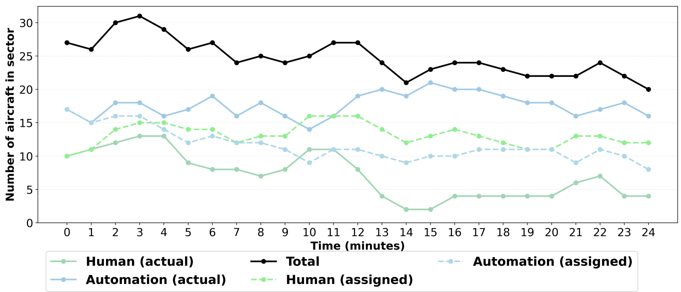

(a) Assigned vs actual aircraft distribution for Allocator A (Participant 1).

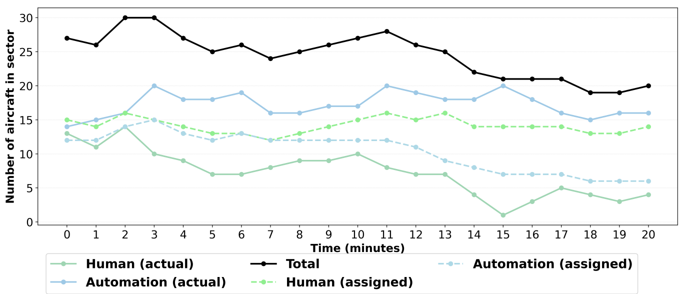

(b) Assigned vs actual aircraft distribution for Allocator B (Participant 1).

Figure 15: Distribution of aircraft in sector every minute for Participant 1 with Allocator A (a) and Allocator B (b).

participant decided to allocate a large number of their aircraft to the digital controller.

The reason for allocation as well as the overall distribution of the experiment scenarios can be found in Figures 17a and 17b. In Figure 17a, the similar distributions of aircraft as in the offline simulation can be seen. The system allocation is not influenced by the number of aircraft assigned to each controller, but only by the predicted number of interactions for the incoming aircraft. Despite a human controller now being present to give their assigned aircraft flight commands compared

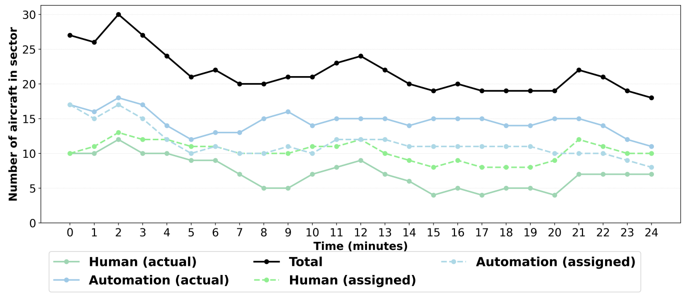

(a) Assigned vs actual aircraft distribution for Allocator A (Participant 2).

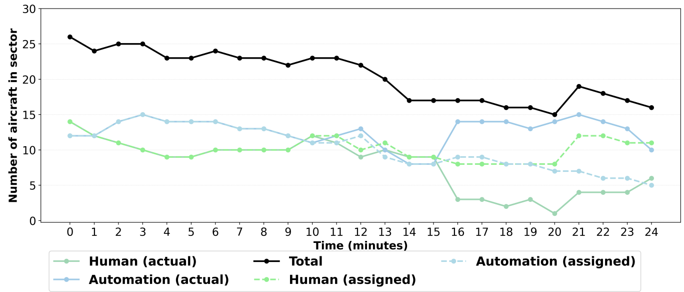

(b) Assigned vs actual aircraft distribution for Allocator B (Participant 1).

Figure 16: Distribution of aircraft in sector every minute for Participant 2 with Allocator A (a) and Allocator B (b).

to the offline simulation, the distribution of aircraft remained relatively similar.

In Figure 17b, it can be seen that the minimum flight distribution requirement comes into play for both participants. Due to their preference of allocating a large number of their aircraft to the automated system, it means that incoming aircraft are automatically allocated to the participant. This was the case for 15% of the total aircraft for Participant 1 and 7% of the aircraft for Participant 2, while this metric was not used in the offline simulation to assign aircraft to the H-ATCO.

Figures 18a and 18b shows the distribution of the number of relevant flights at the moment of allocation for both participants and the offline simulation for allocator A and Allocator B respectively. For allocator A, the distribution remains relatively consistent throughout, with the median value being roughly seven across both participants and the offline simulation. Aircraft assigned to the digital controller had a median value of two across both participant experiments and the offline simulation.

Figure 18b shows larger differences between participants and the offline simulation regarding Allocator B. Participant 1 was generally assigned aircraft that were less complex than those assigned to the digital controller. Participant 2 had a relatively similar distribution to the digital controller, and their median values are closer together than in the offline simulation. The box plots show that Allocator B is more sensitive to the actions of the participant during the simulation while allocator A is relatively consistent between controllers.

Figures 19a and 19b show the number of mixed interactions at allocation moment. Similarities can be seen between the number of mixed interactions in the experiment and in the offline simulation, in which Allocator B actually has a higher number of mixed interactions at allocation.

Investigating the occurrences of the mixed interaction percentage being 50% at allocation moment reveals that

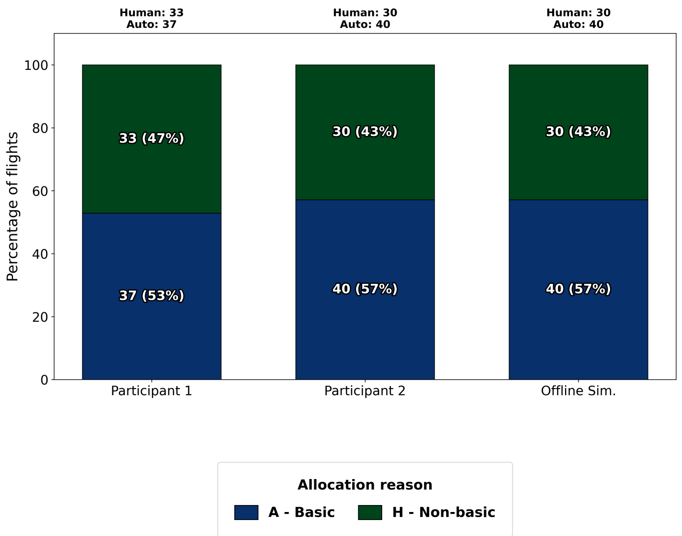

(a) Allocator A.

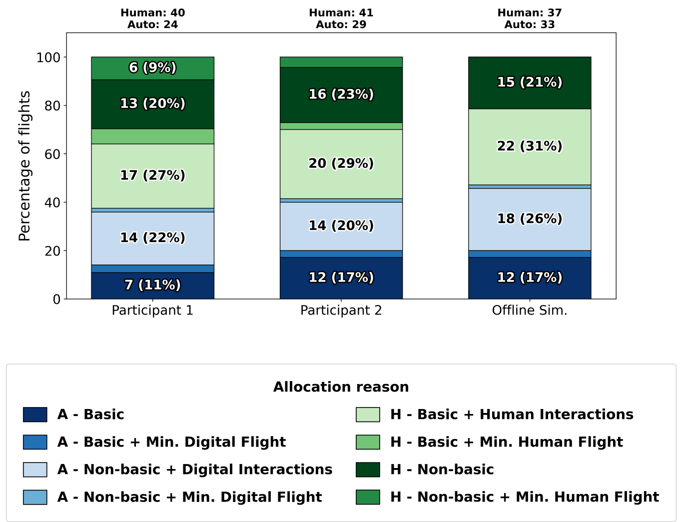

(b) Allocator B.

Figure 17: Distribution of allocation reasons for Allocator A (a) and B (b).

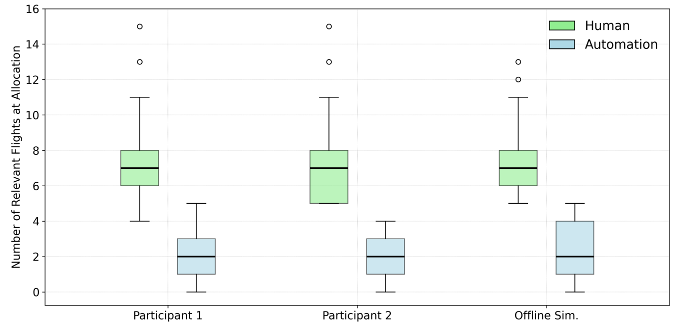

(a) Complexity at allocation moment, Allocator A.

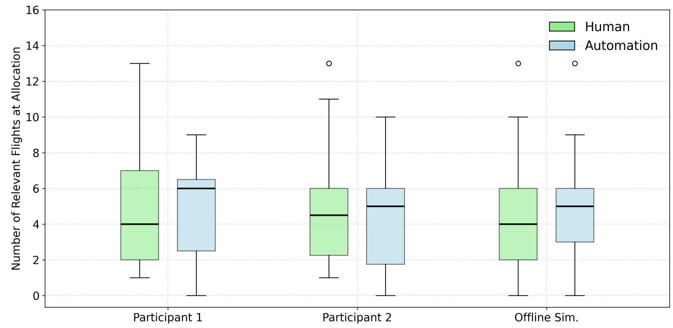

(b) Complexity at allocation moment, Allocator B.

Figure 18: Distribution of number of relevant flights at allocation moment across both allocator A (a) and B (b).

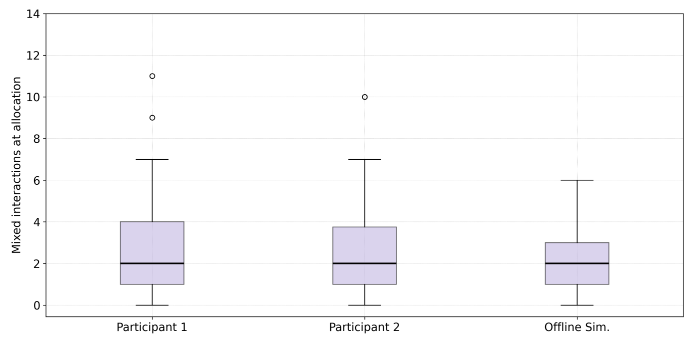

(a) Allocator A.

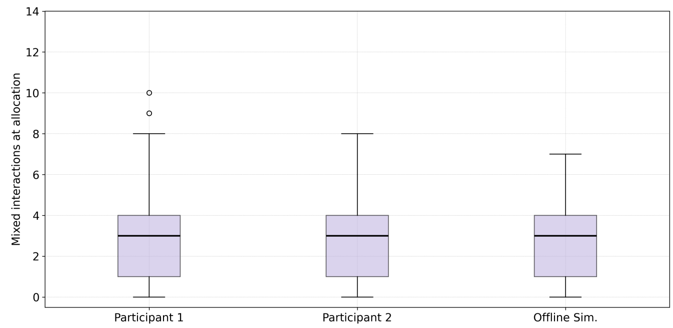

(b) Allocator B.

Figure 19: Distribution of number of mixed interactions at allocation moment across both allocator A (a) and B (b).

Allocator B had a larger number of aircraft at allocation moment with a mixed interaction percentage of exactly 50% as was the case in the offline simulation. Furthermore, the larger number of allocations based on minimum flight distribution could have also contributed to the higher number of mixed interactions seen in Allocator B.

E. Survey Results

Concluding both experiments, the participants were asked to evaluate their experience with the simulation. This included evaluating the allocation schemes, the automation software, and the fidelity of the simulator and its features.

Figure 20 indicates the general opinion of both participants with the allocator models experienced in both experiments. Participant 1 disliked both allocation schemes, and found them quite challenging to deal with. Participant 1 was unable to identify the metrics used in the allocator due to their high workload and focus on managing the traffic inside the sector.

Participant 2 somewhat liked the complexity-based allocator while somewhat disliking the less mixed interactions allocator. Participant 2 said that they found the allocations in the complexity-based allocator to be more logical and easier to understand why aircraft went to a particular controller. They sometimes struggled to understand the allocation choices made by the system in the reduced mixed interaction allocator model.

Participant 2 was able to correctly identify the main metrics driving the allocator choices. This is interesting as Figures 19a and 19b showed that Allocator B failed to reduce the number of mixed interactions overall compared to allocator A. Despite this, the participant was still able to correctly identify that this was the main metric driving the allocations.

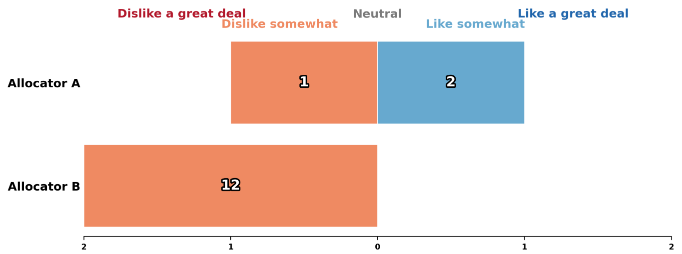

Figure 20: Participants opinion on allocator models.

Figure 21 shows the opinion of both participants regarding certain statements about their experience with the complexity-based allocator experiment. This data was collected directly after the participant completed this section of the experiment. Overall, Participant 1 experienced the complexity-based allocator experiment more negatively than Participant 2. This could be due to the nature of their approach to the experiment. Participant 1 found it very difficult to organize their work the way that they wanted due to the constant re-allocations by the system of flights that they had assumed before the system allocation moment. This made it difficult for them to solve conflicts and use work together with the automation. Participant 2 was more flexible in their approach and adapted to the system. This helped Participant 2 experience automation as a helpful tool rather than a hindrance that increased their workload.

Both participants had a strong distrust in the allocator to make good assignments and both had a strong distrust in the ability of the automation software to handle their assignments. This distrust is consistent with the high number of manual re-allocations carried out within the first 2 minutes of system allocation. Even though Participant 2 was able to adapt their approach to better deal with the allocator and automation features, both participants indicated that the allocator and automation could use further refinement to be useful. Participant 1 noted that the more insight into how the automation intends to handle flights laterally and vertically would help their overview of the sector.

Figure 22 shows the participants opinions regarding certain statements while working with the reduced mixed interactions allocator. As with the complexity-based allocator, Participant 1 had an overall more negative experience with the allocator than Participant 2. Participant 1 felt that the automation was able to handle their assignments much better with the reduced mixed interactions allocator over the complexity-based allocator. Participant 1 also noted that the allocation suggestions made by the system seemed to make more sense than in the complexity-based allocator, although this is not reflected when comparing the two Figures 21 and 22. Both participants indicated a distrust in the allocator to make good assignments, showing that neither allocation system was trusted by either participant.

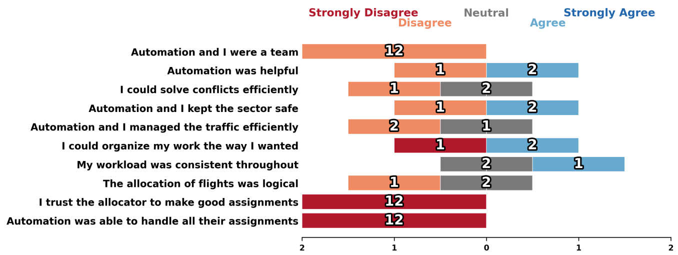

Figure 21: Participants evaluation of the complexity-based allocator.

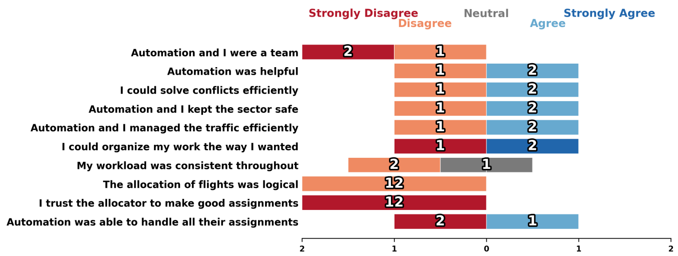

Figure 22: Participants evaluation of the reduced mixed interactions based allocator.

Figure 23 shows the results of the post experiment survey. This survey was carried out after the participant had completed experiments with each allocator and evaluated them. Participant 1 showed no preference regarding either allocator system, only noting that the reduced mixed interactions allocator had a slightly lower workload for Participant 1. Participant 2 had a general preference for the complexity-based allocator, though noting that the reduced mixed interaction allocator experiment yielded more success in managing traffic together with automation, and organization of their work.

This could suggest that the reduction in the number of mixed interactions experienced with Allocator B meant that the two controllers worked more independently alongside each other, making it easier to manage for Participant 2. Mixed interactions could cause frustrations with the human controller, with them having to adapt their plans to the outputs of the automation software, rather than the other way around. However, generally speaking, Participant 2 preferred the logic and consistency of the complexity-based allocator, also noting that they experienced a reduced workload in this scenario. Participant 2 indicated that they were more in agreement with the allocation decisions made by the complexity-based allocator, and was able to understand why certain allocations were made, with this not being the case for the reduced mixed interaction allocator.

Both participants indicated a lower workload experienced with their second experienced allocation algorithm. This could be due to several factors, such as familiarity of the traffic sample, as identical scenarios were used for both simulation runs. Other reasons for the experienced reduction in workload could be due to familiarization with automation software allowing the participant to more easily and consistently predict the commands provided by the automation.

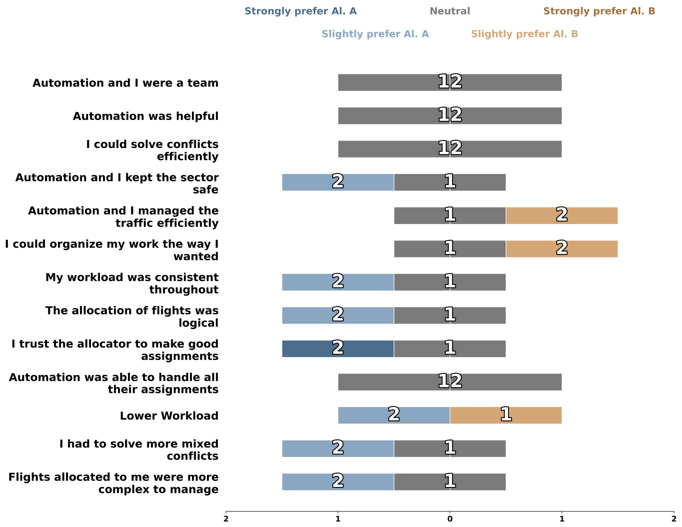

Figure 23: Survey responses comparing both allocators. Al. A is the complexity-only allocator model while Al. B is the reduced mixed interactions allocator model.

Figure 24 provides insight into the opinions of both participants on the automation software. Both participants indicated that automation was unreliable. Participant 2 found that the automation was somewhat predictable while Participant 1 did not find the automation software predictable at all. Participant 1 found the automation very inconsistent while Participant 2 had no strong opinions on the consistency of the automation software.

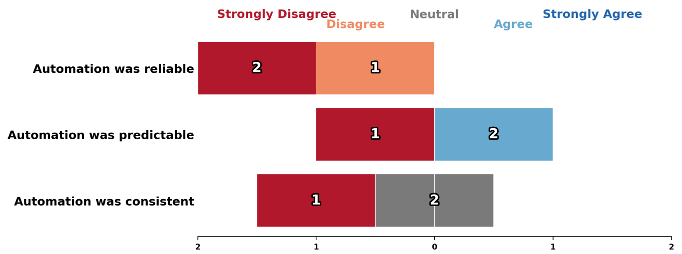

Figure 24: Participants opinions on the automation software.

F. Open Feedback

Both participants were asked to provide insight on their ideal design of an automatic flight allocator. Both participants mentioned that the number of interactions should be considered in the allocator, though noting that these interactions should be weighted. Participant 1 noted that simple conflicts are not reason enough to allocate a flight to the human controller, but that the conflict must also be complex. Participant 2 stated that the interactions should be weighted based on their complexity. Furthermore, Participant 2 indicated that priority of tasks should be considered, such as solving conflicts, inbound restrictions and climbing departures. The tasks required for each aircraft should be weighted by their priority and then allocated accordingly.

The flight filtering tool was used by both allocators to determine the number of relevant flights for an incoming flight. The tool could also be used by participants during the experiment as an aid to help them identify relevant flights when considering a particular flight maneuver. Participant 1 noted the following: "Some pairs were hidden from the filtering while they were clear potential conflicts while others seemed irrelevant which made me stop using it after a few minutes". Participant 2 used the tool a great deal stating: "It's a very nice and useful tool, but some of the highlights did not make sense (tens of miles in trail and still highlighted) but it was generally a net benefit to have". Both participants highlighted that the flight filtering tool did not always do a good job of representing the relevant flights. Additionally, both participants indicated that the VERA tool was not always accurate in providing conflict geometries or minimum distances.

Furthermore, both participants indicated that the capabilities of the automation software should be considered when providing allocations. This means looking at which solutions the automation software is able to apply to a conflict and whether those solutions are adequate to solve the conflict. Participant 1 also noted that aircraft should never be re-allocated to automation without permission after they have already assumed the aircraft.

VIII. DISCUSSION

Out of the two participants, one finished the experiment successfully to completion while the other experienced a crash in the simulation software 4 minutes before the end of their last run. This resulted in no data for the last 4 minutes of that allocator simulation run. The traffic density towards the end of the run was relatively low, and the participant had indicated that most commands had already been issued, so the experiment was not repeated. However, this lack of data for the last four minutes should be noted when looking at the results of Participant 1 as certain measured metrics could be influenced. The results of the experiment showed many weaknesses in the design of the allocator, as well as weaknesses in the features available in the simulator environment. The allocator proved to perform relatively consistently between the offline and online simulation runs, and the allocator designs were clear enough that one of the participants was able to correctly identify the design metrics of both allocators.

A. Experiment sample size

Only two participants carried out the experiment testing out the allocator models in SectorX. The sample size was intentionally small to gain initial feedback regarding the experiment setup and the allocation model. The main purpose of the experiment was to gain insight on the behaviors of the allocator model in a live traffic sample, with both controllers managing their traffic. Given the small sample size, no statistical inference is possible and all findings must be interpreted as exploratory.

Both participants highlighted very different approaches to the experiment, which may have led to them both having contrasting experiences despite using the same traffic samples. When faced with difficulties, particularly regarding behavior of automation software or the allocation algorithm, Participant 1 chose to 'fight' the system to try and organize the work and sector in the way they wanted. When Participant 2 faced similar difficulties, they chose to accept these difficulties and try to learn how to adapt their work style to best cooperate with the system.

With only two participants, it is difficult to say whether these two approaches resulted in the differences in participant experience, however, it is worth considering in the future design of allocation algorithms. Future studies could look at personalizing the allocator to the particular controller. This can take on many forms, such as changing the allocator model according to user preference, changing allocator metrics based on controller/sector characteristics, or adapting the design of the allocation system based on the preferred approach of the controller.

B. Allocation and Automation Cooperation

The flight allocation system was found to behave similarly in the offline simulation and during the experiment with ATCOs. Slight differences in the overall distribution are present between participants and the offline simulation, but the general behavior of the allocation system is stable and consistent throughout despite the vastly different approaches of both participants. The difference in allocator models was noticeable by one of the participants, indicating that the different allocator model designs are significant enough such that they can be observed by controllers.

One flaw of the allocator model design was that the system allocation was always done 60 seconds before sector entry regardless of whether or not the aircraft had already been assumed. Both participants particularly disliked when the system would allocate an aircraft to automation that they had already assumed. This would often disrupt their flight plans. This was not an intended feature of the system and was an oversight in the design. Furthermore, both participants noted that 60 seconds before sector entry was a very late moment for the allocation to occur and that it should ideally occur much earlier. This is something to consider in future designs of an allocator algorithm.

Making the moment of assignment too early can also have potential drawbacks. With this allocator design, the allocation assignment is dynamic and is dependent on the flight plan and movements of aircraft already in the sector. In cases of too early assignment, an allocation made far before the aircraft enters the sector could no longer be valid at the moment the aircraft enters the sector due to commands provided to aircraft already in the sector by either controller. While both controllers agreed that the moment of assignment was too late, the difference in the amount of early assumes could suggest that the ideal time to assign an aircraft is dependent on the preferences of the controller.

Despite the allocator producing consistent assignments based on the model designs, both controllers did not trust the allocation model to make good assignments, nor did they think that the allocation of flights was particularly logical in either experiment. Both the allocation models and automation software were not trusted by either participant, and both participants felt that there was not good cooperation between them and the automation. The distrust of both allocator and automation systems is likely correlated, and design flaws in either negatively affect the performance of the other.

A design flaw of the allocator is that the allocator and the automation are treated as separate entities that do not communicate with each other. The allocation system determines which controller to assign an incoming aircraft to based only on a number of calculations and thresholds set in the allocator. This means that when an aircraft is allocated to automation and it includes a conflict that cannot be resolved in the vertical plane, the human is expected to step in and take over responsibilities of the aircraft. While the algorithm is making a correct assignment based on the metrics, it fails to consider whether or not the assignments it makes can be handled by the digital controller. Furthermore, the digital controller fails to indicate when human assistance is needed to resolve a conflict.

This means that the human controller needs to monitor all aircraft of the digital controller to identify conflicts that are not being resolved. This also means that any allocations to the automation cannot be trusted, and the controller must double check to ensure the assignments are appropriate. This makes it difficult for controllers to critically analyze the allocator design metrics that were the focus of this experiment, due to the large flaw in the allocation system as a whole.

Although participants indicated that they generally did not trust automation, and that it was not very reliable and consistent, both participants made use of the automation software and transferred some of their assigned aircraft to the digital controller. Both participants said that they chose to transfer their aircraft to the automation when the complex parts of the flight plan were solved. This highlights that the distrust in the automation primarily comes from its inability to solve particular conflicts and the lack of communication to the human controller surrounding this.

While this paper uses a flight-based allocation approach, the behavior of both participants suggests that function-based allocation could be more desirable for participants, particularly when the automation is not capable of handling complex situations. In such a scenario, the automation and human work together and give commands to all aircraft in the sector with humans doing the more complex tasks and automation doing the simple tasks. Such behavior was already displayed in this experiment and could be further investigated as an alternative approach to flight-based allocation.

Future studies should look at incorporating the automation capabilities into the allocation algorithm. Before any other factors are considered, the algorithm should first look at any potential conflicts together with the automation software. If the automation software cannot find any suitable solutions, the aircraft should always be assigned to the human controller, regardless of any other metrics. Furthermore, should an automated flight become unmanageable for any reason after allocation, the system should clearly indicate to the human controller that assistance is required. This allows for more trust in the automation system, allowing the allocation algorithm to be more critically evaluated.

A further design flaw of the allocation algorithm lies in the flight filtering algorithm. Both participants indicated that some of the flights included in the flight filtering were irrelevant to the FOI, while other flights that they expected to be included due to clear potential conflicts were missing. As the allocation algorithm uses this as the basis for calculating the number of interactions, which both controllers mentioned to be an important characteristic in their ideal allocator design, it is important that the filtering provides accurate data to avoid incorrect assignments. Especially when combining the allocation algorithm with the automation capabilities, it is essential that every aircraft with potential conflict is considered to ensure that allocations are manageable by the allocated controller.

C. Allocator Metrics

This paper considered three metrics to be used in the design of the allocator that determine whether an aircraft is assigned to the H-ATCO or the D-ATCO. The experiment with only the complexity metric, calculating the number of relevant flights using the flight filtering model of Kumbhar [10] was found to be less variant to H-ATCO choices and actions. A relatively consistent distribution of flights was found across both participants and the offline simulation, and the distribution of the complexity of assigned aircraft was found to be consistent as well. A purely complexity-based allocator model produced the most stable and consistent results.

Possible improvements to the complexity-based allocator suggested by the participants were to give interactions weighting. Rather than using just the number of relevant aircraft, looking deeper into the nature of the interactions, potential conflicts and the required maneuvers to solve the conflicts can provide a better estimate of the complexity of an incoming aircraft. Future studies could further investigate the interactions and construct a weighting for certain interactions to produce a more detailed estimate of the complexity.

One of the allocator models in the experiment primarily allocated flights to reduce the number of mixed interaction between aircraft. In this allocator model, complexity and the minimum distribution of flights to each controller were also considered. Participant 2 noted in informal conversation that they were very confused by some of the allocations made by the system during the experiment with this allocation model. However, Participant 2 also indicated that during the experiment with Allocator B, they were better able to organize their work and manage the traffic efficiently together with automation. The usefulness of reducing mixed interactions as a metric is still unclear and could be further investigated in future studies.

Allocator B was more sensitive to the behavior of the participant, particularly due to the presence of the minimum flight distribution metric. Participant 1 had 25% of the flights assigned to them according to the minimum flight distribution metric, due to them allocating their own aircraft to automation when the complex parts of the flight path had been solved. Participant 2 also displayed a similar strategy, primarily in the last 10 minutes of the experiment run.

While the minimum flight distribution metric was barely present in the offline simulation, it played a significant role with the participant experiments due to their behaviors. This is potentially an unwanted interaction, as aircraft are being allocated less by the main 'desired' metrics. Furthermore, the behavior of the participants as well as their comments indicate that they wish to carry out the complex tasks related to flights and are happy to leave the automation with the simple tasks. The minimum flight distribution metric means that the H-ATCO now receives any incoming flight, whether it be simple or complex. This leads to the H-ATCO possibly having to re-allocate the received simple flights to the D-ATCO, essentially turning the H-ATCO into the allocator.

Possible improvements to this metric could be to change the way in which minimum workload is measured. In this paper, a minimum workload is established purely through the number of flights a controller is managing. Estimating workload through other factors, such as the complexity of flights being managed by each controller, or looking at the required commands of the aircraft being managed by each controller could be a better metric at determining minimum workload. In such a way, even though the H-ATCO has relatively few aircraft, if the workload of these aircraft is significant enough, then incoming aircraft do not need to be assigned to the H-ATCO to simply provide more work but can instead be assigned based on other metrics.

D. Human Task Analysis

Despite the drawbacks in the automation software, both controllers managed to safely handle all air traffic in the sector with the exception of one aircraft pair. This conflict involved a mixed pair aircraft, both approaching the same altitude and lateral position. The reason for this conflict not being resolved by the participant was not made clear. One of the possible reasons for the low number of STCA/LOS alerts is due to the lack of trust in the automation and allocation systems. Both participants indicated a preference for assuming automated aircraft, solving the complex part of their flight path and then handing the aircraft back to the automation. The human controller was closely monitoring flights given to automation due to the lack of trust. This means that the participant kept an overview of the sector as a whole and not only the aircraft assigned to them.

As automation capabilities advance and the assignments made to the automation become more reliable and trustworthy, it is possible that the human starts to ignore automated flights more often, trusting the automation to handle their allocated aircraft. This could lead to the human controller providing commands that cause mixed conflicts, due to not considering the flights being managed by the digital controller. This is not something that could be well investigated in this experiment due to the relatively basic level of automation present.

No general conclusion can be made regarding which allocator model provides a lower workload for controllers. Both controllers indicated experiencing a lower workload during their second experiment scenario. As identical traffic scenarios were used across allocation conditions, learning and expectancy effects cannot be excluded and may have influenced workload ratings. Future studies with a larger sample size of participants could better investigate whether a general experienced workload discrepancy can be found across different allocator models.

IX. CONCLUSION

This paper investigates the possible design metrics of a flight allocator in an en-route human-automation shared traffic environment. Several different metrics were tested after which two models were designed, one focusing on complexity-based allocation, in the form of the number of relevant flights, while the second design focused on reducing the number of mixed interactions.

Experiments with ATCOs revealed that both participants did not trust their automation companion nor did they trust the allocator to make good allocations. For this reason, a large number of allocations were overridden within the first 2 minutes of an aircraft being assigned by the system. The two allocator models were shown to be distinguishable and recognizable, as one of the participants was able to correctly identify the allocation logic of both models.

The data did not reveal any clear insights into the perceived workload of participants across the allocator models. The absence of automation incapabilities being communicated to the human controller as well as these limitations not being considered by the allocation model proved detrimental to the experience of the participants. The results of the experiment did reveal a pattern in both participants to assume flights, solve the complex parts and then transfer to automation to handle the remaining tasks. Participants demonstrated a function-based allocation approach, which may be more appropriate and should be further investigated.

Further studies should develop and refine the complexity-based allocator method to include weighting of tasks and interactions in the complexity calculation. Automation capabilities should also be considered in the decision making process of the allocator model. These approaches will provide more insight in future experiments as to the usefulness of the allocator in distributing traffic to controllers in an en-route human-automation shared traffic environment.

REFERENCES

[1] EUROCONTROL Maastricht Upper Area Control Centre, "Eurocontrol forecast update 2024-2030 – spring update | eurocontrol," Tech. Rep., Feb. 2024. [Online]. Available: https://www.eurocontrol.int/publication/eurocontrol-forecast-2024-2030.

[2] IFATCA. "Staff shortage survey – eur region." (2022), [Online]. Available: https://ifatca.org/staff-shortage-survey-eur-region/.

[3] EUROCONTROL, "European organisation for the safety of air navigation (eurocontrol)," *Reference Works*, Jun. 2023, Institution: Koninklijke Brill NV. DOI: 10.1163/1570-6664_iyb_SIM_org_39214. [Online]. Available: https://referenceworks.brill.com/doi/10.1163/1570-6664_iyb_SIM_org_39214.

[4] SESAR, *SESAR Joint Undertaking | European ATM Master Plan*, 2025th ed. 2024. [Online]. Available: https://www.sesarju.eu/masterplan.

[5] G. de Rooij, A. B. Tisza, and C. Borst, "Flight-Based Control Allocation: Towards Human–Autonomy Teaming in Air Traffic Control," en, *Aerospace*, vol. 11, no. 11, p. 919, Nov. 2024, Number: 11 Publisher: Multidisciplinary Digital Publishing Institute, ISSN: 2226-4310. DOI: 10.3390/aerospace11110919. [Online]. Available: https://www.mdpi.com/2226-4310/11/11/919 (visited on 11/19/2024).

[6] N. H. Mackworth, *Researches on the measurement of human performance (Med. Res. Council, Special Rep. Ser. No. 268)*. Oxford, England: His Majesty's Stationery Office, 1950, p. 156.

[7] M. Endsley and E. Kiris, "The out-of-the-loop performance problem and level of control in automation," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 37, Jun. 1995. DOI: 10.1518/001872095779064555.

[8] A. Stienstra, "Relating air traffic controller perceived complexity to characteristics of individual en-route flights," *TU Delft - Aerospace Engineering*, 2022. [Online]. Available: https://repository.tudelft.nl/record/uuid:ee7d8281-19c9-4611-9e5f-0b4c0f0533f6.

[9] G. de Rooij, A. Stienstra, C. Borst, A. B. Tisza, M. M. van Paassen, and M. Mulder, "Contributing factors to flight-centric complexity in en-route air traffic control," in *Proceedings of the 15th USA/Europe Air Traffic Management Research and Development Seminar (ATM2023)*, Jun. 2023. [Online]. Available: https://www.researchgate.net/publication/371991465_Contributing_Factors_to_Flight-Centric_Complexity_in_En-Route_Air_Traffic_Control.

[10] A. Kumbhar, W. Lyu, and C. Borst, "Determining flight complexity and relevance: Flight-centric filtering for air traffic control: 14th sesar innovation days, sids 2024," *14th SESAR Innovation Days, SIDS 2024*, pp. 1–8, 2024. [Online]. Available: https://www.sesarju.eu/sites/default/files/documents/sid/2024/papers/SIDs_2024_paper_067%20final.pdf.

[11] *ARGOS Factsheet | EUROCONTROL*, Mar. 2023. [Online]. Available: https://www.eurocontrol.int/publication/argos-factsheet.

[12] EUROCONTROL Maastricht Upper Area Control Centre, "Eurocontrol maastricht upper area control centre ops and automation strategy," Tech. Rep., 2019. [Online]. Available: https://skybrary.aero/%20bookshelf/books/5341.pdf.

[13] T. Finck, C. S. Kluenker, and M. Weber, "Conceptual analysis of allocation strategies for air traffic control concepts without conventional sector boundaries," in *2023 Integrated Communication, Navigation and Surveillance Conference (ICNS)*, Apr. 2023, pp. 1–11. DOI: 10.1109/ICNS58246.2023.10124289. [Online]. Available: https://ieeexplore.ieee.org/document/10124289.

[14] A. R. Schmitt, C. Edinger, and B. Korn, "Balancing controller workload within a sectorless atm concept," *CEAS Aeronautical Journal*, vol. 2, no. 1-4, pp. 35–41, 2011. [Online]. Available: http://link.springer.com/10.1007/s13272-011-0031-7.

[15] B. Birkmeier, S. Tittel, and B. Korn, "Controller team possibilities for sectorless air traffic management," in *2016 Integrated Communications Navigation and Surveillance (ICNS)*, IEEE, Apr. 2016, pp. 6C3–1–6C3–10. [Online]. Available: http://ieeexplore.ieee.org/document/7486362/.

[16] T. Finck, M.-C. Névir, and C. Kluenker, "Design validation of a flight centric workload model including atc task change considering influencing factors," in *33rd Congress of the International Council of the Aeronautical Sciences (ICAS)*, Sep. 2022.

[17] U. Metzger and R. Parasuraman, "The Role of the Air Traffic Controller in Future Air Traffic Management: An Empirical Study of Active Control versus Passive Monitoring," en, *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 43, no. 4, pp. 519–528, Dec. 2001, ISSN: 0018-7208, 1547-8181. DOI: 10.1518/001872001775870421. [Online]. Available: https://journals.sagepub.com/doi/10.1518/001872001775870421 (visited on 09/26/2025).

[18] G. de Rooij, C. Borst, M. van Paassen, and M. Mulder, "Flight Allocation in Shared Human-Automation En-Route Air Traffic Control," *21st International Symposium on Aviation Psychology*, pp. 172–177, 2021. [Online]. Available: https://doi.org/10.5399/osu/1148 (visited on 12/16/2025).

[19] S. Loft, P. Sanderson, A. Neal, and M. Mooij, "Modeling and Predicting Mental Workload in En Route Air Traffic Control: Critical Review and Broader Implications," *Human factors*, vol. 49, pp. 376–99, Jul. 2007. DOI: 10.1518/001872007X197017.

[20] R. Mogford, J. Guttman, S. Morrow, and P. Kopardekar, "The Complexity Construct in Air Traffic Control: A Review and Synthesis of the Literature," *Storming Media*, p. 32, Jul. 1995.

[21] E. L. Kaplan and P. Meier, "Nonparametric Estimation from Incomplete Observations," *Journal of the American Statistical Association*, vol. 53, no. 282, pp. 457–481, 1958, Publisher: [American Statistical Association, Taylor & Francis, Ltd.], ISSN: 0162-1459. DOI: 10.2307/2281868. [Online]. Available: https://www.jstor.org/stable/2281868 (visited on 12/18/2025).

Part II

Preliminary Report

1 Introduction

Air traffic globally has been consistently increasing over the last few years since the end of the COVID pandemic [1]. It is expected that by 2026, air traffic demand in Europe will return to the level measured in 2019 and demand is only expected to increase further [1]. This increase in demand puts pressure on the Air Traffic Control (ATC) systems in place, which are currently reaching capacity limits.

In its current form, European airspace is divided into sectors based on state boundaries, international agreements and geographical characteristics. These sectors are managed by Air Navigation Service Providers (ANSP) that monitor and provide air traffic services to flights in these sectors. Each individual sector is managed by two air traffic control officers (ATCO), an executive controller and a planner controller. Both have different responsibilities but work together to manage the air traffic in a sector.

In 2022, 10.4% of all flights were delayed by en-route air traffic flow management (ATFM) regulations, where 36% of this was attributed to ATC capacity [2]. With capacity already being pushed to the limits, a further increase in workload for ATCOs is not possible without consequences to the efficiency and safety of air traffic operations. Along with the increasing demand for air traffic, ANSPs are struggling to employ enough controllers to provide ATC services in their airspace [3]. A survey led by the International Federation of Air Traffic Controllers' Associations (IFACTA) found that an overwhelming majority of ANSPs in Europe had concerns regarding their experienced controller shortages and the effects of this on ATC operations. For these reasons, a restructuring of the operations of Air Traffic Management is required to meet demands.

This is the basis for the formation of the Single European Sky ATM Research Project (SESAR), which aims to develop and deploy solutions to create a fully scalable ATM system that can handle the growing air traffic while remaining safe and secure [4]. One of the specific goals of SESAR is to improve Air Navigation Service productivity, through increased levels of automation support. This means that controllers will handle less repetitive manual tasks, which will then be done by automation, and will be able to focus on more complex tasks [4].

Maastricht Upper Area Control Centre (MUAC) is an Area Control Centre (ACC) responsible for managing traffic in upper airspace (24,500ft-66,000ft) over the Netherlands, Belgium, Luxembourg and north-west Germany. In alignment with the SESAR goals, MUAC is developing a system called the ATC Real Ground-breaking Operational System (ARGOS). ARGOS currently aids ATCOs in detecting and solving potential conflicts, and the goal of MUAC is to eventually have ARGOS autonomously handle traffic without need for ATCO supervision.

The next step in the development of ARGOS, is to have ATCOs and ARGOS working together to control the flights in a given sector. ARGOS would control all the ”basic” routine flights while the ATCO would handle the ”non-basic” more complex flights. Such a division of tasks means that the ATCO is much more likely to stay engaged due to the increased presence of cognitively challenging tasks. Another benefit of this division is that the ARGOS system in its current state is able to handle basic traffic situations, so the division of work would be within the capabilities of the current system [5].

In a study carried out by Stienstra [6], factors that influence the perceived complexity of a flight were identified and used to design a complexity metric. Stienstra's study looked at a solely human controlled environment, and did not take partial automation into account when determining complexity. Future research was recommended to incorporate the effect of partial automation on the complexity of a flight [6]. A different study by de Rooij et al. investigated human-automation teamwork in a sector to manage air traffic [7]. Specifically, the study looked at allocation strategies to allocate flights to either the human or a digital controller. Simple allocation strategies were chosen for this study, and it was concluded that human controllers generally ignored the allocation strategies proposed. It was suggested that follow up studies look at refining the allocation strategy based on complexity of individual flights. By creating a suitable allocator that the human controller generally accepts, the number of interactions between the human and the allocation system can be reduced.

Extending the previous work of these studies, this research project will use the identified complexity metrics to create a flight allocator algorithm that assigns flights to either a human or an automated controller based on the complexity of the flight.

1.1. Research Objective

The research objective of this project is:

To design and implement a flight-based allocator model that allocates flights to either a human controller or a digital computer controller based on the calculated complexity of the flight and evaluate the parameters of the allocator with a simulator experiment.

This research is limited to en-route traffic in upper airspace meaning that any traffic moving to and from an airport will not be included. To conduct the experiment, a Java-based simulator created by the TU Delft will be used. For this research project, in the simulation it is assumed that the positions of every flight can be accurately predicted throughout the full duration of the scenario. Currently this information is not available to ANSPs. This project assumes that the information contained in the extended projected profile (EPP) will be available to obtain very accurate FMS route predictions by the time such a model can actually be tested in a live scenario. This will then provide sufficient information to be able to accurately make these predictions.

1.2. Research Question

The main research question is as follows:

How do the design parameters of a flight allocator model affect the traffic responsibilities of human and digital computer controllers, and how does this affect the efficiency and safety of ATC operations?

To answer this research question, several sub-questions have been created. The first set of sub-questions will be explored through the literature study and will provide a basis for the research. The second set of sub-questions are related to design and testing of a model in which the results of an experiment are to be analyzed.

1. **How can the complexity of a flight be determined?**
    (a) What factors affect the perceived complexity of a flight for ATCOs?
    (b) What is the level of complexity at which ATCOs would let flights be controlled by a computer?
    (c) How does the dynamic environment of a sector affect the complexity of a flight as it travels through the sector?

2. How effective is the traffic allocation model in assigning flights to the correct controller?

(a) What factors contribute to the ATCO claiming a computer-controlled flight?
(b) How does the traffic allocation model affect the ATCO's workload?
(c) What impact does the collaboration of computer controlled and human controlled flights in a sector have on the safety of that sector?

1.3. Report Outline

The first part of the thesis project is the literature study. The literature study consists of two main research topics. The first topic, discussed in Chapter 2, looks at the ATC work domain. In this chapter, the tasks and responsibilities of an ATCO are discussed, and the structure of airspace is explained. Furthermore, common conflict detection and resolution strategies are explained, as well as newer ATC concepts that change the way ATCOs operate.

Chapter 3 deals with defining complexity in air traffic control, with a focus on flight-based complexity. This chapter highlights metrics that have been used to define flight-based complexity. Then, a flight-filtering concept which determines the relevant flights that contribute to the complexity of an individual flight is explained. The concept of basic and non-basic flights in terms of complexity is also discussed. In Chapter 4, automation concepts in air traffic control are discussed. In this chapter, automation practices being employed at EUROCONTROL are explained, as well as their future plans for automation in ATC. Furthermore, challenges that arise due to automation systems in human-in-the-loop systems are mentioned.

In Chapter 5, metrics prevalent in the design of the flight allocator are discussed. The choice of allocator metrics are motivated and explained, and the decision making process of the flight allocator is demonstrated. In Chapter 6 the experiment design setup is discussed, also explaining the variables that are being considered. The findings of the literature study are concluded in Chapter 7.

2

ATC Work Domain

Air traffic controllers are primarily responsible for coordinating the movement of aircraft in such a way that safe distances are maintained between them, while minimizing delays. ATCOs also provide services to flight under their supervision, such as communication updates regarding weather, providing clearances and giving navigational instructions. ATCOs must also foresee potential situations that lead to violations of separation standards. In these cases, ATCOs also provide resolutions in a timely manner to avoid separation violations and ensure safety of all flights in the sector.

Developments in technologies have led to new innovative tools that are used to assist ATCOs in their tasks, thereby decreasing their workload. In order to design useful support tools for controllers, a deep understanding of the knowledge, logic and skills ATCOs use to accomplish their tasks is required. These tools can then be designed in such a way that it complements the work process of controllers, thereby increasing the chance that the tool will be accepted by the controller. For an allocator design, the metrics involved need to be designed in such a way that the allocation of aircraft is accepted by the human controller in order for it to be useful. Understanding the work flow and tasks performed by controllers that contribute to their workload can help determine which metrics are relevant for the allocator.

2.1. Airspace Structure

Airspace is defined as the part of the atmosphere that is above a particular geographical area and is subject to the laws of a particular country or controlling authority [8]. The largest regular division of airspace is called a flight information region (FIR), which is usually under control of a single authority. Within an FIR, flight information and alerting services are provided. These regions are then subdivided by the controlling authority into smaller geographical areas called sectors. Occasionally, FIRs may also be split vertically into a lower and an upper section. The lower section is referred to as an FIR while the upper portion is referred to as the upper flight information region (UIR). The Maastricht Upper Area Control Centre (MUAC) is an example of a UIR, although this FIR spans multiple country borders, which is not always the case.

Airspace classification can be broken down into four different types: controlled, uncontrolled, special use and other airspace. Controlled airspace is a generic term encompassing several different classifications of airspace and is characterized by the fact that ATC services are provided in accordance with the airspace classification. In uncontrolled airspace, ATC has no responsibility or authority to control air traffic. Special use airspace is the designation for airspace in which certain activites must be confined. Such areas are used for things like military operations, prohibited areas or controlled firing areas. Other airspace areas is a general term encompassing the majority of the remaining airspace, and includes military training routes and parachute jump aircraft operations among others.
Controlled airspace is handled by air traffic controllers and can be further subdivided into different airspace regions. These subdivisions are the Control Zone (CTR), Terminal Control Area (TMA), Control Area (CTA) and the Upper Control Area (UTA). The first three subdivisions all deal with lower airspace while the UTA deals with upper airspace. The CTR is the area directly around an airport, and is controlled by tower control. The TMA deals with incoming and outgoing flights between the CTR and the CTA and is handled by Approach/Departure Control. The CTA covers general ATC for flights up to a flight level FL 245 in The Netherlands. Flights in the CTA are managed by the Area Control Centre.

In The Netherlands, Luchtverkeersleiding Nederland (LVNL) is the Air Navigation Service Provider (ANSP) responsible for managing air traffic services for both civil and military aircraft in these three lower airspace sectors. Flights flying between flight levels of FL 245 and FL 660 are in UTA airspace, often referred to as the en-route airspace. MUAC is the ANSP responsible for providing civil-military air traffic services to flights in the UTA for North-Western Europe, covering all en-route traffic flying over The Netherlands.

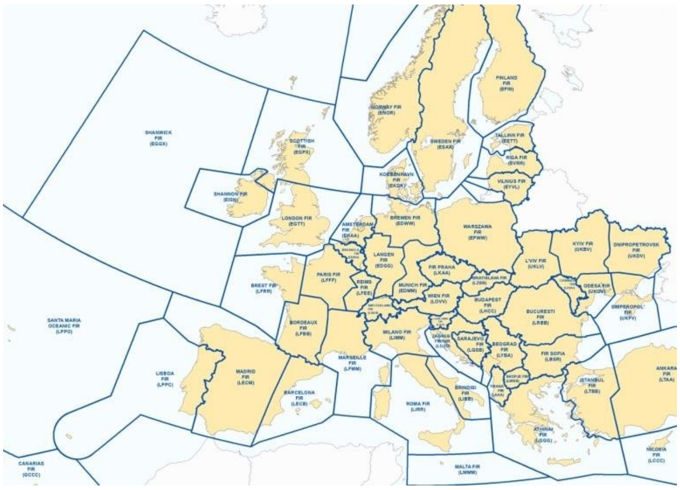

Figure 2.1: Flight Information Regions in Europe [9].

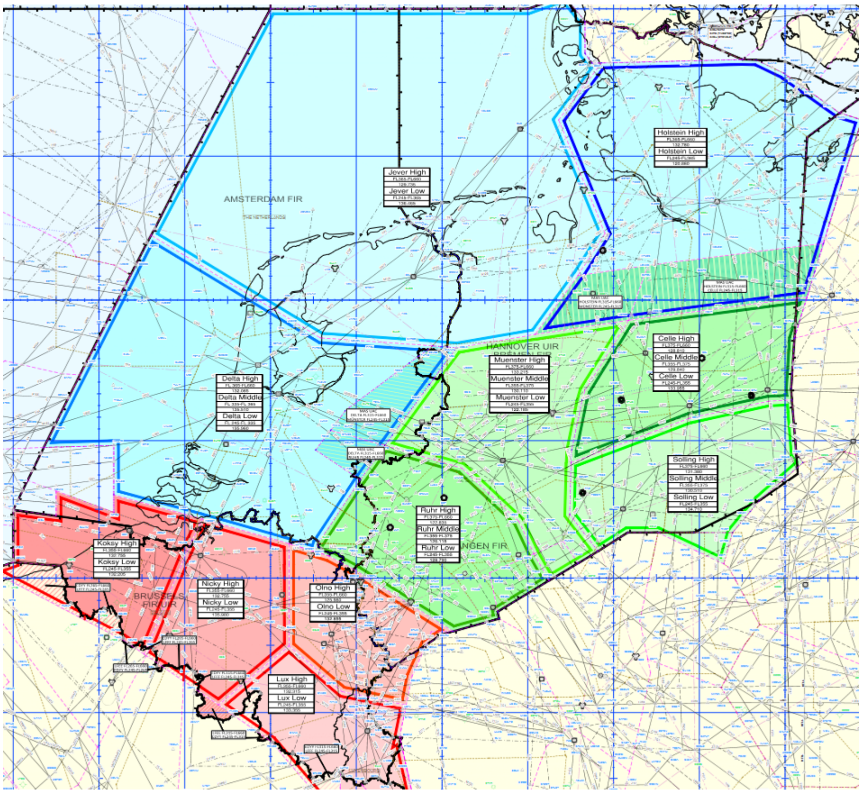

Figure 2.2: Sector division for the upper airspace controlled by MUAC [10].

2.2. ATC Tasks and Responsibilities

2.2.1. Ensuring Safety Standards

The primary responsibility of an air traffic controller is to ensure the safe and expedient flow of air traffic within their sector. In order to ensure safe operations, certain separation minima are defined, which indicate the minimum distance between air traffic that needs to be upheld at all times. There are three types of separation standards that need to be met: Vertical Separation, Lateral Separation and Longitudinal Separation. While the exact minima vary slightly depending on the type of flight and the equipment available, in general, the separation minima are as follows [11]:

- **Vertical Separation:** Minimum vertical separation between aircraft must be 1000ft below FL290, and 2000ft at or above FL290, unless reduced vertical minima separation apply, in which case the minimum vertical separation remains 1000ft above FL290. Vertical minima depend on the accuracy of the altimetry system.

- **Lateral Separation:** Aircraft are required to fly on specified tracks which diverge by a specified amount depending on the navigation aid being used. Minimum lateral separation for aircraft should be at least 5 nautical miles.

- **Longitudinal Separation:** For aircraft flying along same/intersecting tracks, a period of time exists where lateral separation does not exist. Instead, these flights are separated by a time interval of 15 minutes. The time at which flights pass through waypoints are noted, and this is used to ensure longitudinal separation is maintained.

Based on these separation standards, a cylindrical area can be constructed around an aircraft [12], as seen in Figure 2.3. If another aircraft enters the cylindrical area, then the controller responsible for the flight has failed to maintain the separation standards. In such a situation, a loss of separation (LOS) is said to have occurred. LOS situations can lead to collisions between aircraft or dangerous last minute evasion maneuvers. A prediction of a LOS situation is called a conflict, and ATCOs aim to foresee potential conflicts and prevent them before the LOS occurs in order to ensure safe operations within their sector.

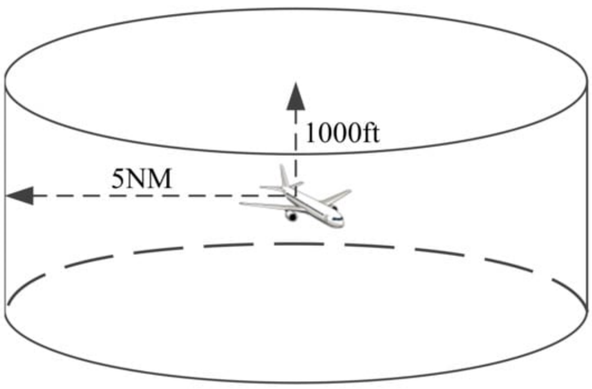

Figure 2.3: Safety Cylinder around aircraft [12].

2.2.2. Work Environment

In most cases, a pair of air traffic controllers are responsible for a sector of airspace. The two controllers take on different roles, one being the executive controller and the other being the planner controller. The executive controller ensures safe operations within a sector through communication with the pilot over radio. The planner controller is responsible for coordinating the handovers and acceptances of flights to or from other sectors. This is done because each sector has its own individual set of standard flight levels, and controllers generally only accept incoming flights that are within a certain flight-level threshold. This ensures a smooth vertical alignment between adjacent flights in a sector.

The teamwork between an executive and planner controller to manage the sector together is crucial in ensuring safe and expedient flow of air traffic. In a study carried out by Svensson [13], ATCOs were asked to rank the most important factors required for successful teamwork. For En-route ATCOs, the most important factors were found to be Adaptability, Mutual Trust and Backup behaviour. Backup behaviour refers to the ability to be able to call over one more person when the workload increases to help out. Furthermore, the planner and executive can execute certain tasks for one another to relieve workload [13]. When designing support tools, the tool must adhere incorporate these factors in order for it to be useful and relevant for ATCOs.

ATCOs infer flight information and sector attributes from a radar display (also called a plan view display (PVD)). An example of a section of a radar display can be seen in Figure 2.5. Flights under control of the ATCO monitoring the sector are highlighted in green. Flights not under the control of the ATCO tend to be flights passing through different sectors and are highlighted in light grey. From the flight labels, the following flight information can be obtained by the ATCO, also visible in Figure 2.4:

1. **Aircraft Callsign:** A group of alphanumeric characters used to identify the flight.
2. **Current Flight Level:** The current altitude of the flight.
3. **Arrow:** Indicates whether the aircraft is currently ascending or descending. No arrow means the aircraft is cruising.
4. **Cleared Flight Level:** The altitude for which the flight is currently cleared.
5. **Heading:** The direction in which the aircraft is currently heading.
6. **Exit Flight Level:** The (required) target altitude the flight must obtain while exiting the sector.
7. **Exit Point:** The specific waypoint at which the flight must exit the current sector and be handed over to the adjacent sector.

Figure 2.4: Example of a flight label.

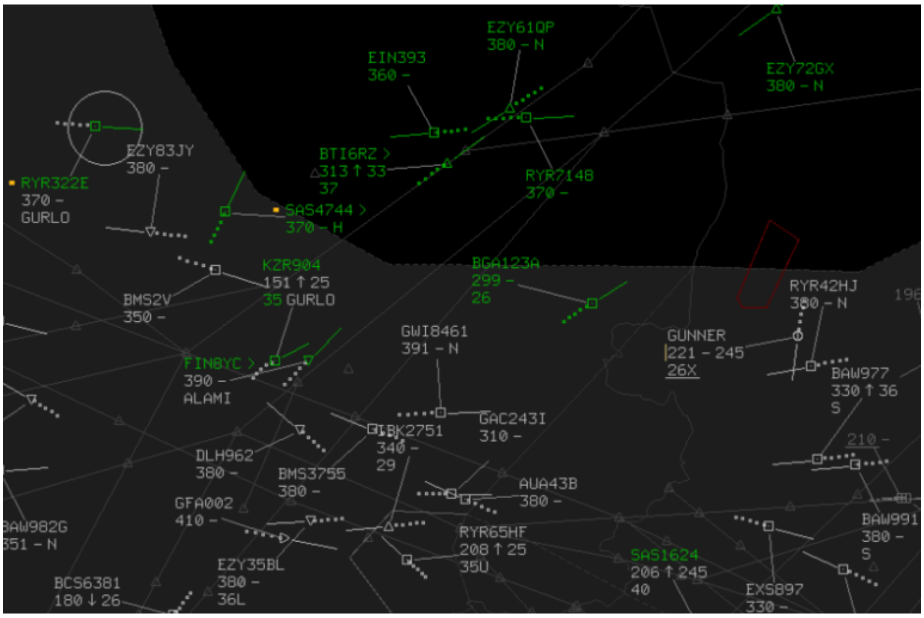

Figure 2.5: Screenshot of a section of the radar display operated at MUAC.

With the information provided on the radar display, the ATCO has a list of tasks they carry out to ensure safe and expedient flow. A summary of control tasks carried out is listed below [14]:

- **Assume Control:** Assume control over new aircraft that enter the sector.
- **Clear to target flight levels:** Give aircraft clearances to their target altitudes.
- **Routing to exit waypoints:** Guide aircraft toward destination while minimizing travel delays.
- **Conflict Detection:** Monitoring aircraft in the sector to predict possible separation violations.
- **Conflict Resolution:** For any detected conflicts, issue reroutes through altitude, route or heading clearances. ATCOs can decide to directly solve a conflict, or wait and monitor the situation, depending on the urgency of the conflict.
- **Conformance monitoring:** Making sure an aircraft is following its cleared course. Monitoring for any deviations of the aircraft from the instructions/clearances given.
- **Hand over control:** Transfer control of aircraft to adjacent sector.

Often, the tasks carried out by an ATCO are interlinked with one another. At the core of an ATCO's work is monitoring. Conformance monitoring and conflict detection are both monitoring tasks by which the ATCO builds a mental image of the sector currently, as well as what the sector will look like in several minutes. Based on new information, or clearances/instructions issued, the mental model of the sector is updated by the ATCO, and each command given is evaluated to ensure it does not lead to conflicts in the future. Besides monitoring tasks, the ATCO also has to issue commands to aircraft to safely guide them through the sector. The ATCO looks at the sector as a whole, and tries to develop an optimal strategy to clear flights to their target flight levels and to route aircraft to their exit waypoints. For this, the ATCO must take into account the current flight plan of the aircraft it is issuing commands for, as well as the flight plans of other (nearby) aircraft in the sector, to avoid issuing commands that lead to potential conflicts in the future. By accounting for potential conflicts before issuing commands, the ATCO can quickly and efficiently communicate with pilots to reduce time needed for communication.

2.2.3. Conflict Detection

When an ATCO identifies that a possible conflict will occur, this is called conflict detection. ATCOs are responsible for predicting and solving conflicts, through extrapolation of the flight path based on the current state of the flight as well as the intent of the flight. Typically, there are three different levels of conflict detection that are carried out: long range, mid-range and short-range [15].

Long-range conflict detection looks at the entire route of an aircraft (departure to arrival) to see if the planned trajectory of the aircraft interferes with any other planned traffic. Long-range conflict detection is usually performed on a daily basis, looking at flight plans and airline schedules. Medium-range conflict detection looks for potential conflicts up to 20 minutes in the future. Medium-range conflict detection occurs multiple times per hour, and is done en-route, looking at the future waypoints of the aircraft to determine possible conflicts.

Medium-Term Conflict Detection (MTCD) systems are tools that are designed to aid the ATCO in detecting potential conflicts in the sector. The goal of MTCD systems is to move from a more reactive form of ATC to proactive control. A MTCD system must satisfy the three following functions [16]:

- The detection and notification to the controller of probable loss of separation between two aircraft.
- The detection and notification to the controller of aircraft penetrating segregated or restricted airspace.
- The detection and display to the controller of aircraft-to-aircraft encounters where, although the required separation is achieved, each aircraft is blocking airspace that might have been used by the other.

Short-range conflict detection occurs every few seconds, and is carried out by both the ATCO and the onboard equipment on the aircraft by the aircraft collision avoidance system (ACAS). Short Term Conflict Alert (STCA) is a ground based safety net with the intent to aid the controller in preventing collisions between aircraft, by generating an alert within a short look-ahead time (2 minutes) to the controller. STCA is a safety net, with the sole purpose to enhance safety.

A study carried out by Eyferth et al. [17] looked at the relationship between look-ahead time and conflict detection. In this experiment, it was found that shorter look-ahead times led to more accurate conflict detection by ATCOs. Furthermore, it was found that conflicts involving climbing/descending aircraft were more easily detectable by ATCOs than horizontal conflicts. Another study by Rantanen [18] looked at how ATCOs determined whether there was a conflict detected. Rantanen theorized that controllers judgment of conflict-detection adhered to a hierarchical strategy. This theory was supported by the results, with Rantanen concluding that ATCOs first looked at altitude information when determining potential conflicts. Flights with no converging vertical profile were very quickly discarded as no-conflicts. For aircraft with the same-altitude, the strategy to determine conflicts becomes significantly more complicated, relying on a combination of trajectory factors including convergence angle, speeds, distance and time to the point of intersection [18]. This study suggests that controllers employ hierarchical structure for gathering information on potential conflicts for flight pairs. The hierarchical structure allows for a reduction in cognitive demand, as controllers could often immediately determine whether an aircraft pair needed further inspection by only considering a specific set of flight characteristics. The hierarchical structure is explained below:

1. **Altitude Overlap:** The first check for controllers is to determine whether there is altitude overlap. If a flight pair has sufficient vertical separation, then it is sufficient to determine that no conflict will occur between the flight pair. In such a case, the horizontal profiles of the flight pair do not need to be considered. In the experiment, no-conflict decisions were made twice as fast when aircraft were at different altitudes, due to the cognitive ease at which this can be evaluated. For flight pairs with changing vertical profiles, extrapolation of the trajectories is required to determine if a potential conflict is possible.

2. **Heading:** If a flight pair are at the same altitude, then the next factor to consider is the heading of the flight. Three types of headings interactions are considered; diverging, converging and parallel tracks. Diverging tracks are tracks in which the flight pair are moving further away from each other, which is an inherently conflict-free interaction. Conversely, with converging tracks, the flight pair are getting closer and closer to each other. Extrapolation is then needed to determine the point of intersection, as well as the positions of the flight pair along the track.

Higher cognitive demand is required to extrapolate the paths of the flight pair, and to determine the minimum distance between the flight pair to see if separation minima are violated. For parallel tracks, if the tracks are separated by a minimum distance greater than the separation minimum, then no conflict is to occur. If the tracks are closer than the minimum distance, then speed information is needed to extrapolate the possibility of a conflict.

3. **Speed:** Lastly, speeds of the flight pair are considered to determine whether a conflict will occur. With flights at the same altitude, and also on horizontally converging tracks, speed is needed to extrapolate the positions of both aircraft over time. By predicting when the first flight will reach the interaction point, and the position of the second flight at this time, the distance between the two flights is predicted to see whether a conflict will occur. Looking at speed to determine the presence of the conflict falls at the bottom of the hierarchy due to its very demanding cognitive effort. Speed-distance calculations of the flight profiles of the aircraft pair must be carried out to determine their closest point of approach and whether or not this violates separation principles.

Along with the hierarchical structure, ATCOs also make use of support tools to help monitor the sector for potential conflicts. Tools like STCA and MTCD systems are used to warn controllers of violations of separation minima in short-term and medium-term conflicts respectively. Other tools such as the VErification and Resolution Advisory tool (VERA) developed by EUROCONTROL can aid the ATCO in detecting potential conflicts. VERA extrapolates flight positions of flight pairs based on radar data, providing ATCOs with a predicted position based on the current trajectories. This can help the ATCO in creating a mental image of the extrapolated profiles of flights, to determine whether or not an aircraft pair will be in conflict in the future.

2.2.4. Conflict Resolution

Having detected a potential conflict between an aircraft pair in the sector, the ATCO is then responsible for solving the conflict by providing a resolution maneuver. The resolution maneuver involves giving a command to change either the altitude, heading or speed of one of the two flights. Conflict resolution can take many forms, and resolutions can be given to either aircraft of the conflicting pair (or even both) depending on what the controller deems best for the traffic flow.

Similarly to the hierarchical structure employed to identify potential conflicts, the same factors are looked at when controllers create resolution maneuvers for aircraft to avoid the conflicts. An ATCO will change the trajectory of one of the aircraft in the flight pair by changing either the altitude, the heading, or the speed of the aircraft [19]. An ATCO decides the resolution maneuver and the aircraft to give the command to based on temporal, situational and technical parameters. [19]. In general, vertical separation maneuvers are the easiest resolution maneuvers to implement, and require the least amount of attention and vigilance from controllers, making it the preferred option[20] [21]. Furthermore, providing an altitude or speed command has no influence on the projection of the flight course depicted on the radar screen, whereas a heading change makes the earlier projection of the flight plan obsolete. A heading command adds additional mental effort for the controller, as the previously available data like the planned time to fly over a specific waypoint is no longer valid [19].

Alternatively, instead of changing a single flight parameter, sometimes a ’direct’ command is given, a command which allows the aircraft to head directly to a distant waypoint. While the direct also causes a heading change, the advantage of a direct command is that it does not require further instructions to get the aircraft back on to the planned route. In the case of a heading command, the aircraft will deviate from its planned flight path, but must eventually require another command to bring it back to the flight path and pass the required waypoints. An overview of each type of resolution maneuver is given below:

1. **Vertical maneuvers:** Vertical maneuvers are carried out by ATCOs by clearing a flight for a new altitude. Vertical maneuvers require little attention from the controller, as a flight pair that maintains their flight levels will not be in conflict [18]. For flights climbing or descending, the flight is likely to cross the flight levels of other aircraft in the sector. This may lead to conflicts with other aircraft during the transition of flight levels, so controllers must consider flights along the altitude band to ensure separation standards are being met. Altitude maneuvers are given by either clearing a flight to an available flight level, or by initiating a step climb/descent whereby the flight amends its flight level in 'steps' of 1000ft to resolve conflicts. [21].

2. **Horizontal maneuvers:** A horizontal maneuver can also be initiated by an ATCO by clearing a flight for a new heading. A 'direct' command can be given to a distant waypoint/exit point, otherwise controllers must give multiple heading commands to the aircraft. The first heading command provides a new heading that will resolve the conflict, while a second heading command issued later is necessary to bring the flight back to its flight path. Horizontal maneuvers require an update of the mental picture of the sector, as projections of the flight path are no longer relevant when a flight deviates from its initial trajectory. Similarly to with vertical maneuvers, care must be taken to consider surrounding traffic when giving heading commands to ensure no new conflicts are created. Horizontal maneuvers include parallel tracks, whereby a new trajectory is assigned which is parallel to the current route of the aircraft. Another horizontal maneuver is to assign a new trajectory 5 miles to the left or right of the original trajectory, then assign it back to the original trajectory at a later time to ensure separation [21]. The availability of lateral maneuvers is also largely dependent on the shape of the sector. For particularly narrow sectors, lateral maneuvers are difficult to execute without causing additional conflicts.

3. **Speed maneuvers:** Speed maneuvers can be initiated by ATCOs to solve conflicts. Speed commands are particularly favourable for cases of same-heading conflicts or overtaking-level conflicts. [22]. Controllers implement speed maneuvers by adjusting the trajectory of the slower aircraft to be behind the faster aircraft to solve the conflict [21]. Speed maneuvers are less preferred than vertical and horizontal maneuvers, particularly at higher altitudes, due to their slower profiles and narrowness[22][23].

Another part of the conflict resolution strategy employed by ATCOs is to determine which aircraft of the aircraft pair should have its course altered to avoid the conflict. In a study carried out by Guleria et al. [24], experiments were performed with ATCOs to determine the strategies employed in conflict resolution. From this study, it was concluded that most ATCOs prefer to maneuver trailing aircraft, reasoning that extending the path of the farther aircraft is effective because it always ensures separation [24]. In cases where it was difficult to determine which aircraft was closer to the conflict point, both aircraft were equally likely to be chosen for the resolution maneuver.

Using one of the three types of maneuvers mentioned above, or a combination of the three, controllers work to resolve conflicts. When giving resolution maneuvers, ATCOs also consider the impact of the maneuver not only on the flight pair, but also on the directly surrounding traffic in the sector, to ensure that an immediate follow-up conflict does not appear. It is not always possible to foresee how conflict resolution maneuvers will impact the rest of the sector over longer time periods. Certain tools can aid ATCOs in determining optimal conflict resolution maneuvers. One such tool is the Lateral Obstacle and Resolution Display (LORD) developed by EUROCONTROL. LORD displays all possibilities and limitations of the current environment to allow the ATCO to select the most preferrable clearances. LORD scans heading and flight level combinations to identify possible conflicts in the lateral and vertical plane, taking into account the predicted trajectories of surrounding traffic (Figure 2.7). VERA is another tool previously mentioned that can also be used to detect conflicts and provide resolution maneuvers to solve the conflicts (Figure 2.6). Such tools can ease the workload of ATCOs in their evaluation of the impact of their resolution maneuvers on the sector.

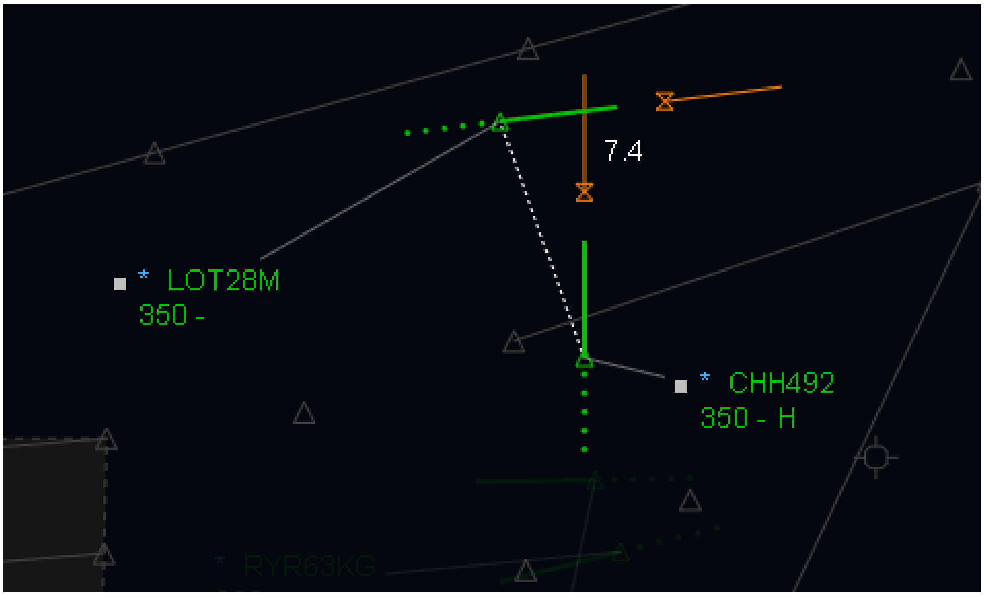

Figure 2.6: Example of the VERA tool in the simulator environment.

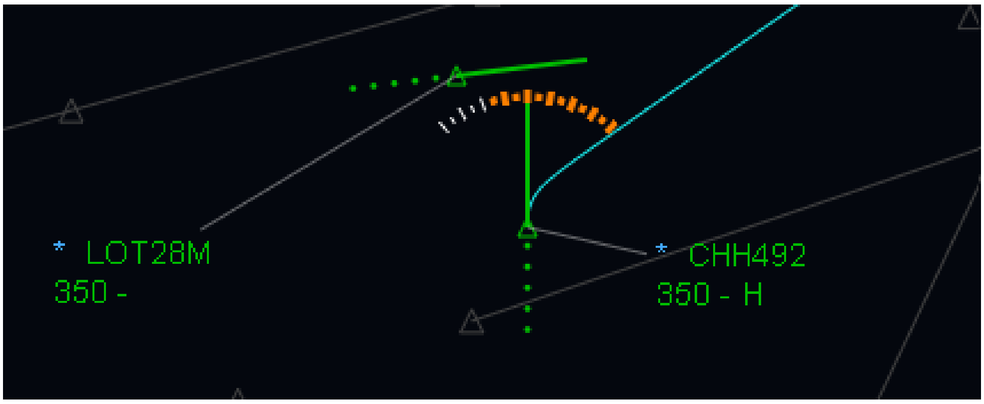

Figure 2.7: Example of the LORD tool in the simulator environment.

2.3. Alternative Concepts: Flight-Centric ATC

Flight-Centric Air Traffic Control, sometimes referred to as Sectorless ATC, is an alternative approach to ATC. Traditionally, an executive and planner controller pair work together to manage a sector. If the traffic in a sector increases, sectors are often subdivided into smaller sectors in order to keep the workload at a manageable level for the ATCOs [25]. However, splitting up sectors into smaller sub-sectors has a limit at which further splitting is no longer beneficial. A large amount of sectors means that the number of coordination activities between sectors will also rise. Furthermore, the smaller the size of the sector, the more difficult it becomes to implement tactical and strategic control [25] as well as conflict resolution of aircraft, due to the limited space available. As an alternative solution to the growing air traffic problem, the concept of Flight-Centric ATC has been introduced.

The concept of Flight-Centric ATC is to have an aircraft controlled by only a single ATCO during its entire flight. The concept is described as ”instead of having two controllers controlling one sector containing *n* aircraft, one controller will be responsible for a limited number of *m* aircraft, from departure to arrival terminal area”[25, p.1]. In Flight-Centric ATC, the controllers only communicate with the aircraft assigned to them, and surrounding traffic is controlled by other controllers. Figures 2.8 and 2.9 show examples of conventional and Flight-Centric ATC divisions.

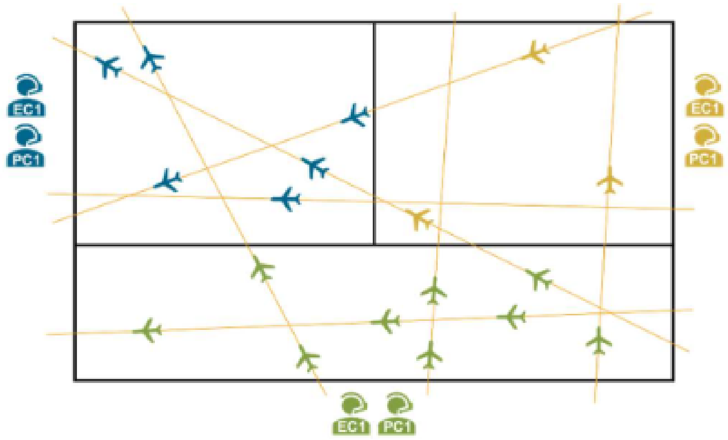

Figure 2.8: Example of conventional ATC [26].

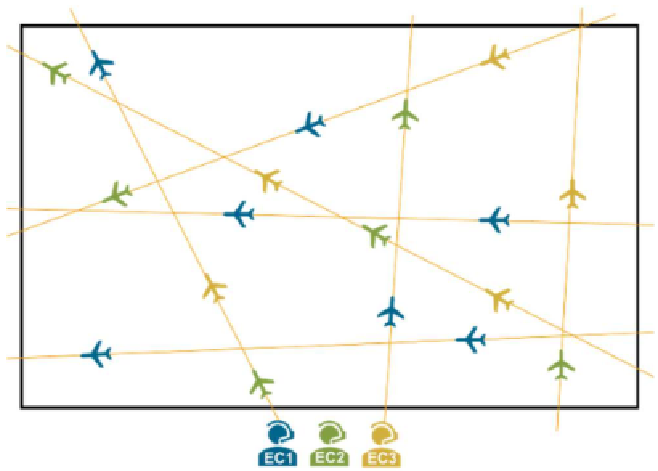

Figure 2.9: Example of Flight-Centric ATC [26].

In the case of a conflict between two aircraft controlled by different controllers, clear rules should determine which ATCO is responsible for solving the conflict. A draft of a set of rules which could be used to determine who is responsible for solving conflicts can be found in [27]. Along with changes in conflict detection and resolution methods, the way in which information is displayed will also undergo changes with Flight-Centric ATC. As controllers are now responsible for a number of aircraft which can be separated by much greater distances than with traditional ATC, a PVD may no longer be suitable to provide the controller with a good overview of the traffic situation. Birkmeier et al. [28] offers insight into possible alternative display methods that could be used with sectorless ATC to accommodate controllers. One such an example of a new display format is called a tiled display, whereby a controller has one traffic display for each aircraft under their control, and the controller is only responsible for one aircraft in each of these displays. [28].

2.3.1. Allocation Strategies

While the implementation of Flight-Centric ATC would bring about several changes to the work domain of an ATCO, another important task that would need to change is the strategy employed for the allocation of flights. With sector boundaries no longer existing, flights must be assigned to an ATCO based on other criteria. Schmitt et al. introduced four initial simple concepts that could be used for aircraft allocation [25]:

- **Uniform allocation of aircraft:** The only relevant factor in this allocation strategy is the number of aircraft assigned to each controller. Each controller receives the same number of aircraft to monitor, irrespective of the complexity of the individual traffic.

- **Uniform allocation of ATCO workload:** The workload of an ATCO is considered when assigning aircraft, with the goal to uniformly distribute air traffic based on workload. For this to be implemented, a workload model would need to be implemented that can measure the amount of workload experienced by a controller. This becomes difficult as not all controller activites can be recognized by sensors.

- **Allocation based on similar trajectories:** Assigns aircraft based on some simililarity criterion. Trajectories can be similar in space, or in space and time. This can allow for the ATCO to become familiar with the surrounding space of the trajectories, similar to becoming familiar with a given sector. However, dealing with similar trajectories may also be monotonous which can affect the ATCOs performance [25].

- **Allocation by conflicting aircraft:** Assigns aircraft based on predicted conflicts. aircraft are assigned such that conflicting aircraft pairs are being controlled by one controller. This would then solve the problem arising when a conflict occurs between aircraft controlled by different ATCOs, as no discussion is needed to decide how the conflict should be resolved. This strategy does not provide insight on how to assign conflict-free aircraft, and is best used as a supplement for a basic strategy.

Further allocation strategies are described in [26]. In this paper, different types of allocators are also discussed. Three types of allocators are mentioned, Automatic, manual and semi-manual allocation. In automatic allocation, aircraft are assigned to controllers automatically based on the allocation strategy and rules defined. In automatic allocation, no human verification is done, meaning no extra human personnel is required. Automatic allocation fails to account for the individual work experience of a controller on duty [26]. In manual allocation, a human is responsible for assigning aircraft based on personal work experience. In this way, individual situations can be better catered for in the allocation process. Lastly, in the semi-manual allocation, a combination of automatic and manual are used. This can be approached in several different ways. The human allocator can use the allocation system as a support tool to obtain suggestions for allocations. Alternatively, the automated allocation system can have control, with the human being able to override when they see fit. The third approach is to allow the human controller to dynamically adjust the parameters of the automatic allocation system, thus altering the allocation strategy slightly depending on the situation [26].

2.4. Conclusion

In current ATC operations, controllers work in pairs with an executive controller and a planner controller. Together, they are responsible for managing the safe and expedient flow of air traffic in a sector of airspace. The primary responsibility of an air traffic controller is to ensure safety standards are being upheld, through separation standards that should be maintained. Furthermore, controllers are responsible for guiding aircraft through their sector to their exit points, where they will be handed off to a different controller in a different sector.

ATCOs make use of decision support tools in order to aid them with their responsibilities. Systems such as VERA or STCA and MTCD provide aid in detecting conflicts between flight pairs in the sector. Other systems like LORD can assist ATCOs in determining the best course of action for conflict resolution maneuvers. The use of decision support tools in ATCO tasks can help reduce their workload.

Alternative approaches to ATC operations are being researched. One such emerging concept is that of Flight-Centric ATC. Flight-Centric ATC has controllers manage aircraft from departure to arrival terminal area. In this way, sectors are abolished, leading to a new ATC work-domain in which conflicts between different controller aircraft pairs occur. Several strategies exist for the allocation of flights and for rules regarding conflict resolution. Alternative operations for ATC are being investigated to find solutions to the growing problem of air traffic. Reducing workload through concepts like Flight-Centric ATC or decision support tools can help increase the capacity of airspace.

Flight-Centric ATC is a solution that focuses primarily on evolving the role of human controllers to be more efficient to allow for an increase in airspace capacity. Other approaches, such as ARGOS, focus on more automation-based innovations, reducing the need for human controllers. This is discussed in more detail in the following chapter.

3

Automation in Air Traffic Control

With demand for air traffic expected to grow significantly, and with current practices being limited by controller workload, research into automation has been a key topic of interest in the ATM community. More specifically, research is focusing on human-automation interaction in order to shift some of the workload to automation, increasing the capacity of the human controller and the sector being monitored. However, introducing automation in the control loop can have negative consequences for the human controller.

This chapter focuses on the current state of automation at EUROCONTROL, as well as the vision for SESAR to further implement automation. Furthermore, the potential risks of introducing automation in the human-control loop are discussed. The chapter is concluded by discussing task divisions in human-automation shared control environments.

3.1. Automation at EUROCONTROL

With the expected increase in air traffic in the coming years, aviation authorities around the world have introduced projects such as NextGen in America and SESAR in Europe. The goal of these projects are to modernize aviation practices through the introduction of automation in ATC systems to increase the airspace capacity. According to the European ATM Master Plan, "*The SESAR vision aims to deliver a resilient and fully scalable ATM system capable of handling growing air traffic made up of a diverse range of manned and unmanned air vehicles in all classes of airspace, in a safe, secure, sustainable manner*" [4]. One of the primary goals of SESAR is to implement the concept of trajectory-based operations (TBO). TBO enables airspace users to plan and carry out their own preferred flight trajectories so that destinations can be reached on time and cost-effectively. TBO introduces a more flight-centric approach to ATC, providing ATCOs with exact trajectories in 4D which can be used as a reference for the expected position as well as the expected time to reach a waypoint.

In order to achieve TBO, ATM modernization must look at the flight as a whole rather than in segmented portions [4]. One of the ways that SESAR aims to modernize ATM is through improving the productivity of air navigation services (ANS). The first steps taken by EUROCONTROL start with increased levels of automation support in ATC, where controllers will perform less manual and repetitive tasks as these will be delegated to automation. Ultimately, the goal is to reach full automation for ATC which would allow for the defragmentation of European Airspace, a necessity for implementing TBO [4].

Figure 3.1 shows an automation model from the SESAR master plan indicating the general plan for the implementation of automation. In this figure, four tasks are defined, Information Acquisition and Exchange, Information Analysis, Decision and Action Selection, and Action Implementation. For each of these tasks, the degree of automation support is given for each of the five levels of automation. Currently in ATC, only lower levels of automation are present, whereby action can only be initiated by the human controller. Reaching level three and beyond enables actions to be initiated by automation and can be used to take over some of the tasks traditionally carried out by an ATCO thereby reducing their workload. The vision in the SESAR master plan is to have ATM in Europe at automation level two by 2035.

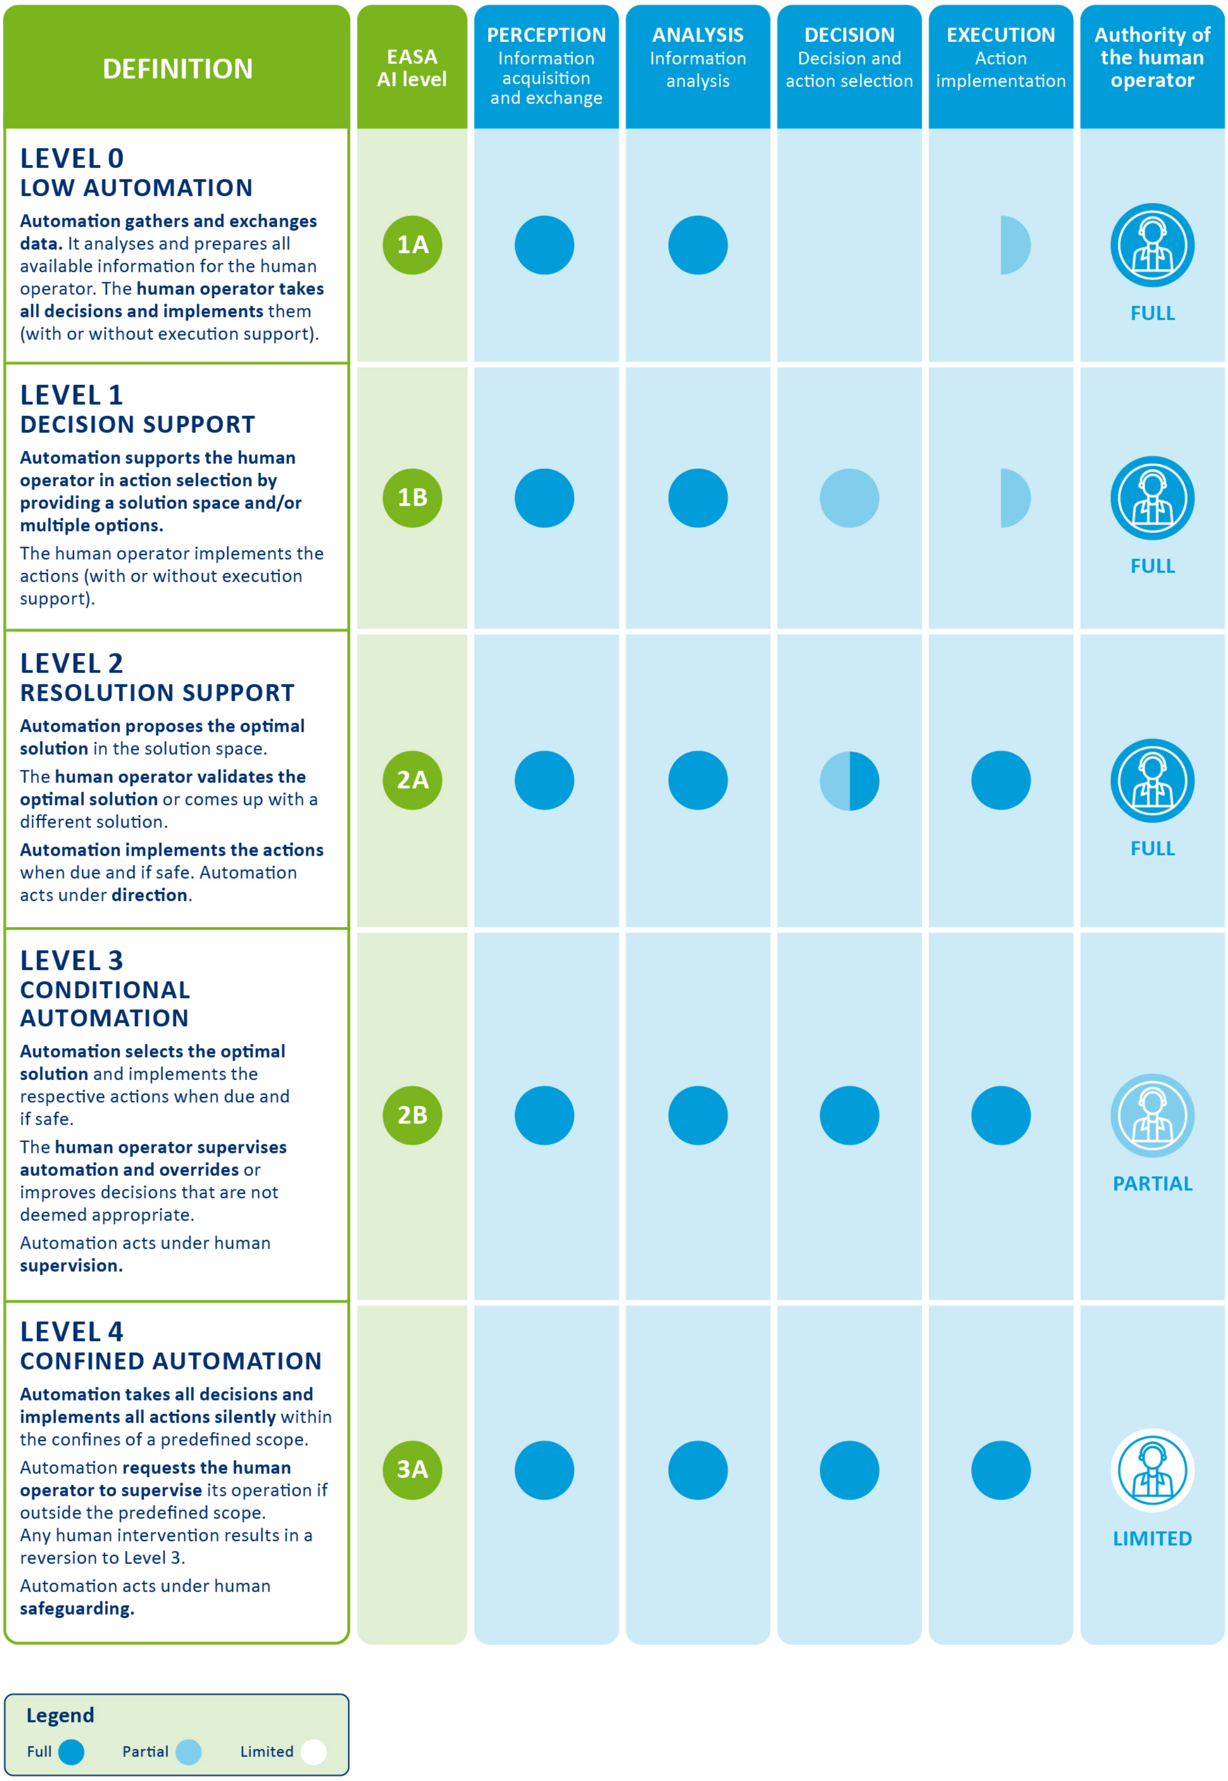

Figure 3.1: Levels of Automation in the SESAR master plan [4].

3.1.1. ARGOS

The latest software being developed by MUAC to achieve the automation goals presented in the SESAR vision is ARGOS. According to a MUAC environmental report ARGOS is quickly becoming the most intelligent assistant for ATCOs. The goal of ARGOS is to "Let ATCOs focus on the real, challenging work, to do what they are best at, and leave the routine work to the machine." [29] With the development of ARGOS, the vision is to be able to delegate basic traffic to automation and have ATCOs control complex flights, resulting in reduced workload for the controller. MUAC defined three objectives that will aid in the implementation of the vision [29].

Objective 1: Fully automated pre-tactical phase

Objective 2: Automated decision making and execution support for complex traffic

Objective 3: Fully automated separation assurance in the basic traffic

Similarly to Figure 3.1, a 10 level model for ARGOS has been created to describe the different levels of automation. This model ranges from L0-L9 where L0 is ARGOS doing nothing, and L9 is ARGOS solely managing all flights [29]. The remaining levels progress from basic support provided by automation to full automation where ATCOs can still choose to take flights away from ARGOS. These levels are used as operational guidelines for the implementation of ARGOS.

To achieve objectives 2 and 3, a three-step plan is used to introduce the automation aspects into the ATCO environment. Each of these three steps corresponds to a different level of automation in the 10 level model, namely L3, L5 and L8. In L3, ARGOS will only suggest actions to the ATCO. When an aircraft enters MUAC airspace, ARGOS suggests a flight plan to the ATCO. The ATCO has full control and can decide whether or not to accept the proposed plan. In L5, control over a sector is divided between ARGOS and the ATCO. ARGOS manages certain flights, whereby for each flight the flight plan is presented to the ATCO and then executed. The ATCO controls the remainder of the flights in the sector, and can opt to take control over any flights being managed by ARGOS. In L8, control has shifted from the ATCO to ARGOS, with ARGOS managing all flights. The primary role of the ATCO is then to monitor flights that ARGOS deems outside its comfort zone. In this situation, the ATCO can then decide to take control of flights managed by ARGOS. Doing so would ultimately degrade the level of automation back to L5.

Through these three steps, MUAC aims to transition from human-centric ATC to a more automation-centric ATC, while still keeping a human ATCO involved in both. In the lower levels, L3 and L5, responsibility of all traffic in the sector remains with the ATCO, even if some flights are being managed by ARGOS. The ATCO is responsible for monitoring those flights and ensuring that ARGOS is behaving as expected. Initially ARGOS can only handle easier flights and simple scenarios, whereby the logic used can be tested comfortably with the frameworks decribed in L3 and L5 levels of automation. As ARGOS is further developed and improved upon, the automation system can start to handle more complex scenarios, enabling automation to take more control of the sector. This leads to the implementation of L8, where the focus of ATCO changes completely from controlling to monitoring. The ATCO should be ready to intervene should ARGOS not be able to comfortably manage the sector. For this, MUAC claims that the ATCO does not need to be in the loop for every aircraft at all times, but that ARGOS will be able to help get the controller back in the loop if extra monitoring or an intervention is deemed necessary.

3.2. Automation Challenges

While an increase in automation in ATC will have many potential benefits, it is also important to consider the drawbacks that come with the implementation of automation. In a paper published by Bainbridge, the so called ”Ironies of Automation” are introduced. Bainbridge states that the purpose of introducing automation is to replace a human operator that may be seen as unreliable and inefficient [30]. However, in doing so, the design of the automation system leaves the human operator to complete tasks that were deemed too difficult to automate. Ultimately, the system is then reliant on the human operator to control difficult tasks when the purpose was to completely eliminate them from the control loop. Furthermore, the design premise of these automation systems is to compensate for unreliable and inefficient human operators, yet these automation systems are also designed and created by humans, leading to questions on the reliability of the automation system.

Another irony of automation highlighted by Bainbridge has to do with the deterioration of cognitive skills. For an operator monitoring an automated system where interference is required, the operator must use knowledge to resolve the situation. The efficient retrieval of knowledge from long-term memory is dependent on the frequency of its use [30]. Assuming a reliable automated system, the knowledge is rarely accessed which may cause difficulties for the operator to recall the knowledge during critical moments. Furthermore, Bainbridge explains that this type of knowledge is only developed through use and feedback about its effectiveness, which is also negatively affected by the infrequent use.

Introducing higher levels of automation as is planned for the ARGOS system changes the role of the human controller to a monitoring role. In this case, the human is expected to intervene when the automatic systems are not behaving correctly. According to several vigilance studies [31], it was found extremely difficult for a human to maintain focus and attention towards a source of information where very little happens for more than half an hour. This study suggests that the human operator is therefore not well suited for the monitoring of unlikely abnormalities, as they will be unable to retain focus. Further studies [32] have also shown that monitoring tasks ended up having negative consequences on both the situational awareness of the opeartor as well as the decision time needed to address a system breakdown. This, coupled with the fact that no reduction in workload was measured brings in to question the benefits of introducing human-automation systems with humans in the supervisory position.

According to Endsley and Kiris [32], a loss of situational awareness is said to be one of the largest contributors to the out-of-the-loop (OOTL) performance problem. This OOTL problem discusses the diminished abilities of human operators working with automation to detect system errors and subsequently perform tasks manually in the face of automation failures. Endsley and Keris conclude that the lack of situational awareness stems from three major areas: vigilance, complacency and monitoring.

Vigilance and monitoring have already been discussed, where studies have shown that humans monitoring a system have difficulties maintaining their vigilance, especially in systems with high levels of reliability, where intervention only occurs in rare cases. Lack of attention is sometimes attributed to the task being too simple which leaves the human feeling bored, however, a study by Parasuraman found that these vigilance problems persist even when dealing with complex tasks. Parasuraman concludes that ”vigilance effects can be found in complex monitoring and that humans may be poor passive monitors of an automated system, irrespective of the complexity of events being monitored.” [33]

In a study conducted by Bergeron, it was found that increased levels of automation in an autopilot caused the pilot to be more likely to lose track of where they were [34]. Endsley and Kiris claim this was due to complacency, which occurs when the human has too much trust in the automated system, leaving them to become absorbed in other tasks and thus no longer maintaining a sufficient level of situational awareness [32]. In cases where automation systems become very reliable, the result is that the humans end up relying on the automation as well, possibly making it difficult for the human to carry out the tasks should automation failure occur. Further research by Parasuraman found that performance effects as a result of complacency were more likely to occur in a multitask environment [33]. It was found that when the operator’s sole task was to monitor the automation, complacency had no effect on the performance of the controller.

Endsley summarizes all the automation challenges that affect the human-automation relationship as the "automation conundrum". As the level of automation is increased in a system, and the automation becomes more reliable and robust, the less likely that the human operator overseeing the automation system will be aware of critical information and be able to take over manual control if necessary [35]. The relationship between factors involved in the automation conundrum are depicted in the human-autonomy system oversight (HASO) model. An overview of the model can be found in Figure 3.2.

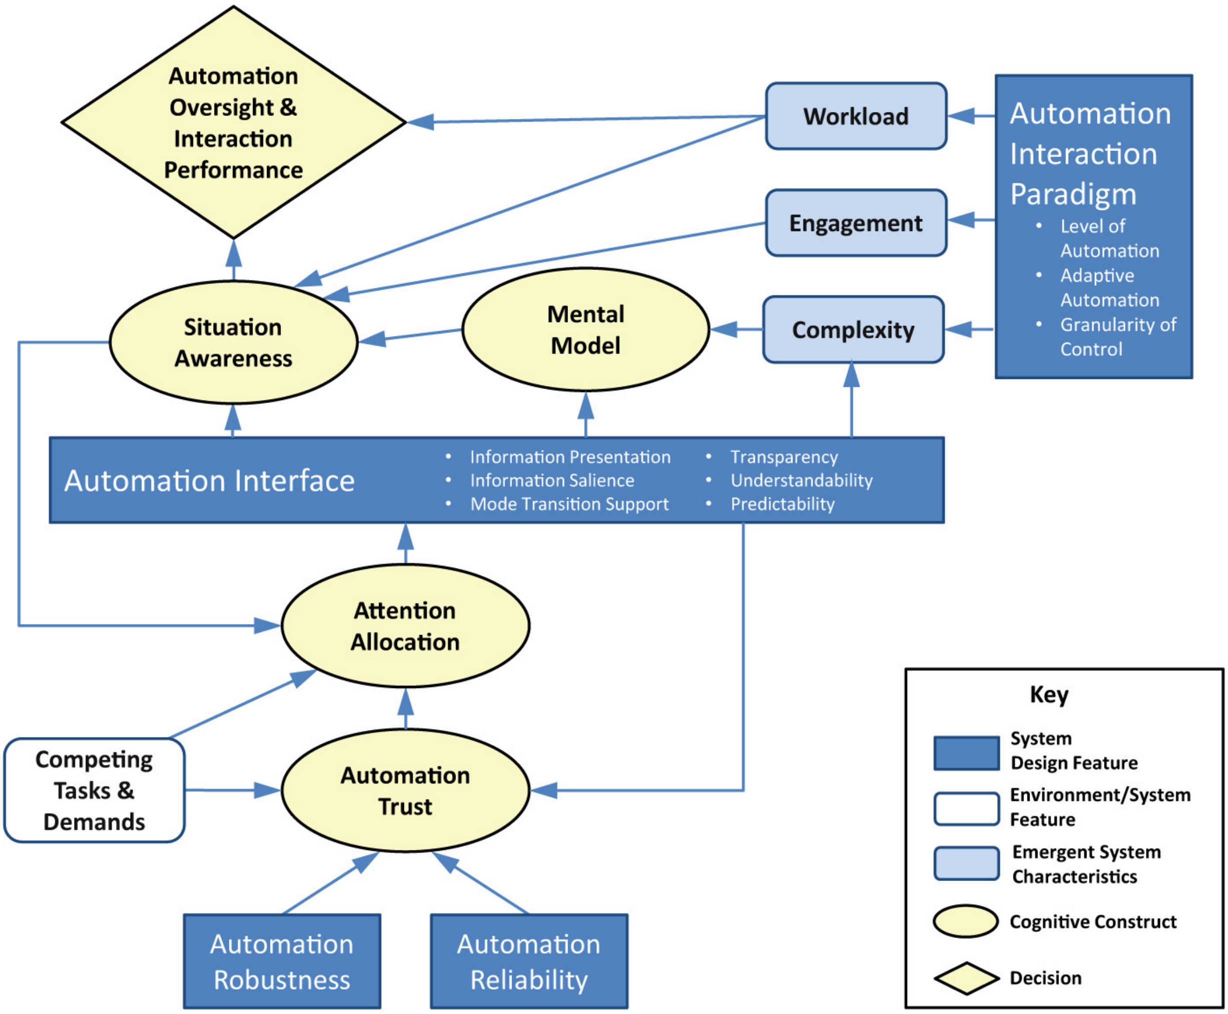

Figure 3.2: HASO model depiciting the factors influencing the automation conundrum [35].

This model shows that the performance of the human operator in the supervision and intervention of automation tasks depends on the situational awareness and the workload of the human operator. The operator must have the situational awareness to realize that the automation is performing incorrectly or that the situation is outside the capabilities of the automation. Furthermore, the operator must have sufficient mental capacity, meaning the operator must have sufficient time and resources to be able to handle the issue.

In Figure 3.2, there are four important automation design features mentioned that influence the oversight and performance of the operator. Increasing automation robustness and automation reliability leads to an increase in automation trust from the operator. This model shows that increased automation trust will influence the allocation of attention of the operator, with more trust increasing the likelihood that the operator not paying enough attention to the automation systems.

The automation interface contains all the information presented by the automation system. This includes information that helps the operator understand the logic automation systems use to complete their tasks, with good clarity leading to automation transparency and predictability for the operator. The design of the automation interface influences the apparent complexity of the system and the ability of the operator to create and maintain an accurate mental model. When the automation interface is transparent and understandable, the operator in turn places more trust in the automation system [35].

Lastly, the automation interaction paradigm is an automation design feature that influences the success of the operator in managing the system. The automation interaction paradigm consists of three different automation paradigms. The granularity of control can range from high (manual control) to low (goal-based control) whereby only a high-level goal needs to be provided to the system. As the control granularity decreases, the workload of the operator should also decrease. However, lower granularity could also lead to less situational awareness due to the increased amount of automation tasks [35]. Adaptive automation is another paradigm in the HASO model which attempts to increase operator engagement through use of intentional periods of manual control for otherwise automated tasks. The trigger for periods of manual control is flexible and can be based on parameters such as set time periods, critical events or drops in human performance. The purpose of adaptive automation is to positively impact human performance through reduction of workload. The influence of the different levels of automation on the operator have already been discussed in this chapter.

It is evident that the introduction of automation in human-automation systems provides many challenges that must be overcome. While automation may reduce workload and increase the overall efficiency and capacity of the system, care must be taken to ensure that the human operator remains a relevant and engaged part of the control loop. This is important for systems whereby control tasks are shared between the operator and automation as well as for systems whereby the operator acts as a supervisor, intervening only where automation fails. The human-automation system must therefore be carefully designed in such a way that the potential negative effects of human-automation cooperation are mitigated.

An approach to achieve increased levels of automation is through a parallel team of multiple agents working autonomously and interdependently to achieve shared goals [7]. It is hypothesized that this approach avoids the automation challenges raised above [36]. With ARGOS, MUAC aims to create a parallel system in which a human ATCO and a digital (computer) ATCO work independently to control and manage the sector [7]. In this way, the automation system can reduce the workload of the human, while the human remains engaged and involved with their tasks.

3.3. Human-Automation Shared Traffic Responsibilities

On the way to full automation in ATC, there are several stages whereby the operator and automation must share control of the sector. In this section, the findings of previous studies regarding the task division of these shared human-automation systems are discussed.

In 2012, NASA published a paper discussing a series of studies carried out aimed at investigating the introduction of automated separation assurance as a solution to the growing capacity problem in US airspace. It was concluded that with the introduction of automated systems, a higher level of air traffic could be safely managed [37]. In one of these studies, an experiment was conducted in which controllers operated in a comprehensive work environment. ATCOs were responsible for completing a wide range of ATC tasks, working together with automation that was also performing control tasks. In a survey conducted after the experiment, participants were asked to give feedback on what tasks could be allocated to automation. The results of the survey can be found in Figure 3.3.

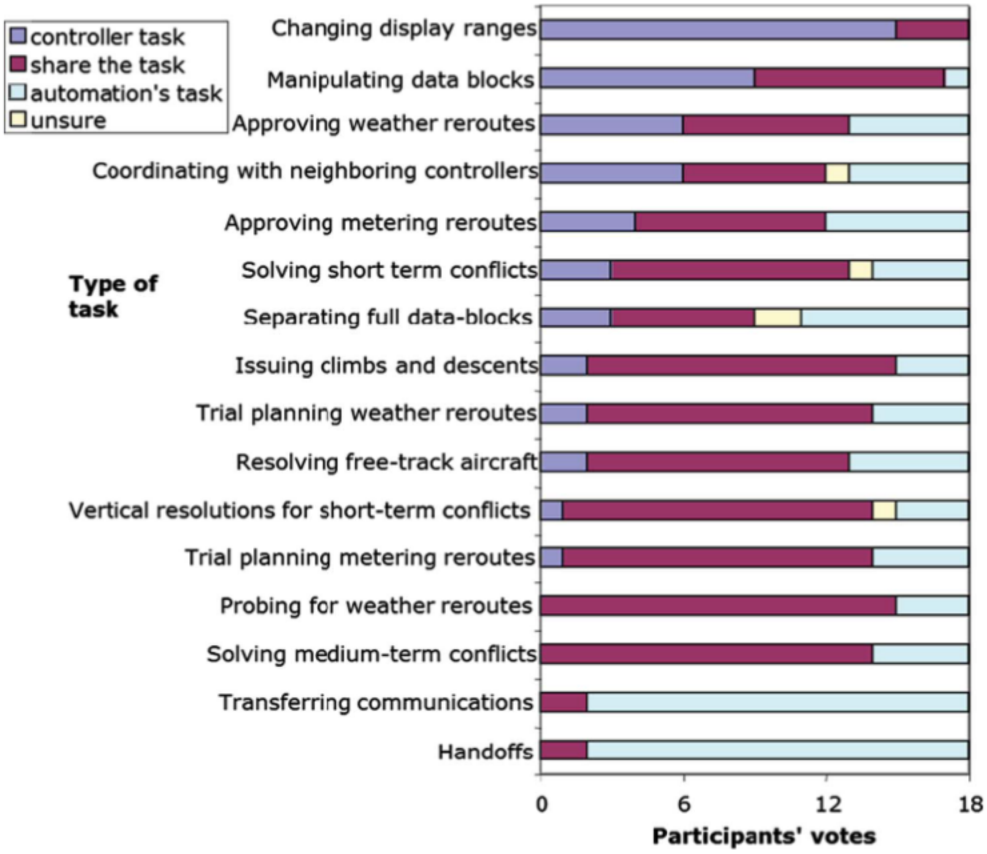

Figure 3.3: Task allocation as given by ATCOs [37].

Looking at the results, the only tasks that ATCOs had a strong preference for doing themselves was changing the display ranges and manipulating data blocks. These tasks both affect the mental model and situational awareness that the ATCO has over the system, so they prefer to be in control of this. On the other side, handoffs and transferring communications were mostly preferred to be given to automation to handle. For these individual routine tasks, ATCOs had a strong preference to let either themselves or the automation be fully responsible for that particular task. The remaining tasks were all decision-making tasks, where ATCOs showed their preferred response was to share the task with automation, indicating the willingness of ATCOs to incorporate increased levels of automation into the ATC work environment.

An experiment carried out by de Rooij et al. examined how ATCOs allocated flights in a shared human-automation en-route airspace. In this experiment, the ATCO had full control over which flights were given to automation, and were able to switch who had control of a particular flight at any point in its trajectory [38]. In a survey similar to the one carried out by NASA, the preferred task allocation between human and automation was collected. The results can be found in Figure 3.4. These results show similar findings as in Figure 3.3 with ATCOs showing a preference for controlling short-term tactical actions and leaving automation to (partially) control more strategic long-term tasks.

Along with the task allocation survey, the factors influencing the decision of the ATCO to assume or give control to automation for a particular flight was also surveyed. The results of this can be found in Figure 3.5. From this survey, the factors deemed important by ATCOs were largely those that looked at the potential interactions of the flight of interest, as well as considering the capabilities of the automation. Another finding of the study was that all ATCOs unanimously agreed that complex flights, which were defined in this experiment to be flights that require more than 2000ft level change, should be handled manually, while basic flights were preferred to be delegated to automation.

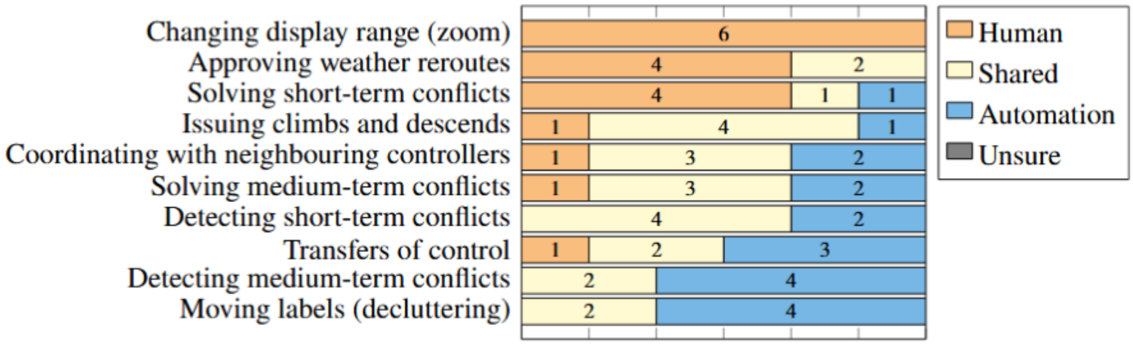

Figure 3.4: Task allocation as given by ATCOs [38].

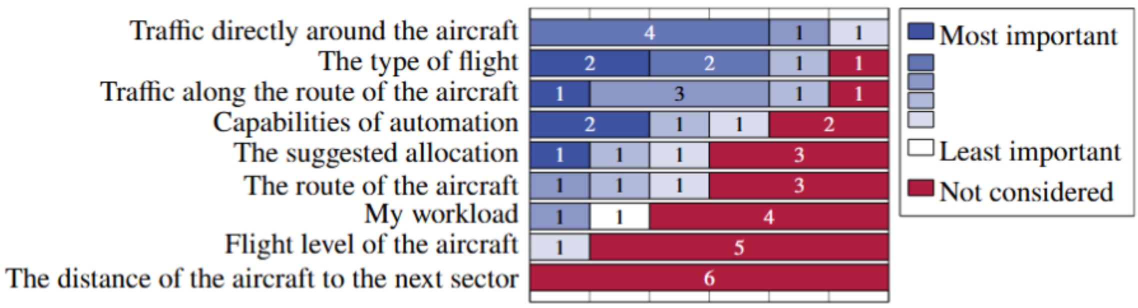

Figure 3.5: Driving factors for ATCO allocation strategy [38].

3.4. Conclusion

In ATC, automation is a tool that can be used to shift parts of the work from humans to automated systems. The implementation of automation tools can vary depending on the desired support level, ranging from decision support tools to fully controlled automation in a sector. At EUROCONTROL, modernization projects to include higher levels of automation in ATC are already well underway. The latest tool being developed is ARGOS, an intelligent assistant for ATCOs aimed at controlling basic traffic, to allow for controllers to deal with more complex work. One of the steps in the journey to implementing the MUAC vision for ARGOS is to split traffic within a sector between ARGOS and an ATCO. ARGOS would be in charge of the simpler, more basic traffic while the ATCO would be responsible for managing the more complex traffic.

While several different levels of automation support exist, each with their own benefits, the introduction of automation in the ATC work domain also brings challenges that must be considered. Through increased levels of automation, the human ATCO can experience a decline in their situational awareness and focus, as their tasks shift from more active control tasks to monitoring. Therefore, when introducing automation systems in ATC, it is extremely important that the potential effects on the ATCOs performance are accounted for. Keeping workload and engagement of the ATCO at an optimal level is important to keep them aware and functioning at a high level.

Studies carried out looking at human-automation shared traffic responsibilities show a strong preference from ATCOs to allow automation to handle very simple tasks. ATCOs preferred to be in charge of tasks that affected their mental model and situational awareness, while more routine tasks like handoffs were preferred to be delegated to automation. Through these studies, a better understanding is built of the preferences of ATCOs when it comes to sharing responsibility of a sector with automation systems.

4 Complexity in Air Traffic Control

Air traffic complexity is defined as the difficulty of monitoring and managing a specific traffic situation [39]. ATC complexity has been the subject of many studies related to ATC operations, which can be divided into two categories. The first is sector complexity, which looks at factors influencing the complexity of a given sector as a whole. The second is flight-centric complexity, in which complexity is determined for all flights individually, rather than looking at the influence of all flights on the sector.

The goal of determining the complexity in both cases is to be able to measure the relationship between complexity and controller workload. Limitations in the capacity of the controller workload are often the limiting factor when it comes to increasing airspace capacity. By measuring workload through complexity, solutions can ultimately be derived to increase airspace capacity.

4.1. Controller Mental Workload

To be able to understand the limits of the controller workload capacity, several studies have been carried out to model mental workload and determine its relationship with air traffic complexity. Hilburn [40] provides a generalized definition of mental workload applicable to ATC, defining mental workload as a representation of an interaction between a task and an operator. The mental workload looks at the difficulty of a task, as well as the required effort to achieve the task, which can vary for different task-operator combinations. In Hilburn's research paper, a distinction is made between taskload and workload, where taskload is defined as the objective demands of the task and the workload is the subjective demand experienced by the controller while performing the task.

Difficulties arise when creating a model for mental workload due to its subjective nature and the reliance on studying how the brain processes and handles information. Several different strategies have been used to model the control workload. Schmitt [25] looked at modeling workload through the measurement of observable controller actions, creating a control difficulty index, a weighted sum of the expected frequency of occurrence of events that affect controller workload. The tasks were weighted according to the time required to execute. Antulov et al. [39] argued that such an approach was limiting, as only observable actions are considered.

Other methods used to model mental workload include subjective workload measurement techniques, of which there are many. One such method is the Instantaneous Self-Assessment (ISA) in which workload is measured through immediate subjective ratings of work demands as an air traffic controller performs a task. In a study evaluating ISA as a valid measure of workload, Tattersall and Foord [41] concluded that ISA ratings were correlated with post-task ratings of workload, and that ISA ratings were sensitive to changes in task difficulty. The main drawback found was that ISA responses negatively affected the performance of the primary task.

Although the modeling of mental workload has many challenges due to its subjective nature, Hilburn's literature review [40] stated that traffic density is the most cited, studied and evaluated traffic characteristic contributing to controller workload. Traffic density was found to have an impact on communication time and manual performance time.

An increase in traffic density makes overseeing and managing traffic in the sector more difficult, as well as limiting the solution space for conflict resolution maneuvers. A study by Hah et al.[42] found that traffic density had a significant linear relationship with the controller workload. Furthermore, increased levels of controller workload have a negative and significant effect on controller performance [43]. Controller workload thus play a significant role in the capacity of airspace. Solutions must be found that reduce the controller workload while still facilitating the growing demand for air traffic.

The introduction of automation as a digital controller allows for some aircraft in the sector to not be handled by the human controller. In this way, the workload of the human controller can be reduced without reducing the traffic density. This is the idea behind ARGOS, an automation system developed by MUAC which is further explained in Section 3. However, due to limitations in automation software, not all aircraft can currently be handled by automation. At this time ARGOS can only handle simpler aircraft, due to the difficulty of modeling more complex scenarios. An understanding of complexity metrics are therefore needed to be able to assign digital controllers to simple aircraft, and reduce the workload of the human controller.

4.2. Complexity Metrics

In order to implement automation to control traffic alongside a human controller, a distinction needs to be made between basic and non-basic flights. This distinction is made by looking at the complexity of individual flights and establishing a complexity threshold for which all flights that are more complex are considered non-basic. To achieve this, complexity metrics for flights need to be defined so that the complexity of an individual flight can be measured.

4.2.1. Sector Complexity

Previously, most research focused on determining the complexity of a sector as a whole. For this, several metrics were used. The simplest of which is the density of the airspace, in which the complexity of the sector scales with the number of aircraft present in this sector. In European airspace, flow management staff use the aircraft density to determine when to merge or split sectors [44]. Dynamic density is another indicator used to determine the complexity of a sector as an extension of the airspace density method. The dynamic density is determined through the structural and flow characteristics of the airspace [45] [40].

Structural characteristics of the airspace are static, including factors like terrain configuration, number of airways, and airway crossings. Flow characteristics are dynamic, and are subject to change with time. This includes factors like number of aircraft, weather, and flow restrictions. The interaction of all these factors produces a nonlinear way to determine air traffic complexity. Several dynamic density models have been developed, and they tend to be sector dependent, due to sector geometries as well as the subjective choice by ATCOs in determining relevant complexity factors. This makes the dynamic density model a sector-dependent metric to determine the complexity of the sector [44].

4.2.2. Flight Centric Complexity

While many frameworks exist for determining the complexity of a given sector, a less-studied concept is looking at the complexity of individual flights. Stienstra [6] carried out an experiment in order to identify characteristics of flights that influenced the perceived complexity of the flight. To do this, five complexity metrics were derived from sector complexity metrics. These derived metrics looked at describing the complexity of individual flights, taking into account the trajectory and exit flight level (XFL) of the flight [6]. The five metrics Stienstra derived are as follows:

1. **Number of potential loss of separation (LOS) situations on the trajectory -** Defined as a situation in which two flights come within 7 NM of each other without sufficient vertical separation. The output is the number of predicted LOS situations along the trajectory.

2. **Number of crossings on the trajectory -** When a flight comes within 10 NM of the considered flight in the lateral plane, a crossing has occurred. The output is the predicted number of crossing along the trajectory.

3. **Closest point of approach (dCPA) -** Defined as the minimum distance in the lateral plane that will occur between the considered flight and any other flight in the sector. ATCOs always take action if the predicted minimum distance is less than 7 NM, otherwise they allocate additional attention to flights for which the dCPA is between 7 and 10 NM [6]. The dCPA is given in NM.

4. **Time to the loss of separation moment (tLOS) -** Defined as the time remaining until loss of separation occurs between the considered flight and any other flight in the sector. The boundary for when tLOS requires action is dependent on the sector geometry, though an unsolved tLOS of 3 minutes or less is generally said to make an ATCO nervous [6]. The tLOS is given in seconds.

5. **Required flight level change -** Defined as the difference between the current flight level and the required flight level of the considered flight. The required flight level change is given in feet.

After developing the five complexity metrics for individual flights, an experiment was performed to determine the ability of these individual metrics to predict the experienced complexity given by an ATCO [6]. The experiment was carried out by focusing on a flight of interest (FOI) and determining the effort required to bring the FOI from its current location to the required sector exit point. All flights from the background traffic that were considered in the complexity assessment were also indicated by the ATCO. Stienstra concluded that the experimental setup was limiting and by only considering single traffic characteristics in the experiment scenarios, nothing significant could be concluded about the prediction of experienced complexity. It was also mentioned that future work should look at a combination of factors to determine complexity, and focus should be placed on the correlation between different metrics that could influence the complexity to be able to understand what the dominant factors affecting complexity are.

Despite the results of the experiment being inconclusive, a questionnaire was also given at the end of the experiment, asking ATCOs to indicate which factors contributed to the complexity of the flights. There were six factors provided for the ATCOs to choose from, and multiple could be indicated per flight. The six factors are listed below. The list is in descending order from the most recurring factor to the least recurring factor.

1. XFL of the FOI
2. Possible loss of separation between FOI and included flight
3. Time remaining before loss of separation
4. Flight level of FOI
5. The fact that the FOI is climbing or descending
6. Solution space available for the FOI

In a follow up study by de Rooij et al. [46] an experiment was conducted with fifteen professional ATCOs from MUAC. The purpose of the experiment was to determine which background flights influenced the complexity assessment of a selected flight for a controller. The reason for selecting a relevant background flight was recorded in a questionnaire. The results were then used to interpret the relationship between the number of flights considered and the perceived complexity of the flight of interest. The results of the study can be found in Figure 4.1.

Findings from the study noted that the model performance was weak. However, some conclusions were still drawn from this study. From the results, it was concluded that altitude overlap was by far the most important feature affecting perceived complexity, which is closely related to the number of included background flights. It was noted that flights closer to their exit point (XCOP) were also deemed important, due to the limited solution space to resolve any conflicts.

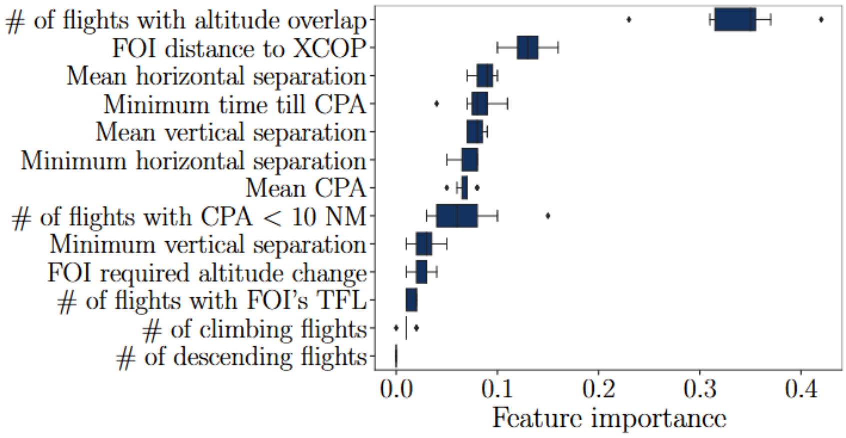

Figure 4.1: Feature importance of regression model [46].

Limitations of the study were that the experiment was carried out in a static scenario. In order to determine whether the results are generalizable, dynamic scenarios need to be tested. Furthermore, de Rooij et al. recommend that future work look at the determination of a complexity threshold, as well as carrying out an experiment to validate the operational applicability of a complexity model in a shared human-automation airspace.

4.3. Flight Filtering

Building off the flight-centric complexity study conducted by de Rooij et al., a follow-up study by Kumbhar et al. [23] looked at determining the relevance of background flights for a particular flight of interest. The information was then used to produce a flight-centric filtering mechanism for the flight of interest, fading all flights not deemed relevant to the flight of interest. The filter parameters were determined and the filtering mechanism was then tested using the data collected by de Rooij et al.[46].

The filter was created by mapping out the decisions carried out by an ATCO during enroute traffic. The ATCOs first scan radar display to identify flights that require clearance. A clearance can be issued to due to the need to resolve conflicts, guiding flights to their transfer flight levels or to optimize the flight trajectory towards the sector exit.[23]. When a flight has been selected, the ATCO performs a conflict detection task, checking the impact of the clearance on the safety of the surrounding traffic. The flowchart for a typical ATCO's conflict detection strategy can be found in Figure 4.2.

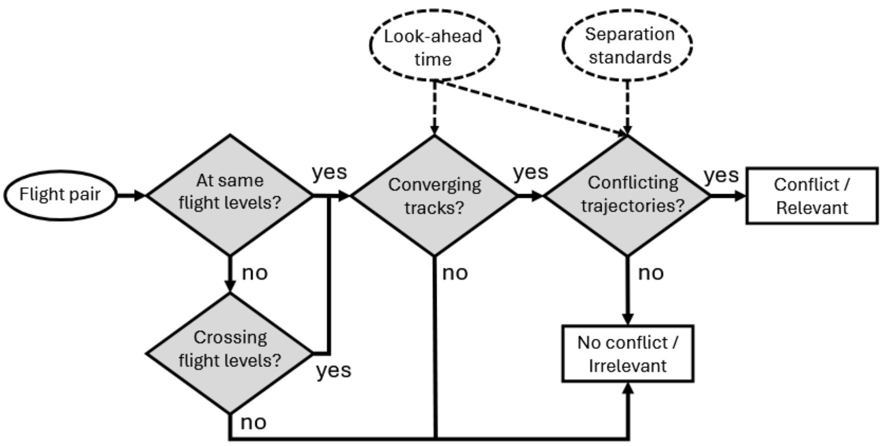

Figure 4.2: ATCO conflict detection flowchart (Dashed-lines are dependent on the ATCO) [23].

The filtering mechanism was then modeled using the flow chart, following a 3 step process to determine the relevance of a particular flight for a flight of interest. The first step is to check for altitude overlap, where the initial cognitive effort to check for overlapping altitudes is relatively low. If there is no altitude overlap, then further examination of other flight characteristics are irrelevant for determining flights of interest.

If two flights have overlapping altitudes, then the next step is to determine the spatial overlap of the flights. In order to ensure horizontal separation, the ATCO determines if the flight trajectories are parallel, converging, or diverging from the flight of interest. Should the flight have spatial overlap, or if the flights are likely to share a common waypoint or cross paths, then the flight remains of interest.

Should there be both altitude and spatial overlap, then the final step is to check for temporal overlap. By extrapolating the positions of flights using speed-distance calculations, the time and closest distance at which flights reach potential interaction points is calculated, giving an indication of the immediate threat level of the interaction. If a flight satisfies all three criteria, then this flight is determined to be a potential interacting flight relevant to the flight of interest. All other flights are deemed to be irrelevant to the flight of interest.

Having determined all the relevant criteria for the filtering mechanism, the filter was then constructed. Two different types of filters were created accounting for the differences in look-ahead time used to predict future positions. Look-ahead time is dependent on the ATCO as shown in Figure 4.2 and is influenced by the traffic scenario. During heavy traffic, shorter look-ahead times are generally preferred by ATCOs, and in low traffic, longer look-ahead times are used [23]. This led to the creation of two filtering concepts, state-based and intent-based extrapolation.

In state-based extrapolation, flights profiles are projected forward based on their current positions, headings and speeds. These are then used to determine the closest point of approach with other flights and do not account for the intended flight plans. State-based extrapolation is said to be a more 'reactive' approach, and is much more accurate for a shorter time scale.[23]. In intent-based extrapolation, the position of a flight is extrapolated using the planned altitude and speed profiles. The path between two planned waypoints is a line segment which is used to look for spatial overlaps while the planned time of arrival at waypoints is used for the temporal overlap. Where spatial and temporal overlap exists, the CPA is calculated by predicting the state of the flight pairs at the start of the critical route segment [23]. A visual representation of the geometry of the state and intent based filters can be found in Figure 4.3.

Following the creation of a state-based and intent-based filter, a third filter was also constructed. The consolidated filter was developed, combining the state-based and intent-based filtering parameters. All three filters were then tested by comparing the relevant flights selected by each of the filters for a particular flight of interest with the relevant flights selected by an ATCO in the same traffic scenario. The consolidated filter performed well in correctly predicting relevant flights from all actual relevant flights, with a success rate of 90.46%. An example of the results of the filtering mechanisms developed can be found in Figure 4.4.

Conclusions of the study note that the success of the filter is highly dependent on the ability to robustly predict interactions between flights while also noting that real-world uncertainties were not considered. Kumbhar et al. also note that the flight-filtering mechanism can be used in future studies as a basis for a flight-allocation algorithm. The consolidated filter could be used in combination with automation systems to compute complexity of incoming flights based on the amount of relevant flights for the incoming flights. This information could then be used to assign a flight as 'basic' or 'non-basic' and the flight is then assigned to either a human controller or automation [23].

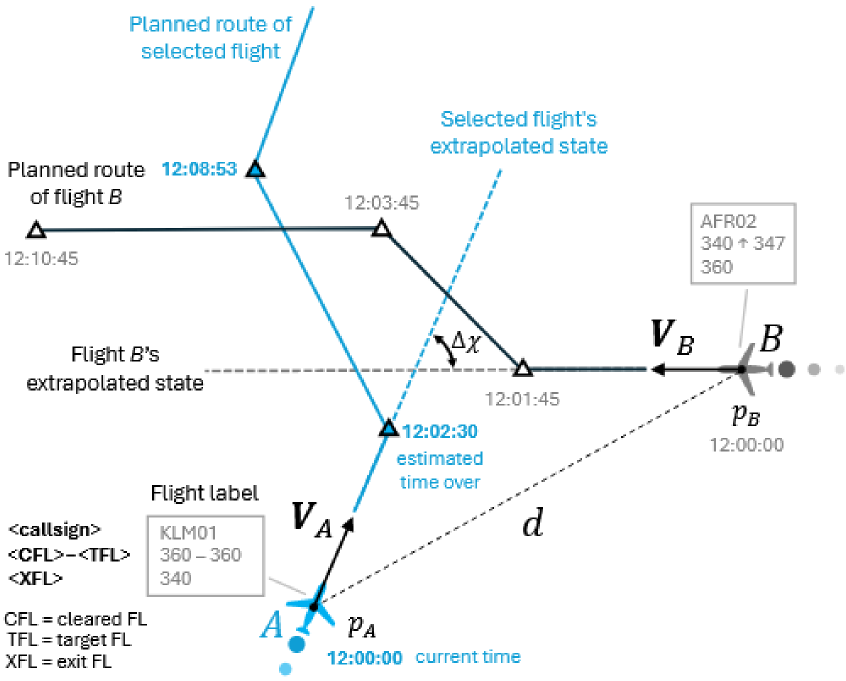

Figure 4.3: Geometry of state and intent based extrapolation [23].

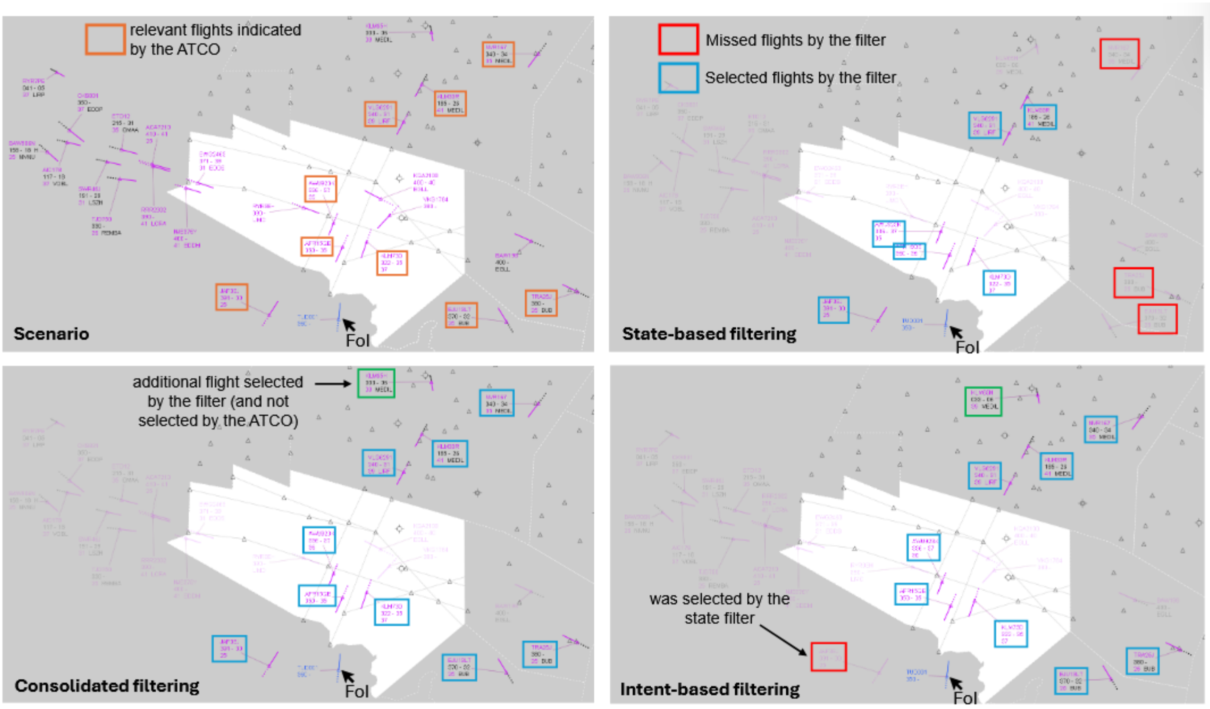

Figure 4.4: An example of the filter results [23].

4.4. Basic vs. Non-Basic Flights

In a human-automation ATC system, flights need to be designated to either the human controller or to automation. In order to decide the division of flights, a distinction can be made between basic flights and non-basic flights, whereby basic flights are given to automation and non-basic flights are given to the human operator. In general, basic flights are flights with little interactions and are relatively simple to manage in terms of workload. Non-basic flights have more interactions and require greater attention thus having a higher workload than basic flights. There are two reasons for the chosen strategy regarding the division of flights. The first reason is based on limitations of the capabilities of the automation software. The necessary maneuvers and strategies to manage basic traffic are well within the capabilities of the automation system, ensuring that the system is more reliable. The second reason for this assignment strategy is that the human ATCO will remain more engaged and kept in the loop due to the higher cognitive demand. This concept is further explained in Section 3.2.

Dividing the flights between basic and non-basic flights means that the human controller can give more attention to the difficult to manage flights, while leaving the simple flights to be controlled by the automation system. In theory, this allows for a potential increase in the sector capacity, as some of the workload the human operator spends monitoring and managing basic flights is now available. In order to maximize the benefits of such a shared control system, a threshold boundary needs to be determined for basic and non-basic flights.

An experiment was conducted by MUAC where an algorithm was developed to determine whether a flight was considered basic or non-basic. A flow chart describing the rough logic of the allocation system can be found in Figure 4.5. At a time X before the flight enters the sector, the allocator decides whether the flight should be allocated to the controller or an automated system. The allocation is fixed for the full crossing duration. The first check is to determine whether the aircraft has Controller Pilot Data Link Communications (CPDLC) capabilities. CPDLC allows for the exchange of communication between controller and pilot without the use of voice transmission. If the aircraft does not have CPDLC capabilities, then it is immediately allocated as non-basic.

Next, an overview of historical data is collected for flights with the same 'flow' parameters, i.e. flights with the same flight plan characteristics. These parameters include the Aerodrome of Departure (ADEP), the Aerodrome of Destination (ADES) and the Exit Coordination Point (XCOP). For flights with the same flow parameters, historical data is used to collect any information on Direct (DCT) commands given to the flight. If the flight had been given a DCT in historical data, then it would be allocated as non-basic, otherwise the algorithm would continue. This criterion was eventually dropped, as more than 95% of the flights had received a DCT at least once in historical data for the BOL sector.

Lastly, the amount of predicted vertical interactions with surrounding flights is computed. This is computed using the vertical profile inferred from the flight plan. Any potential interaction occurs when flights are within 8NM of each other and within 10FL of each other. The interactions are calculated using the flight plans as well as historical data of DCTs. This generates "what-if"s for flights, where historical DCTs are used to linearly interpolate a new trajectory. Should the flight have a potential interaction with the new "what-if" trajectory, this is recorded. If the flight of interest has two or more interactions with other flights, then it is considered non-basic. If the flight has less than two interactions with other flights during its time in the sector, then the flight is considered basic.

An employee of MUAC explained that a further expansion of this method was also used, but its logic was not added to the flow chart. For flights that would be considered basic, a final check is done by looking at flights with a singular possible interaction. For these flights, if the interaction consists of another flight that is considered non-basic, then the initially basic flight is then also non-basic. This is because it would solve any potential conflicts regarding how the interaction is resolved and by which controller. This logic was then used to test and monitor the resultant percentage of basic and non-basic flights allocated in a sector, although the results were unavailable.

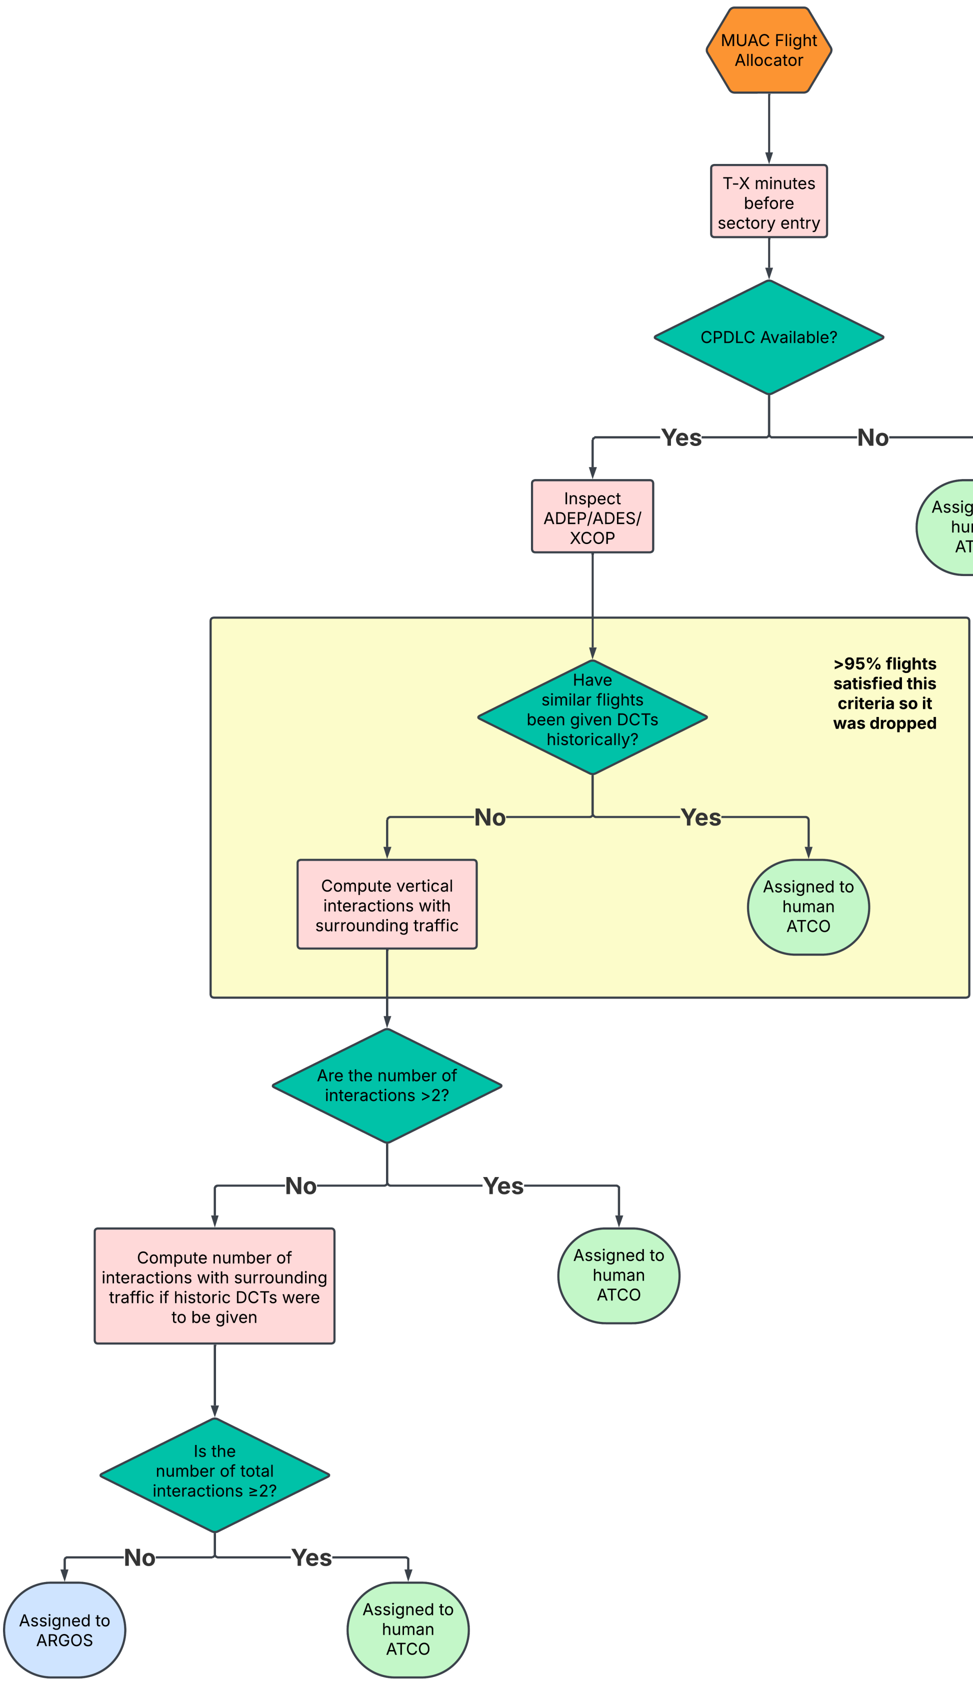

Figure 4.5: MUAC flow chart for determining controller assignments.

4.5. Conclusion

Complexity in ATC can have various definitions. Traditionally, complexity metrics are used to determine the complexity of a sector as a whole. More recent research aims to measure the complexity of individual flights within a sector. The relationship between individual flights and their respective complexity can be used as the basis for flight allocation methods. Using these individual flight complexity metrics, as well as other methods such as flight filtering, can help separate air traffic into different categories.

Through this implementation, a distinction can be made between less complex, 'basic' flights and more complex 'non-basic' flights. Using this as a basis for a traffic allocator can split the workload in a given sector between several human/automation controllers. This reduction in controller workload gives room for improvements to be made in regards to airspace capacity.

5 Flight Allocator Design

Based on the findings of the literature study as well as the research objective, stated again below, a flight allocator has been designed. The flight allocator is to allocate basic flights to the computer and non-basic flights to the human controller.

To design and implement a flight-based allocator model that allocates flights to either a human controller or a digital computer controller based on the calculated complexity of the flight and evaluate the parameters of the allocator with a simulator experiment.

5.1. Allocator Metrics

In order to design a flight-based allocator based on flight complexity, a clarification needs to be made on the metrics used to calculate complexity. As found in the literature review, there are a large number of factors that influence the complexity of a flight. Although some complexity characteristics have generally been deemed more important than others, no clear consensus on the contribution of each complexity metric to the total complexity of the flight exists. Furthermore, it was found that the influence of certain complexity metrics is dependent on the controller, making it more difficult to produce a general complexity model that can be applied for all controllers. The initial design for the flight allocator logic can be found in Figure 5.1. In this figure, the values n,x, and y are placeholder values that have yet to be defined.

5.1.1. Flight Filtering Mechanism

For these reasons, the basis of the flight allocator is the flight filtering mechanism designed by Kumbhar [47] [23]. The flight filtering mechanism allows a controller to see all relevant aircraft for a particular flight of interest, essentially looking for all surrounding traffic that could have a possible interaction with the FOI. The flight filtering mechanisms filters based on factors directly related to flight complexity, looking at metrics such as separations and altitude overlaps. In this way, flight complexity is defined as an integer value, corresponding to the number of relevant flights for the flight of interest(the amount of flights that are highlighted by the filter). This provides a general metric that can be used to calculate complexity of flights, from which the allocator can allocate a flight of interest to either a human controller or automation based on the number of relevant flights.

5.1.2. Human-Automation Controller Interactions

While the flight filtering mechanism acts as an estimate of the initial complexity of the FOI, several other factors that influence flight complexity not covered by the flight filter are also used in the flight allocator model. If a flight pair has a detected conflict, but the flights are both being controlled by different controllers, several issues can arise. Firstly, in such situations where controllers are surrounded by air traffic that is controlled by other controllers, whether that be other human controllers or automation, the mental picture of the sector can be deteriorated. Controllers risk being less aware of the intentions and activities of surrounding traffic not being controlled by them.

This can lead to surprise conflicts, where the controller has less time to react to the conflict and figure out a sensible resolution. Furthermore, conflicts between traffic controlled by different controllers can require an extra layer of communication. The controllers must communicate to decide who is responsible for solving the conflict, which adds an extra layer of complexity and extra time needed to resolve.

For this reason, an additional input for the flight allocator is to check which controller is responsible for every aircraft relevant to the flight of interest. Although the initial allocation method of a flight is only dependent on the complexity threshold of the number of relevant flights, exceptions to the method can be made based on the controller responsible for the relevant flights. In situations where a flight would be considered basic, due to a lower number of interactions with other flights, the flight allocator would initially assign the flight to the computer. However, if the interactions of the FOI are all with flights controlled by the human controller, then assigning the FOI to automation may actually cause an unwanted increase in workload. Several human-automation controller interactions would exist, complicating conflict resolution for the human controller. Therefore, in such a situation, it could be useful to assign the FOI to the human controller, despite the FOI being considered basic by the flight filtering complexity calculation.

5.1.3. Workload of Human Controller

Another factor that will influence the allocation model is an estimate of the workload of the human controller. While workload is subjective to the controller and difficult to measure, the drawbacks of having too much or too little workload have been well-established. Workloads outside of appropriate ranges can deteriorate the ability of the controller to carry out their tasks, either through stress and overwhelming them in the case of too much workload, and boredom and lack of focus in cases of too little workload. Introducing limits to the number of flights that can be assigned to the human controller can help strike a balance in workload to allow optimal performance.

In cases of low workload, flights that are usually assigned to automation are assigned to the human controller to give them more tasks to do. In cases of high workload, flights that would usually be assigned to the human controller and instead assigned to the automation to relieve the controller of additional workload. The implementation of the workload metric in the flight allocator can take two forms.

The first method for the workload metric in the flight allocator design is to have fixed thresholds set in the allocator model that determine when workload is at an unacceptable level. This would include a simplified workload model, whereby the contribution to the workload of each flight is estimated based on its calculated complexity. The sum of the workload calculations for all flights controlled by the human controller would then give an indication of the workload of the controller. Drawbacks of this method are that the workload calculation may not be an accurate representation of the actual workload of the controller, as it does not account for external factors that could influence the workload. Furthermore, workload is a subjective measure that is dependent on the controller, and applying a general model to all controllers may lead to problems where the thresholds are poorly chosen for certain controllers.

Alternatively, workload thresholds could be indicated by the human controller which the allocator model then uses to determine the allocation of the flight. Controllers can indicate through a button on the screen if their experienced workload is either too low/too high. This information can then be used by the allocator model to adjust future flight allocations depending on the preferences of the controller. This allows controllers to tailor the allocation algorithm to their own experienced workload.

Possible drawbacks of the workload threshold as a metric in the allocator model is that this metric essentially overrides all previously mentioned metrics. This can lead to suboptimal allocations being made, whereby complex flights could be assigned to automation, even though the automation is ill-equipped to deal with the complex flight. Furthermore, if a human controller deems the workload too high, and future assignments are done to automation, an increase in workload for the human controller is still possible. If the flight now assigned to the automated system has a large number of interactions with flights already being controlled by the human controller, the human controller may still experience an increase in workload although they were not assigned the flight. Whether such a situation would have a positive or negative influence on the perceived workload of the controller is unknown at this stage.

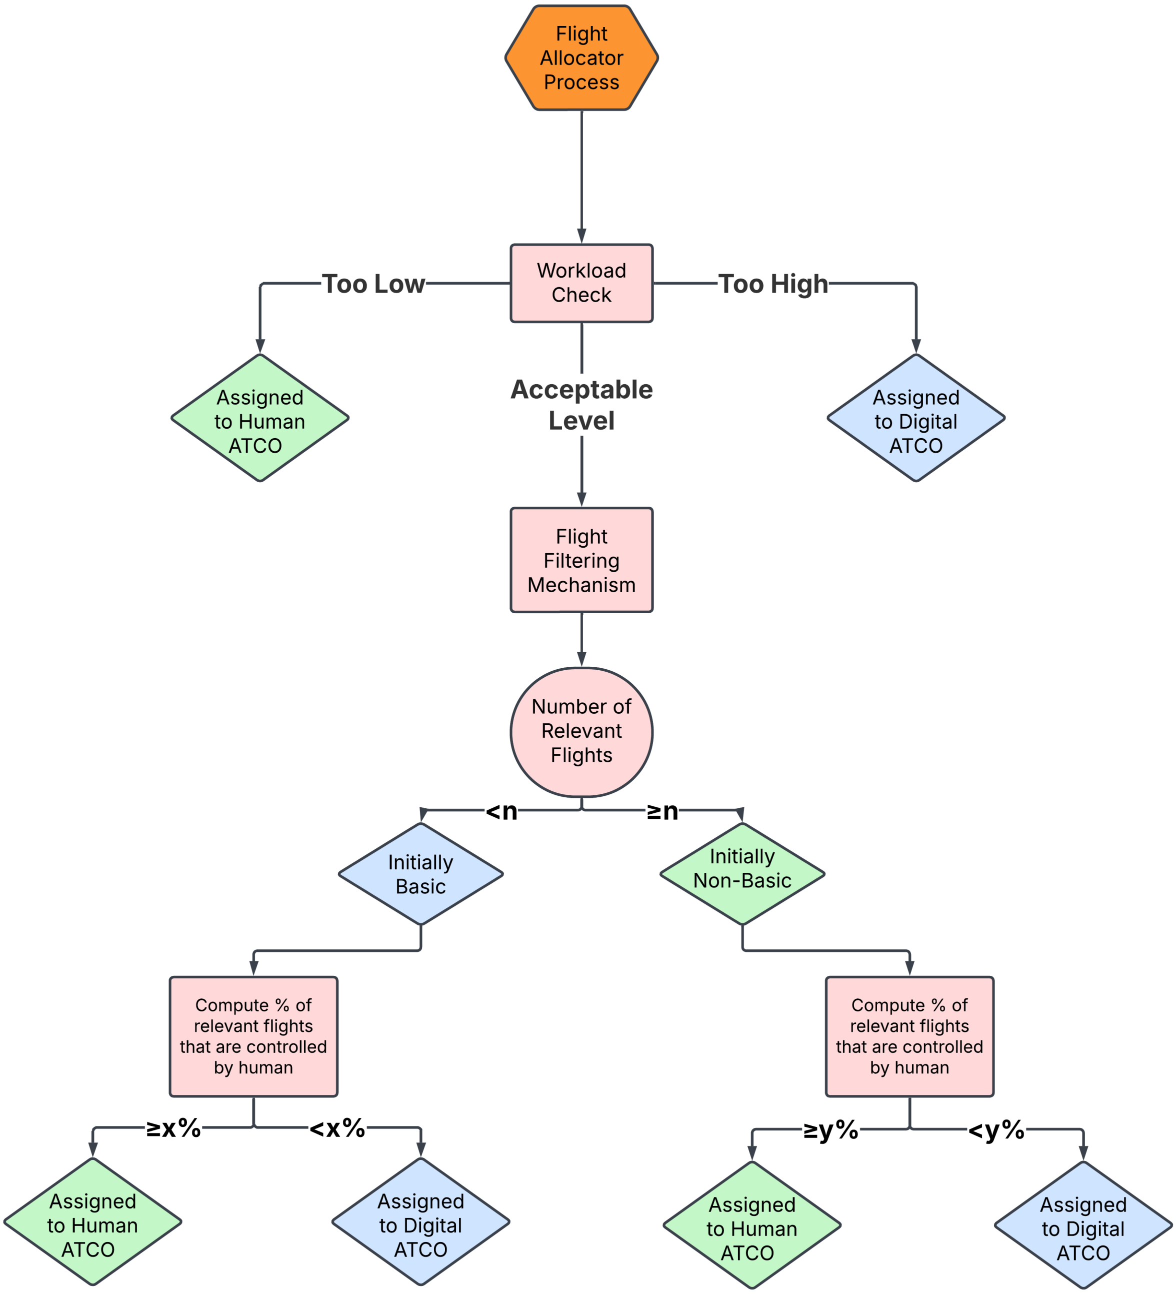

Figure 5.1: Designed flight allocator flow chart for determining controller assignments.

5.2. Further Considerations

Figure 5.1 shows the preliminary design of the flight allocator for this research project. The preliminary design raises some questions that need to be addressed as the design of the allocator is developed further. Three main questions were raised during the initial design of the flight allocator:

1. **Is the number of relevant flights consistent?**

2. **What is the limit of the digital ATCO capabilities?**

3. **What happens when the digital ATCO is no longer able to handle an aircraft?**

The first question looks at the consistency of the number of relevant flights for a flight of interest. The number of relevant flights is determined by the amount of flights that are caught by the flight filtering mechanism. The flight filtering mechanism looks at all flights currently in the sector to determine the number of relevant flights. This produces an initial value, which will vary as the aircraft travels through the sector. New aircraft entering the sector at a later point in time could have a future interaction with the flight of interest, but this is not picked up by the flight filtering mechanism. Furthermore, as conflicts pass and flights leave the sector, the number of relevant flights can also reduce. This leads to two considerations that need to be investigated.

1. Can we consider traffic that has not yet entered the sector?
2. How will new aircraft entering the sector affect the complexity of existing air traffic?

These points should be considered to determine whether the number of relevant flights is somewhat consistent as the flight of interest travels through the sector. If the number of relevant flights is found to be inconsistent and unstable, then a reevaluation of how the flight allocator uses this metric to allocate flights may be necessary. A brief initial investigation into the stability of the number of interactions over time was carried out. For this, an existing sample of traffic was simulated and five aircraft were randomly chosen to see how the number of relevant flights varied over the duration of the simulation. A screenshot of the traffic sample used with the investigated flights highlighted can be found in Figure 5.2. The results of the investigation can be found in Figure 5.3.

Although Figure 5.3 shows some fluctuations in the number of interactions over time, it is still unknown whether these fluctuations can be considered small enough such that the complexity of the flight is not drastically altered. Flight NJE418R is a high-flying overhead flight at a flight level of FL450, suggesting that the higher flying aircraft have less interactions. Flights KLM42U, KLM26V and CAI19FQ travel through the center of the sector, at lower altitude levels, representing a more ordinary flight. This comes with larger fluctuations in the number of relevant flights, likely due to the more congested traffic situation with lots of aircraft entering and leaving the center of the sector. The flight plan of UAE70M stays closer to the outside of the sector, a possible explanation for the lower number of relevant flights and fewer variations. These five aircraft represent different types of traffic we may expect to see in a sector. Further investigation should look at the average fluctuation for all flights in the sector, as well as to determine what an acceptable range of fluctuation is such that the initial allocation is still valid.

The capabilities of the digital ATCO must also be considered. In the current allocator design, allocations are done based on three metrics, none of which currently account for the capabilities of the digital ATCO. In the current iteration of the automation software, the digital ATCO is only able to solve conflicts vertically, and is unable to solve conflicts in the horizontal plane. This is a factor that needs to be considered in the allocation process. The allocator must only assign aircraft to the digital ATCO that the digital ATCO can resolve potential conflicts for. Assigning aircraft that the digital ATCO is unable to handle will result in the human ATCO having to take control, which means that the initial assignment was inappropriate.

Furthermore, a solution needs to be created for the situation in which an aircraft entering the sector causes an earlier allocation assignment to become invalid. If a flight is initially assigned to the digital ATCO, but at a later moment in time, a new flight enters the sector which has a conflict with the initial flight that can only be resolved through horizontal maneuvers, then the digital ATCO is no longer able to handle the situation. For this, the allocation of the aircraft must be switched to the human ATCO, or the human ATCO must take control of the new flight and be responsible for solving the conflict between the human and digitally controlled aircraft.

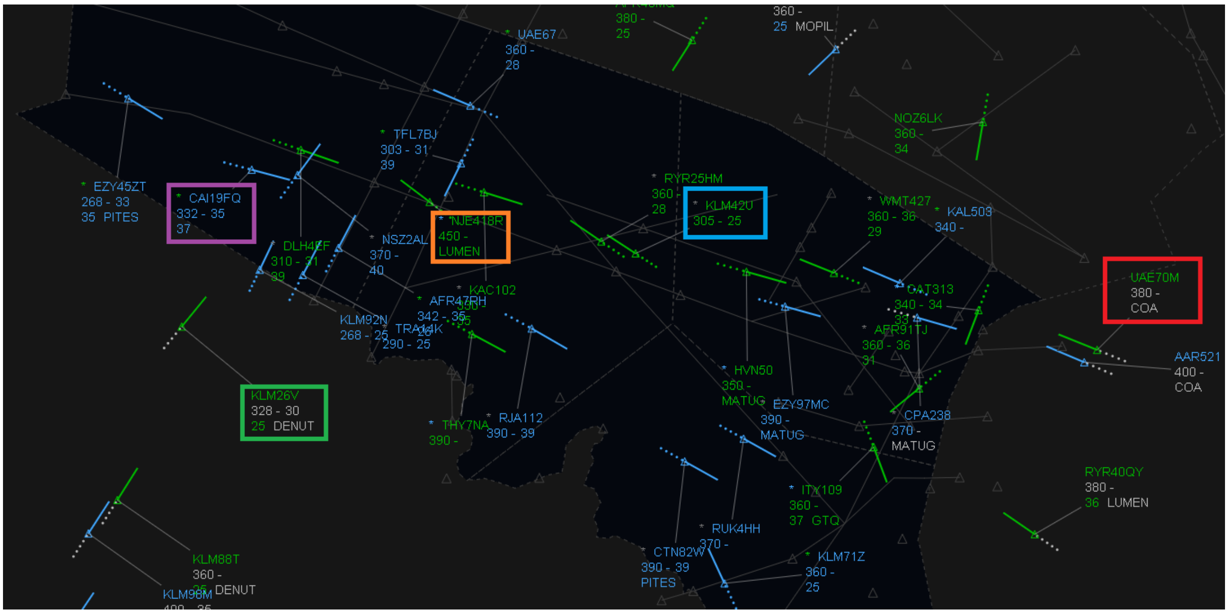

Figure 5.2: Screenshot of traffic sample used.

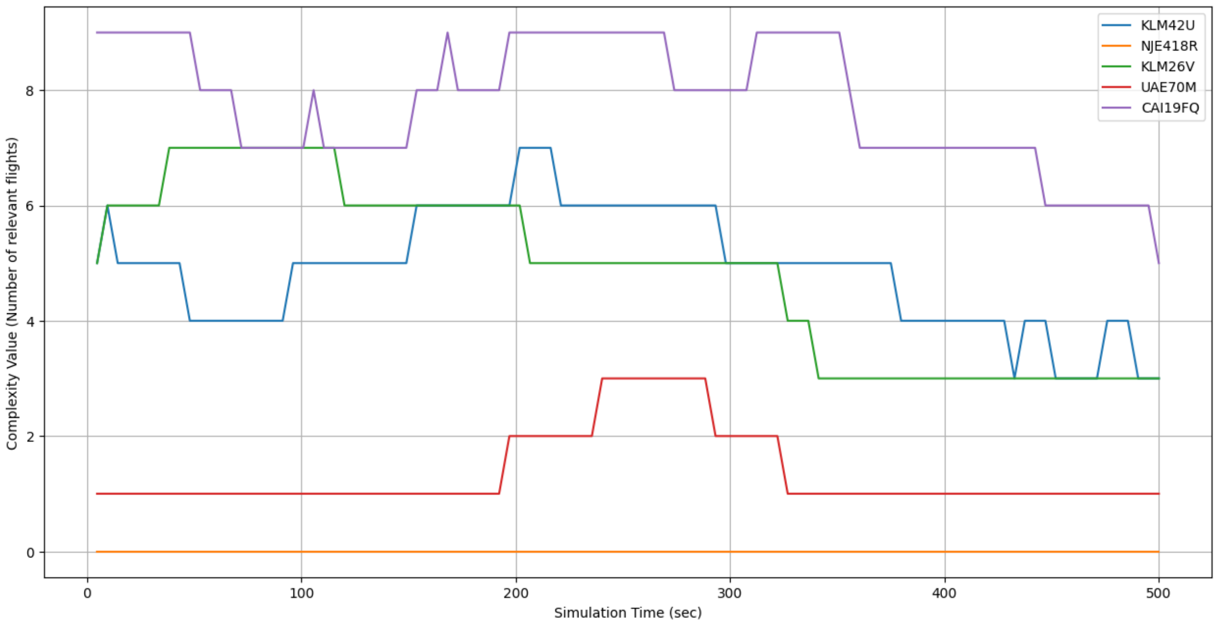

Figure 5.3: Number of relevant flights over time for five selected aircraft.

5.3. Conclusion

The three metrics proposed that make up the logic of the flight allocator are the number of relevant flights, the percentage of human-automation controller interactions and the perceived workload of the human controller. The exact threshold for all metrics will need to be determined and evaluated. The design of the flight allocator is still in the preliminary stage, and there are some questions and concerns that need to be addressed in the design process.

To evaluate the usefulness of the metrics, an experiment has to be performed with ATCOs to determine the effectiveness of the allocator and the metrics chosen. Using the data collected in the experiment, the metrics can be further fine-tuned, and the next steps to enhance the quality of the flight allocator model can be determined.

6 Preliminary Experiment Design

In the previous chapter, the design of the flight allocator model is discussed. The flight allocator will be able to assign flights to either a human controller or automation depending on metrics calculated in the flight allocator. The objective of the experiment is to evaluate the usefulness of the allocator for ATCOs and to answer the following research question:

How do the design parameters of a flight allocator model affect the traffic responsibilities of human and digital computer controllers, and how does this affect the efficiency and safety of ATC operations?

In the experiment, MUAC ATCOs will be asked to evaluate their experience with the flight allocator. Different scenarios are designed to determine the influence of flight allocator metrics on controller operations. The data and evaluations collected in the experiment will be used to evaluate the feasibility of human-automation teamwork in a sector and the usefulness of the flight allocator.

6.1. Participants and Apparatus

The experiment participants will be MUAC ATCOs. The standard radar screen used at MUAC will be mimicked with SectorX, a TU Delft built Java-based simulator. This will reduce the time required for participants to familiarize themselves with the interface. Different to real ATCO operations, all work-domain clearances will have to be issued via CPDLC, as no radar connection is present for this experiment. Participants will be equipped with a standard mouse and keyboard used to navigate the simulation interface.

6.2. Experiment Procedure and Tasks

The experiment will begin with a briefing and training phase. In the briefing phase, an introduction is given to the setup and purpose of the experiment. The logic and operations of the automated controller will be explained to ensure participants are familiar with the strategies employed by the automation. In the training phase, participants undergo training, in which the information and concepts explained in the briefing phase can be practiced. Furthermore, participants will be familiarizing themselves with the simulation software which will help them adjust to the differences between the experiment environment and their real work-domain.

Next is the experiment phase, in which ATCOs handle multiple traffic scenarios, with varying flight allocator designs. In each scenario, ATCOs are tasked with managing and overseeing the traffic in their sector. This involves solving conflicts between flight pairs alongside flights being controlled by automation. Any human-automation controller conflicts should be solved by the human controller. After the experiment has concluded, participants will be asked to fill in a questionnaire to provide more feedback on their experiences. This will be used to gain more insight on factors such as their trust in the automation, trust in the allocator assignments, the upkeep of their mental model, and strategies they employed.

6.3. Experiment Variables

In this experiment, the dependent variables will be measured through manipulation of the independent variables. This will be done while considering the control variables to alleviate possible biases.

6.3.1. Independent Variables

The independent variables are variables that are manipulated during the experiment to determine their significance by measuring the dependent variables. In this experiment, the independent variables are the design aspects of the flight allocator. Through manipulation of the flight allocator thresholds, the allocation distribution of flights to human or automation can be altered. By modifying the thresholds, the impact on the dependent variables will be measured. There are two independent variables in the flight allocator model that can be modified:

- **Threshold for Number of Relevant Flights:** The number of relevant flights for a flight of interest is calculated by the flight filtering mechanism. The flight of interest is then classified as initially basic or non-basic based on this criteria. The threshold for number of relevant flights can be altered by either increasing or decreasing the value required for a flight to be considered non-basic. The threshold will be manipulated to study the effect of this allocator parameter on the sector and human operations. There will be three levels with which the threshold will be varied. The exact placement of the thresholds is still to be determined, but the three levels will consist of a low, medium, and high threshold.

- **Threshold for Human-Automation Controller Interactions:** The percentage of relevant flights already being controlled by the human controller is calculated for a flight of interest. A threshold percentage then defines which controller will be responsible for the flight of interest. The threshold can be modified to change the allocation distribution of flights to the human controller or automation. The threshold will be manipulated to study the effect on sector and human operations. For this variable, there will be two levels, a low and a high threshold.

Participants will be dealing with only one of the independent variables. For their experimental scenarios, one of the independent variables will vary during the experiment while the other will remain fixed. When changing one of the two independent variables, the other must be controlled to make sure that the effect of the independent variable on the dependent measures is studied.

6.3.2. Dependent Measures

Dependent measures represent the variables measured during an experiment. These dependent measures are influenced by the manipulation of the independent variables. The dependent measures for this research project have been categorized and are listed below.

Efficiency:

- Average heading deviation from the flight trajectory
- Additional miles covered by flights

Safety:

- Average conflict resolution time
- Number of conflict alerts
- Duration of loss of separation
- Minimum separation between flights in the sector

Workload:

- Number of flights cleared to target waypoints
- Number of flights cleared to target altitudes
- Subjective Workload Ratings

Acceptance of allocator:

- Amount of manual overrides of allocation
- Time duration between allocation and manual override

The research will measure the dependent variables through manipulation of the independent variables. The findings from this will be used to answer the research question and draw a conclusion. In addition to the dependent variables mentioned above, a questionnaire will provide insight into the participants' experience with respect to automation trust, trust in the allocation model, and other feedback related to the experiment.

6.3.3. Control Variables

Control variables are variables that are kept constant throughout the experiment, which is done to avoid the influence of these variables on the results. The control variables for this experiment are the sector shape, the traffic sample and the automation logic. In future experiments, these control variables could become independent variables to further study the strength and versatility of the allocation algorithm in different situations.

- **Sector Shape:** The sector shape is kept constant throughout the duration of the experiment. Look-ahead time and conflict resolution strategies can be different depending on the shape of the sector. For example, in narrower sectors, look-ahead time is reduced and conflict resolution maneuvers in the horizontal plane become significantly more difficult. To prevent this from influencing the dependent measures, the sector shape is kept constant.

- **Traffic Sample:** The traffic sample is maintained constant throughout the duration of the experiment. In high-density traffic, look-ahead time is reduced compared to lower-density traffic situations. Differences in the traffic density or traffic distribution can influence the workload experienced by the controller. The traffic sample that will be used will be the same sample used by de Rooij in their experiments [7].

- **Automation Logic:** The logic employed by the automation software to solve conflicts will be kept constant throughout the whole experiment. Keeping the automation conflict resolution tactics consistent throughout allows the human controller to be better able to predict what the automation will do. Differences in automation logic can influence the traffic behavior as well as the mental model of the human controller and may alter their experienced workload.

7 Conclusion

With the growing demand for air traffic, research focuses on increasing the capacity of airspace. One of the methods being investigated is increasing levels of automation, which will change how ATC operations are carried out. Automation support can take many different forms, and emerging concepts suggest different levels of automation assistance. This research project primarily focuses on the concept of shared responsibility of handling air traffic between an ATCO and automation. Recent studies have demonstrated the feasibility of a shared responsibility concept. These studies divide traffic based on the complexity of the flight in which ”basic” traffic is to be allocated to the automation system, while having ”non-basic” traffic be allocated to the ATCO [37][38][6].

In order to share the traffic responsibilities between ATCOs and automation in a sector based on complexity, a flight allocator needs to be designed which can allocate traffic to the correct controller. To achieve this, a literature study was conducted to identify characteristics that influence flight complexity. Findings from the literature study noted that flight complexity is largely influenced by the both the amount and nature of interactions of the flight of interest with surrounding air traffic. A flight-filtering concept, which identifies air traffic that may have a possible interaction with the flight of interest was found to be useful as a basis for determining flight complexity.

From the literature study, a design for a flight allocator was produced. In this flight allocator, three different metrics are used in tandem to determine the allocation of a flight to an ATCO or automation. The metrics involved include the number of interactions, the controllers involved in an interaction, and the ATCO workload.

The next stage of the research project is to produce the flight allocator, and then perform experiments to determine whether the tool is useful, and if the design metrics chosen were appropriate. In this experiment, ATCOs will be asked to control traffic in a simulated air-traffic environment alongside automation. The flight allocator would be responsible for assigning traffic to either controller. The results of this experiment should provide insight into the feasibility of sharing traffic responsibilities in a sector through use of a flight allocator. Furthermore, insight should be gained on the impact of changing the thresholds of flight allocator metrics on the performance of the human controller and the sector operations.

References

[1] EUROCONTROL Maastricht Upper Area Control Centre. *EUROCONTROL Forecast Update 2024-2030 - Spring Update | EUROCONTROL*. Tech. rep. Feb. 2024. URL: https://www.eurocontrol.int/publication/eurocontrol-forecast-2024-2030.

[2] "European Organisation for the Safety of Air Navigation (EUROCONTROL)" (June 2023). Institution: Koninklijke Brill NV. DOI: 10.1163/1570-6664_iyb_SIM_org_39214. URL: https://referenceworks.brill.com/doi/10.1163/1570-6664_iyb_SIM_org_39214.

[3] IFATCA. *Staff Shortage Survey – EUR Region*. IFATCA. Dec. 26, 2022. URL: https://ifatca.org/staff-shortage-survey-eur-region/ (visited on 03/07/2025).

[4] SESAR. *SESAR Joint Undertaking | European ATM Master Plan*. 2025th ed. Dec. 12, 2024. 90 pp. URL: https://www.sesarju.eu/masterplan (visited on 03/20/2025).

[5] *ARGOS Factsheet | EUROCONTROL*. Mar. 2023. URL: https://www.eurocontrol.int/publication/argos-factsheet.

[6] Amber Stienstra. "Relating Air Traffic Controller Perceived Complexity to Characteristics of Individual En-Route Flights". en. *TU Delft - Aerospace Engineering* (2022). URL: https://repository.tudelft.nl/record/uuid:ee7d8281-19c9-4611-9e5f-0b4c0f0533f6 (visited on 06/27/2025).

[7] Gijs de Rooij et al. "Flight-Based Control Allocation: Towards Human–Autonomy Teaming in Air Traffic Control". *Aerospace* 11.11 (Nov. 2024). Number: 11 Publisher: Multidisciplinary Digital Publishing Institute, p. 919. DOI: 10.3390/aerospace11110919. URL: https://www.mdpi.com/2226-4310/11/11/919.

[8] David Crocker. *Dictionary of Aviation*. English. 2nd ed. London: A&C Black Publishers Ltd, 2007. URL: https://www.academia.edu/18758790/Dictionary_of_Aviation (visited on 02/18/2025).

[9] Thomas Standfuß et al. "Performance Assessment of European Air Navigation Service Providers". Sept. 2018. DOI: 10.1109/DASC.2018.8569839.

[10] EUROCONTROL. *Maastricht Upper Area Control Centre (MUAC) sectorisation chart | EUROCONTROL*. en. Jan. 2025. URL: https://www.eurocontrol.int/publication/maastricht-upper-area-control-centre-muac-sectorisation-chart (visited on 02/18/2025).

[11] INTERNATIONAL CIVIL AVIATION ORGANIZATION. *Procedures for Air Navigation Services - Air Traffic Management*. 16th ed. Doc 4444. INTERNATIONAL CIVIL AVIATION ORGANIZATION, 2016.

[12] Zhuang Wang et al. "Review of Deep Reinforcement Learning Approaches for Conflict Resolution in Air Traffic Control". en. *Aerospace* 9.6 (June 2022). Number: 6 Publisher: Multidisciplinary Digital Publishing Institute, p. 294. DOI: 10.3390/aerospace9060294. URL: https://www.mdpi.com/2226-4310/9/6/294 (visited on 02/20/2025).

[13] Åsa Svensson et al. "Design implications for teamwork in ATC". en. *Cognition, Technology & Work* 22.2 (May 2020). Company: Springer Distributor: Springer Institution: Springer Label: Springer Number: 2 Publisher: Springer London, pp. 409–426. DOI: 10.1007/s10111-019-00579-y. URL: https://link.springer.com/article/10.1007/s10111-019-00579-y (visited on 02/20/2025).

[14] Clark Borst. *E-UI design doc & demonstrator*. English. Tech. rep. D4.1. SESAR, Sept. 2021, p. 38. URL: https://mahalo.lr.tudelft.nl/wp-content/uploads/2021/10/D4.1-E-UI-design-doc-demonstrator_v02.pdf (visited on 02/21/2025).

[15] Kevin Richard et al. "Scenario-Driven Development and Testing of ATC Conflict Detection". en. *AIAA Scitech 2019 Forum*. San Diego, California: American Institute of Aeronautics and Astronautics,

Jan. 2019. DOI: 10.2514/6.2019-1481. URL: https://arc.aiaa.org/doi/10.2514/6.2019-1481 (visited on 02/25/2025).

[16] EUROCONTROL. *EUROCONTROL Specification for Monitoring Aids*. Tech. rep. 2. Mar. 2017. URL: https://www.eurocontrol.int/sites/default/files/publication/files/EUROCONTROL-SPEC-0139%20MTCD%20Ed%202.0.pdf (visited on 02/25/2025).

[17] Klaus Eyferth et al. "A model of air traffic controllers' conflict detection and conflict resolution". *Aerospace Science and Technology* 7.6 (Sept. 2003), pp. 409–416. DOI: 10.1016/S1270-9638(03)00064-6. URL: https://www.sciencedirect.com/science/article/pii/S1270963803000646 (visited on 02/25/2025).

[18] Esa M. Rantanen et al. "Hierarchical Conflict Detection in Air Traffic Control". *The International Journal of Aviation Psychology* 15.4 (Oct. 2005). Publisher: Taylor & Francis _eprint: https://doi.org/10.1207/s15327108ijap1504_3, pp. 339–362. DOI: 10.1207/s15327108ijap1504_3. URL: https://doi.org/10.1207/s15327108ijap1504_3 (visited on 02/25/2025).

[19] Oliver Späth et al. "Conflict Resolution in En Route Traffic - A Draft Concept for an Assistance System Compatible with Solutions of Air Traffic Controllers". *MMI Interaktiv* 1.05 (Aug. 2001). URL: https://dl.gi.de/items/bba1d3f2-c384-40c9-83e5-87504fb32a89.

[20] Esa M. Rantanen et al. "Conflict Resolution Maneuvers in Air Traffic Control: Investigation of Operational Data". *The International Journal of Aviation Psychology* 22.3 (July 2012). Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/10508414.2012.691048, pp. 266–281. DOI: 10.1080/10508414.2012.691048. URL: https://doi.org/10.1080/10508414.2012.691048 (visited on 02/26/2025).

[21] Selina Fothergill et al. "Conflict-Resolution Heuristics for En Route Air Traffic Management". en. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 57.1 (Sept. 2013), pp. 71–75. DOI: 10.1177/1541931213571018. URL: https://journals.sagepub.com/doi/10.1177/1541931213571018 (visited on 02/26/2025).

[22] Fitri Trapsilawati et al. "Why do Controllers Choose the Conflict Resolution Maneuvers that They Do?" *The International Journal of Aerospace Psychology* 32.1 (Jan. 2022). Publisher: Routledge _eprint: https://doi.org/10.1080/24721840.2021.1925119, pp. 24–38. DOI: 10.1080/24721840.2021.1925119. URL: https://doi.org/10.1080/24721840.2021.1925119 (visited on 02/26/2025).

[23] Ajay Kumbhar et al. "Determining Flight Complexity and Relevance: Flight-Centric Filtering for Air Traffic Control: 14th SESAR Innovation Days, SIDS 2024". *14th SESAR Innovation Days, SIDS 2024* (2024), pp. 1–8. URL: https://www.sesarju.eu/sites/default/files/documents/sid/2024/papers/SIDs_2024_paper_067%20final.pdf (visited on 06/27/2025).

[24] Yash Guleria et al. "Towards conformal automation in air traffic control: Learning conflict resolution strategies through behavior cloning". *Advanced Engineering Informatics* 59 (Jan. 2024), p. 102273. DOI: 10.1016/j.aei.2023.102273. URL: https://www.sciencedirect.com/science/article/pii/S1474034623004019 (visited on 02/26/2025).

[25] Angela R. Schmitt et al. "Balancing controller workload within a sectorless ATM concept". *CEAS Aeronautical Journal* 2.1-4 (Dec. 2011), pp. 35–41. DOI: 10.1007/s13272-011-0031-7. URL: http://link.springer.com/10.1007/s13272-011-0031-7.

[26] Tobias Finck et al. "Conceptual Analysis of Allocation Strategies for Air Traffic Control Concepts without Conventional Sector Boundaries". *2023 Integrated Communication, Navigation and Surveillance Conference (ICNS)*. ISSN: 2155-4951. Apr. 2023, pp. 1–11. DOI: 10.1109/ICNS58246.2023.10124289. URL: https://ieeexplore.ieee.org/document/10124289 (visited on 11/11/2024).

[27] Bettina Birkmeier et al. "Sectorless ATM and advanced SESAR concepts: Complement not contradiction". *29th Digital Avionics Systems Conference*. ISSN: 2155-7209. Oct. 2010, pp. 2.D.5–1–2.D.5–8. DOI: 10.1109/DASC.2010.5655475. URL: https://ieeexplore.ieee.org/document/5655475/?arnumber=5655475 (visited on 02/27/2025).

[28] Bettina Birkmeier et al. "Controller team possibilities for sectorless air traffic management". en. *2016 Integrated Communications Navigation and Surveillance (ICNS)*. Herndon, VA, USA: IEEE, Apr.

2016, pp. 6C3–1–6C3–10. DOI: 10.1109/ICNSURV.2016.7486362. URL: http://ieeexplore.ieee.org/document/7486362/ (visited on 10/22/2024).

[29] EUROCONTROL Maastricht Upper Area Control Centre. *EUROCONTROL Maastricht Upper Area Control Centre OPS and Automation Strategy*. Tech. rep. 2019. URL: https://skybrary.aero/%20bookshelf/books/5341.pdf.

[30] Lisanne Bainbridge. “Ironies of automation”. *Automatica* 19.6 (Nov. 1, 1983), pp. 775–779. DOI: 10.1016/0005-1098(83)90046-8. URL: https://www.sciencedirect.com/science/article/pii/0005109883900468 (visited on 01/17/2025).

[31] N. H. Mackworth. *Researches on the measurement of human performance. (Med. Res. Council, Special Rep. Ser. No. 268.)* Researches on the measurement of human performance. (Med. Res. Council, Special Rep. Ser. No. 268.) Pages: 156. Oxford, England: His Majesty”s Stationery Office, 1950. 156 pp.

[32] Mica Endsley et al. “The Out-of-the-Loop Performance Problem and Level of Control in Automation”. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 37 (June 1, 1995). DOI: 10.1518/001872095779064555.

[33] Raja Parasuraman et al. “Performance consequences of automation-induced ’complacency”’. *International Journal of Aviation Psychology* 3.1 (Jan. 1, 1993). NTRS Author Affiliations: NASA Headquarters, Catholic Univ. NTRS Report/Patent Number: ISSN: 1050-8414 NTRS Document ID: 19930055574 NTRS Research Center: Legacy CDMS (CDMS). URL: https://ntrs.nasa.gov/citations/19930055574 (visited on 01/23/2025).

[34] Hugh P. Bergeron. “Single Pilot IFR Autopilot Complexity/Benefit Tradeoff Study”. *Journal of Aircraft* 18.9 (1981). Publisher: American Institute of Aeronautics and Astronautics _eprint: https://doi.org/10.2514/3.57549, pp. 705–706. DOI: 10.2514/3.57549. URL: https://doi.org/10.2514/3.57549 (visited on 01/23/2025).

[35] Mica R. Endsley. “From Here to Autonomy”. *Human Factors* (Dec. 15, 2016). Publisher: SAGE PublicationsSage CA: Los Angeles, CA. DOI: 10.1177/0018720816681350. URL: https://journals.sagepub.com/doi/full/10.1177/0018720816681350 (visited on 01/23/2025).

[36] Joseph B. Lyons et al. “Human–Autonomy Teaming: Definitions, Debates, and Directions”. *Frontiers in Psychology* 12 (May 28, 2021). Publisher: Frontiers. DOI: 10.3389/fpsyg.2021.589585. URL: https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2021.589585/full (visited on 03/20/2025).

[37] Thomas Prevot et al. “Toward Automated Air Traffic Control—Investigating a Fundamental Paradigm Shift in Human/Systems Interaction”. *International Journal of Human-Computer Interaction* 28.2 (Feb. 2012), pp. 77–98. DOI: 10.1080/10447318.2012.634756. URL: http://www.tandfonline.com/doi/abs/10.1080/10447318.2012.634756 (visited on 10/29/2024).

[38] de Rooij et al. “Flight Allocation in Shared Human-Automation En-Route Air Traffic Control”. 2021. URL: https://www.semanticscholar.org/paper/Flight-Allocation-in-Shared-Human-Automation-Air-Rooij-Borst/6bc3a64d1528a65af74fc0b866b97dd29272a1e4 (visited on 01/23/2025).

[39] Bruno Antulov-Fantulin et al. “Determining Air Traffic Complexity – Challenges and Future Development”. *Promet - Traffic&Transportation* 32 (July 2020), pp. 475–485. DOI: 10.7307/ptt.v32i4.3401.

[40] Brian Hilburn. “Cognitive complexity in air traffic control: a literature review”. *Tech. rep. EUROCONTROL* (Jan. 2004). URL: https://www.eurocontrol.int/node/9861.

[41] ANDREW J. TATTERSALL et al. “An experimental evaluation of instantaneous self-assessment as a measure of workload”. *Ergonomics* 39.5 (May 1996). Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00140139608964495, pp. 740–748. DOI: 10.1080/00140139608964495. URL: https://doi.org/10.1080/00140139608964495.

[42] Sehchang Hah et al. “The Effect of Air Traffic Increase on Controller Workload”. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50 (Oct. 1, 2006), pp. 50–54. DOI: 10.1177/154193120605000111.

[43] Haksa Ewanda Alvin et al. “The Effect Of Workload On Air Traffic Controller Performance At Perum LPPNPI Kupang Branch”. *International Journal of Progressive Sciences and Technologies* 35.1 (Oct. 2022).

[44] Maria Prandini et al. “Toward Air Traffic Complexity Assessment in New Generation Air Traffic Management Systems”. *IEEE Transactions on Intelligent Transportation Systems* 12.3 (Sept. 2011), pp. 809–818. DOI: 10.1109/TITS.2011.2113175. URL: http://ieeexplore.ieee.org/document/5723748/ (visited on 10/22/2024).

[45] Richard Mogford et al. “The Complexity Construct in Air Traffic Control: A Review and Synthesis of the Literature”. *Storming Media* (July 1, 1995), p. 32.

[46] Gijs de Rooij et al. “Contributing Factors to Flight-Centric Complexity in En-Route Air Traffic Control”. en. *Proceedings of the 15th USA/Europe Air Traffic Management Research and Development Seminar (ATM2023)*. Savannah, GA, USA, June 2023. URL: https://www.researchgate.net/publication/371991465_Contributing_Factors_to_Flight-Centric_Complexity_in_En-Route_Air_Traffic_Control.

[47] A.K.Vijay Kumbhar. “Guiding Visual Attention to Relevant Flights in Supporting Air Traffic Controller Decision Making”. en. *TU Delft - Aerospace Engineering* (2024). URL: https://repository.tudelft.nl/record/uuid:b6f6e40b-c61e-4718-ae0b-baf731ad9809 (visited on 06/27/2025).

A

Informed Consent Form

Designing a complexity-based flight allocator for a shared human-automation air traffic control environment

You are being invited to participate in research titled “Designing a complexity-based flight allocator for a shared human-automation air traffic control environment.” This study is being done by Thijs Verkade from the TU Delft under the supervision of Dr. Ir. Clark Borst.

The purpose of this research study is to investigate different design models for automatic flight allocators in a shared traffic environment. The research study aims to investigate the impact of these different flight allocators on the performance, safety and efficiency of the sector as well as the workload experienced by the Air Traffic Controller. The experiment will take you approximately 110 minutes to complete. The experiment is structured as follows:

Experiment Procedure - 110 Minutes					
30 minutes	Briefing and Training				
	25 minutes	Experiment 1			
		10 minutes			
			Questionnaire 1		
			25 minutes	Experiment 2	
				20 minutes	Questionnaire 2

The data will be used to analyze the usefulness of different flight allocator design metrics. The results will be published as part of an Aerospace Engineering Master Thesis. Your task will be to expedite the traffic in the Upper Control Area (UTA) in the simulation environment called SectorX. Two experiments will be conducted, after each a short questionnaire will be provided which you are kindly asked to take part in. A final questionnaire at the end will ask you to compare the two experiment scenarios.

Your participation in this study is entirely voluntary, and you can withdraw at any time. You are free to refrain from answering any questions. Since a copy of this informed consent form will be stored by the responsible researcher, the participant has the right to be forgotten.

Thank you for your participation in this experiment. If you have any questions or concerns, do not hesitate to contact us.

Contact Information

Researcher

Thijs Verkade
m.m.verkade@student.tudelft.nl
+31 6 27227785

Contact Information

Supervisor

Dr. Ir. Clark Borst
c.borst@tudelft.nl
+31 15 2789099

Consent Form

PLEASE TICK THE APPROPRIATE BOXES		
A: GENERAL AGREEMENT – RESEARCH GOALS, PARTICPANT TASKS AND VOLUNTARY PARTICIPATION	**Yes**	**No**
1. I have read and understood the study information dated [*xx/11/25*], or it has been read to me. I have been able to ask questions about the study and my questions have been answered to my satisfaction.	☐	☐
2. I consent voluntarily to be a participant in this study and understand that I can refuse to answer questions and I can withdraw from the study at any time, without having to give a reason.	☐	☐
3. I understand that taking part in the study involves having simulation data automatically stored in an anonymous manner when completing the experiment.	☐	☐
4. I understand that taking part in the study involves me answering questions in questionnaires.	☐	☐
5. I understand that taking part in the study involves taking notes about the things I do or say during the experiment.	☐	☐
B: Use of information in the study	**Yes**	**No**
10. I understand that personal information collected about me that can identify me, such as name and signature, will not be shared beyond the study team.	☐	☐
11. I agree that my simulation data can be quoted in research outputs on an anonymous basis.	☐	☐
12. I agree that my responses, views or other input can be quoted anonymously in research outputs.	☐	☐
11. I understand that information I provide will be used for analysis and scientific publications on an anonymous basis.	☐	☐

C: Future use and reuse of information by others	Yes	No
12. I give permission for the recorded simulation data and answers to questionnaires that I provide, to be archived in secure folders, so that it can be used for future research and learning. All data is stored anonymously. Access is safeguarded and not to be used for commercial use.	☐	☐

Signatures

____________________________ ___________________________ _________

Name of participant Signature Date

I, as researcher, have accurately read out the information sheet to the potential participant and, to the best of my ability, ensured that the participant understands to what they are freely consenting.

__________________________ ____________________ _________

Name of researcher Signature Date

Study contact details for further information:
Thijs Verkade

m.m.verkade@student.tudelft.nl

+31 6 27227785

B

Experiment Briefing

Experiment Briefing

Analysis of Flight Allocation Methods in a Shared Airspace

Responsible Researchers: Thijs Verkade (m.m.verkade@student.tudelft.nl) and Clark Borst (C.Borst@tudelft.nl)

Introduction

Dear participant,

Thank you very much for taking the time to participate in this research for my master Thesis! With this simulation experiment, we would like to gain insight into a future concept of air traffic control operations. In this experiment flights in your airspace will be delegated to either the computer or to you based on different allocation schemes. Generally, the 'basic' flights in your airspace are allocated to the computer, leaving you with the more interesting and complex flights. The purpose of this experiment is to gain insight into controller experiences with the different allocation schemes.

The duration of the experiment is approximately 1.5 hours.

Kind regards,

Thijs Verkade

Master Student in Aerospace Engineering

Experiment set-up and procedure

Simulator

The simulator that we use in this experiment is developed by the TU Delft. It has been designed to mimic the MUAC interface. You will have access to VERA to assess potential conflicts. For the sake of the experiment, there will be several differences that we will let you experience at the start of the experiment. Most importantly, you only need to input clearances in the system (simulating CPDLC) as there is no voice R/T.

Scenarios

You will be presented with two scenarios of ca. 25 minutes each, based on a radar sample. You will be responsible for the entire Brussels sector group. Approximately half of the flights in this airspace are allocated to the computer, meaning that you have no control over them. These 'automated flights' are coloured blue. The flights that have been allocated to you are coloured green. In either scenario, the allocation of incoming flights is based on different rules, which you will be asked to evaluate after the experiment. The exact allocation scheme will not be revealed until after the experiment.

Your Task

It is your task to issue any clearances that you deem necessary to ensure that all flights can safely reach their XCOP at the correct TFL. You can give altitude, heading or route/direct clearances. Speed is at pilot discretion and cannot be modified by you. You can preview the planned trajectories and use VERA to assess potential conflicts. Blue flights will automatically be assumed, controlled and transferred by the computer.

After each scenario, you will be asked to briefly evaluate your experience with managing traffic in shared airspace, as well as evaluating the factors you think contributed to allocation. Following both scenarios, you will be asked to compare the two scenarios and their allocation schemes.

Automation

The computer will climb blue flights as early as possible, descend them as late as possible and send them on direct routes to the XCOP when able. It has a lookahead time of 8 minutes and will take care of automatically preventing and/or solving conflicts between blue flights in this time frame. It will not steer blue flights into conflict with green flights, but if such conflicts do occur (after 8 minutes, or because you modified a green flight), the computer will not solve them. You are responsible for solving these ‘mixed’ conflicts, by issuing a clearance to the green flight. Automation will notify you when such conflicts are predicted, 8 minutes before the LOS, by marking the conflicting pair with VERA.

Simulator Features

Manual Reallocation

You may choose to reallocate flights at any time should you wish. The allocator can provide possible candidates for reallocation based on the allocator metrics. A potential desirable candidate for reallocation is indicated by a star next to the callsign. An example of this can be found in Figure 1. If you do not wish to reallocate flights, you may ignore any stars next to aircraft callsigns.

Flight Filtering

You can make use of the flight filtering model present in the simulation software to display flights that may be relevant to a particular flight that you are interested in. The flight filtering algorithm checks for any altitude overlap, spatial overlap (checking for common waypoints or crossing paths) and temporal overlap (estimating the position of flights in time) between two flights. In the case that such an overlap is present, the flight is considered relevant and is labeled as a possible interaction. When clicking down on the cleared waypoint of a flight, all other background flights that are not relevant will be faded. A visualization of this can be seen in Figures 2 and 3. Figure 3 shows the effect of using the flight filtering.

Figure 1:
Simulator Interface

Figure 2:
Example of flight filtering (Before using filter)

CAI19FQ is the flight of interest

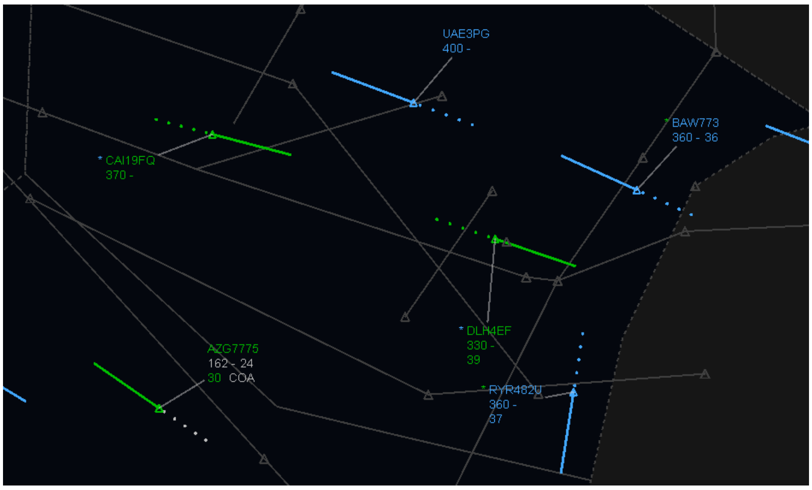

Figure 3:
Example of flight filtering (After using filter)

CAI19FQ is the flight of interest

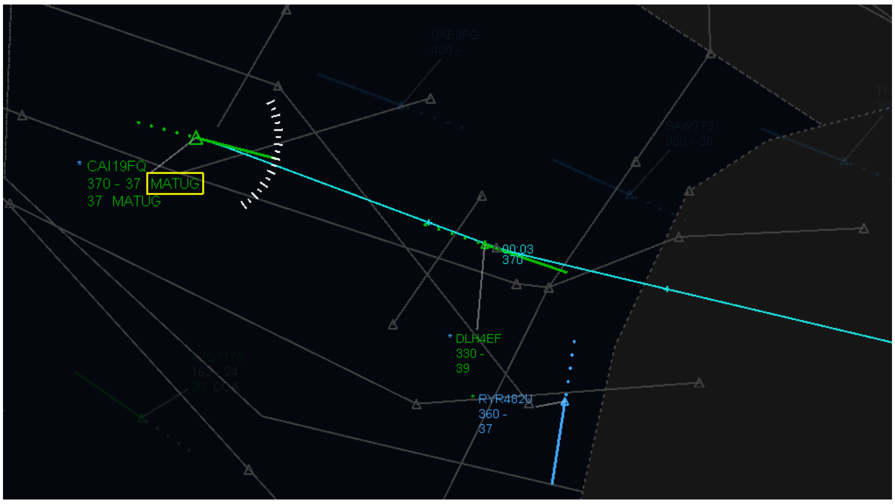

Your Rights

Participating in the experiment is on voluntary basis. You may cancel your participation at any moment for any reason, even during the experiment. This has no consequences for you. Your personal performance is not part of the study, there is no wrong or right, and there is no competition between you and your colleagues.

Data that are collected during the experiment will be anonymised before they are stored by assigning a random identifier to you. Only the principal researcher (Thijs Verkade) and his daily supervisor (Clark Borst) can link the data to you personally. This relation will never be made available to others (neither to MUAC).

By participating, you give consent to publishing the data in anonymised form. At the start of the experiment you will be asked to sign a consent form to make sure that you have read and understood what participating in the experiment means.

C
Post Scenario Questionnaire

Post-Scenario Questionnaire

* This form will record your name, please fill your name.

1. Flight Allocation Algorithms:
In this experiment you experienced a flight allocation algorithm that decided which flights should go to the computer and which should be assigned to you. Please give an indication of how you rate your experience with the allocation algorithm.

	Dislike a great deal	Dislike somewhat	Neutral	Like somewhat	Like a great deal
Flight allocation algorithm	■	■	■	■	■

2. Which factors did you think contributed to the allocation decisions made by the flight allocator?
(Select as many as you think were relevant)

- ■ Amount of potential interactions with other aircraft
- ■ Minimizing amount of interactions between human and computer controlled flights
- ■ Ensuring a minimum workload for human controller
- ■ Ensuring a minimum workload for computer controller
- ■ I am not sure what factors contributed to flights being allocated to either me or the computer

3. Please indicate how much you agree with the following statements:

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
Automation and I were a team	■	■	■	■	■
Automation was helpful	■	■	■	■	■
I could solve conflicts efficiently	■	■	■	■	■
Automation and I kept the sector safe	■	■	■	■	■
Automation and I managed the traffic efficiently	■	■	■	■	■
I could organize my work the way I wanted	■	■	■	■	■
My workload was consistent throughout	■	■	■	■	■
The allocation of flights was logical	■	■	■	■	■
I trust the allocator to make good assignments	■	■	■	■	■
Automation was able to handle all their assignments	■	■	■	■	■

4. Please give an indication of your overall workload

	Extremely low	Slightly low	Acceptable	Slightly high	Extremely high
Experienced workload	■	■	■	■	■

5. If you manually re-allocated any flights, could you explain why you chose to re-allocate these flights?

6. Were there any things that you particularly liked or disliked about the flight allocation in this scenario?

7. Is there anything else that you would like to share about the allocation algorithm?

D

Post Experiment Questionnaire

Post-Experiment Questionnaire

* This form will record your name, please fill your name.

1. Please indicate which algorithm you preferred according to the following statements.
If your experience was the same in both scenarios you can choose neutral.

	Strongly prefer first algorithm	Slightly prefer first algorithm	Neutral	Slightly prefer second algorithm	Strongly prefer second algorithm
Automation and I were a team	■	■	■	■	■
Automation was helpful	■	■	■	■	■
I could solve conflicts efficiently	■	■	■	■	■
Automation and I kept the sector safe	■	■	■	■	■
Automation and I managed the traffic efficiently	■	■	■	■	■
I could organize my work the way I wanted	■	■	■	■	■
Consistent workload	■	■	■	■	■
The allocation of flights was logical	■	■	■	■	■
I trust the allocator to make good assignments	■	■	■	■	■
Automation was able to handle all their assignments	■	■	■	■	■
Lower workload	■	■	■	■	■
I had to solve more mixed conflicts	■	■	■	■	■
Flights allocated to me were more complex to manage	■	■	■	■	■

2. If you were to design an automatic algorithm for allocating flights to either a computer or a human ATCO, what factors would you consider when determining allocation?
These can be factors used in the experiment or completely new factors.

3. Automation

The following questions are related to the functioning of the (automated) computer controlled flights, which was independent of the flight allocation algorithm.

	None at all	A little	A moderate amount	A lot	A great deal
How much trust did you have in the automation?	■	■	■	■	■

4. Please indicate to what extent you agree or disagree with the following statements about the automation.

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
It was reliable	■	■	■	■	■
It was predictable	■	■	■	■	■
It was consistent	■	■	■	■	■

5. Is there anything else that you would like to share about the automation?

6. Simulator Fidelity

How would you rate the realism of the following simulation aspects?

	Very Unrealistic	Slightly unrealistic	Acceptable	Very realistic	Perfect
Interface (look and feel)	■	■	■	■	■
Aircraft behaviour (vertical rates, turns)	■	■	■	■	■
Traffic scenario (density, routes)	■	■	■	■	■

7. Do you have any comments about the simulation?

8. Any closing thoughts about the experiment?

E

Questionnaire Responses

Following are the responses of the participants while answering the subjective responses and open-feedback questions related to the allocation algorithm. P1 first carried out the experiment with Allocator A, and then the experiment with Allocator B. P2 first carried out the experiment with Allocator B, followed by the experiment with Allocator A.

E.1. Allocator A

1. Please give an indication of how you rate your experience with the allocation algorithm.

Dislike a great deal	Dislike somewhat	Neutral	Like somewhat	Like a great deal
	P1		P2	

2. Which factors did you think contributed to the allocation decisions made by the flight allocator? (Select as many as you think were relevant)

Identified factor	Participant(s)
Amount of potential interactions with other aircraft	P2
Minimizing amount of interactions between human- and computer-controlled flights	
Ensuring a minimum workload for the human controller	
Ensuring a minimum workload for the computer controller	
I am not sure what factors contributed to flights being allocated to either me or the computer	P1
Other	

3. Please indicate how much you agree with the following statements:

Table E.1: Participant responses to Likert-scale statements regarding interaction with automation

Statement	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
Automation & I were a team		P1, P2			
Automation was helpful		P1		P2	
I could solve conflicts efficiently		P1	P2		
Automation & I kept the sector safe		P1		P2	
Automation & I managed the traffic efficiently		P2	P1		
I could organize my work the way I wanted	P1			P2	
My workload was consistent throughout			P2	P1	
The allocation of flights was logical		P1	P2		
I trust the allocator to make good assignments	P1, P2				
Automation was able to handle all of their assignments	P1, P2				

4. **Please give an indication of your experienced workload?**

Extremely low	Slightly low	Acceptable	Slightly high	Extremely high
	P2		P1	

5. **If you manually re-allocated any flights, could you explain why you chose to re-allocate these flights?**

P1: I reallocated flights to auto after solving the complex part of their flight path. When the system randomly reallocated some assumed flights to auto I also reallocated the associated conflicting flights to auto to find out how the system would deal with it.
P2: Once I had solved the conflicts of certain flights and they became low workload, i would transfer them to the automation (departures reaching cruise level).

6. **Were there any things that you particularly liked or disliked about the flight allocation in this scenario?**

P1: Disliked: working on my own flights which suddenly got reallocated to auto. System not recognizing that a flight SHALL be reassigned to me to ensure meeting exit conditions. Not being able to visualize how the auto flights will be handled laterally and vertically (with the flight leg for instance).
P2: The allocation seemed more consistent than in the first run.

7. **Is there anything else that you would like to share about the allocation algorithm?**

P1: Slightly too busy exercise to be able to fully understand the logic behind the allocation algorithm.

E.2. Allocator B

1. Please give an indication of how you rate your experience with the allocation algorithm.

Dislike a great deal	Dislike somewhat	Neutral	Like somewhat	Like a great deal
	P1,P2			

2. Which factors did you think contributed to the allocation decisions made by the flight allocator? (Select as many as you think were relevant)

Identified factor	Participant(s)
Amount of potential interactions with other aircraft	
Minimizing amount of interactions between human- and computer-controlled flights	P2
Ensuring a minimum workload for the human controller	
Ensuring a minimum workload for the computer controller	
I am not sure what factors contributed to flights being allocated to either me or the computer	P1
Other	

3. Please indicate how much you agree with the following statements:

Statement	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
Automation & I were a team	P2	P1			
Automation was helpful		P1		P2	
I could solve conflicts efficiently		P1		P2	
Automation & I kept the sector safe		P1		P2	
Automation & I managed the traffic efficiently		P1		P2	
I could organize my work the way I wanted	P1				P2
My workload was consistent throughout		P2	P1		
The allocation of flights was logical		P1,P2			
I trust the allocator to make good assignments	P1, P2				
Automation was able to handle all of their assignments	P2			P1	

4. Please give an indication of your experienced workload?

Extremely low	Slightly low	Acceptable	Slightly high	Extremely high
		P1,P2		

5. If you manually re-allocated any flights, could you explain why you chose to re-allocate these flights?

P1: Tried to be a bit more proactive than in the first exercice to simulate a more realistic situation. The

suggestions made by the system to reallocate seemed to make more sense.
P2: Transfer low-workload flights to automation , assume flights that were relevant for each other (succeeding departures).

6. Were there any things that you particularly liked or disliked about the flight allocation in this scenario?
P1: Disliked: working on my own flights which suddenly got reallocated to auto. System not recognising that a flight SHALL be reassigned to me to ensure meeting exit conditions. Not being able to visualise how the auto flights will be handled laterally and vertically (with the flight leg for instance).
P2: It took some of the difficult flights.

7. Is there anything else that you would like to share about the allocation algorithm?
P1: Slightly too busy exercise to be able to fully understand the logic behind the allocation algorithm.

E.3. Comparison Questionnaire

1. Please indicate which algorithm you preferred according to the following statements.

Statement	Strongly prefer Al. A	Slightly prefer Al. A	Neutral	Slightly prefer Al. B	Strongly prefer Al. B
Automation and I were a team			P1,P2		
Automation was helpful			P1,P2		
I could solve conflicts efficiently			P1,P2		
Automation and I kept the sector safe		P2	P1		
Automation and I managed the traffic efficiently			P1	P2	
I could organize my work the way I wanted			P1	P2	
My workload was consistent throughout		P2	P1		
The allocation of flights was logical		P2	P1		
I trust the allocator to make good assignments	P2		P1		
Automation was able to handle all their assignments			P1,P2		
Lower Workload		P2		P1	
I had to solve more mixed conflicts		P2	P1		
Flights allocated to me were more complex to manage		P2	P1		

2. If you were to design an automatic algorithm for allocating flights to either a computer or a human ATCO, what factors would you consider when determining allocation? These can be factors used in the experiment or completely new factors.
P1: Complexity (a simple conflict not being complex enough in my opinion to be allocated to the human but for that you need the ability to issue HDGs). No auto reallocation without the human opinion. Correct auto-VERA to indicate conflicts. A clear ”I can't do it anymore” warning from the computer.
P2: The amount of interactions with an aircraft , paired with the complexity of those interactions. The priority setting of certain tasks (solve conflicts, meet inbound restrictions, climb departures). The capabilities of the automation, which solutions is it able to apply (can it solve a level conflict with headings, or will it require to change level?). Are the flights standard vs non-standard flights or situations?

3. How much did you use the flight filtering?

Not at all	A little	A moderate amount	A lot	A great deal
	P1			P2

4. Is there anything else that you would like to share about the flight filtering?

P1: Indication of minimum distance on each pair would be a plus. Small buffers as well (8NM laterally) for instance. Some pairs were hidden from the filtering while they were clear potential conflicts while others seemed irrelevant which made me stop using it after a few minutes.

P2: Its a very nice and useful tool, but some of the highlights did not make sense (tens of miles in trail and still highlighted) . Generally, a net benefit to have.

5. How much trust did you have in the automation?

None at all	A little	A moderate amount	A lot	A great deal
P1	P2			

6. Please indicate to what extent you agree or disagree with the following statements about the automation.

Statement	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
It was reliable	P2	P1			
It was predictable	P1			P2	
It was consistent	P1		P2		

7. Is there anything else that you would like to share about the automation?

P1: The inability to find out future automatic clearances (i.e. where the system will give what clearance) makes it difficult to work in a hybrid automated/human environment. The system's inability to tell me "I can't do it anymore, take that flight back" means monitoring automated flights for potential exit issues instead of focusing on own traffic.

P2: Since it did not know the inbound restrictions, it did not meet any of them.

8. How would you rate the realism of the following simulation aspects? Please compare the simulation to the actual MUAC working position and live traffic. Try to exclude de-tails from your judgment that might be missing but were not relevant for the current experiment.

Simulation Aspect	Very un-realistic	Slightly unrealis-tic	Accept-able	Very realistic	Perfect
Interface (look and feel)				P1, P2	
Aircraft behavior (vertical rates, turns)		P1	P2		
Traffic scenario (density, routes)			P1	P2	

9. Do you have any comments about the simulation? Were there features missing? Were there things that you really liked about the simulator? What are possible points of improvement for future experiments?

P1: Label deconfliction would help. System respecting real LOA condition (entry/exit) and not only using COP as a reference would also help.
P2: The VERA did often times provide incorrect conflict geometries or minimum distances.

F

Additional Results

This chapter shows instances of STCA and LOS infractions during the experiment. Furthermore, all raw results from the offline simulation are also included in this chapter.

F.1. Safety Violations

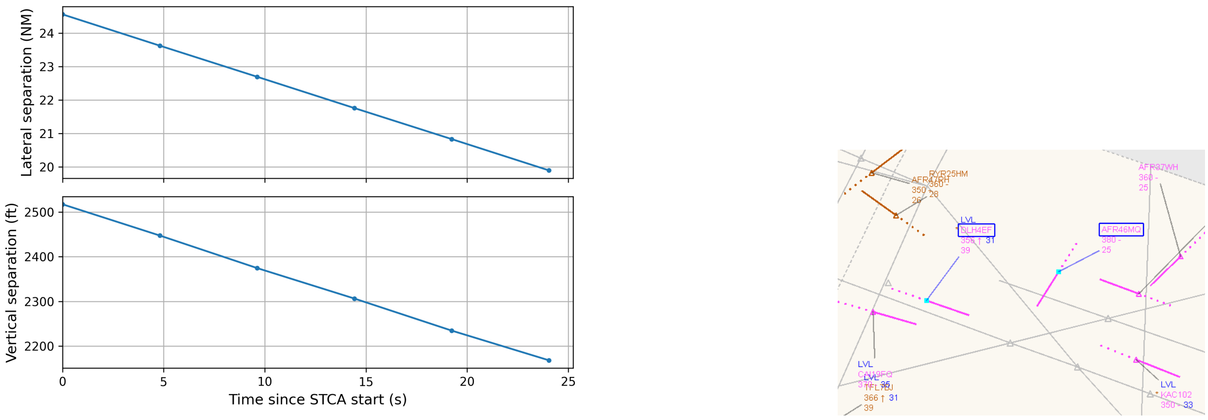

(a) Participant 2, Allocator A: STCA Duration

(b) Participant 2, Allocator A: Conflict Geometry (Inverted colors)

Figure F.1: STCA duration and conflict geometry for Participant 2 and Allocator A. Shaded red areas indicate the time period for which a LOS has occurred. Pink aircraft are controlled by the H-ATCO and the auburn aircraft are controlled by the D-ATCO.

Figure F.1b shows the conflict geometry during an STCA provided between two aircraft controlled by the H-ATCO. This STCA occurred during Participant 2's run with the complexity based allocator. The DLH4EF aircraft was set to climb to FL390 but was commanded to level off at FL360 by Participant 2 as soon as the STCA triggered. This solved the potential conflict and a LOS was avoided. Figure F.1a shows that the STCA lasted for less than 25 seconds, and was resolved quickly by Participant 2.

Figure F.2b shows the conflict geometry during an STCA provided between a mixed pair of aircraft. This is during Participant 2's experiment run with the reduced mixed interaction allocator. The aircraft are not converging, but are separated by 1,000ft, the minimum requirement for vertical separation. As the aircraft pass each other, Participant 2 issues an altitude command to NOZ6LK to lower to their exit flight level. The aircraft begins descending causing the vertical separation to be breached. However, as the aircraft are no longer converging on each other, the alert is quickly resolved. This is not a particularly dangerous situation and the descend command was issued when the aircraft were diverging so no potential conflict was possible. The duration of this STCA and the LOS was less than five seconds.

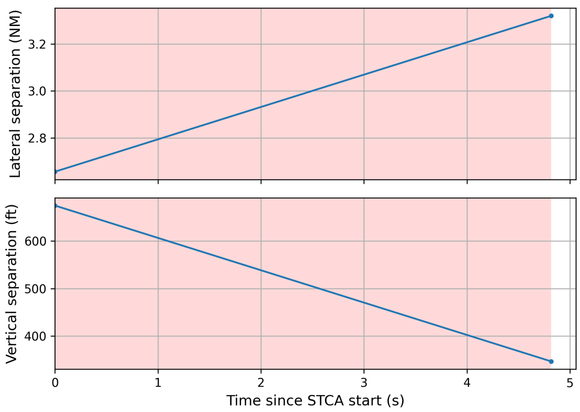

(a) Participant 2, Allocator B: STCA Duration

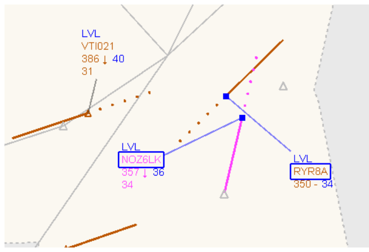

(b) Participant 2, Allocator B: Conflict Geometry (Inverted colors)

Figure F.2: STCA duration and conflict geometry for Participant 2 and Allocator B. Shaded red areas indicate the time period for which a LOS has occurred. Pink aircraft are controlled by the H-ATCO and the auburn aircraft are controlled by the D-ATCO.

F.2. Offline Simulation Results

F.2.1. Scenario 1

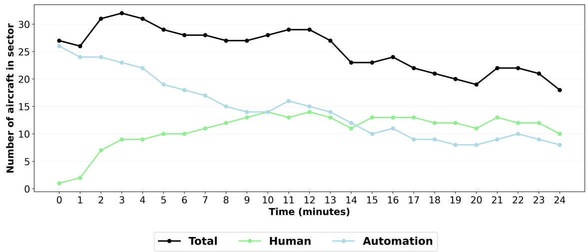

(a) Automated start.

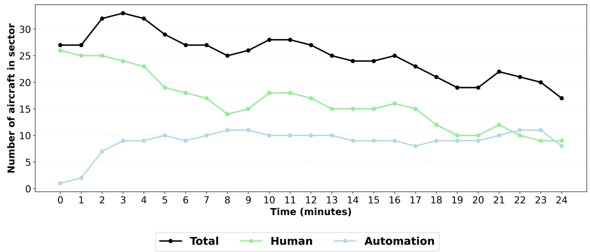

(b) Human start.

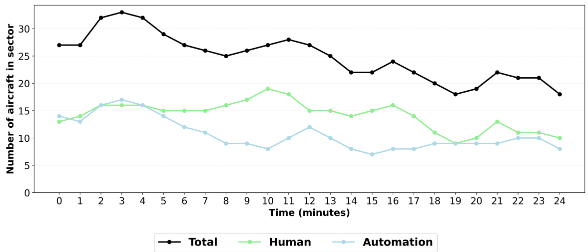

(c) No initial allocation.

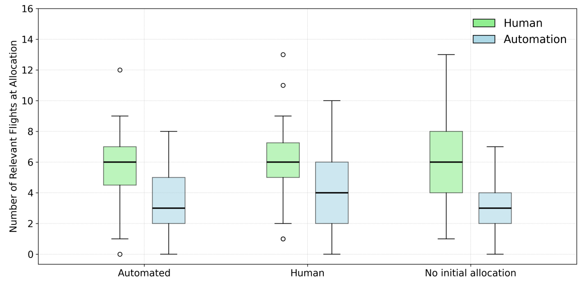

(d) Number of relevant flights at allocation.

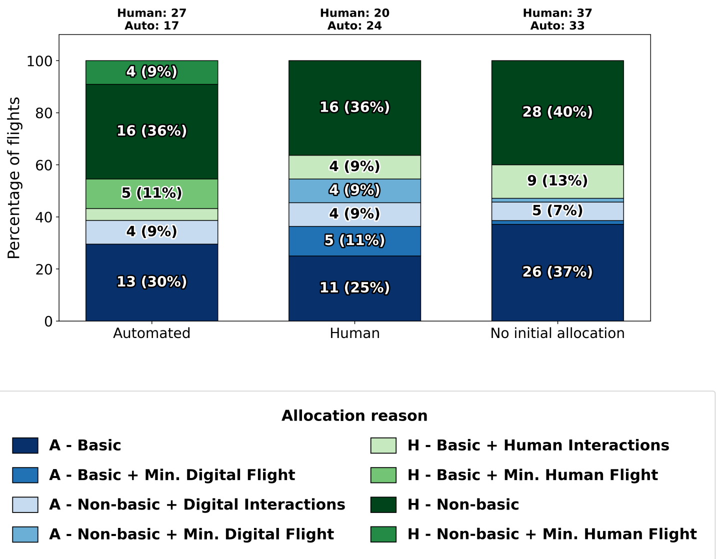

(e) Allocation distribution.

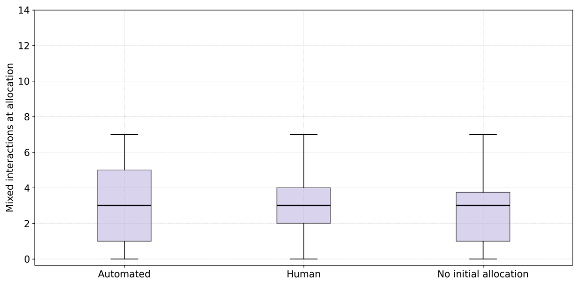

(f) Number of mixed interactions at allocation.

Figure F.3: Scenario 1 Default parameter values (Int. Threshold: 5, Min. Threshold: 25%, Mixed Int. Threshold: 70%).

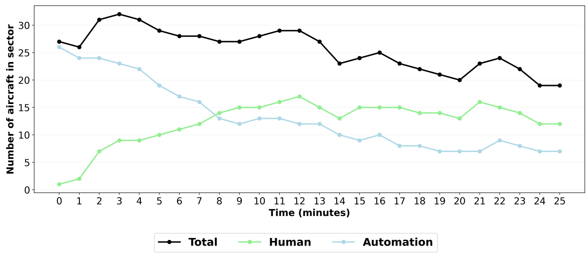

(a) Automated start.

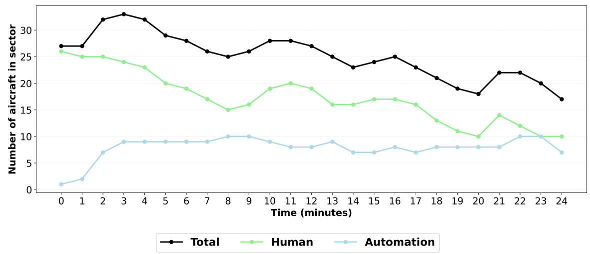

(b) Human start.

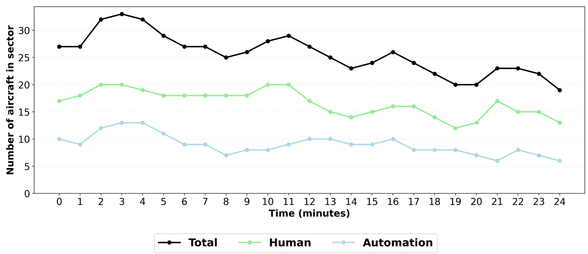

(c) No initial allocation.

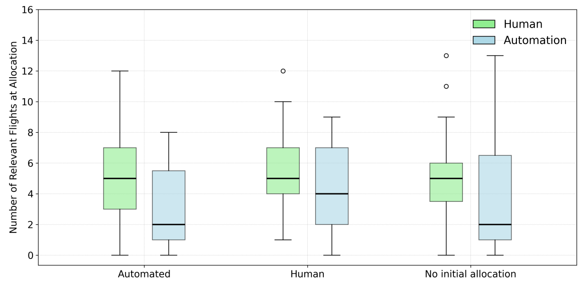

(d) Number of relevant flights at allocation.

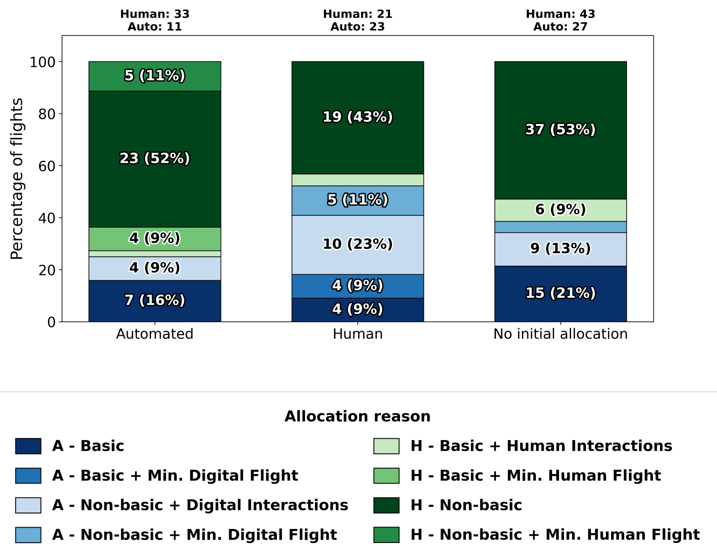

(e) Allocation distribution.

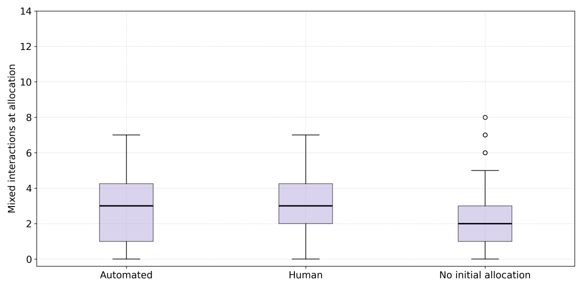

(f) Number of mixed interactions at allocation.

Figure F.4: Scenario 1 Parameter Values: (Int. Threshold: 3, Min. Threshold: 25%, Mixed Int. Threshold: 70%).

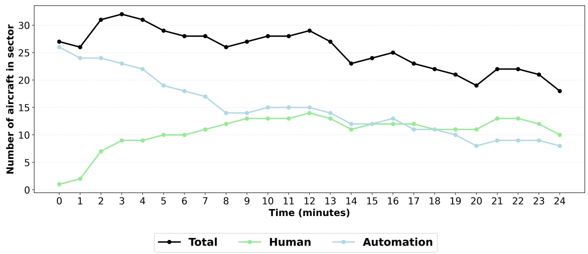

(a) Automated start.

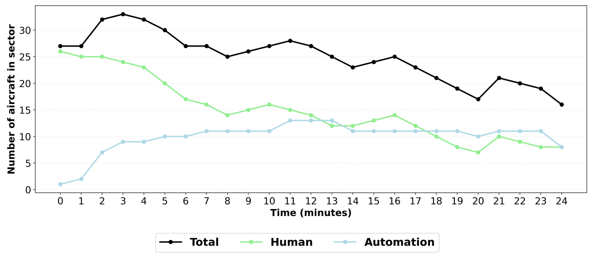

(b) Human start.

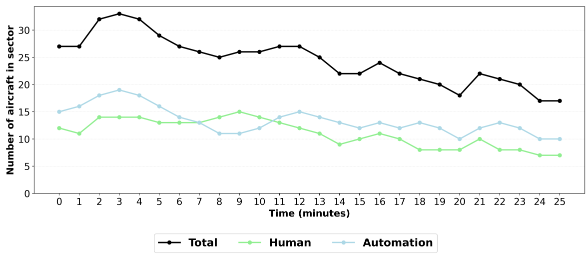

(c) No initial allocation.

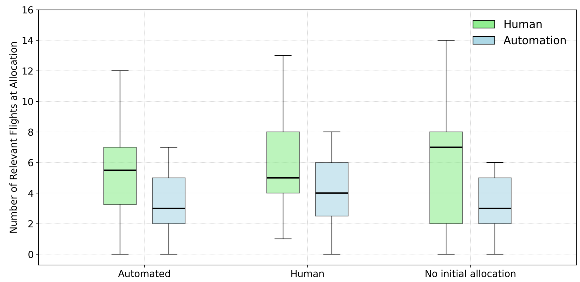

(d) Number of relevant flights at allocation.

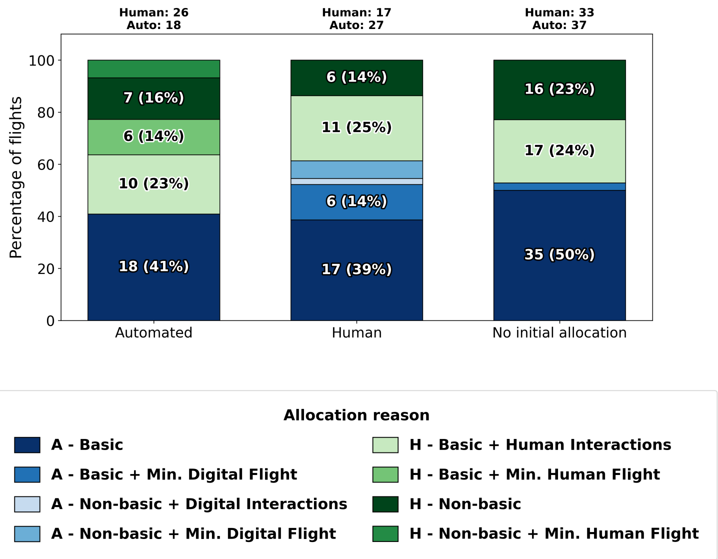

(e) Allocation distribution.

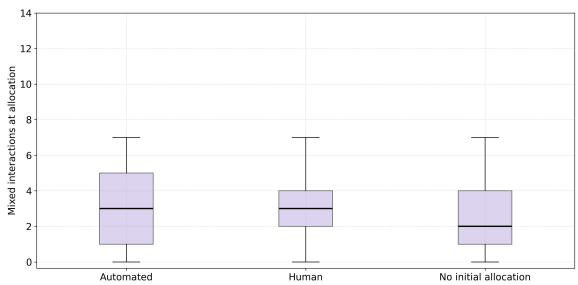

(f) Number of mixed interactions at allocation.

Figure F.5: Scenario 1 Parameter Values: (Int. Threshold: 7, Min. Threshold: 25%, Mixed Int. Threshold: 70%).

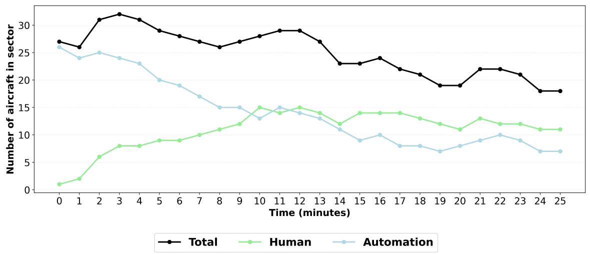

(a) Automated start.

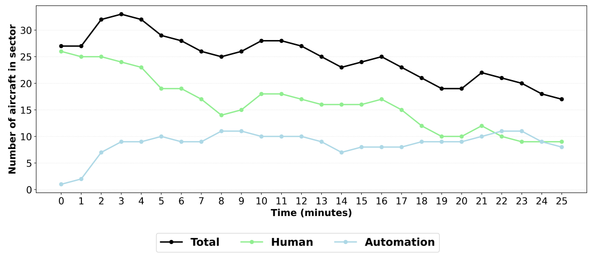

(b) Human start.

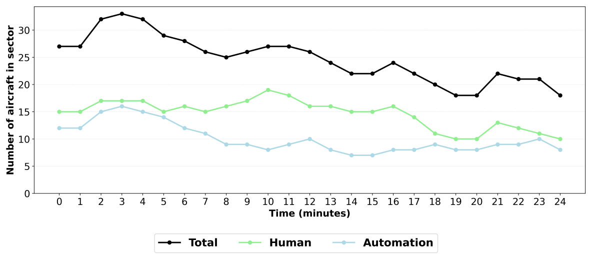

(c) No initial allocation.

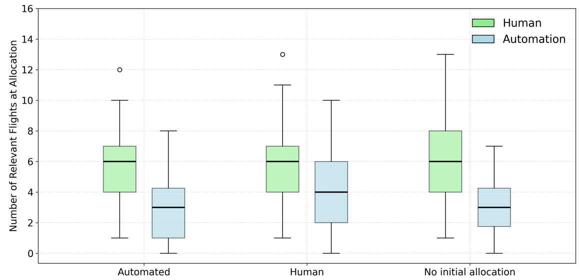

(d) Number of relevant flights at allocation.

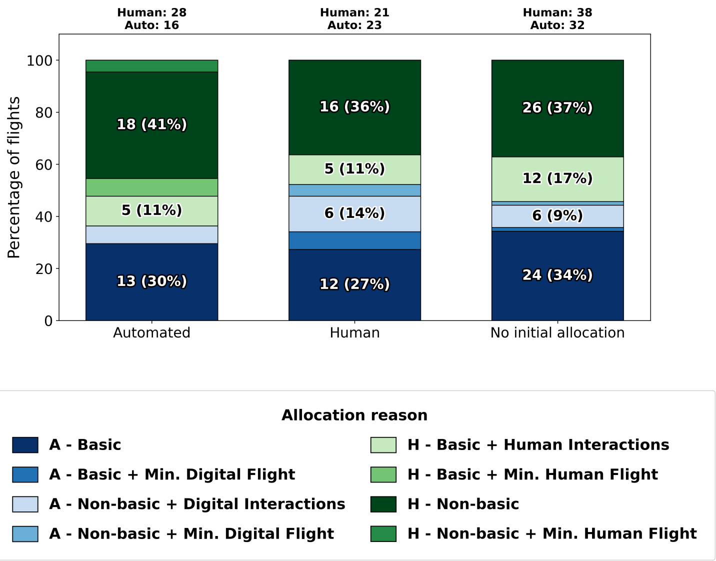

(e) Allocation distribution.

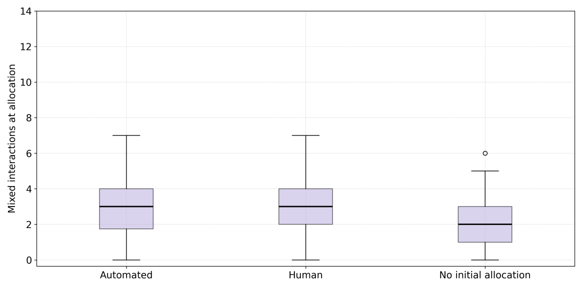

(f) Number of mixed interactions at allocation.

Figure F.6: Scenario 1 Parameter Values: (Int. Threshold: 5, Min. Threshold: 15%, Mixed Int. Threshold: 70%).

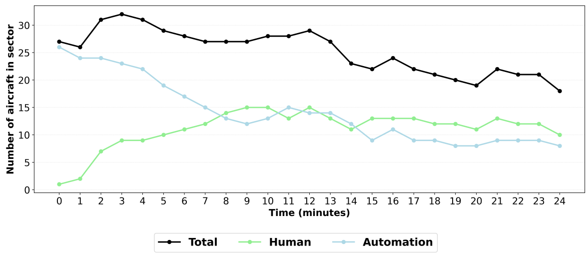

(a) Automated start.

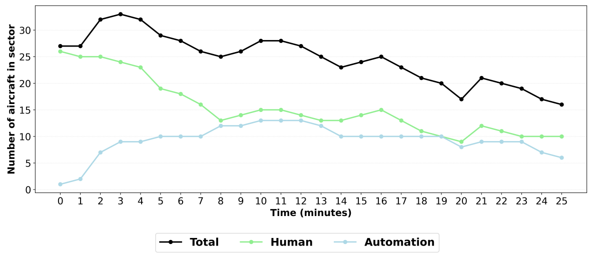

(b) Human start.

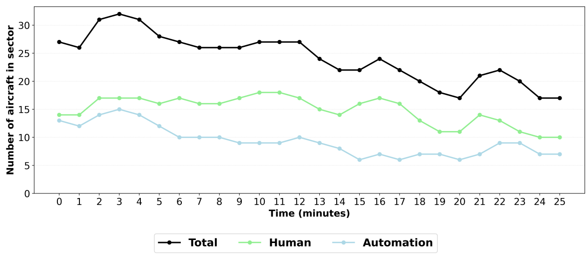

(c) No initial allocation.

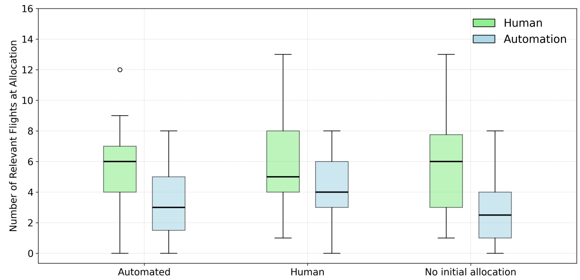

(d) Number of relevant flights at allocation.

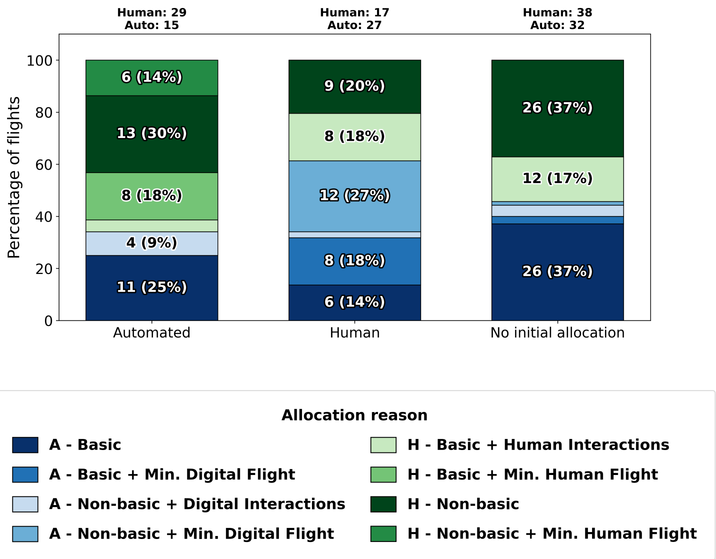

(e) Allocation distribution.

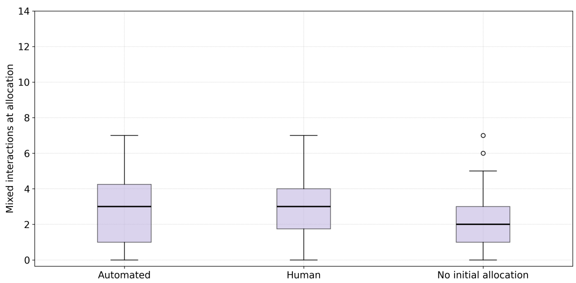

(f) Number of mixed interactions at allocation.

Figure F.7: Scenario 1 Parameter Values: (Int. Threshold: 5, Min. Threshold: 35%, Mixed Int. Threshold: 70%).

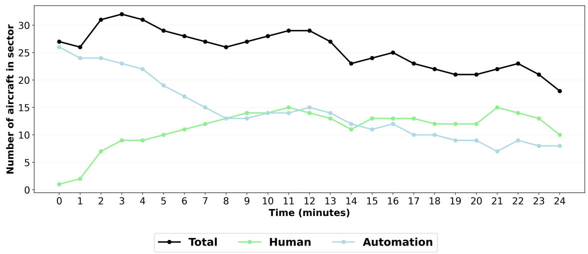

(a) Automated start.

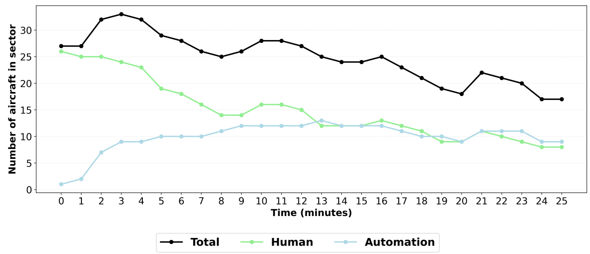

(b) Human start.

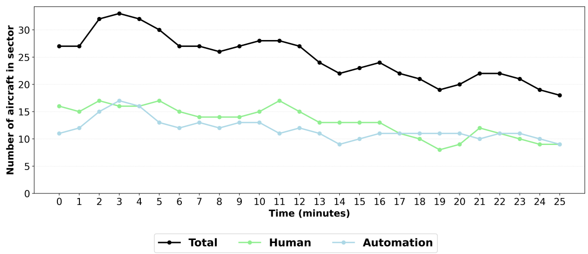

(c) No initial allocation.

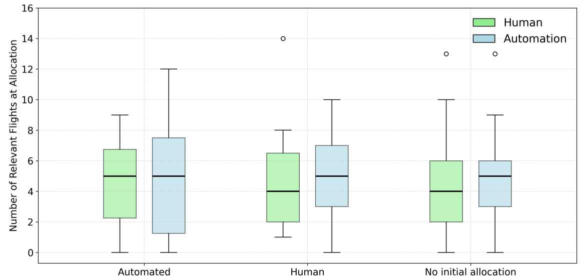

(d) Number of relevant flights at allocation.

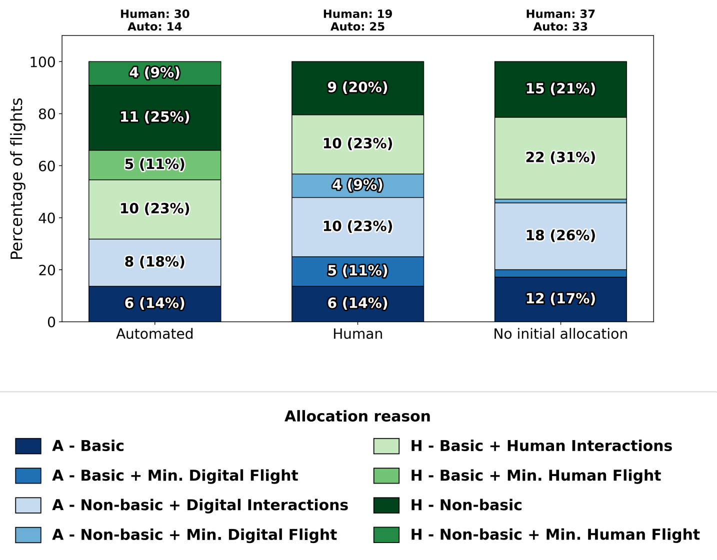

(e) Allocation distribution.

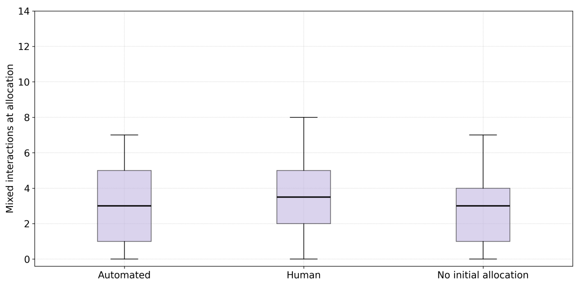

(f) Number of mixed interactions at allocation.

Figure F.8: Scenario 1 Parameter Values: (Int. Threshold: 5, Min. Threshold: 25%, Mixed Int. Threshold: 50%).

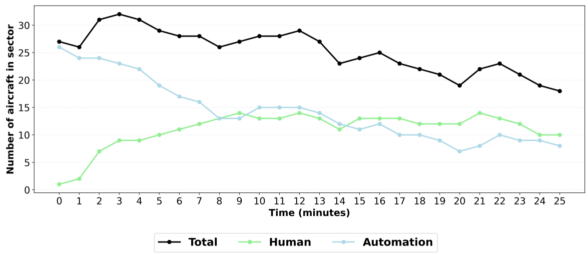

(a) Automated start.

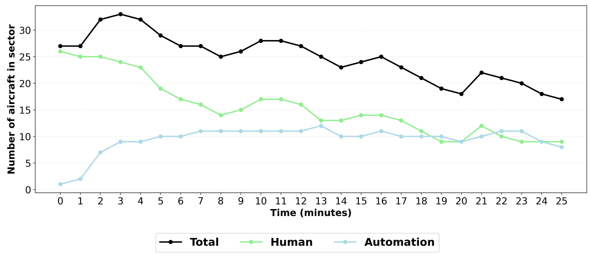

(b) Human start.

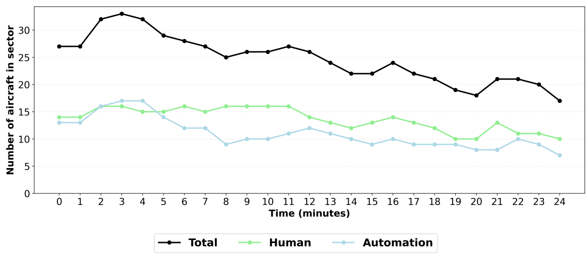

(c) No initial allocation.

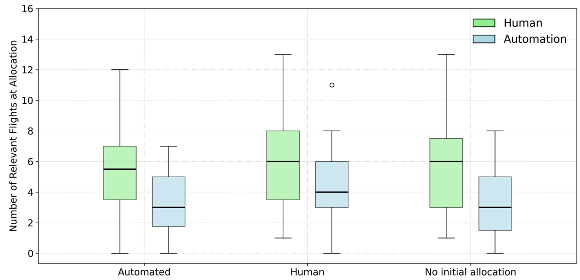

(d) Number of relevant flights at allocation.

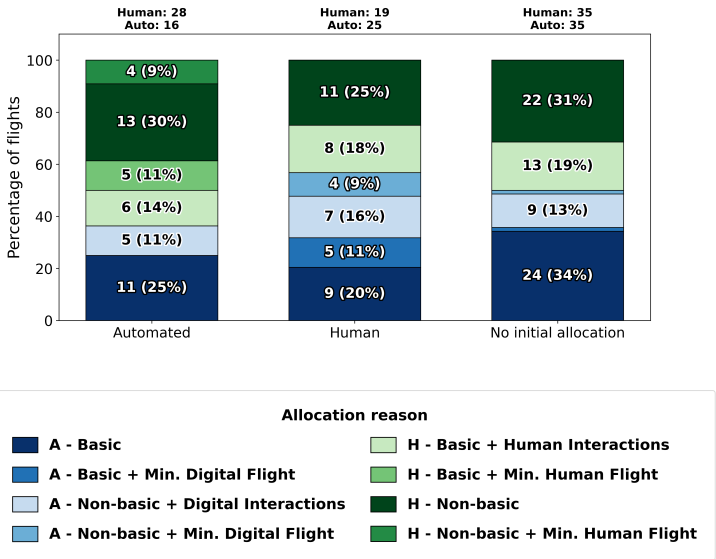

(e) Allocation distribution.

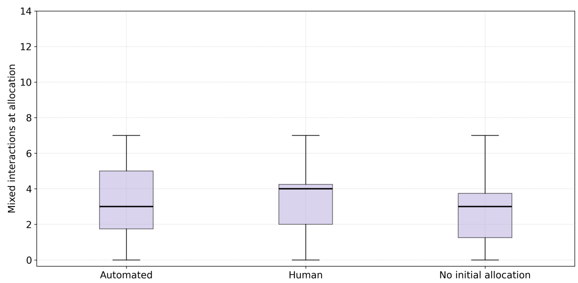

(f) Number of mixed interactions at allocation.

Figure F.9: Scenario 1 Parameter Values: (Int. Threshold: 5, Min. Threshold: 25%, Mixed Int. Threshold: 60%).

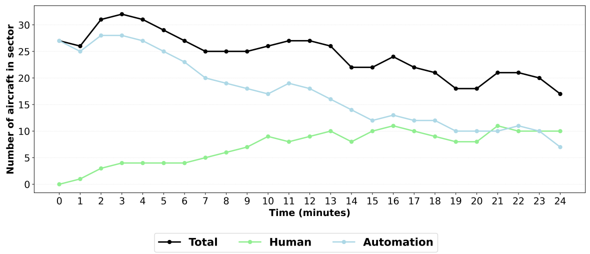

(a) Automated start.

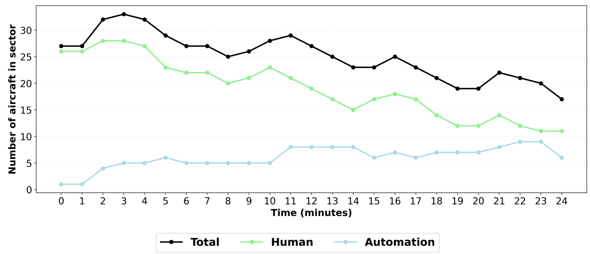

(b) Human start.

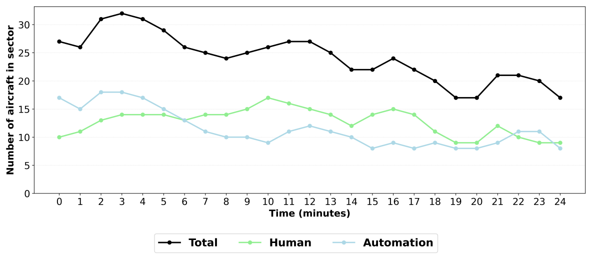

(c) No initial allocation.

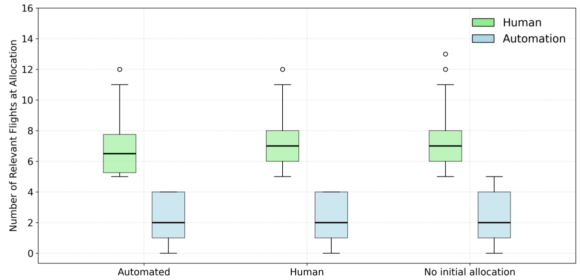

(d) Number of relevant flights at allocation.

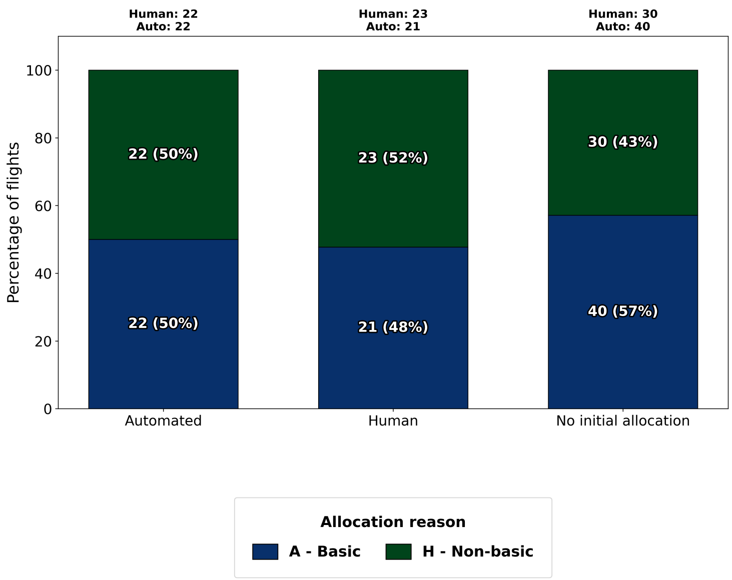

(e) Allocation distribution.

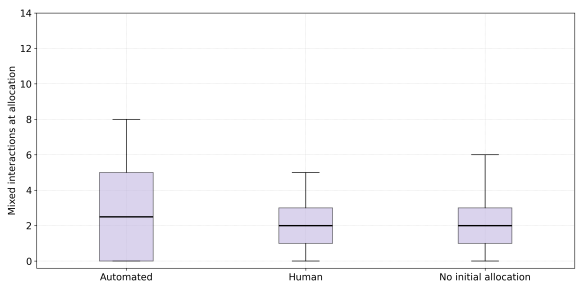

(f) Number of mixed interactions at allocation.

Figure F.10: Scenario 1 Complexity Only: (Int. Threshold: 5, Min. Threshold: N/A, Mixed Int. Threshold: N/A).

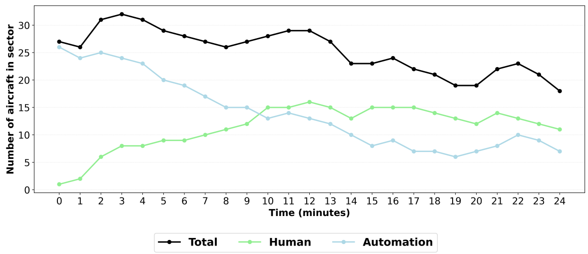

(a) Automated start.

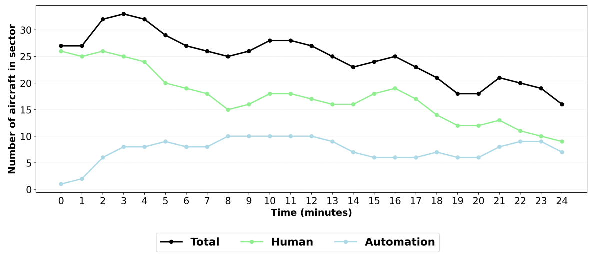

(b) Human start.

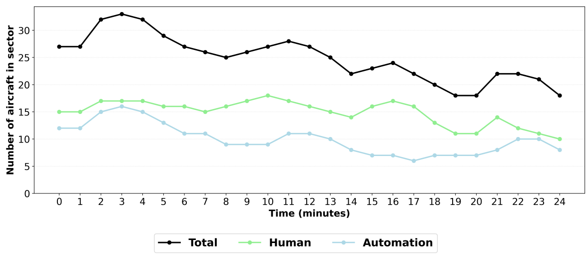

(c) No initial allocation.

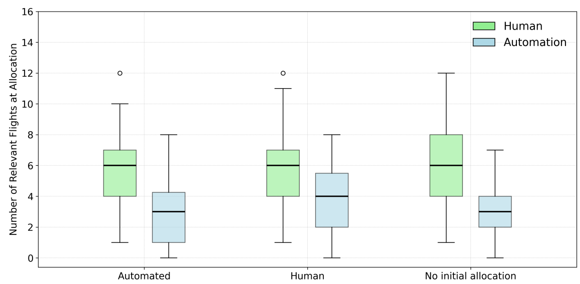

(d) Number of relevant flights at allocation.

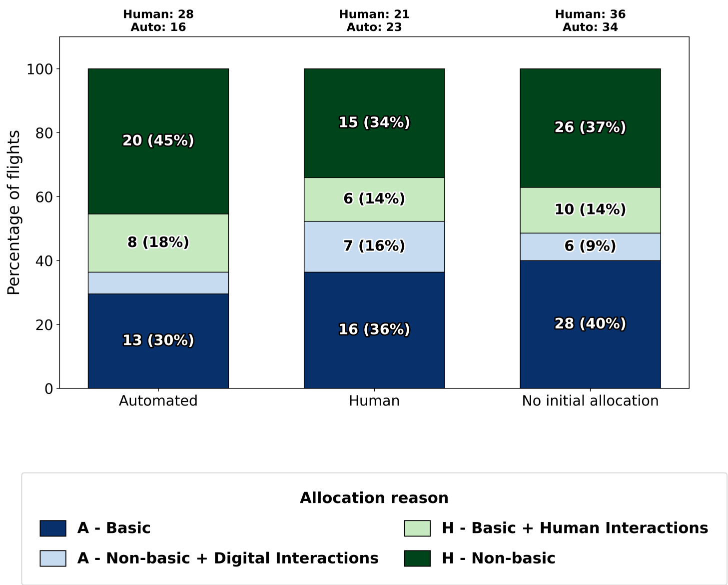

(e) Allocation distribution.

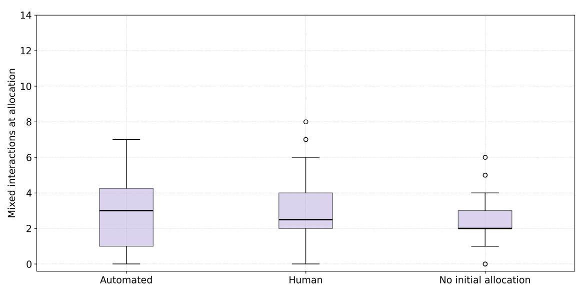

(f) Number of mixed interactions at allocation.

Figure F.11: Scenario 1 No Minimum Flight Threshold: (Int. Threshold: 5, Min. Threshold: N/A, Mixed Int. Threshold: 70).

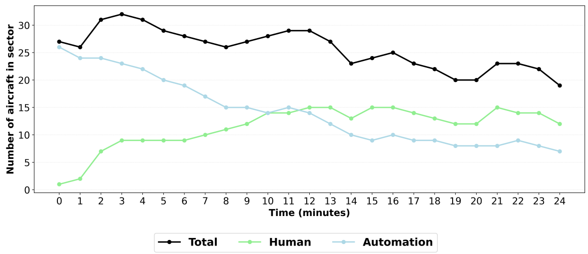

(a) Automated start.

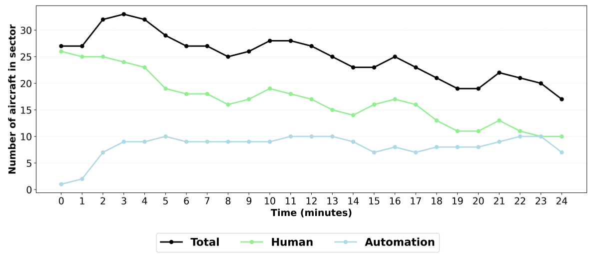

(b) Human start.

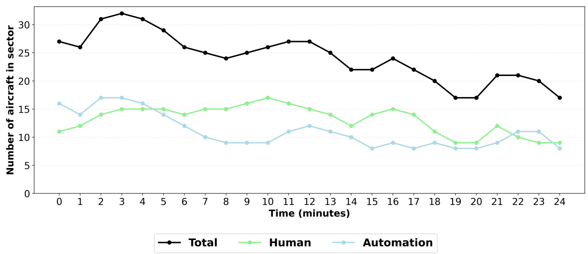

(c) No initial allocation.

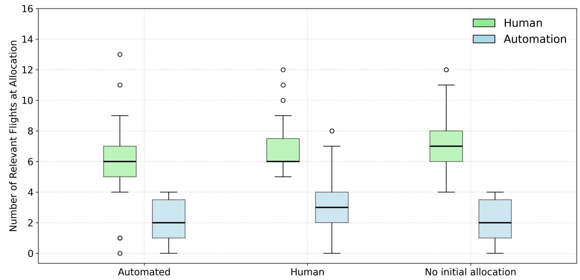

(d) Number of relevant flights at allocation.

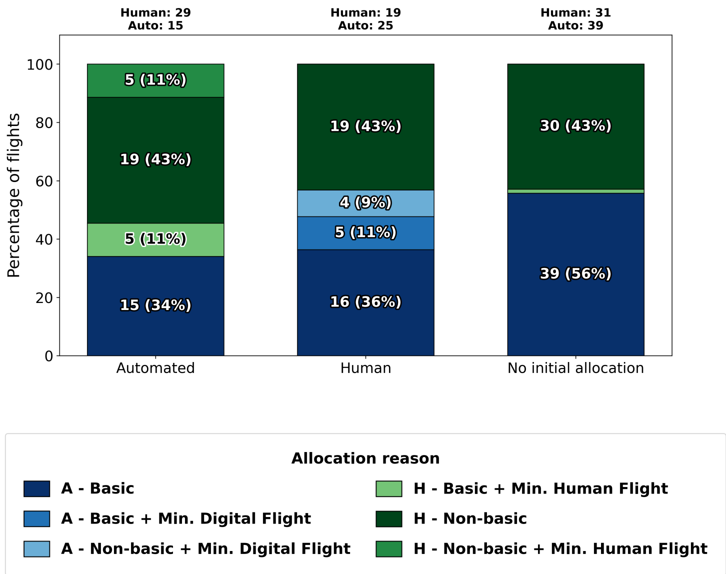

(e) Allocation distribution.

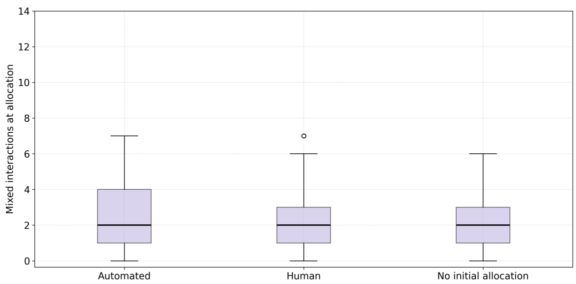

(f) Number of mixed interactions at allocation.

Figure F.12: Scenario 1 No Mixed Interaction Threshold: (Int. Threshold: 5, Min. Threshold: 25, Mixed Int. Threshold: N/A).

F.2.2. Scenario 2

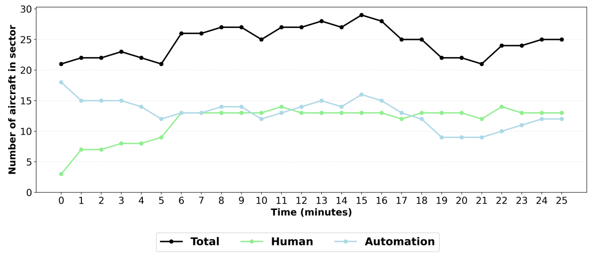

(a) Automated start.

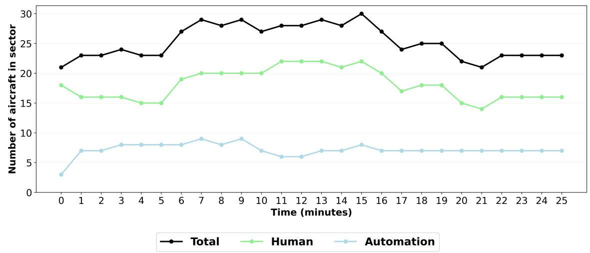

(b) Human start.

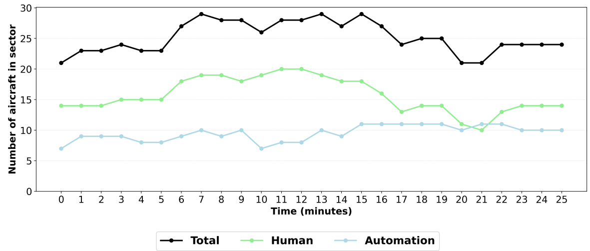

(c) No initial allocation.

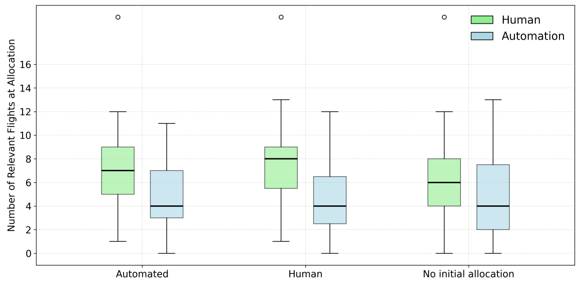

(d) Number of relevant flights at allocation.

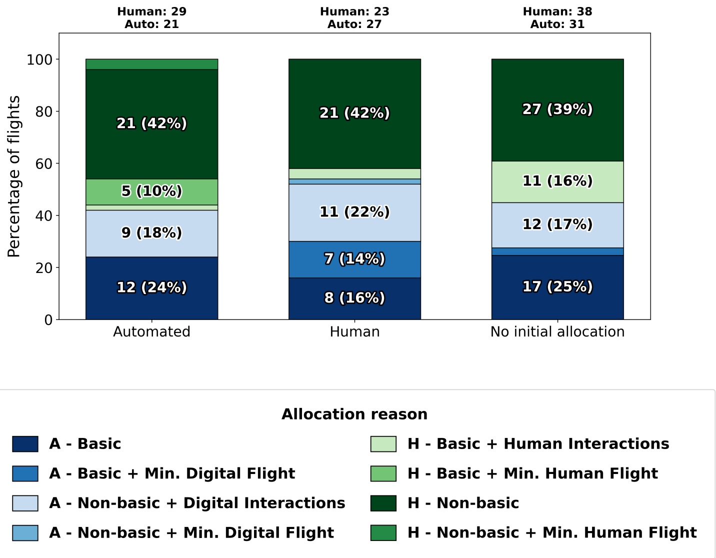

(e) Allocation distribution.

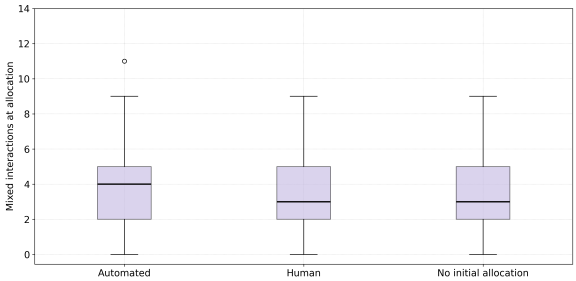

(f) Number of mixed interactions at allocation.

Figure F.13: Scenario 2 Default parameter values (Int. Threshold: 5, Min. Threshold: 25%, Mixed Int. Threshold: 70%).

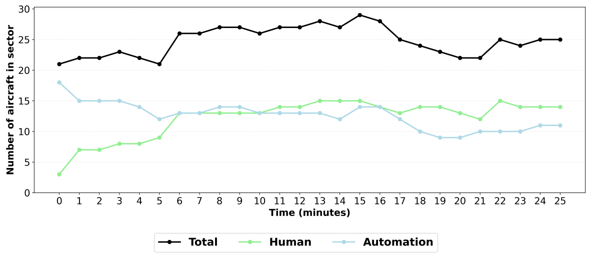

(a) Automated start.

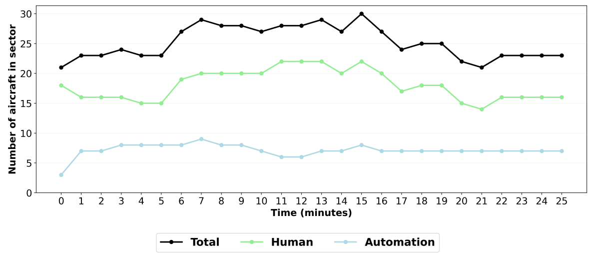

(b) Human start.

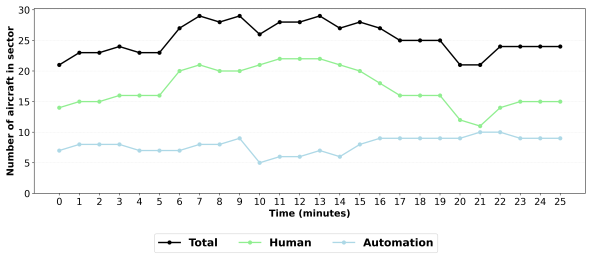

(c) No initial allocation.

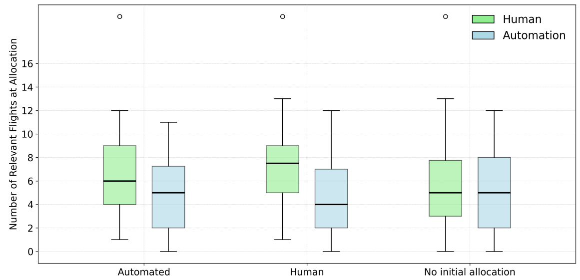

(d) Number of relevant flights at allocation.

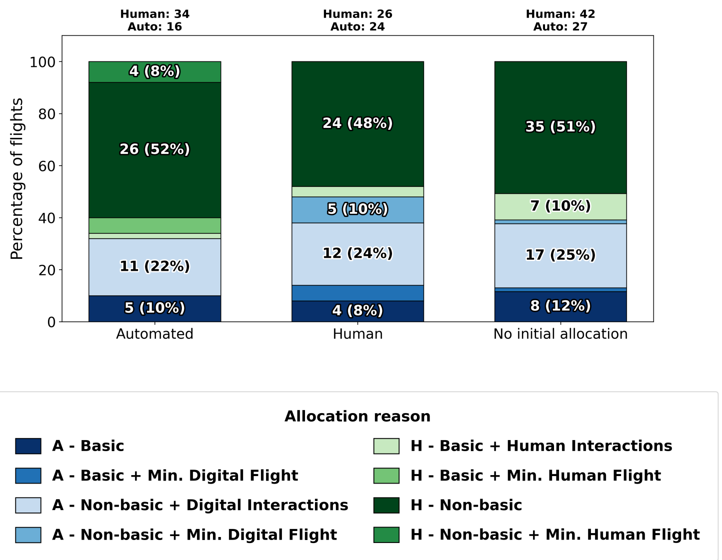

(e) Allocation distribution.

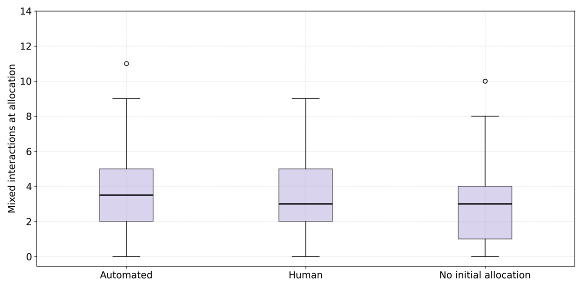

(f) Number of mixed interactions at allocation.

Figure F.14: Scenario 2 Parameter Values: (Int. Threshold: 3, Min. Threshold: 25%, Mixed Int. Threshold: 70%).

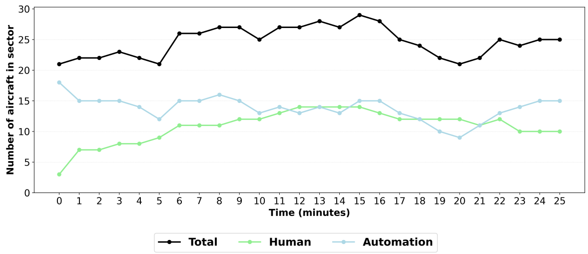

(a) Automated start.

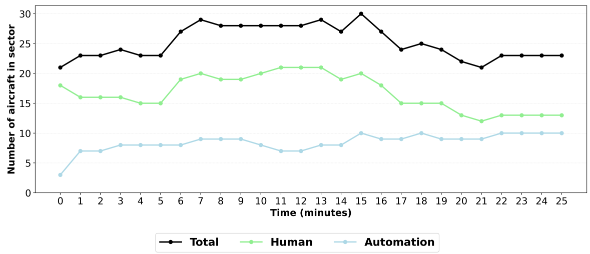

(b) Human start.

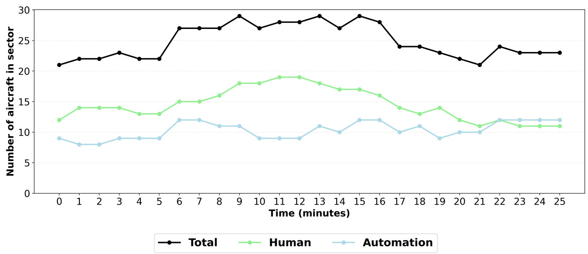

(c) No initial allocation.

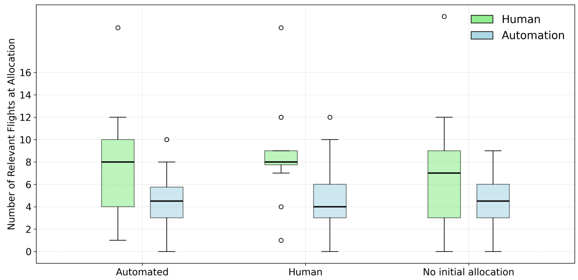

(d) Number of relevant flights at allocation.

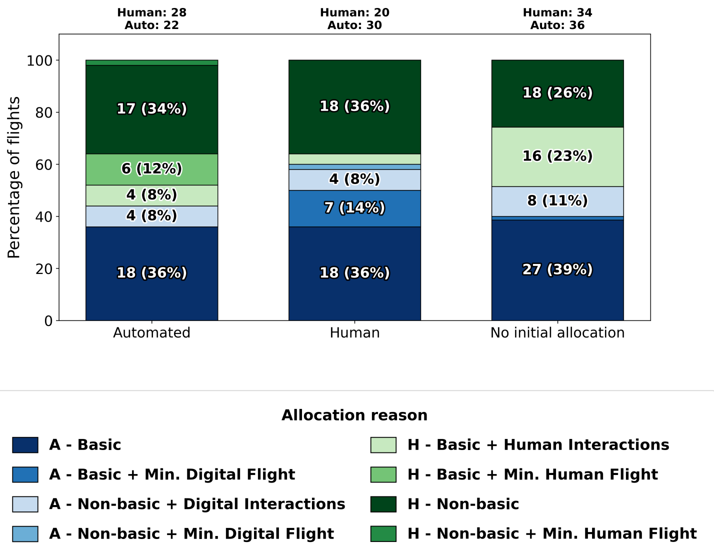

(e) Allocation distribution.

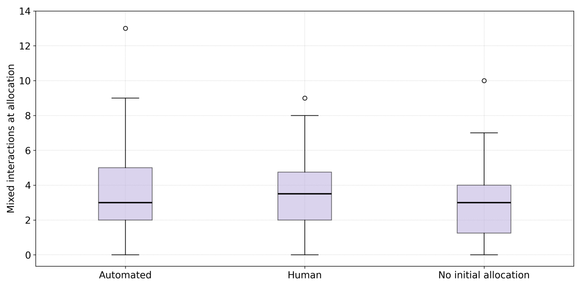

(f) Number of mixed interactions at allocation.

Figure F.15: Scenario 2 Parameter Values: (Int. Threshold: 7, Min. Threshold: 25%, Mixed Int. Threshold: 70%).

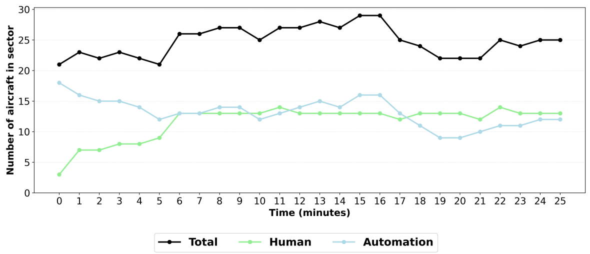

(a) Automated start.

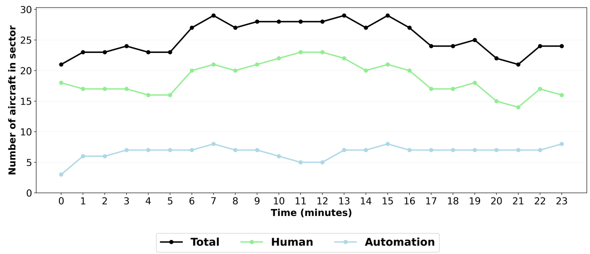

(b) Human start.

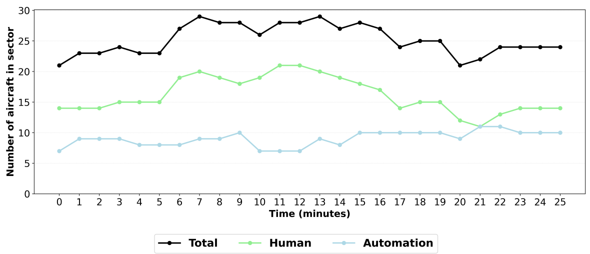

(c) No initial allocation.

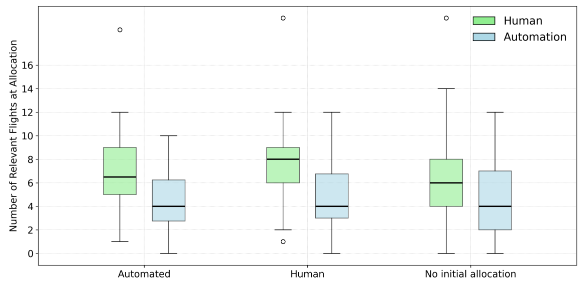

(d) Number of relevant flights at allocation.

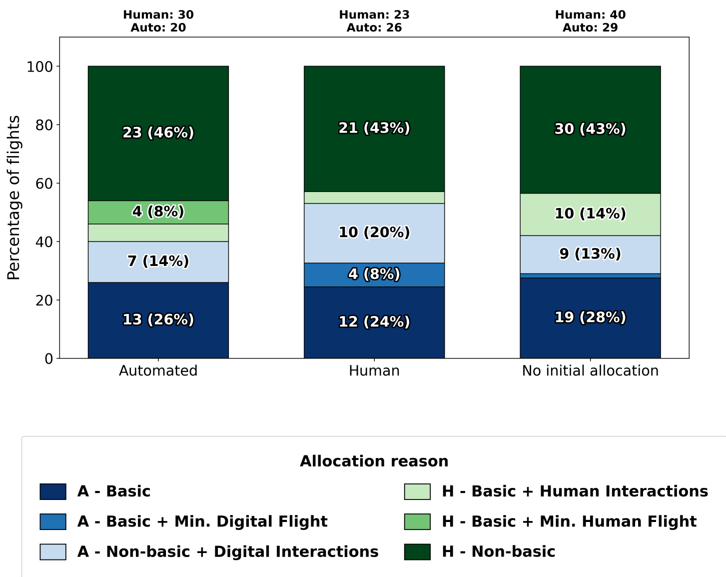

(e) Allocation distribution.

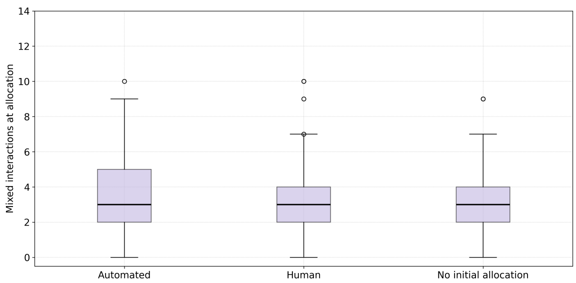

(f) Number of mixed interactions at allocation.

Figure F.16: Scenario 2 Parameter Values: (Int. Threshold: 5, Min. Threshold: 15%, Mixed Int. Threshold: 70%).

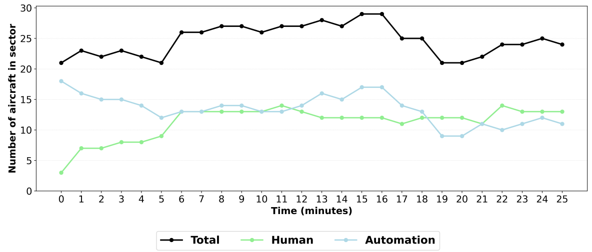

(a) Automated start.

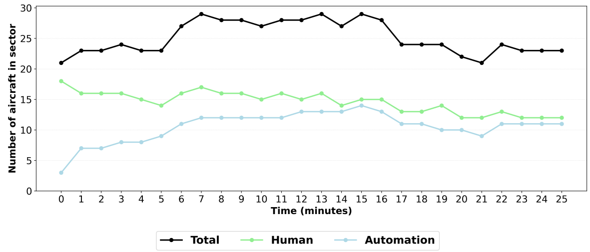

(b) Human start.

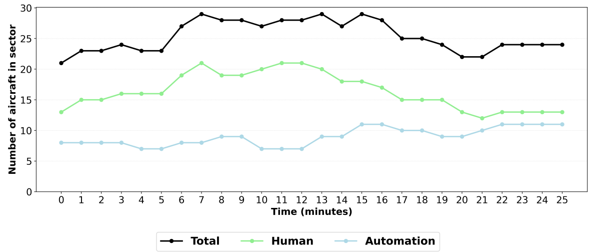

(c) No initial allocation.

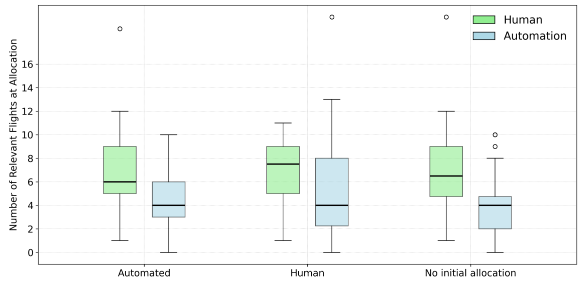

(d) Number of relevant flights at allocation.

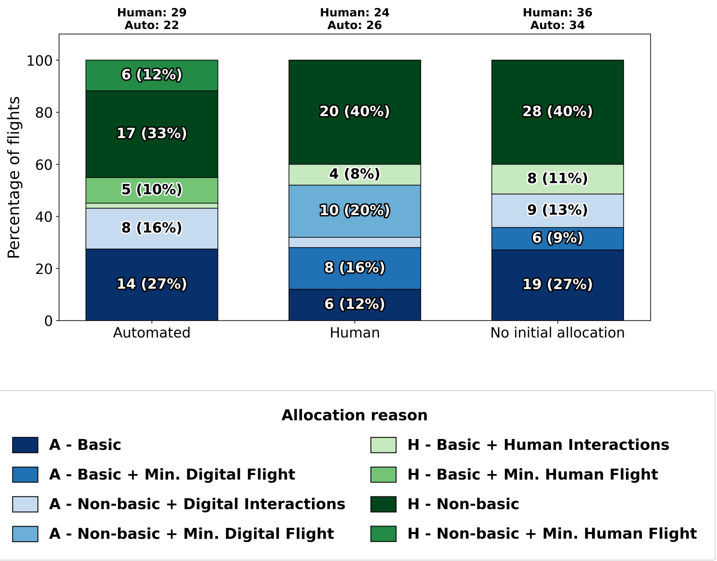

(e) Allocation distribution.

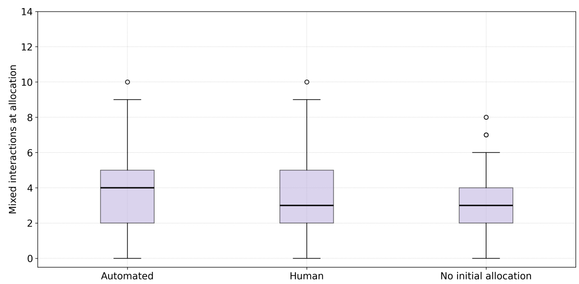

(f) Number of mixed interactions at allocation.

Figure F.17: Scenario 2 Parameter Values: (Int. Threshold: 5, Min. Threshold: 35%, Mixed Int. Threshold: 70%).

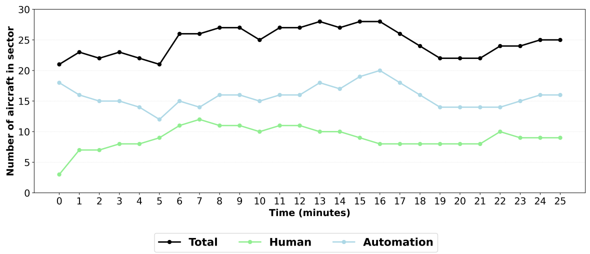

(a) Automated start.

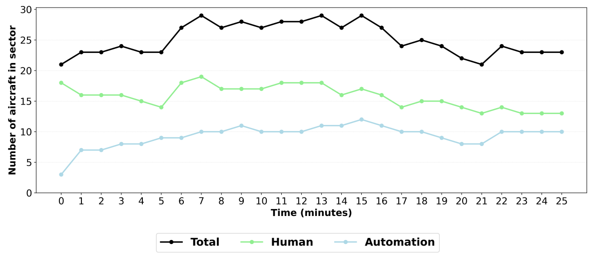

(b) Human start.

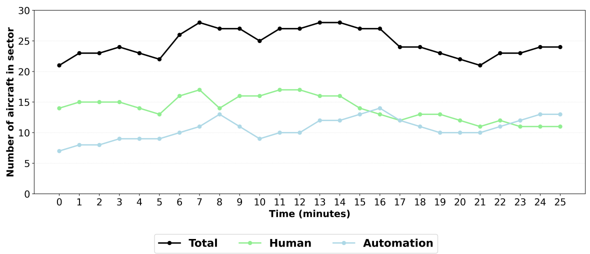

(c) No initial allocation.

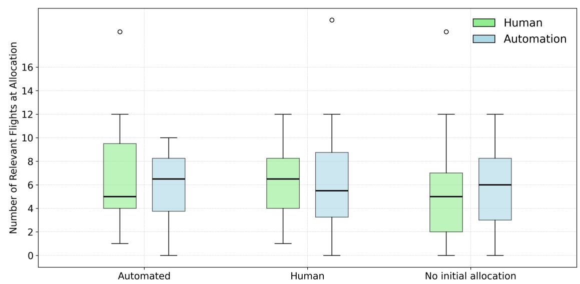

(d) Number of relevant flights at allocation.

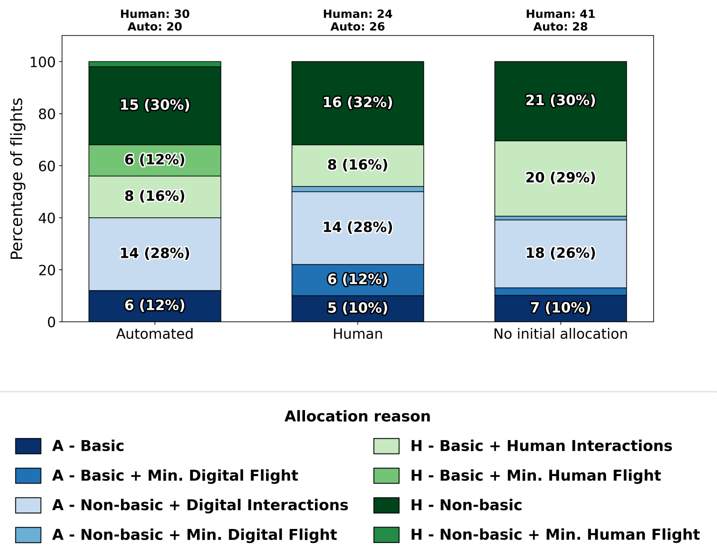

(e) Allocation distribution.

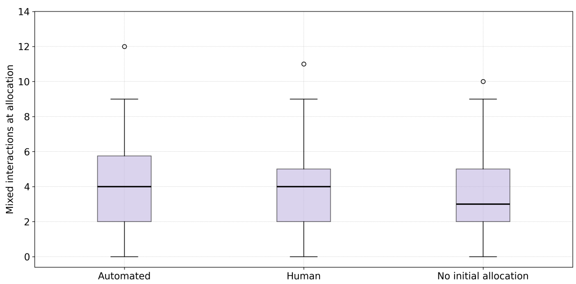

(f) Number of mixed interactions at allocation.

Figure F.18: Scenario 2 Parameter Values: (Int. Threshold: 5, Min. Threshold: 25%, Mixed Int. Threshold: 50%).

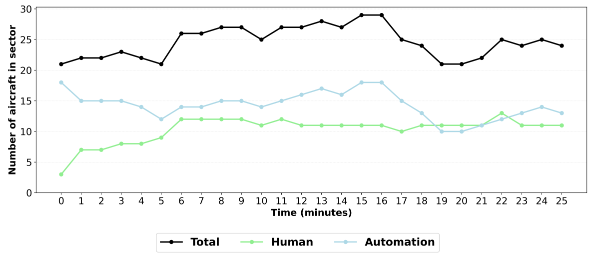

(a) Automated start.

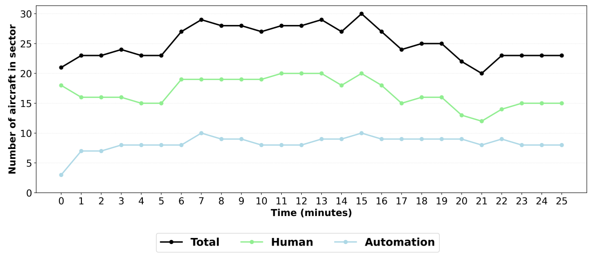

(b) Human start.

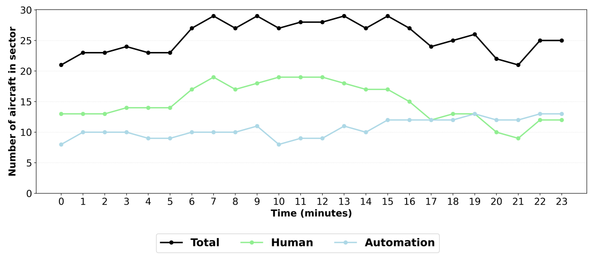

(c) No initial allocation.

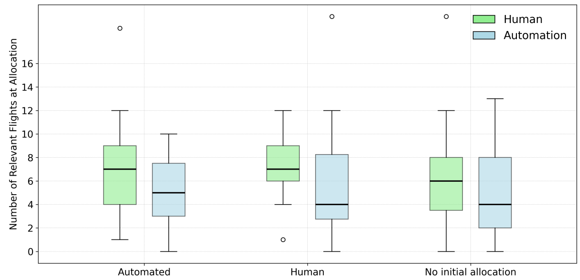

(d) Number of relevant flights at allocation.

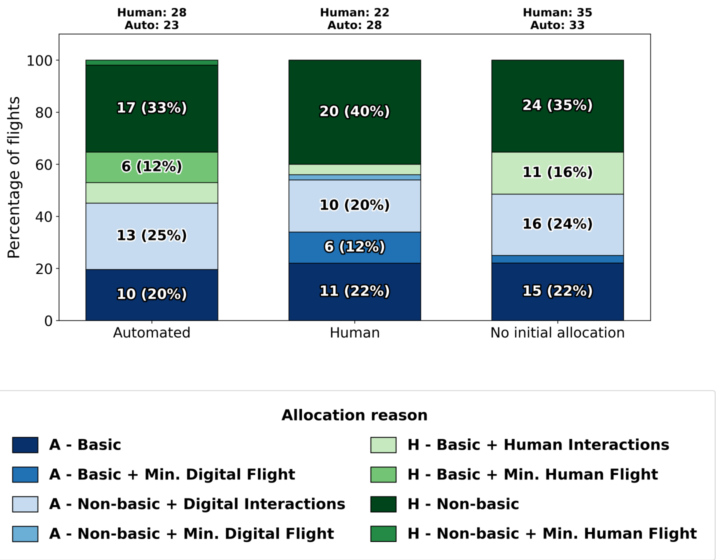

(e) Allocation distribution.

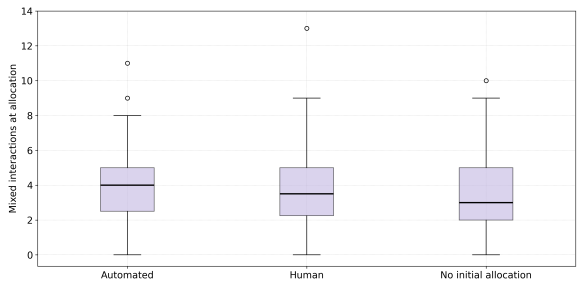

(f) Number of mixed interactions at allocation.

Figure F.19: Scenario 2 Parameter Values: (Int. Threshold: 5, Min. Threshold: 25%, Mixed Int. Threshold: 60%).