Solution Space Concept for Trajectory-Based Air Traffic Control
An Ecological Approach

Klomp, R.E.

DOI
10.4233/uuid:c1a8a6ee-b589-4243-8e2f-e021dfafe1fa

Publication date
2023

Document Version
Final published version

Citation (APA)
Klomp, R. E. (2023). *Solution Space Concept for Trajectory-Based Air Traffic Control: An Ecological Approach*. [Dissertation (TU Delft), Delft University of Technology]. https://doi.org/10.4233/uuid:c1a8a6ee-b589-4243-8e2f-e021dfafe1fa

Solution Space Concept for Trajectory-Based Air Traffic Control

An Ecological Approach

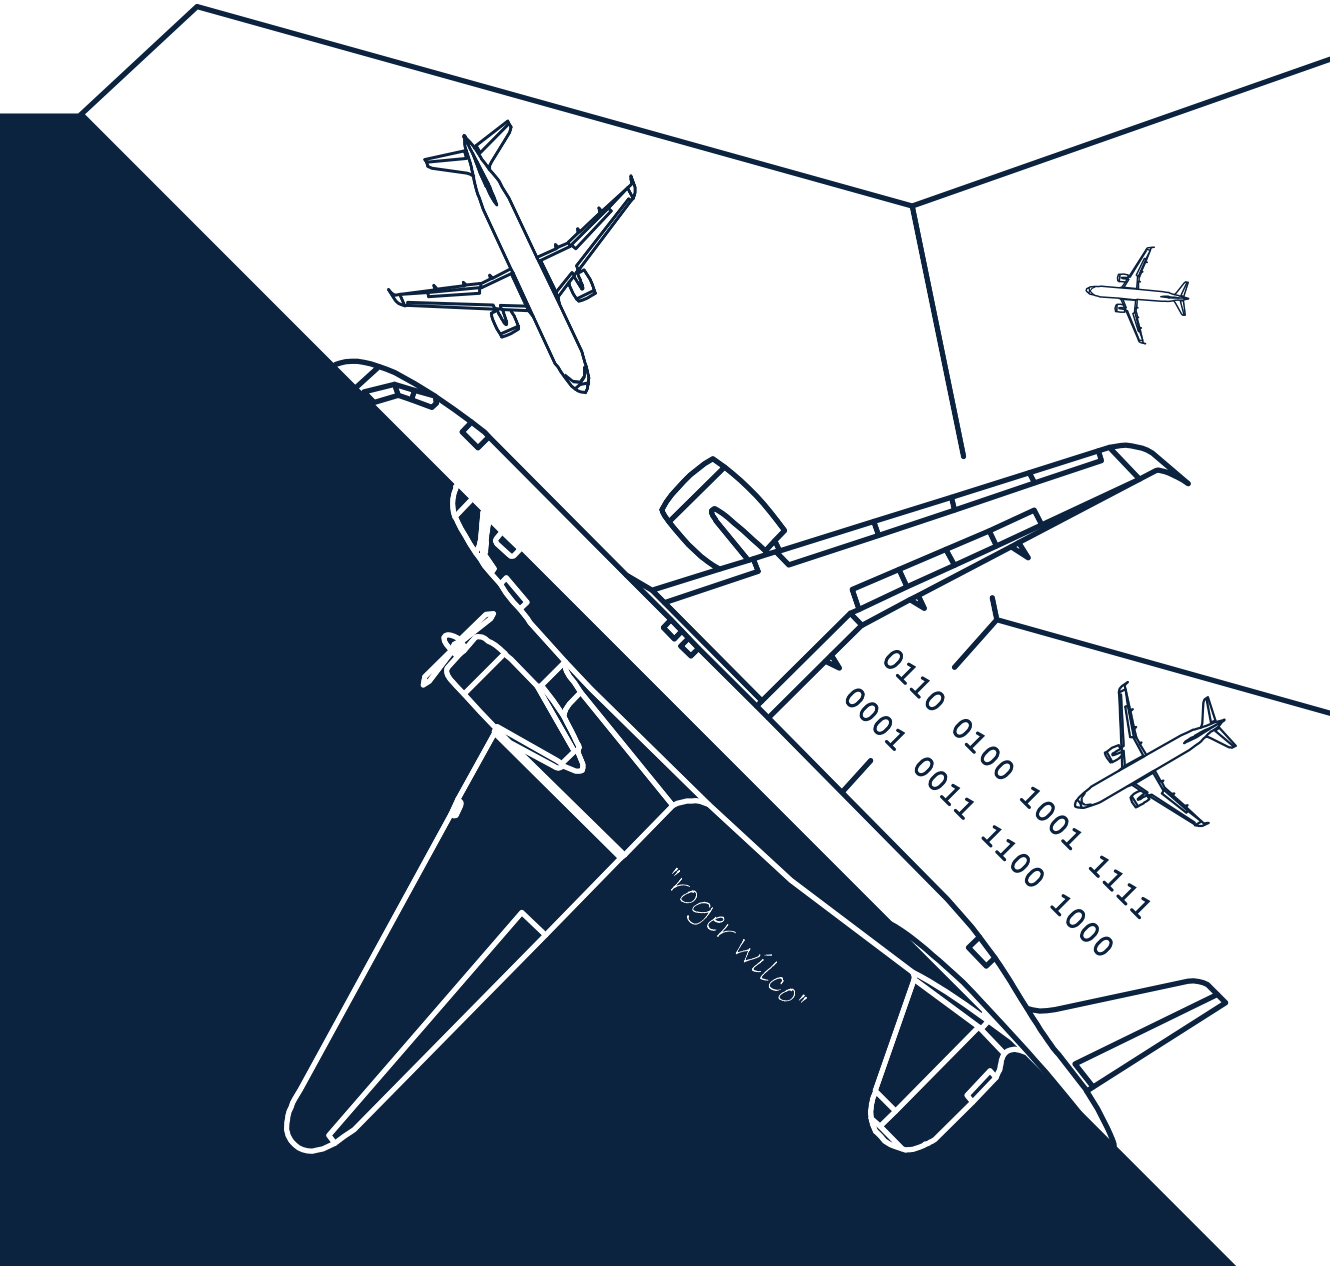

Rolf Klomp

Solution Space Concept for Trajectory-Based Air Traffic Control

An Ecological Approach

Solution Space Concept for Trajectory-Based Air Traffic Control

An Ecological Approach

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen
chair of the Board for Doctorates
to be defended publicly on
Tuesday 2 May 2023 at 10:00

by

Rolf Edwin KLOMP

Master of Science in Aerospace Engineering
Delft University of Technology, the Netherlands

born in Voorburg, the Netherlands

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Prof. dr. ir. M. Mulder,	Delft University of Technology, promotor
Dr. ir. C. Borst,	Delft University of Technology, co-promotor
Independent members:	
Prof. dr. D. Kaber	University of Florida
Dr. D. Schaefer	EUROCONTROL
Prof. dr. J. M. Schraagen	TNO
Prof. dr. ir. J. C. F. de Winter	Delft University of Technology
Prof. dr. ir. J. M. Hoekstra	Delft University of Technology

The work described in this thesis has been carried out at the Control and Simulation section at Delft University of Technology. Part of this work was supported by EUROCONTROL acting on behalf of the SESAR Joint Undertaking (SJU) and the European Union as part of SESAR Work Package E project C-SHARE.

Dr. ir. M.M. van Paassen contributed to a significant extent in the realization of this thesis.

Keywords: Air Traffic Control, Trajectory-based Operations, Solution Space, Human-centered Automation
Cover: “Contrast” by R. E. Klomp
Printed by: Ipskamp Printing

ISBN: 978-94-6473-071-5

An electronic version of this dissertation is available at
`http://repository.tudelft.nl/`.

SUMMARY

Solution Space Concept for Trajectory-Based Air Traffic Control

THE increasing demand in worldwide air travel is foreseen to push the limits of capacity of the current Air Traffic Management (ATM) system within the coming decade. Therefore, two major international programs have been initiated to fundamentally restructure the way in which Air Traffic Control (ATC) is performed. A key pillar in both programs is the introduction of Trajectory-based Operations (TBO), in which highly accurate gate-to-gate defined four-dimensional (4D) trajectories will form the basis of the work of future Air Traffic Controllers (ATCos) and pilots alike. There is a consensus that human controllers and not automation should remain ultimately responsible for the safety of operations. However, the precise task of the ATCo, and the extent of automation autonomy and authority are not clear.

The work in this thesis explored a human-centered—*ecological*—approach to the design of decision-support interfaces for tactical 4D trajectory-based control. This approach focuses on the use of automation as a means to enhance the capabilities of human controllers, and (in this work) enables them to retain their current role and responsibilities in the more complex operational environment of TBO. Automation is used to integrate and directly visualize the boundaries of feasible trajectory-based control actions, overlaid with go-and no-go areas resulting from separation requirements to other traffic. The resulting *Solution Space* concept empowers controllers to perform effective (manual) decision making and action implementation.

The research hinged around answering three main questions. First, what are the underlying system goals, functions, relationships, and constraints inherent to the work domain of tactical 4D trajectory-based air traffic control, and how do they translate into Solution Spaces? Second, in what visual form can these Solution Spaces best be mapped to a decision-support interface such that they support the work and strategies of trajectory-based control? Third, how do human controllers with different levels of expertise use interactive implementations of the resulting concept, what strategies do they apply, and what lessons can be learned?

To answer the first question, the frameworks of Cognitive Work Analysis (CWA) and Ecological Interface Design (EID) provided guidance in identifying the elements that underlie the ecology of the trajectory-based control task. As part of CWA, Rasmussen's Abstraction Hierarchy (AH) was used to explore the work domain on various levels of abstraction; from the higher level (abstract) system goals such as safety and efficiency, to the lower level (physical) elements such as sector shape and size and aircraft intent. Together with a set of hierarchical means-end links that connect the elements in a top-down manner, the AH encompasses the functions, relationships, and constraints inherent to the work domain, and defines the playing field for trajectory-based control.

Using the AH as a basis, and resulting from the integration of task-relevant information on the higher levels of abstraction, the construct of the Solution Space is introduced. Solution Space is conceptualized as the comprehensive set of—fully four-dimensional—control space, constrained by both the internal constraints of a controlled aircraft, and by external constraints within the air traffic control work domain. That is, the instantaneous Solution Space of an aircraft represents the physically feasible control space to manipulate its trajectory (i.e., within the performance envelope). Avoidance zones caused by rule-based—intentional—separation requirements with respect to other dynamic traffic are then mapped on top of the control spaces in the form of go and no-go areas. The resulting Solution Spaces then represent all available trajectory-based control actions, regardless of their optimality.

Implementations of the concept focused on the en-route control task, and led to the development of interactive, real-time visual representations of horizontal and time-based Solution Spaces in prototype decision-support interfaces. Horizontal and time-based Solution Spaces show the available and safe control spaces for lateral- and time-based (i.e., speed) trajectory revisions respectively. The design and refinement of the interfaces followed an iterative approach by applying the lessons learned from various human-in-the-loop experiments. The levels of automation in this work were limited to information automation only; decision selection and action implementation were fully at the discretion of the human controller. The concept, however, does not rule out the addition of higher levels of action automation acting upon the same definition of the Solution Space, and thereby either partially or completely taking over the tasks of the human controller.

Three human-in-the-loop experiments with iterations of the decision-support prototypes are presented and show that controllers with varying levels of expertise can all successfully use Solution Space interfaces to safely control traffic, albeit in a different way. In the first experiment that only featured horizontal Solution Spaces, novices were found to perform more risky and/or less efficient strategies. They often used the Solution Spaces as their primary means of decision selection, and frequently chose solutions close to the borders of safe and unsafe actions. Experts performed more robust and pro-active control, typically using the Solution Spaces to validate their premeditated strategies. However, they also indicated to have lower trust in the presented information. In the second iteration of the prototype, additional intentional constraints were added and visualized within the horizontal Solution Spaces in the form of larger separation buffers and areas that lead to inefficient solutions. The second human-in-the-loop experiment—a repetition of the first experiment, but with the addition of the intentional constraints—showed that the control *robustness* of novice participants improved to the same level of the experts in the prior experiment. This indicates that the experts already knew how to avoid wandering into 'tight' solutions, and likely explains their lower trust in the Solution Spaces that identified such solutions as safe. The third human-in-the-loop experiment featured both horizontal and time-based Solution Spaces, and focused on control strategies. Results showed that the applied strategies, and thereby the resulting control efficiency, varied greatly between individual participants. On the one hand, this indicates that the interface is individual sensitive, and allows enough freedom for controllers to formulate and implement their own unique strategies. On the other hand, this emphasizes that ecological interfaces, by themselves, do not remove or replace the advantages of expertise and deep domain knowledge gained by extensive training and experience.

The limited scope and research limitations in this thesis led to initial Solution Space decision-support concepts that—although promising—were governed by fully deterministic constraints, and did not encompass the full four-dimensionality of trajectory-based operations. These are areas that need to be addressed in future work in order to gain more traction, and thereby acceptance within the operational communities. First, uncertainty and fuzzy constraints should explicitly be taken into account in the elements and functions in the AH. These probabilistic traits should then become an inherent part of the definition of Solution Spaces themselves. Second, the Solution Spaces should be extended to encompass all possible dimensions of trajectory control. The absence of control in the vertical dimension should especially be addressed. And finally, the above should be translated into useful partial Solution Space-based decision-support interfaces that can be perceived in an intuitive way, without ignoring or simplifying the underlying work domain complexities. This, however, is neither a trivial nor a straightforward endeavor and will require overcoming many challenges.

CONTENTS

CHAPTER 1

INTRODUCTION

This chapter provides an introduction and background to the challenges in designing effective air traffic controller decision-support tools for 4D trajectory-based control. The research problem and the main research questions addressed in this thesis are introduced. An overview is provided of the research approach and the used methodologies, together with the definition of the scope and the research assumptions. The chapter concludes with an outline of the main chapters, and a short description of their position and importance within the context of the combined work.

1.1 Background

THE Air Traffic Management (ATM) system is responsible for providing the means to coordinate safe and efficient global air traffic movements. The system encompasses all functions that assist aircraft to successfully plan and execute their flight, including Air Traffic Control (ATC), Air Traffic Flow and Capacity Management (ATFCM), Air Traffic Services (ATS), aeronautical meteorology, and air navigation systems. These functions are provided by a total of 83 Air Navigation Service Providers (ANSPs), and are governed by a set of unified standards and recommended practices that are issued by the International Civil Aviation Organization (ICAO). In 2019, the worldwide ATM system was responsible for handling an average of 105.000 commercial flights daily [1]. Although the recent COVID-19 pandemic temporarily reduced the number of traffic movements by 50% in Europe in 2020, the number of movements is foreseen to be back at pre-pandemic levels as early as 2022 [2]. In the long-term forecast, the ATM system is foreseen to reach the limits of its capacity within the next decades [3].

1.1.1 The Future of Air Traffic Management

There is a consensus by the operational ATM community that fundamental *organizational* and *technological* changes in the current system are required to remain effective and sustainable in future operations [4]. Three main bottlenecks can be identified which prevent a capacity growth in the current system. Firstly, most systems operated by the ANSPs have been developed by individual organizations, which has resulted in different levels of capabilities and services around the world. Therefore, procedures and transitions between the airspace of different ANSPs are not well harmonized [5], [6]. Secondly, as the majority of ANSPs operate on a national level, the size and shape of aeronautical sectors is mainly determined by geographic borders. The fragmentation of the airspace and airways, especially in Europe, prevents aircraft to fly their optimal point-to-point trajectories. Finally, the technologies supported by the current ATM systems do not fully utilize the capabilities of modern aircraft. Most commercial aircraft have extensive digital data-link capabilities and are able to execute their flight plans with a high degree of precision. However, the ATC community has, traditionally, been conservative with adopting new technologies.

In order to overcome these bottlenecks, two major projects have been initiated to fundamentally restructure the current ATM system. The Single European Sky ATM Research Joint Undertaking (SESAR-JU), an initiative of the European Union, and the Next Generation Air Transportation System (NextGen) program, initiated by the U.S. Federal Aviation Administration (FAA) [4], [7]. Both programs lean heavily on the introduction of *higher levels of automation*, a transition to *digital communications* between the ground segment and aircraft (voice by exception), and a *thorough restructuring of the airspace and ATM procedures*. In Europe, this transition is foreseen to facilitate a three-fold increase in airspace capacity, a factor ten improvement in safety, and a fifty percent reduction in ATM service cost by 2040 [8].

A paradigm shift in both programs is the transition towards a unified system of *Trajectory-based Operations* (TBO). Highly accurate gate-to-gate defined 4D (i.e., space and time) trajectories, stored in automated support tools, and shared between airspace and ground actors via digital data-link, will form the basis of the work of the future Air Traffic Controllers (ATCos) and pilots alike. Trajectories are continuously and collaboratively refined by all stakeholders, from early flow-based planning to balance airspace demand versus capacity, to isolated detailed 4D reference trajectories (4DTs) that are corrected for the latest airspace constraints and weather forecasts. In a mature state, this will enable a much more efficient use of airspace, a higher accuracy of traffic movements, and more predictability within the system.

Although many technical aspects of TBO have already been addressed, the *role of the human*, and the extent of human-automation coordination is not yet clear. Longer time-horizon planning tasks such as airspace allocation, route optimization, and capacity balancing typically lend themselves well to be highly automated with minimal human input. At this stage, uncertainty is not yet taken into account, allowing automation to solve a mainly deterministic puzzle. During the *execution phase* of the 4DTs, however, deviations from the planned trajectories, originating from a myriad of possible foreseen and unforeseen events, are unavoidable. Although various tasks such as conformance monitoring and limited trajectory adjustments can be performed by automation, automation is not (yet) able to outperform human control in terms of creativity and out-of-the-box problem solving. Therefore—given that the stakes are high—it is crucial that the human controller should remain in control of monitoring the automation, and should be able to step in when needed.

The added complexity of trajectory-based operations, and the added abstraction of introducing *time* as an explicit control variable, will result in a scenario in which a future ATCo will not be able to perform his/her work without the help of extensive automation aides. Therefore, the development of effective decision-support tools is essential to enable controllers to monitor conformance at a glance, provide the means to effectively reason about the functioning of automation, and to allow for intervention when necessary.

1.2 Problem Definition

The envisioned future ATM system, centered around 4D trajectory-based operations, will inevitably bring a significant shift in the work of the air traffic controller. Given the more complex nature of data-link driven 4D trajectory management, the need for more automation support is clear. However, the ATM community acknowledges that fully automating the ATM system is nearly impossible, leading to the consensus that humans should remain responsible for operational safety [4], [7]. A *central role* is foreseen for the human controller whose work will be oriented more towards higher level airspace management and decision making, while routine tasks will be taken over by automation [4]. The success of this *human-centered* system hinges on three crucial challenges that need to be addressed.

The first challenge is determining *how far* the next generation ATM system *should and can be automated*. Increasing the level of automation in itself is not good or bad, however, in many other complex socio-technical domains this has often introduced new problems, problems that are harder to tackle than those intended to be solved in the first place [9], [10]. Breakdowns in human-automation coordination can result in numerous issues such as a poor understanding of the way automation works, skill degradation, complacency, and over-reliance. The question then is, at what levels of autonomy and authority should the automation be allowed to operate?

This leads to the second challenge; is it possible to exploit the advantages of automation whilst *maintaining a competent and skilled work force*? Currently, air traffic controllers perform a high level of *manual, mostly tactical, control*, often with very limited automation support. Their situation awareness (SA) is built up by perception of the low level elements, comprehension of the current situation, and prediction of the future status [11], [12]. Studies show that implementing automation that replaces manual control can lead to reduced vigilance and out-of-the-loop performance problems [10], [13]. How then should the task allocation between the human controller and automation be structured, and can automation be designed in such a way that the controllers remain fully in-the-loop, allowing them to effectively identify problems and step in when intervention is required?

Third and finally, the ATM community has built up a reputation for being highly conservative in embracing new technologies, tools, and automation. Numerous initiatives have been undertaken to design and develop automated controller decision-support tools, but however failed to gain traction within the operational community [14]–[22]. In many cases, (the lack of) *controller ac-*

ceptance played a crucial role in the rejection of these technologies. Any new technology aimed at taking over (parts of) work from trained professionals will *need to be accepted by the users* as articulated in previous studies [23]–[26].

Given that the shift towards trajectory-based ATM will rely heavily on automation support, it is clear that the above issues need to be addressed in an *early design phase*. The overarching problem statement that forms the basis of this thesis is formulated as follows:

> Problem statement
>
> What information should be made salient to the human controller in a decision-support tool for tactical 4D trajectory-based control, in which form should this information be presented, and what levels of human-automation interaction should be supported?

1.3 Research Approach

The ultimate goal of this work has been to design prototype *ecological* decision-support interfaces for real-time, tactical, 4D trajectory-based air traffic control. Previous research in this field mainly explored *automation-centered* approaches, focusing on new conflict detection and resolution (CD&R) algorithms and/or highlighting (computed) problem areas and avoidance zones to controllers [14]–[22]. However, many of these efforts were met with mixed reactions by the operational community. The ecological approach sets itself apart as it focuses on translating the inherent physical and intentional work domain structure into *action-relevant control spaces* [27], [28]. In essence, providing a view of the ecology-based *Solution Space* for control, irrespective of how, or by whom subsequent *action-implementation* is performed.

The initial part of the work was undertaken within the SESAR Long-Term and Innovative Research (WP-E) project C-SHARE, that focused on identifying a common framework that can be used to underlie the design of both decision-support interfaces and automated tools. The research progressed to the design of the proposed decision-support interfaces named the *Solution Space* concept. Focus was on enabling effective human-automation collaboration by marrying the strong points of both; the unmatched computational power and accuracy of automation, combined with the unmatched creativeness and problem-solving capabilities of the human operator.

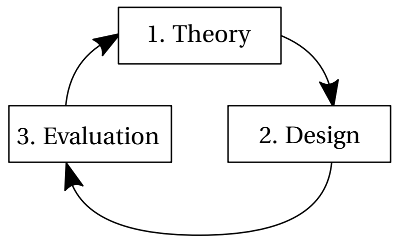

Figure 1.1: Schematic of the research cycle.

The approach taken in this research has been a highly iterative process, and followed the design cycle as illustrated in Fig. 1.1. Each step in the cycle is aimed at investigating and (partially) answering one of the three main research questions. The research questions are focused on the theoretical underpinnings of the Solution Space concept, the design and development of the interface concepts themselves, and experimental evaluations aimed at investigating various aspects of human performance and strategies. The results obtained and lessons learned in each consecutive cycle—either good or bad—allowed further refinement, modification, and expansion of the framework and interface concepts.

1.3.1 Research questions

Finding a common ground. The first part of this research is focused on identifying a *common ground* for tactical trajectory-based air traffic control that can support the mental model and decision making of both human controllers and automation support. Different from previous work on decision-support interfaces for trajectory-based ATC such as, most noteworthy, the PHARE GHMI project that focused on visualizing problem areas *resulting* from control actions [17], the approach in this thesis aims to describe the *action-relevant* Solution Space that underlies all feasible and safe control actions. The same definition of Solution Space can then be used as a basis for *shared cognition* between human operators and automation support alike [29], [30]. The approach follows the frameworks of *Cognitive Work Analysis* (CWA [31], [32]) and *Ecological Interface Design* (EID [33], [34]). Here, instead of taking any existing and/or available automation as the starting point for systems design, a top-down decomposition of the work domain is made, initially independent of any specific solution or division of roles between humans and automated actors.

In CWA/EID, the functional decomposition is performed in a *Work Domain Analysis* (WDA [32], [34]). As part of the WDA, a so-called *Abstraction Hierarchy*

(AH [35], [36]) is often used as a useful template to lay-out the work domain on various levels of abstraction; from the higher level (abstract) system goals such as *safety* and *efficiency*, to the lower level (physical) elements such as *sector shape and size* and *aircraft intent*. Together with a set of hierarchical *means-end links* that connect the elements in a top-down manner, the AH encompasses the functions, relationships, and constraints inherent to the work domain, and defines the joint playing field that bounds both human and automated actions.

The elements and interrelations in the Abstraction Hierarchy can then be integrated and translated into *action-relevant control spaces*. Fig. 1.2 shows a schematic representation of such an *action space* that bounds all actions from an initial state to reach a desired target state. The action space is ultimately constrained by a *physical* boundary that can or may not be crossed, and is governed by constraints innate to the work domain and the laws of physics (e.g., terrain and other traffic, aircraft performance envelope). *Intentional* constraints following from rules, regulations, and design choices further constrain this space by adding additional safety margins (e.g., separation buffers, additional margins within the performance envelope). The resulting action space then represents the control space within which any control action—either human or automated—will lead to a safe and feasible solution, regardless of optimality.

Based upon this common ground, CWA and EID then progresses to consider the human user and the human cognitive processes. By making the elements of the action space salient to the operator, and using *the same* elements

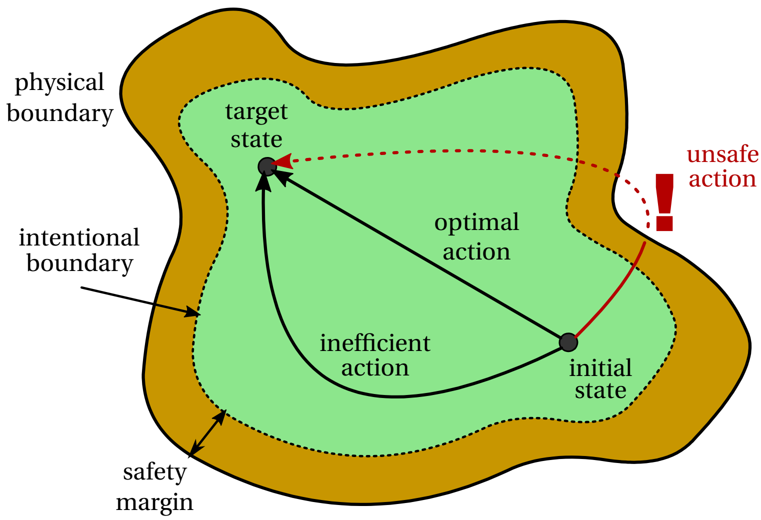

Figure 1.2: Schematic representation of an action space.

as a basis for automation support, build-up of situation awareness is promoted, and human-automation mismatches (i.e., what is it doing?) can be mitigated in the early design phase. Emphasis is put on supporting various levels of decision making, from low level skill-based behavior to more complex knowledge-based reasoning. In many cases, this approach leads to new insights, novel forms of automation, and a workflow that better adheres to the cognitive affordances of the human operator.

The ecological approach has been explored and applied successfully in the design of decision-support interfaces in a variety of other complex socio-technical domains (process control [32], health care [37], and aviation [38]–[41]). Most relevant to this research was the design of an ecological decision-support interface for 4D-trajectory re-planning on the flight-deck [42]. That research focused on a similar concept of Solution Space as the action space for pilots to re-plan their *own* trajectory in order to meet their own *aircraft-centric* goals. However, unlike previous studies that focused on *self-organizing, distributed* control, the work in this thesis aims to investigate *if* and *how* the ecological approach can effectively be applied in a *centralized* environment; the explored concepts focus on the Solution Space of individual aircraft *within* the constraints of the centralized ATC environment. The question then is, will controllers be able observe and integrate the Solution Spaces of the individual aircraft to reason upon, and formulate their control strategies within the broader context of the centralized environment?

Given the fundamentally different nature of the centralized control task, and in order to find the common ground for 4D trajectory-based air traffic control, the first research question has been formulated as follows:

> Research question 1
>
> What are the underlying system goals, functions, relationships, and constraints inherent to the work domain of tactical 4D trajectory-based air traffic control, and what are the implications on interface design?

Designing an ecological interface. The second part of this research is focused on the visualization of the Solution Spaces and the design of ecological decision-support interfaces for tactical trajectory-based air traffic control. The concepts presented in this work are inspired by EID, which aims to make work domain constraints (e.g., laws of physics, rules and regulations) salient on an interface in such a way that operators can directly perceive and act within the *action-relevant control space*. Here, the WDA provides guidance in identifying what functions and constraints should somehow be made visible, however, it does not provide a *ready-made recipe* for the shape and form in which that information should be presented.

Ecological interfaces aim to integrate goal-oriented information by adhering to the means-end links between the various levels of abstraction, enabling the operator to monitor overall conformance at a glance, but also to zoom in to lower levels in more detail. Such interfaces do not impose a specific set of actions or strategies, but rather show the complete task-relevant action space bound by the ecology. This allows the operator to remain *actively* in control, acknowledges *individual differences* between humans, and allows them to utilize their own *expertise*, *skills*, and *strategies* which has been found to be crucial for acceptance [43].

Although the Solution Space concepts explored in this research focus on visualizing the control space for manipulating *individual* aircraft trajectories within the constraints of the centralized ATC environment, the task of air traffic controllers is ultimately one of centralized control. Unlike aircraft-centric trajectory re-planning on the flight deck that entails modifying and/or optimizing one's own flight path, the task of air traffic controllers is to manipulate the trajectories of *multiple* aircraft to achieve the overarching system goals (i.e., safe and efficient operations). This will require the controllers to observe and integrate the Solution Spaces of multiple aircraft in order to make trade-offs and formulate strategies. Therefore, it is crucial that the mechanization of the interface supports the controller to reason about the effect of individual control actions within the centralized system.

The translation of Solution Spaces to visual representations and the interface design and mechanization itself, however, remains a creative process in which *subjective* trade-offs and design choices are made. This entails the visual presentation of information, human-automation interaction, and control task support. This leads to the second research question:

> Research question 2
>
> In what visual form can the constraints and functions that govern the tactical 4D trajectory-based control task best be mapped to a decision-support interface such that it supports the work and strategies of air traffic controllers as best as possible?

Evaluating the concept. To close the research cycle, the third part of this work will focus on human-automation interaction with (partial) implementations of the Solution Space decision-support concept. That is, how effective are the ecological Solution Space interface concepts for the centralized task of air traffic control? Do they allow controllers to formulate and implement their own unique control strategies? How do operators with different levels of expertise use the interface? What control strategies can be identified under different airspace configurations?

A total of three experiments focus on human-automation interaction with the Solution Space concept for *en-route* air traffic control. Here, the first experiment will look at how controllers with different levels of *expertise* use the *horizontal* Solution Space concept. Special attention is put on the *robustness* of control in scenarios with different levels of airspace complexity. In other studies that have explored ecological displays on the flight deck, concerns have been articulated that pilots sometimes opted for solutions that were close to the border between safe and unsafe actions [27], [39], [40], [44], [45]. Previous research has also shown that limit-seeking control behavior can be reduced by explicitly visualizing rules, regulations, and procedures on ecological interfaces [46], effectively defining additional safety buffers within the control space itself. The second experiment will investigate how the robustness of control is affected by explicitly visualizing such *intentional constraints*. Finally, the last experiment will look closer at controller *strategies* using an integrated horizontal- and *time-based* Solution Space interface concept. In addition to the previous experiments that only supported horizontal trajectory manipulations, this interface also offers controllers the possibility to manipulate speeds along trajectories, enabling various different approaches to controlling the traffic.

The experiments in this work primarily focus on the differences in human-automation interaction between controllers with different levels of expertise. However, the experiments also aim to identify any issues and concerns with the Solution Space visualizations and interface concepts. The lessons learned in the initial experiments will be addressed in this thesis, however, other lessons will be important to be addressed in any subsequent iterations of the concept in future work. The third research question is:

> **Research question 3**
>
> How do human controllers with different levels of expertise use (partial) implementations of the Solution Space concept for tactical 4D trajectory-based control, what strategies do they apply, and what lessons can be learned?

1.4 Research Scope

The work presented in this thesis is divided into three parts as illustrated in Fig. 1.3. In the first part (❶), the three main time-horizons of airspace organization and refinement in trajectory-based operations are laid out. The implications thereof to both the division of roles between the human and automated actors, and on the design of automation and decision-support tools are discussed. The second part (❷) focuses on the design of the interface and the means of controller-automation interaction. Also, a set of novel metrics have been derived that were used to measure the trajectory-based *robustness* in the evaluation of the concept. The final part (❸) presents a set of three human-in-the-loop experiments that have been performed with partial implementations of the concept.

Due to the fundamental nature of the research, the scope of the work is bound by various assumptions and simplifications. For the design, development, and refinement of the Solution Space concept, the task of real-time, tactical air traffic control by a single ATCo, and in a single sector has been taken as a starting point. Primary focus in this thesis is on the *en-route* ATC environment, however, the underlying concept does not limit its extension to other ATC domains. The main assumptions and simplifications are categorized below.

Airspace and traffic properties. A schematic overview of the work domain of tactical *en-route* 4D air traffic control is illustrated in Fig. 1.4. Traffic is controlled within a fixed sector, bounded by a border that contains multiple fixed entry and exit waypoints. A fully mature scenario of trajectory-based op-

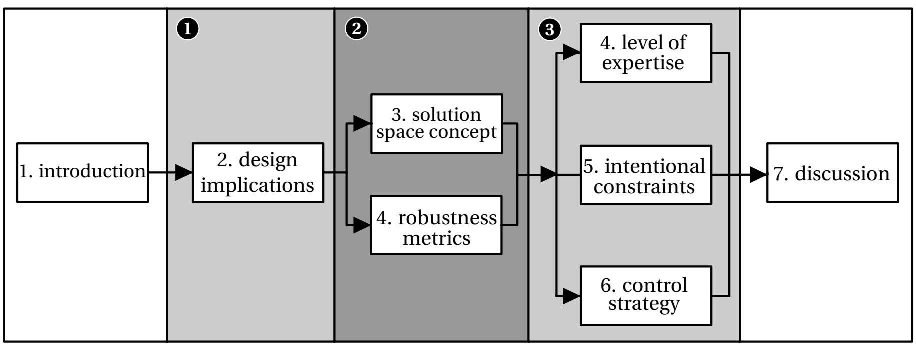

Figure 1.3: Schematic overview of the thesis structure. ❶ Theoretical background and implications on decision-support design. ❷ Solution Space concept and derivation of metrics to evaluate trajectory-based robustness. ❸ Human-in-the-loop experiments.

erations is assumed. That is, all controlled aircraft flying through the sector are assumed to be fully 4D- and data-link capable. A further assumption is that the airspace- and route allocation have been planned, optimized, and deconflicted *a-priori*. Initially, all trajectories are—in principle—conflict free, and all aircraft are expected to enter and exit the sector through a given waypoint pair at pre-planned times (metering fixes). The vertical degree of freedom has been omitted, and only lateral- and speed modifications are considered. Finally, all aircraft are assumed to closely adhere to their cleared trajectories, external environmental factors such as wind are not taken into account (i.e., deterministic, no uncertainty).

Controller task. The main controller task is to monitor the progress of flights, and to identify and revise cases where operational *perturbations* such as delays, hazardous weather, and restricted airspace result in conflicting, unsafe, or infeasible trajectories. For example, the trajectory of the observed aircraft in Fig. 1.4 conflicts with a restricted area inside the sector requiring action. The controller must then manually intervene and re-route that aircraft around the restricted area, whilst maintaining sufficient separation with other traffic and meeting the required sector exit time. In the design of the Solution Space concept, emphasis has been placed on supporting control actions that adhere to the original planning as close as possible to prevent cascading effects in adjacent sectors. That is, decision-support is aimed at making the feasible *control space* salient that leaves the sector exit parameters (i.e., sector hand-off location and time) unchanged.

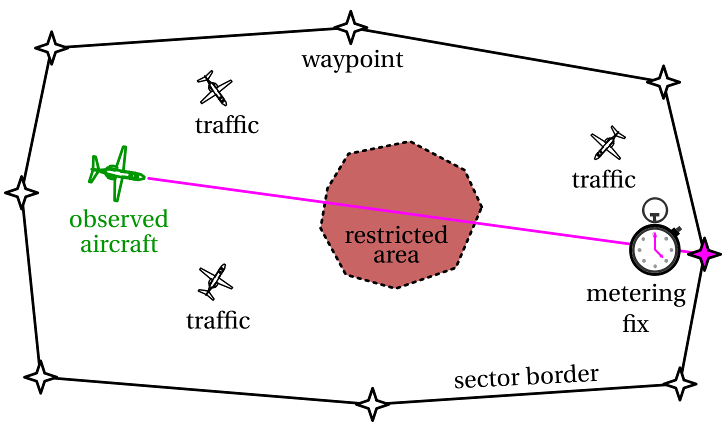

Figure 1.4: Schematic overview of the en-route, tactical, ATC work domain that illustrates the task intended to be supported.

Experiment limitations. The human-in-the-loop experiments presented in this research have been short-term experiments with partial implementations of the Solution Space concept. No longitudinal studies have been performed. The participants consisted of a varying mix of active ATCos, ATC domain experts, scientific staff, and M.Sc. or Ph.D. students. Experiment conditions, sector size and shape, and traffic samples do not directly reflect real-world operations and standards, but have been tailored to elicit decision-making and control by the participants.

1.5 Thesis Outline

This thesis consists of seven chapters that describe the design and refinement of the ecological Solution Space concept for 4D trajectory management. Chapter 2 (Shared Cognition in Air Traffic Management) and Chapter 4 (Expertise, Strategies and Robustness) are adaptations of published, peer-reviewed journal papers. Chapter 3 (Solution Space Concept) and Chapter 5 (Intentional Constraints) are adaptations of peer-reviewed conference papers. Chapter 6 (Evaluating the Integrated Solution Space Concept) presents results of a to-be published human-in-the-loop experiment with the final, integrated display concept. Chapters are introduced with a short introduction of the relevance of the work within the context of the overall thesis, together with original paper title, co-authors and publication details (where relevant).

Chapter 2: Shared Cognition in Air Traffic Management. This chapter outlines a step-wise approach to the definition and refinement of 4D trajectories in Trajectory-Based Operations (TBO), and discusses the implications thereof on automation design and the design of human-centered decision support tools. A Work Domain Analysis (WDA) is presented in the form of an Abstraction Hierarchy (AH) for three different time-horizons of the refinement of the airspace structure and 4D trajectories. The implications on display design and on the division of tasks between the human operator and automation support is highlighted. The chapter acts as the main literature survey and theoretical foundation of this thesis.

Chapter 3: Solution Space Concept. The design of the integrated Solution Space concept for 4D trajectory management is presented in this chapter. A short overview of relevant previous work is given, followed by the theoretical foundations for the elements and interactions incorporated in the interface. Next, the mapping of these elements into action-relevant control spaces is presented. A control task analysis has been performed by using the Skills-, Rules-, and Knowledge (SRK) taxonomy to identify the likely various modes of human-interaction with the interface.

Chapter 4: Expertise, Strategies and Robustness. This chapter presents the first human-in-the-loop experiment with a partial implementation of the Solution Space concept. First, the horizontal Solution Space representation is briefly introduced, followed by an investigation of the various control strategies by means of a decision ladder and an information flow map. A new metric is presented to quantify how robust control actions are to further system disruptions. The set-up of the human-in-the-loop experiment is presented, and the results of the experiment are discussed.

Chapter 5: Intentional Constraints. In this chapter, the human-in-the-loop experiment of the previous chapter is repeated with the addition of 'intentional' constraints to the original horizontal Solution Space. Results of a fast-time simulation are presented that quantify the effect of the conflict angle and 'additional separation' on the robustness of control actions. Further, the design choices of visual mapping of the intentional constraints to the original horizontal solution space display are discussed. The set-up of the second human-in-the-loop experiment is presented, together with the experiment results.

Chapter 6: Control Strategies. This chapter presents the final human-in-the-loop experiment with the *enhanced* Solution Space concept that featured both the horizontal and time-based Solution Space representations. The chapter and experiment are focused on the various control strategies that can be applied with the interface. The chapter starts with an analysis of the effect of both horizontal and time-based conflict resolution strategies on the efficiency of resulting trajectories. Next, an overview of the experiment interface and its elements is provided. The experiment set-up is described in detail, followed by a presentation of the results.

Chapter 7: Discussion and Conclusions. Discussions and conclusions of previous chapters are summarized and combined, to provide an overarching discussion on the design and employment of ecological air traffic controller decision-support tools for 4D trajectory-based operations. Limitations of this study, together with real-world implications and recommendations for future work are provided. The chapter ends with a set of conclusions that address the main research question.

Bibliography

[1] "State of Global Aviation Safety 2020," International Civil Aviation Organization, Montréal, QC, Canada, Report, 2020.

[2] "EUROCONTROL Forecast Update 2021-2027, European Flight Movement and Service Units, Three Scenarios for Recovery from COVID-19, October 2021," EUROCONTROL, Brussels, Belgium, Report, 2021.

[3] "Global Air Navigation Plan 2016-2030," International Civil Aviation Organization, Montréal, QC, Canada, Report, 2016.

[4] "SESAR 2020 Concept of Operations Edition 2017," EUROCONTROL, Brussels, Belgium, Report, 2017.

[5] D. C. Crespo and P. M. De Leon, *Achieving the Single European Sky: Goals and Challenges*. Alphen aan den Rijn, The Netherlands: Kluwer Law International, 2011.

[6] M. Baumgartner and M. Finger, "The Single European Sky gridlock: A difficult 10 year reform process," *Utilities Policy*, vol. 31, pp. 289–301, 2014.

[7] "Concept of Operations for the Next Generation Air Transportation System, Version 3.2," Joint Planning and Development Office, Washington, DC, USA, Report, 2011.

[8] "European ATM Master Plan 2020 Edition: Executive View," SESAR Joint Undertaking, Brussels, Belgium, Report, 2019.

[9] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

[10] M. R. Endsley and E. O. Kiris, "The Out-of-the-loop Performance Problem and Level of Control in Automation," *Human Factors*, vol. 37, no. 2, pp. 381–394, 1995.

[11] M. R. Endsley and M. D. Rodgers, "Situation awareness information requirements for en route air traffic control," DOT/FAA, Lubbock, TX, USA, Report, 1994.

[12] M. R. Endsley and D. B. Kaber, "Level of Automation Effects on Performance, Situation Awareness and Workload in a Dynamic Control Task," *Ergonomics*, vol. 42, no. 3, pp. 462–492, 1999.

[13] C. E. Billings, "Human-Centered Aircraft Automation: A Concept and Guidelines," NASA Ames research Center, Moffet Field, CA, USA, Report, 1991.

[14] M. S. Eby, “A Self-organizational Approach for Resolving Air Traffic Conflicts,” *The Lincoln Laboratory Journal - Special issue on air traffic control*, vol. 7, pp. 239–254, 2 1994.

[15] H. R. Idris, R. Vivona, J.-L. Garcia-Chico, and D. Wing, “Distributed Traffic Complexity Management by Preserving Trajectory Flexibility,” in *Proceedings of the 26th IEEE/AIAA Digital Avionics Systems Conference*, 2007, pp. 1–13.

[16] H. R. Idris, T. El-wakil, and D. J. Wing, “Trajectory Planning by Preserving Flexibility: Metrics and Analysis,” in *Proceedings of the AIAA Guidance, Navigation, and Control (GNC) Conference and Exhibit*, 2008, pp. 1–14.

[17] M. J. Whiteley and I. Wilson, “PHARE Advanced Tools Problem Solver,” EUROCONTROL, Brussels, Belgium, Final Report, 1999.

[18] T. Prevot, P. U. Lee, N. Smith, and E. A. Palmer, “ATC Technologies for Controller-Managed and Autonomous Flight Operations,” in *Proceedings of the AIAA Guidance, Navigation, and Control (GNC) Conference and Exhibit*, 2005, pp. 1–43.

[19] J.-P. Nicolaon, G. Dean, and W. Miller, “An Overview of ARC 2000 Version 3 from the Operational Point of View,” EUROCONTROL experimental center, Bretigny-sur-Orge, France, Tech. Rep. 286, 1995.

[20] L. Goldmuntz, J. T. Kefaliotis, L. A. Kleiman, R. A. Rucker, S. L., and D. Weathers, “The AERA Concept,” DOT/FAA, Washington D.C., USA, Tech. Rep. FAA-EM-81-3, 1981.

[21] W. P. Niedringhaus, “A Mathematical Formulation for Planning Automated Aircraft Separations for AERA 3,” DOT/FAA, Washington D.C., USA, Tech. Rep. DOT/FAA/DS-89/20, 1989.

[22] H. Erzberger, T. A. Lauderdale, and Y. C. Chu, “Automated Conflict Resolution, Arrival Management, and Weather Avoidance for Air Traffic Management,” in *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, Institution of Mechanical Engineers, 2012, pp. 930–949.

[23] M. van Gool and H. Schroter, “PHARE Final Report,” EUROCONTROL, Brussels, Belgium, Tech. Rep. DOC 99-70-09, 1999.

[24] C. A. L. Westin, C. Borst, and B. G. Hilburn, “Strategic Conformance: Overcoming Acceptance Issues of Decision Aiding Automation?” *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 41–52, 2016.

[25] M. Bekier and A. M. Williamson, "Why Air Traffic Controllers Accept or Refuse Automated Technology," in *16th International Symposium on Aviation Psychology*, 2011, pp. 615–620.

[26] M. Bekier, B. Molesworth, and A. M. Williamson, "Tipping point: The narrow path between automation acceptance and rejection in Air Traffic Management," *Safety Science*, vol. 50, no. 2, pp. 259–265, 2012.

[27] C. Borst, J. M. Flach, and J. Ellerbroek, "Beyond Ecological Interface Design: Lessons From Concerns and Misconceptions," *IEEE Transactions on Human-Machine Systems*, vol. 45, no. 2, pp. 164–175, 2015.

[28] M. M. Van Paassen, J. Ellerbroek, M. Mulder, and J. M. Flach, "Ecological Interface Design for Vehicle Locomotion Control," *IEEE Transactions on Human-Machine Systems*, vol. 48, no. 5, pp. 541–555, 2018.

[29] E. Hollnagel and D. D. Woods, "Cognitive systems engineering: new wine in new bottles.," *International journal of humancomputer studies*, vol. 51, no. 2, W. Karwowski, Ed., pp. 339–356, 1999.

[30] E. Hollnagel and D. D. Woods, *Joint Cognitive Systems*, E. Hollnagel and D. D. Woods, Eds. Taylor & Francis Group, LLC, 2005, ch. 6, pp. 113–133.

[31] J. Rasmussen, A. M. Pejtersen, and L. P. Goodstein, *Cognitive Systems Engineering*, ser. Wiley series in systems engineering October. Wiley-Interscience, 1994, p. 396.

[32] K. J. Vicente, *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. Mahwah, NJ: Lawrence Erlbaum Associates, 1999, pp. 1–415.

[33] K. J. Vicente and J. Rasmussen, "Ecological Interface Design: Theoretical Foundations.," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 22, no. 4, pp. 589–606, 1992.

[34] C. M. Burns and J. Hajdukiewicz, *Ecological Interface Design*. Boca Raton, FL, USA: CRC Press, 2004.

[35] J. Rasmussen, "The Role of Hierarchical Knowledge Representation in Decision Making and System Management," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 15, no. 2, pp. 234–243, 1985.

[36] ——, *Information Processing and Human-machine Interaction: An Approach to Cognitive Engineering*, A. P. Sage, Ed., ser. North-Holland series in system science and engineering. New York, NY, USA: Elsevier Science Publishers, 1986, vol. 12.

[37] T. R. McEwen, J. M. Flach, and N. Elder, "Interfaces to Medical Information Systems: Supporting Evidence Based Practice," in *Proceedings of the 2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, San Diego, CA, USA, 2014, pp. 1–6.

[38] M. H. J. Amelink, M. Mulder, M. M. Van Paassen, and J. M. Flach, "Theoretical foundations for a total energy-based perspective flight-path display," *The International Journal of Aviation Psychology*, vol. 15, no. 3, pp. 205–231, 2005.

[39] C. Borst, H. C. H. Suijkerbuijk, M. Mulder, and M. M. Van Paassen, "Ecological Interface Design for terrain awareness," *The International Journal of Aviation Psychology*, vol. 16, no. 4, pp. 375–400, 2006.

[40] S. B. J. van Dam, M. Mulder, and M. M. Van Paassen, "Ecological Interface Design of a tactical airborne separation assistance tool," *IEEE Transactions on Systems, Man and Cybernetics, Part A*, vol. 38, no. 6, pp. 1221–1233, 2008.

[41] J. Ellerbroek, M. Visser, S. B. J. van Dam, M. Mulder, and M. M. Van Paassen, "Design of an Airborne Three-Dimensional Separation Assistance Display," *IEEE Transactions on Systems, Man & Cybernetics, Part A*, vol. 41, no. 5, pp. 863–875, 2011.

[42] B. J. A. Van Marwijk, C. Borst, M. Mulder, M. Mulder, and M. M. Van Paassen, "Supporting 4D Trajectory Revisions on the Flight Deck: Design of a Human-Machine Interface," *The International Journal of Aviation Psychology*, vol. 21, no. 1, pp. 35–61, 2011.

[43] B. G. Hilburn, C. A. L. Westin, and C. Borst, "Will Controllers Accept a Machine That Thinks Like They Think? The Role of Strategic Conformance in Decision Aiding Automation," *Air Traffic Control Quarterly*, vol. 22, no. 2, pp. 115–136, 2014.

[44] C. Borst, M. Mulder, and M. M. Van Paassen, "Design and Simulator Evaluation of an Ecological Synthetic Vision Display," *Journal of Guidance, Control, and Dynamics*, vol. 33, no. 5, pp. 1577–1591, 2010.

[45] J. Ellerbroek, K. C. R. Brantegem, M. M. Van Paassen, N. de Gelder, and M. Mulder, "Experimental Evaluation of a Coplanar Airborne Separation Display," *Human-Machine Systems, IEEE Transactions on*, vol. 43, no. 3, pp. 290–301, 2013.

[46] J. Comans, "Visualizing Rules, Regulations, and Procedures in Ecological Information Systems," Ph.D. Dissertation, Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands, 2017.

Chapter 2

Shared Cognition in Air Traffic Management

As machines become more and more efficient and perfect, so it will become clear that imperfection is the greatness of man. —*Ernst Fischer (Journalist, Writer, Politician, 1899–1972)*

This chapter outlines a step-wise approach to the definition and refinement of 4D trajectories in three distinct phases of Trajectory-Based Operations (TBO). The implications on the division of tasks between the human operator and automation support are discussed, and the implications on automation design and the design of human-centered decision support tools are highlighted. The concept of Travel Space is introduced that formed the basis for the Solution Space construct. The chapter acts as the main literature survey and theoretical foundation of this thesis.

Paper Title	Adapted from: Designing for Shared Cognition in Air Traffic Management
Authors	M.M. van Paassen, C. Borst, R.E. Klomp, M. Mulder, P. van Leeuwen, and M. Mooij
Published in	Journal of Aerospace Operations, Vol. 2, Issue.1-2, 2013, pp. 39-51.

ABSTRACT

It is to be expected that the task of an air traffic controller will change with the introduction of four-dimensional (space and time) trajectories for aircraft, as can be seen in ongoing developments in ATM systems in Europe (SESAR) and the US (NextGen). It is clear that higher levels of automation will need to be developed to support the management of four-dimensional trajectories, but a definite concept on a distribution of the roles of automation and human users has not yet been well defined. This chapter presents one approach to the design of a shared representation for 4D trajectory management. The design is based on the Cognitive Systems Engineering framework and by using a formative approach in the analysis of the work domain, a step-wise refinement in the planning and execution of 4D trajectories is proposed. The design is described in three Abstraction Hierarchies, one for each phase in the refinement. The ultimate goal is to design a shared representation that underlies both the design of the human-machine interface and the rationale that guides the automation. It is foreseen that such a shared representation will greatly benefit the shared cognition in ATM and allows shifting back and forth across various levels of automation. Further work will focus on the refinement of the work domain analysis for the tactical phase of 4D trajectory management, and the development and validation of a joint cognitive decision-support interface.

2.1 Introduction

CURRENTLY, air traffic controllers (ATCos) perform a sector-based, tactical form of control. They are responsible for planning and managing traffic within their assigned airspace, often with little help from automated tools [1]. In the coming decades, the task of an air traffic controller is predicted to undergo a large transformation. The pull for transformation comes from the increasing demands which are placed on the air traffic management (ATM)-system [2]–[5]. A push is provided by technological advances on the air- and ground side of the ATM-system, which make a new form of air traffic control (ATC) possible [6], [7]. This is expected to result in a situation where aircraft four-dimensional (4D, i.e., space and time) trajectories stored in automated support tools form the basis for the ATCos work [8]–[11].

A fundamental difference between current practice and future air traffic management is the explicit use of a 4D definition of aircraft trajectories (4DT) as a shared representation between air and ground segments. In SESAR, this 4DT has been defined as a Reference Business Trajectory (RBT) [4]. Supported by a communications network, the System Wide Information Management (SWIM), the definition of the 4DT is to be shared, such that all parties involved have access to relevant and the most up-to-date flight information. This will enable more accurate planning, prediction and monitoring of traffic movements, resulting in a higher capacity and more efficient use of the available airspace.

Given that the stakes are high, and that the ATM work domain inevitably has too many unforeseen situations to create a fully automated solution (i.e., the work domain can be characterized as "open", [12]), human users will have to remain actively involved in the system. The future air traffic controller will not, as he or she is doing currently, provide hands-on instructions to the aircraft, in essence creating the aircraft's 4D trajectory in real time. Rather, controllers will work on (refining) a pre-planned definition of the aircraft 4D trajectory (4D flight plan), visualized and represented to them by automated tools. Future or modernized aircraft will have the capability to receive 4D trajectory updates on the flight deck through data-link, and implement their flight according to this trajectory with a high degree of precision.

Although considerable research has been devoted to exploring this future approach to air traffic control with 4D trajectory support [13]–[17], a definite concept on a distribution of the roles of automation and human users has not yet been defined. It is clear that higher levels of automation will need to be developed, but the 'central role' of the human operator is not well defined yet. Many other complex socio-technical domains have shown that the introduction of higher levels of automation often introduce new problems, problems that are often harder to resolve than those intended to be solved in the first place [18]. Increasing the level of automation is not good or bad in itself, but with more automation, greater coordination between people and technology will be required[19]. Breakdowns in coordination may result in humans having difficulty getting the automation to do what they want and, conversely, a poor understanding of the way automation works [20]. To mitigate breakdowns in coordination between human and automated agents, it is imperative to create new tools for coordination so as to make automated systems 'team players'.

This chapter explores one possible design approach for facilitating a successful human-automation collaboration in a future 4DT-based ATM system. In the SESAR WP-E project 'C-SHARE' a Cognitive Work Analysis (CWA) approach is adopted [21], [22]. CWA starts from an analysis of the work domain, identifying goals and functions in this work domain, and in a design, it is possible to start top-down, initially independent of the chosen solutions for the system. The ultimate goal of the project is to find a "common ground" or shared representation for manipulating 4DTs that can be used to develop the automated tools as well as the human-machine interfaces. It is foreseen that a shared representation for automation and humans will greatly benefit the shared cognition in ATM. The chapter outlines a step-wise approach to the definition and refinement of 4DTs and discusses some of the implications on automation and display design.

2.2 The Structure and Function of Airspace

The re-engineering of the air traffic management system is a design process, which will be approached here following the paradigm of CWA. That means that the first step in CWA, the Work Domain Analysis (WDA), will be started in a top-down fashion. Part of the WDA will reflect the constraints innate in the work domain itself, for example the fact that aircraft need to have sufficient clearance from terrain and other aircraft (separation). However, other functions in the WDA are influenced by the design choices, both for the current system and for the envisaged new ATM systems.

2.2.1 Work domain analysis

The work domain analysis will be done by constructing an Abstraction Hierarchy (AH) [20]. In constructing an AH for a new domain, the main challenge is to select the proper choice for the abstract functions in this domain. In process technology and energy generation systems, where CWA originated, the abstract functions that describe the energy and mass balances in a system form an appropriate choice [12]. In the description of a single vehicle, the WDA at this level focuses on locomotion and on (potential and kinetic) energy balances [23]. For the case of ATM, the principal functions at this level are proposed to be identified as locomotion, localization, communication, separation, and a principle we refer to as "travel space".

Locomotion is a function of the moving elements in the ATM system, realized by flight for aircraft, and drifting for weather. *Localization* is the function that determines the position of these moving elements, either on-board, by the navigation system, or on the ground, by the ATM surveillance systems. *Communication* supports localization and decision making in the system by sharing intentions, plans, and localization results. *Separation* is the principal means for safety in the ATM system: at all times a proper clearance to other aircraft, terrain and hazardous weather must be maintained.

The identification of *travel space* as a separate functionality in this analysis warrants additional explanation. We define travel space as the function offered by the air and infrastructure to the moving elements in the ATM system – the aircraft – to implement their locomotion. Other elements in the system, such as weather, terrain and including other aircraft, affect the possibility to use the available airspace in certain ways [24], [25]. That is, for a given aircraft fixed (terrain) and moving (other aircraft, weather) obstacles limit the maneuvering possibilities for that aircraft both in space and time. Travel space is by definition four-dimensional, and the way its constraints can be managed stands at the heart of (current and future) ATM.

2.2.2 Travel space function constraints

Identifying the possibilities for *travel* as a function in our analysis enables us to investigate the implications of 4D operations in our design of new ATM systems. Many constraints in this function are unavoidable; removing them would require removing terrain or other traffic. However, the solutions chosen for our ATM system, such as the communication and navigation systems, the legal infrastructure and the way in which we plan and coordinate trajectories, affect the shape and characteristics of the travel space function.

Communication limitations. Current ATM mainly uses voice communication. To enable efficient use of this communication channel which is limited in bandwidth – on the other hand, it is extremely flexible – the actors taking part in this communication need to have agreed on extensive background information. This makes it possible to only use pre-defined and published waypoints and discrete altitude levels for defining flight plans. Digital data-link will allow the actors to exchange fully-defined flight plans with custom 4D waypoints, allowing for more flexible use of the airspace.

Navigation systems. Traditionally, limitations of navigation systems provided constraints on where flight was possible. In the early days of commercial aviation, railroads and other landmarks formed the basis for the air structure. Later, radio navigation aids, such as the four-course radio range, VOR and NDB beacons largely determined the use of airspace. The navigation aids thus determined which parts of the airspace are usable as travel space, and how these can be used. Much of these restrictions have been removed as aircraft are increasingly able to perform satellite-based Area Navigation (RNAV), meaning that navigation can be performed without the direct need for physical land-based nav-aids.

Legal infrastructure. A further constraint on the locomotion is provided by the administrative organization of airspace. The (current) division in airspace sectors imposes restrictions on the paths of aircraft, basically because the handling and the transition of an aircraft from one sector to another requires a buffer zone between the sectors, and effort from the controllers and pilots. Aircraft trajectories are effectively constrained to transitions between sectors with more or less perpendicular angle to the sector boundaries. Short paths through sectors, such as the perpendicular traversal of narrow sectors, or passing through a corner of a sector, are difficult to manage and therefore uncommon.

Planning and coordination. Currently, the control of the traffic within an airspace sector is normally the job of a single ATCo, or of a small team of two to five. Support by automated tools is fairly limited, and the extent to which a 4D trajectory is known ahead of time is very limited. This forces an ATCo to impose additional structure on the use of airspace.

The technological advances in navigation systems and communications foreseen in SESAR and NextGen can remove part of the constraints on the travel space function, opening the way for more economical—direct—routes.

2.3 Operational Concept

2.3.1 Overview

This chapter sketches an operational concept for the future ATM system that largely uses the functionality foreseen in the ATM master plan [2]. In particular, the functionalities provided by 4D trajectory management and information-exchange with a SWIM system are combined in a concept that assumes a central role for the human actors in the system.

A chronological step-wise refinement of the 4D aircraft trajectories is proposed. Given the central role reserved for the human actors in this process, to enable them to contribute in the planning and management of 4D-based operations, proper support for visualizing, evaluating and modifying trajectories and airspace allocation is required. Also, the interdependencies and the amount of work in defining 4DTs for all aircraft in the ATM system is very large, and thus require automated systems to support the human users in this task.

Task division between humans and automation is often approached as an ‘allocation problem’; either the human actor *or* the automation is selected for a task. Prime examples for this can be found within the aircraft themselves;

the task of stabilizing the aircraft is normally allocated to the autopilot, and the navigation along a trajectory is performed by a combination of the autopilot following a flight plan entered by the pilot in the Flight Management System.

The first guiding principles in task allocation have been laid out in what is now known as Fitts' list [26]. In this concept, part of the tasks in the foreseen ATM concept are indeed *assigned* with these principles, such as the tactical monitoring for deviation between actual flights and the agreed 4D trajectories. However, other tasks are foreseen to be performed *jointly* by automation and humans, and some tasks can be done *in parallel* by automation and humans.

In most complex systems, however, many tasks are too ill-defined to be handled by automation. Such 'fuzzy' tasks are typically assigned to human operators. To support operators in those cases, a proper visualization of the problem space can help. Examples of such visualizations are the Ecological Interface Design for the example process system DURESS [27], or, more recently, visualizations for airborne traffic avoidance [28], [29], and 4D trajectory manipulations on the flight deck [30], [31]. These can be seen as automation support, where algorithms are used to visualize the work domain constraints in such a way that operators can implement appropriate control strategies. This leads to a task that can be performed *jointly* by automation and humans. The resulting cognition can be seen as a *joint* effort of the automation, and in particular the visualization of the problem, and the human user [32].

Within the SESAR overall operational concept, several stages in the refinement of 4D trajectories are foreseen. The design presented in this chapter focuses on three stages: (i) Short-term planning, performed 24 to 12 hours in advance of the flight; (ii) Pre-tactical planning, conducted several hours to 30 minutes in advance of the flight, and (iii) a Tactical monitoring phase (30 min-

(a) Short-term planning (b) Pre-tactical planning (c) Tactical monitoring

Figure 2.1: Summary of the stages in refinement and implementation of 4DTs. Only for the tactical control the actual aircraft flight data are used (radar symbol). For the pre-tactical and tactical control, assistance from automated agents is foreseen.

utes to now). In contrast, the full SESAR design starts with seasonal planning. A summary of the foreseen phases is given in Figure 2.1. The interaction foreseen in these three phases between users, their display and support tools and automated agents is discussed in the following subsections.

2.3.2 Short-term planning

Short-term planning – corresponding with SESAR terminology – takes place approximately 24 to 12 hours in advance. This phase starts with an inventory of intended flights, initially designed as the shortest and most economical route to the destination. A visualization will be used to show the use of airspace, including "hot spots", with high concentrations of traffic. The human planners use this representation to create a global structuring of the airspace (e.g., restricting the number of flights in certain areas, reserving altitudes for certain headings, making sure that there is 'spare' airspace to handle unforeseen disturbances or to re-structure the flows to be able to handle a change in runway at an airport, etc.). The function of the automation in this stage is mainly to provide visualization and identification of hot-spots.

The result of this stage is a planning of the airspace 'structure', i.e., the travel space will be partly pre-allocated. NextGen flow corridors [5] might be an example of this. The 4D trajectories are then modified by automated algorithms to conform to this structure resulting in an indirect de-confliction (e.g., to adhere to capacity limits defined for the airspace). However, overlapping conflicts that may exist between the 4D trajectories are not identified or resolved, since the 4D trajectories are not yet (or can not yet be) sufficiently defined to perform this step.

Part of the work domain analysis is given in the Abstraction Hierarchy in Figure 2.2. The work domain analysis describes the functionality and constraints inherent to the work domain, in this case of Air Traffic Management. An Abstraction Hierarchy describes one and the same system or work domain at different levels of abstraction [20]. The top level is the *functional purpose (FP)* level, containing the ultimate goals identified for the system. The *abstract function (AF)* level describes the basic principles and processes in the work domain that enable the realization of these goals. In this case, the basic mechanisms at work are obstruction (e.g. by weather) and allocation of space. The *generalized function (GF)* level further specifies this in terms of "systems solutions". Normally, an AH has two further – more detailed – levels, that are not yet specified for this study [33].

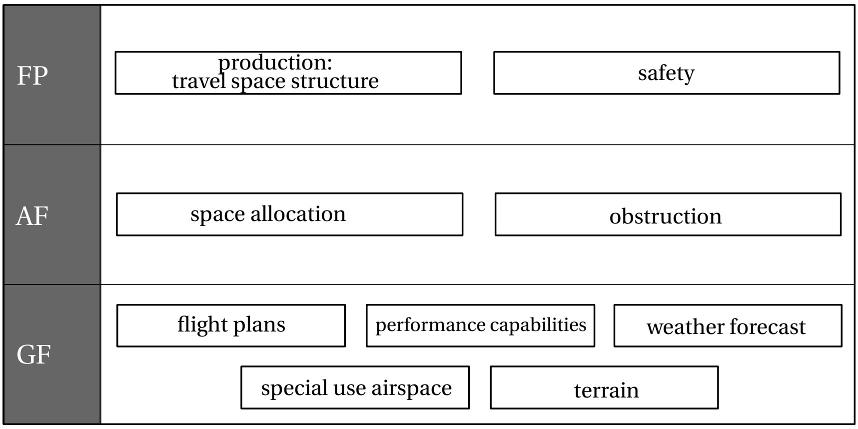

Figure 2.2: Short-term planning stage, top three levels of the AH.

The product that comes out of the short-term planning step is a "structure" for the travel space; choices are made to reduce traffic at places where large volumes of traffic are expected, and additional capacity is reserved where needed, for example as a contingency for weather phenomena. This planned structure should achieve the goals identified at the top level in the AH.

2.3.3 Pre-tactical planning

This takes place from several hours up to approximately half an hour in advance of current time. By *refining* the 4D trajectories, taking into account aircraft performance and weather, the 4DTs are defined in more detail to be "in principle" conflict free. The adjustments to 4DTs can be performed by human operators and automated agents in parallel. A proper visualization of the travel space functions are used as a template for the cognitive process; human operators can use this visualization to directly perceive the effects of path and speed manipulation. It also serves as a shared memory, offering a workspace to automation and human agents alike. The result of this stage is that the 4DTs are de-conflicted and in accordance with the airspace structure defined in the previous step.

Part of the work domain analysis for this stage is given in Figure 2.3. The airspace structure generated in the previous stage is now a generalized function. It defines a rough plan for the generation and modification of trajectories for individual flights, and it acts as an additional constraint in this analysis, imposing limits on flights but providing an overall means to simplify the planning process, analogous to the way the current airway structure is used to shape

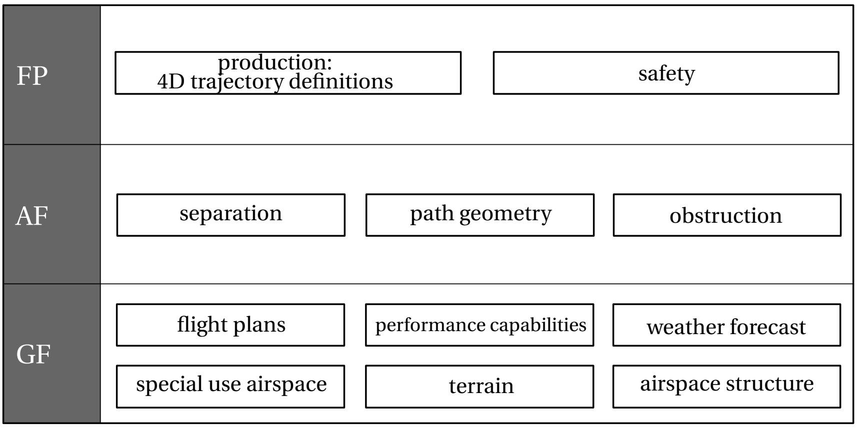

Figure 2.3: Pre-tactical planning stage, top three levels of the AH.

air traffic. Observing the required separation, possible obstruction (by terrain, weather, etc.) and the geometric constraints of each flight's path, this stage results in refined definitions for the 4DTs.

2.3.4 Tactical monitoring

At this stage the *planned* 4DTs of the different flights are conflict free. However, in the execution of flights, small deviations from these planned trajectories are expected to be unavoidable. Automated agents monitor the execution of the trajectories and provide limited solutions (e.g., speed and minor path adjustments) to keep the flights conflict free. The visualization now serves to inform the human users of the progress of the flights and of the actions of the automated monitoring agents. The situation awareness thus built up permits the human user to perform the higher – system – level monitoring, and to intervene when unforeseen circumstances make this necessary.

Part of the work domain analysis for this stage is given in Figure 2.4. At this stage, the physical function level (the next level below the generalized function) will be formed by the functionality related to the aircraft and physical devices in the ATM system. While the previous stages mainly involved planning flights (first globally, resulting in the travel space structure, then in more detail), the result from this stage are the actual flights. Real-time communication therefore becomes an important function in the tactical monitoring phase.

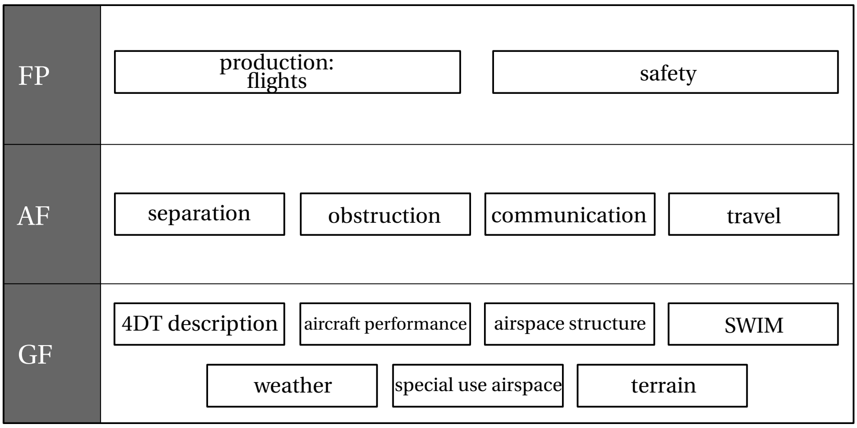

Figure 2.4: Final tactical monitoring stage, top three levels of the AH.

2.3.5 Joint cognitive representation

The joint cognition by the human users and automation differs in these three stages. In the short-term planning, the contribution by the automation is mainly the visualization. The human users are primarily responsible for the planning result. In the pre-tactical planning, the automation and human users contribute on a more or less equal basis. Here, the visualization serves as the representation of the commonly used (space) resources. The tactical monitoring phase most closely resembles current high levels of automation, with a large contribution of automated agents, utilizing a probabilistic road-map method, to the final solution.

The work domain analysis, which in current approaches to Ecological Interface Design (EID) serves as an input to the display design process only, will be used for both the design of the automation and the displays in this project (C-SHARE). That is, the constraints on the 4D trajectories as identified in the WDA are transformed into a representation that underlies both the design of the human-machine interface as well as ground- and air-automated tools.

By virtue of being based on a ‘shared’ representation, the results of the work done by automated tools will be compatible with the human’s representation of the work domain, and can be visualized in a human-centric fashion. With automation and humans working on a shared cognitive representation, this also allows shifting back and forth across various levels of automation, from fully manual control to fully automated control, and in principle also supports the transition between SESAR’s unmanaged and managed airspace.

2.4 Implications for Display Design

In EID, the analysis of the work domain is a primary input for the actual design of the display. However, the design of a display presentation is still a creative step, the WDA does not result in a "recipe" for how the display is to be created, it only provides guidance in determining what functions should somehow be made visible in the display [31]. The following first inventory of the important elements, and the way they might be visualized, is given here:

Short term. For the short-term planning of the travel space structure, obstruction and space allocation are considered primary functions at the abstract function level (Figure 2.2). The product of this stage should be the travel space structure for the next day, indicating how airspace will be allocated for flights, and where disruptions are expected and thus buffers are reserved. The input to this work is the set of flight plans as filed by airline companies and airspace availability. Important aspects of the visualization will be the obstruction, by weather cells, terrain, or temporary restricted airspace. A global visualization is needed of the *traffic flow* (not per 4D trajectory, but as a whole), as is a visualization of the means to modify this flow by structuring the travel space.

Pre-tactical. Pre-tactical planning should result in initial conflict-free 4DTs. The visualization should show the travel space structure created in the previous step. Since the planning is done in parallel by automated agents and human users, communication between the agents and humans on the ongoing work, and allocation of (space) resources is important (Figure 2.3). The result is mainly the path geometry of the individual 4DTs. At this stage, the constraints imposed by the aircraft performance capabilities and separation should be visible. An important feature of the display is the visualization of the relation between the possible modifications to the 4DTs and the effect on separation and performance.

Tactical. In the tactical planning, much of the actual work should be performed by automated agents. Flights that are operating on or near the 4DT defined in the previous stage can be monitored automatically. The visualization for the human user should enable checking of conformance to the 4DT at a glance. At this stage the detection of anomalies is important. Since the actual implementation depends on real-time communication, an indication of "communication health" for all flights should be given (Figure 2.4). Handling flights with problems, that need to be diverted from their route, requires a visualization of separation from other flights and of buffer zones that can be used to safely divert the flight.

2.5 Discussion

The largest change for an ATCo will be to step away from direct, hands-on control of aircraft trajectories, and to implement the desired traffic flow in real time in conformance with an operation in which traffic is planned in detail beforehand. For individual flights, it has proven possible to implement, monitor and manipulate 4D trajectories, usually in the context of all other aircraft being controlled traditionally. The case when *all* aircraft are to be controlled based on their 4DT means a tremendous step, and a real-time visualization of how all trajectories will evolve *in time* is a tremendous challenge for display designers. Whereas dimensionality of the control problem explodes, the visualization and display techniques remain limited by, among others, clutter issues, and physical constraints like screen size and resolution.

The main outcome of project C-SHARE will be a representation for the tactical and strategic manipulation of 4D trajectories. This framework for 'joint cognition' can act as a basis for designing both the automation support and human-machine interfaces, in the air and on the ground, from one and the same perspective. It is very likely that during the development and testing of prototypes, the Work Domain Analysis will need to be augmented and/or partly changed. Several human-in-the-loop experiments are foreseen that will show to be crucial in converging the design and analysis iterations to a representation of 4DT management that can indeed be used for both automation and human-machine interface design.

2.6 Conclusions

This chapter outlines a possible approach for the creation of the new work domain in Air Traffic Management. The envisaged future situation in SESAR and NextGen, in which aircraft will be able to fly four-dimensional (space and time) trajectories, requires planning, monitoring, and if necessary modification of these trajectories. The approach proposed in this paper is based on Cognitive Systems Engineering, and assumes three successive steps in the refining and final implementation of the four-dimensional trajectories. Automation support comes in the form of visualization of the constraints in the planning phase, and collaborating agents in the execution phase. An initial Work Domain Analysis has been performed for these three phases, and critical functions for each of the phases have been identified. Further work will focus on the refinement of the WDA, the creation of actual interface prototypes, and the evaluation of the design by means of human-in-the-loop experiments.

Bibliography

[1] E. A. P. B. Oprins and M. Schuver, "Competentiegericht opleiden en beoordelen bij LVNL [Competence based training and assessment at LVNL]," *HUFAG nieuwsbrief*, no. 6, pp. 2–4, 2003.

[2] O. Dlugi, T. Astheimer, C. Baldoni, J. M. Bara, R. Brown, R. Cato, G. O'Connell, M. Standar, J. Stenson, and G. Tauss, "The ATM Target Concept," SESAR Consortium, Toulouse, France, Deliverable 3, 2007.

[3] O. Dlugi, T. Astheimer, C. Baldoni, R. Brown, R. Cato, M. Erb, D. Hilton, G. O'Connell, A. Sohn, and J. Stenson, "The ATM Deployment Sequence," SESAR Consortium, Toulouse, France, Deliverable 4, 2008.

[4] "SESAR 2020 Concept of Operations Edition 2017," EUROCONTROL, Brussels, Belgium, Report, 2017.

[5] "Concept of Operations for the Next Generation Air Transportation System, Version 3.2," Joint Planning and Development Office, Washington, DC, USA, Report, 2011.

[6] R. A. Slattery, "Examination of Possible Future Terminal Area Procedures," in *Proceedings of the AIAA, Guidance, Navigation, and Control (GNC) Conference and Exhibit*, 1996, pp. 1–9.

[7] M. G. Ballin, J. M. Hoekstra, D. J. Wing, and G. W. Lohr, "NASA Langley and NLR Research of Distributed Air/Ground Traffic Management," in *AIAA's Aircraft Technology, Integration and Operations 2002 Technical*, 2002, pp. 1–13.

[8] M. J. Whiteley and I. Wilson, "PHARE Advanced Tools Problem Solver," EUROCONTROL, Brussels, Belgium, Final Report, 1999.

[9] P. G. A. M. Jorna, D. Pavet, M. Van Blanken, and I. Pichancourt, "PHARE Ground Human Machine Interface (GHMI) project," EUROCONTROL, Brussels, Belgium, Summary Report, 1999.

[10] W. W. M. Heesbeen, J. M. Hoekstra, and M. S. V. Valenti Clari, "Traffic Manager, Traffic Simulation for Validation of Future ATM Concepts," in *AIAA Modeling and Simulation Technologies Conference and Exhibit*, 2003, pp. 1–10.

[11] T. Prevot, P. U. Lee, N. Smith, and E. A. Palmer, "ATC Technologies for Controller-Managed and Autonomous Flight Operations," in *Proceedings of the AIAA Guidance, Navigation, and Control (GNC) Conference and Exhibit*, 2005, pp. 1–43.

[12] K. J. Vicente and J. Rasmussen, "Ecological Interface Design: Theoretical Foundations.," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 22, no. 4, pp. 589–606, 1992.

[13] D. Lindgren and A. Törne, "4D Trajectory Management Decision Support," FOI, Swedish Defense Agency, Linköping, Tech. Rep. FOI-R–2405–SE, 2008.

[14] R. Weber and C. Eva, "Subliminal ATC Utilizing 4D Trajectory Negotiation," in *7th ICNS Conference*, IEEE, Herndon (VA), 2007, pp. 1–9.

[15] M. Jackson, "Role of Avionics in Trajectory-Based Operations," *Aerospace and Electronic Systems Magazine, IEEE*, vol. 25, no. 7, pp. 12–19, 2010, ISSN: 0885-8985.

[16] S. Landry, "Human Centered Design in the Air Traffic Control System," *Journal of Intelligent Manufacturing*, vol. 22, pp. 65–72, 1 2011, ISSN: 0956-5515.

[17] D. McNally and C. Gong, "Concept and Laboratory Analysis of Trajectory-Based Automation for Separation Assurance," in *AIAA-MST-2006*, AIAA, Keystone, CO, 2006, pp. 1–17.

[18] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

[19] K. Christoffersen and D. D. Woods, "How to make automated systems team players," in *Advances in Human Performance and Cognitive Engineering Research*, E. Salas, Ed., 2nd ed., vol. 2, Elsevier Science Ltd, 2002, ch. 1, pp. 1–12.

[20] J. Rasmussen, *Information Processing and Human-machine Interaction: An Approach to Cognitive Engineering*, A. P. Sage, Ed., ser. North-Holland series in system science and engineering. New York, NY, USA: Elsevier Science Publishers, 1986, vol. 12.

[21] J. Rasmussen, A. M. Pejtersen, and L. P. Goodstein, *Cognitive Systems Engineering*, ser. Wiley series in systems engineering October. Wiley-Interscience, 1994, p. 396.

[22] K. J. Vicente, *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. Mahwah, NJ: Lawrence Erlbaum Associates, 1999, pp. 1–415.

[23] F. M. Heylen, S. B. J. Van Dam, M. Mulder, and M. M. Van Paassen, "Design of a Vertical Separation Assistance Display," *Proceedings of the AIAA Guidance, Navigation, and Control (GNC) Conference and Exhibit*, no. AIAA 2008-6969, 2008.

[24] M. M. Van Paassen, "Functions of Space and Travellers," *Proceedings of the XVIII European Annual Conference on Human Decision Making and Manual Control, Loughborough (UK)*, pp. 268–274, 1999.

[25] J. J. Gibson and L. E. Crooks, "A Theoretical Field Analysis of Automobile Driving," *The American Journal of Psychology*, vol. 51, no. 3, pp. 453–471, 1938.

[26] P. M. Fitts, "Human Engineering for an Effective Air-Navigation and Traffic-Control System," Ohio State University Research Foundation, Technical Report, 1951.

[27] K. J. Vicente and J. Rasmussen, "The Ecology of Human-Machine Systems II: Mediating Direct-Perception in Complex Work Domains," *Ecological Psychology*, vol. 2, no. 3, pp. 207–249, 1990.

[28] S. B. J. Van Dam, M. Mulder, and M. M. Van Paassen, "Ecological Interface Design of a Tactical Airborne Separation Assistance Tool," *IEEE Transactions on Systems, Man & Cybernetics, Part A*, vol. 38, no. 6, pp. 1221–1233, 2008.

[29] J. Ellerbroek, M. Visser, S. B. J. van Dam, M. Mulder, and M. M. Van Paassen, "Design of an Airborne Three-Dimensional Separation Assistance Display," *IEEE Transactions on Systems, Man & Cybernetics, Part A*, vol. 41, no. 5, pp. 863–875, 2011.

[30] M. Mulder, R. Winterberg, M. M. Van Paassen, and M. Mulder, "Direct Manipulation Interfaces for In-Flight Four-Dimensional Navigation," *International Journal of Aviation Psychology*, vol. 20, no. 3, pp. 249–268, 2010.

[31] B. J. A. Van Marwijk, C. Borst, M. Mulder, M. Mulder, and M. M. Van Paassen, "Supporting 4D Trajectory Revisions on the Flight Deck: Design of a Human-Machine Interface," *International Journal on Aviation Psychology*, vol. 21, no. 1, pp. 35–61, 2011.

[32] M. M. Van Paassen, M. Mulder, and S. B. J. Van Dam, "Cognitive Work Analysis for Vehicle Locomotion - Aircraft Short-Term Trajectory Planning," *Proceedings of the Nineth IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design and Evaluation of Man-Machine Systems*, vol. 46, p. 5, 2004.

[33] A. M. Bisantz and K. J. Vicente, "Making the Abstraction Hierarchy Concrete," *Int. J. Human-Computer Studies*, vol. 40, pp. 83–117, 1994.

CHAPTER 3

SOLUTION SPACE CONCEPT

The computer is incredibly fast, accurate, and stupid. Man is incredibly slow, inaccurate, and brilliant. The marriage of the two is a force beyond calculation. —*Leo M. Cherne (Economist, 1912–1999)*

The design of the integrated Solution Space concept for 4D trajectory management is presented in this chapter. The notion of Solution Space can be seen as the visualization of travel space as described in the previous chapter. An overview of relevant previous work is given, followed by the theoretical foundations for the elements and interactions incorporated in a human-machine interface. The mapping of these elements into action-relevant control spaces is discussed. Finally, a control task analysis is presented using the Skills-, Rules-, and Knowledge (SRK) taxonomy to identify the likely various modes of human-interaction with the interface.

Paper Title	Adapted from: Solution Space Concept: Human-Machine Interface for 4D Trajectory Management
Authors	R.E. Klomp, R. Riegman, C. Borst, M. Mulder, M.M. van Paassen
Published in	Proceedings of the Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2019).

ABSTRACT

The current evolution of the ATM system, led by the SESAR program in Europe and the NextGen program in the US, is foreseen to bring a paradigm shift to the work of the air traffic controller. In both programs, a key pillar is the shift from hands-on, tactical control by voice towards more strategic, 4D (time and space) trajectory-based management via digital data-link. There is a general consensus that the human operator should remain to play a central role, supported by higher levels of automation and advanced decision-support tools. Previous work in the design of such automated tools has seen promising results, however, also highlighted various critical issues that need to be addressed in order to gain acceptance by the operational community. This chapter presents a concept decision-support tool for 4D trajectory management that aims to overcome these issues. Following an approach based on the framework of ecological interface design, rather than imposing (a set of) discrete solutions, action-relevant Solution Spaces are directly visualized to the controller. These Solution Spaces represent all feasible control actions, regardless of their optimality. This approach leaves the controller fully in-the-loop, and allows them to apply their own—individual sensitive—control strategies.

3.1 Introduction

THE current evolution in the Air Traffic Management (ATM)-system is expected to result in a situation where high-precision, gate-to-gate, four-dimensional (4D, i.e., space and time) trajectories for aircraft, stored in automated support tools, will form the basis for the work of the human controller [1]–[3]. This new form of Air Traffic Control (ATC) implies a fundamental shift in the work of the Air Traffic Controller (ATCo), one that will no longer be possible without higher levels of automation and advanced decision-support tools.

Although a definite breakdown of the distribution of roles and coordination between the human operator and automation is not yet well defined, the controller is foreseen to remain to play a *central role* in monitoring and revising trajectories in the real-time—tactical—phase [3]. In this scenario, airspace and route-allocation will be structured and optimized in the *pre-tactical* (i.e., planning) phase beforehand to achieve optimal system performance in terms of safety, efficiency and productivity. However, system variance due to *unforeseen* delays, hazardous weather, temporarily restricted airspace, etc., will *inevitably* require (small) modifications to the pre-planned trajectories in the real-time execution phase.

In previous work, various prototypes of advanced automation support for ATC have been explored, but have not been embraced by the operational community [4]–[6]. According to Westin *et al.* [4], *controller acceptance* played a critical role, where acceptance is driven by how much the support tool *conforms* or matches with the skills and strategies of humans. Thus the challenge remains how to design a decision-support tool that facilitates 4D trajectory management, whilst accounting for the individual controllers' expertise and strategies.

In this chapter, a novel human-machine interface concept is introduced that supports the controller with the task of tactical 4D trajectory management. Different from other prototypes, the proposed *Solution Space Concept* focuses on portraying solution spaces (instead of problem areas) that bound all feasible control actions, regardless of their optimality. These Solution Spaces—constructed by automation—enable controllers to observe the full range of control actions, and allow them to perform individual sensitive control strategies that are close to current-day practices (i.e., '4D vectoring'). By visualizing action-relevant areas rather than *issuing* specific actions, flexibility and adaptation in control will be facilitated in a way that keeps the controllers *in full control*, and allows them to control traffic in their unique, creative, and self-driven way.

3.2 Previous Work

In previous studies, various decision-support tools for 4D trajectory management by ATC and on the flight deck have been developed and prototyped [7]–[20]. These range from highly automated tools that provide a set of discrete conflict resolution options to advisory-level tools that highlight probable conflicts or conflict areas. Three concepts will be discussed that inspired the design of the Solution Space concept, and also underlined the importance of controller acceptance and human-automation conformance:

1. the Highly Interactive Problem Solver (HIPS [14]) concept that visualizes *avoidance zones* to aid the ATCo to identify problem areas when manipulating 4D trajectories, and

2. an ecological direct manipulation interface that visualizes *action-relevant* control spaces for revising 4D trajectories on the *flight deck* [7], [8], and

3. the state-based *velocity obstacles* concept [9], [10], [20] for ATC that visualizes *both* control spaces and avoidance zones for vectoring an aircraft.

3.2.1 Highly Interactive Problem Solver

In the late 1990's an elaborate study was performed by the PHARE consortium to develop a ground-based planning and de-confliction tool for 4D trajectory management, and resulted in the Highly Interactive Problem Solver (HIPS) concept (Fig. 3.1). Although the initial outset was to provide controllers with highly automated conflict resolution tools, that idea was discontinued at an early stage due to controllers' lack of acceptance [14]. Subsequently, an approach was adopted that allowed the controller to manually modify a trajectory, and show the outcome of the modifications overlaid on the radar screen in the form of "avoidance zones". These zones indicate areas in the sector where conflicts can occur along the trajectory of a selected aircraft. After each trajectory modification (routing, altitude, speed), the avoidance zones would be updated to reflect the impact of that modification.

Although HIPS was well-received by the research community, the final project report highlighted a number of key findings and issues [14]. The avoidance zones did not always correctly represent the controllers' perception of the nature and severity of the problems (i.e., human-automation mismatches). Sometimes there could be inconsistencies between the conflict detection algorithms and the problem-solver tool, resulting in ambiguous information causing controller distrust. The main conclusion was that *controller acceptance* and *interface intuitiveness* were crucial for the operational acceptance of any future ATC decision-support tool.

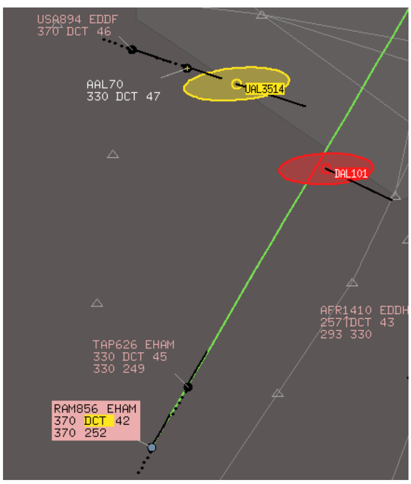

Figure 3.1: PHARE HIPS display showing red and yellow avoidance zones along the planned route of selected aircraft RAM856 (taken from [21]).

3.2.2 4D-Trajectory revisions on the flight deck

Various *ecological direct manipulation* interfaces for 4D-trajectory revisions by aircrews on the flight deck have been developed in previous research at the Delft University of Technology [7], [8]. These prototypes visualize *affordance zones* (i.e., control spaces) for the relocation of a single flight plan waypoint as illustrated in Fig.3.2. Here, the magenta affordance zone highlights the physical airspace where a selected waypoint (UD) can be relocated to navigate around weather, and to meet a time-based constraint at a metering waypoint (FIX).

The interface visualizes the *solution space* for manipulating the trajectory of the *ownship*, however, de-confliction with the trajectories of other aircraft is not taken into account (i.e., *egocentric* control). Further, this concept does not directly visualize the control space for manipulating speed and/or timings at waypoints. Therefore, in this form, the flight deck re-planning prototypes are less applicable (or practical) for the *decentralized* task of air traffic control. The concept of displaying the action relevant control space for trajectory modifications as a physical area inside the airspace, however, has inspired the design of the Solution Space concept presented in this chapter.

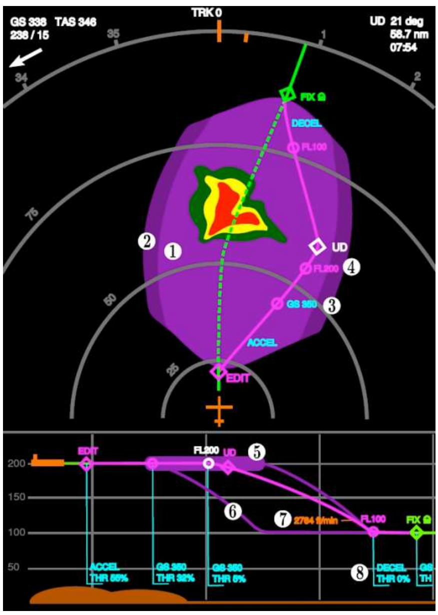

Figure 3.2: The 4D-trajectory flight deck re-planning interface showing the affordance and tolerance zones for re-routing a trajectory around hazardous weather (taken from [8]).

3.2.3 Velocity obstacles

The velocity obstacles method, initially proposed as a collision avoidance tool for ships [22], visualizes state-based conflict information based upon the relative velocities between two (or more) moving objects. Fig. 3.3 shows how tangent lines can be drawn from a controlled aircraft (A_{con}) to the protected zone of an observed aircraft (A_{obs}). When the tip of the resulting *forbidden beam zone* (FBZ) is offset by the velocity vector of the observed aircraft in the velocity plane, a so called state-based solution space can be constructed. By plotting the minimum and maximum speed of the controlled aircraft in the velocity plane, the solution space displays *all* instantaneous combinations of heading and speed that will lead to a conflict with the observed aircraft (i.e., when the velocity vector of the observed aircraft is inside the forbidden beam zone).

The velocity obstacles method visualizes the action relevant control space for de-conflicting aircraft by means of vectoring (i.e., providing discreet speed and heading commands). Though, because of its state-based form, it does not lend itself well for trajectory-based air traffic control. The concept of overlaying avoidance zones *on top of* the available control space, however, combined with the visualization of the control space itself as shown in Fig. 3.2, formed the foundation for the design of the trajectory-based Solution Space concept.

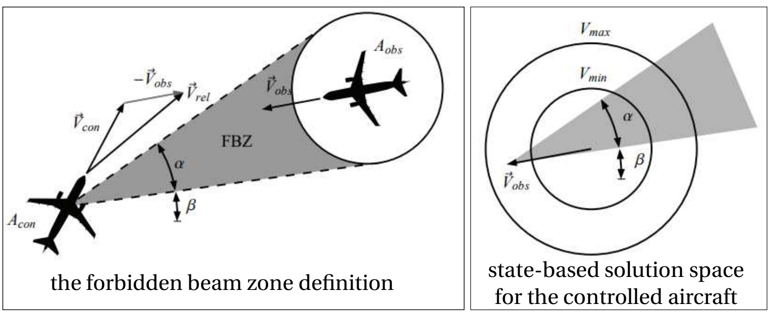

Figure 3.3: Definition of the forbidden beam zone and the state-based solution space of an observed aircraft. (adapted from [20]).

3.3 Theoretical Foundations

The Solution Space concept has been the result of various design iterations and refinements, based upon the framework of Ecological Interface Design (EID [23], [24]). In this section, first, a discussion of the theoretical motivation behind the selected design approach will be given. Next, the scope of the work domain, and limitations herein, for which the Solution Space concept has been designed will be described. Finally, a breakdown of the work domain is provided in the form of an Abstraction Hierarchy (AH) that formed the basis for identifying the elements, functions and interrelations that have been made salient on the interface.

3.3.1 Theoretical motivation

The success of the envisioned future ATM system will hinge on three important challenges that need to be addressed. The first challenge is determining *how far* the next generation ATM system should and *can be automated*. Given the more complex nature of data-link driven 4D trajectory management, the need for more automation support is clear. However, the ATM community acknowledges that fully automating the ATM system is nearly impossible, leading to the consensus that humans should remain the ultimate responsible for operational safety. But at what levels of autonomy and authority should automation mingle in the work of humans? The answer to this question leads to the second challenge; Is it possible to exploit the advantages of automation whilst *maintaining a competent and skilled human workforce*? Third and finally, any new technology aimed at taking over (parts of) work from trained professionals will *need to be accepted by the users* as articulated in PHARE studies and others [4].

The Solution Space concept aims to overcome these challenges by putting emphasis on the design of the support tool(s) on visualizing the physical and intentional *boundaries on control actions* to controllers, rather than having a computer algorithm providing a single, optimized solution, or showing only the *result* of control actions. That is, automation uses flight data (from digital data-links) to calculate and visualize 'Solution Spaces' for control, but lets the operator decide on a specific course of action within the available action space.

This design philosophy is inspired by Ecological Interface Design (EID), which aims to make work domain constraints (e.g., laws of physics) salient on an interface in such a way that people can directly perceive the space of possibilities and act upon it by utilizing their own expertise and skills. The concept is geared towards moving humans closer to the decision and control loops and provide a deeper insight into the constraints governing their work, which all can help to preserve, and perhaps even extend existing expertise. Research on *acceptance* indicates that technology should acknowledge the individual differences between humans, and thus provide support that matches the individual user[25]. The Solution Space concept is indeed individual-sensitive as it does not directly dictate what actions a controller must perform, but rather enables her to solve problems in her own way.

3.3.2 Scope

The design of the Solution Space concept is initially focused on real-time, tactical 4D trajectory re-planning, by a single air traffic controller in a single sector with traffic that is fully 4D- and data-link capable. We assume that the airspace- and route allocation have been planned, optimized, and de-conflicted *a-priori*. The controller task is then to *revise* trajectories in cases where operational perturbations such as delays, hazardous weather, and restricted airspace result in conflicting, unsafe, or infeasible trajectories.

For this concept, only en-route traffic has been considered that passes through the sector from predefined entry and exit waypoints, in zero-wind conditions, and at a single flight level (i.e., no vertical movements). Emphasis has been placed on supporting control actions that adhere to the original planning as much as possible to prevent cascading effects in adjacent sectors. That is, the horizontal and time-based trajectory manipulations have been designed such that the sector exit parameters (i.e., exit location and time) remain unchanged.

3.3.3 Work domain break-down

The outcome of the Work Domain Analysis (WDA) is crucial to understand what and how the various relationships and constraints in the work domain shape the work of the ATCo, and what the control activities of the ATCo will be (i.e. the work that the ATCo will have to execute). Ultimately, the WDA provides insight into the kind of information the ATCo needs to know about the system (the means) to be able to successfully execute the control task (the ends). The WDA, however, does not provide direct guidance about the representational form of the systems information that needs to be displayed to the ATCo.

The WDA in this study has been performed by constructing an Abstraction Hierarchy (AH) for the work domain envisioned in the scope. In the AH, the elements, functions and interrelations apparent in the work domain are laid out on different levels of abstraction. It provides a top-down model of the work domain that reveals the functional demands and the complete set of goal-relevant constraints. The higher levels in the AH represent the *system purpose*, whereas the lower levels represent the elemental data and *physical form*. The different levels are connected by *means-end links*, representing the logical flow and decomposition of the work domain interconnections.

Such a decomposition also aids in coping with system and interface complexity. That is, given that continuously presenting the full complexity of the system to the human operator is impractical (if not impossible), the various levels provide the means for adapting the interface to allow for shifting from higher to lower levels of abstraction. The overall goal (normal system operation) is defined by adhering to the ultimate system goals represented in the top levels. Overall conformance to these goals should be directly visible to the operator at all times. In case of non-conformance (e.g., a predicted loss of separation or restricted area intrusion), the means-end links can be followed to shift to a more detailed view of the problem in order to find, and reason about, a specific detailed solution.

The abstraction hierarchy that has been derived for the tactical 4D-trajectory re-planning task is shown in Fig. 3.4. The elements, functions and constraints are represented on five levels of abstraction, together with the means-end links that connect them; *(FP) Functional Purpose* – What is the purpose of the system? *(AF) Abstract Function* – What are the governing laws and priorities of the work domain? *(GF) Generalized Function* – How can these laws and priorities be achieved and upheld? *(PFu) Physical Function* – What are the components and their capabilities? *(PFo) Physical Form* – What is the physical form of these components? A detailed description is given in the following section of how these elements have been mapped to visual representations and implemented in an interactive decision-support tool.

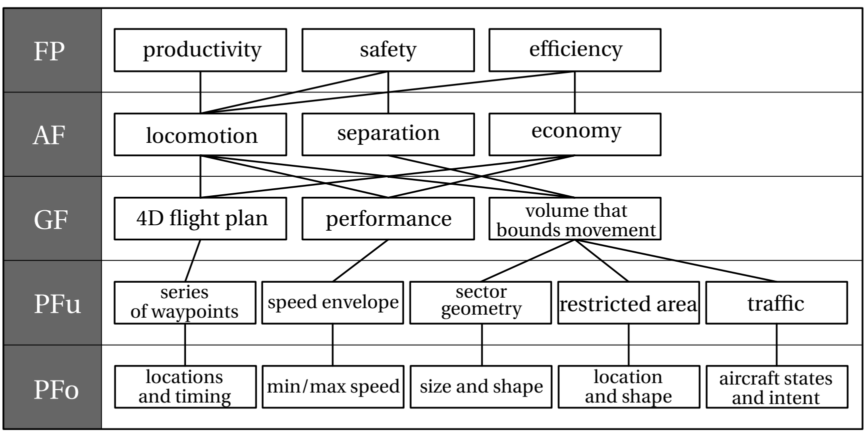

Figure 3.4: The Abstraction Hierarchy as a basis for the design of the Solution Space concept.

3.4 Solution Space Concept

The onset of the Solution Space concept is to facilitate flexibility and adaptation in control, allowing humans to reason about, and come up with creative solutions for challenging problems. The set of feasible control actions in air traffic control is, however, first and foremost bounded by the laws of physics (locomotion). In terms of manipulating a 4D trajectory, while adhering to the imposed metering constraints, flight dynamics and aircraft performance play a crucial role. These are the *internal constraints* related to the aircraft that is being controlled. The presence of other traffic and restricted areas then impose further *external constraints* that can be mapped inside the physical control space in the form of *avoidance zones* [26]–[28].

In this study, vertical trajectory manipulations are not part of the scope, reducing the dimensionality of the overall control space into three dimensions; manipulation of the horizontal flight path, defined by a set of waypoints, and manipulation of the timings at these waypoints. As there are inherent difficulties such as *occlusion of information* and *dimensional warping* when displaying more than two dimensions on a two-dimensional screen, the interface has been split up into two distinct parts. One is the top-down electronic plan view display (PVD) in which the *horizontal Solution Space* (HSS) is shown, and the other is a so-called *time-space diagram* (TSD) on which the *time-based Solution Space* (TSS) is shown.

First, the construction of the horizontal and time-based Solution Spaces bounded by internal aircraft constraints is discussed, followed by the derivation of the overlaid avoidance zones imposed by external and intentional constraints. Next, the integrated display concept is presented, including the proposed methods for interactive trajectory manipulation. Finally, the real-time computer-based implementation and considerations for visualizing the Solution Spaces are discussed.

3.4.1 Internal constraints

Fig. 3.5 shows the causal functions in the work domain that are related to the internal aircraft constraints in tactical 4D-trajectory control. Locomotion is achieved by following a 4D flight plan that must adhere to the physical performance limits of the aircraft. This paragraph shows how these internal constraints can be mapped on the PVD and TSD to form the basic shapes of the horizontal and time-based Solution Spaces.

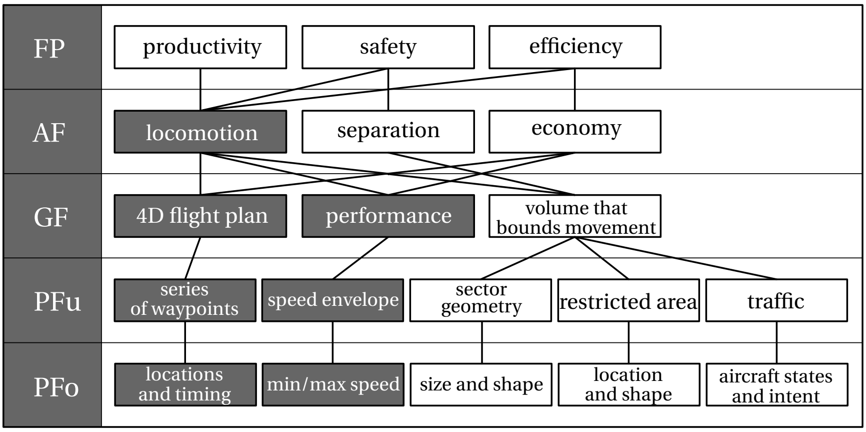

Figure 3.5: Elements in the AH related to the *internal* aircraft constraints.

Horizontal control space

Fig. 3.6 shows an aircraft flying along a straight trajectory towards a time-metered position fix (FIX). If the aircraft deviates from the direct route (i.e., elongates the path by flying a dog-leg), a higher ground speed is required in order to meet the metered time (Required Time of Arrival, RTA) at the fix. The speed envelope within which an aircraft can operate is physically limited by the lower maneuvering speed (V_{min}) and upper maximum operating speed, or Mach buffet limit (V_{max}), which can be converted to ground speed given the current atmospheric conditions. Assuming that the RTA at the fix is a hard constraint, the speed envelope of the aircraft bounds the physical control space through which that aircraft can be rerouted.

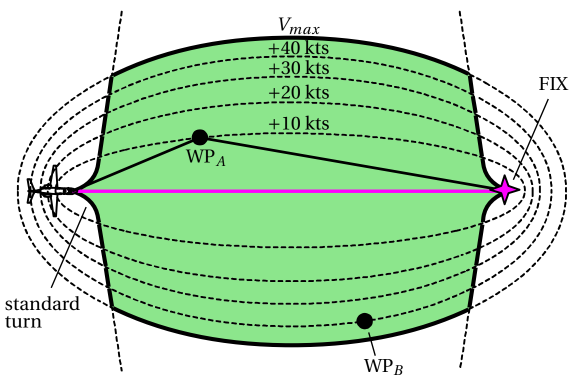

Figure 3.6: Construction of the *horizontal* control space.

Fig. 3.6 shows how the speed envelope affects the possibilities for placing an intermediate waypoint along the trajectory. Without considering turn dynamics and wind, all rerouting options for a given speed increase lie on an ellipse with the aircraft and fix as focal points. For example, the intermediate waypoint WP_A will split the original trajectory into two equal-speed segments that must be flown 10 knots faster in order meet the fix RTA. Intermediate waypoints that require a larger path deviation, such as WP_B, result in a higher required speed. By taking the ellipse that results from flying at V_{max} as the limiting factor, and taking into account a standard turn in the direction toward the metering fix, the green shaded horizontal control space can be constructed as shown in Figure 3.6. Any waypoint that is placed inside the horizontal control space, splitting the trajectory into two equi-speed segments, will thus result in a feasible new trajectory that does not violate the time constraint.

Time-based control space

The control space along the trajectory can also be expressed in the time domain. A useful way to visualize this is by using a *Time-Space Diagram* (TSD) [29], as illustrated in Fig. 3.7. Here, the horizontal axis represents the *distance to go* along the trajectory *towards* the metering fix. The vertical axis represents the corresponding predicted *time* at the points along the trajectory.

The slope of the line is a measure of the ground speed of the aircraft along the trajectory. A faster speed will result in a more shallow slope (the aircraft will arrive earlier at the fix), and vice versa, a slower speed will result in a steeper slope. Any curved section in the time-space trajectory indicates a speed change (i.e., acceleration or deceleration).

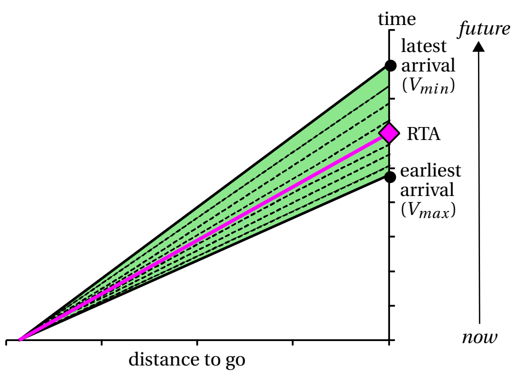

Figure 3.7: Construction of the *time-based* control space.

The time-based control space is also constrained by V_{min} and V_{max}, and results in the latest and earliest possible arrival times at the fix, respectively. Changing the arrival time at the fix will violate the RTA constraint and does not meet the earlier imposed constraints. However, for trajectories that consist of multiple waypoints, the times at individual waypoints can be manipulated separately such that the final RTA is still met.

There is a direct link between the horizontal and time-based control spaces. If the lateral path becomes longer due to a deviation, the margin between the earliest and latest arrival time becomes smaller, and the slope of the time-space trajectory becomes shallower. Vice versa, if the fix RTA is changed in the TSD, the horizontal control space will widen or narrow accordingly.

3.4.2 External constraints

Fig. 3.8 shows the elements in the work domain that are related to the external constraints acting on the trajectory of a controlled aircraft from the centralized perspective of an ATCo. The volume that bounds movement is primarily shaped by the geometry of the controlled sector. Inside the sector, movement can be further bound by the presence of restricted areas (e.g., due to weather, temporary restricted airspace, etc.), and by the separation requirements with respect to other traffic. In this study, restricted areas are considered to be static, and visually re-routing aircraft around these zones is considered to be a trivial task. Therefore, only the mapping of separation constraints caused by other dynamic traffic will be discussed.

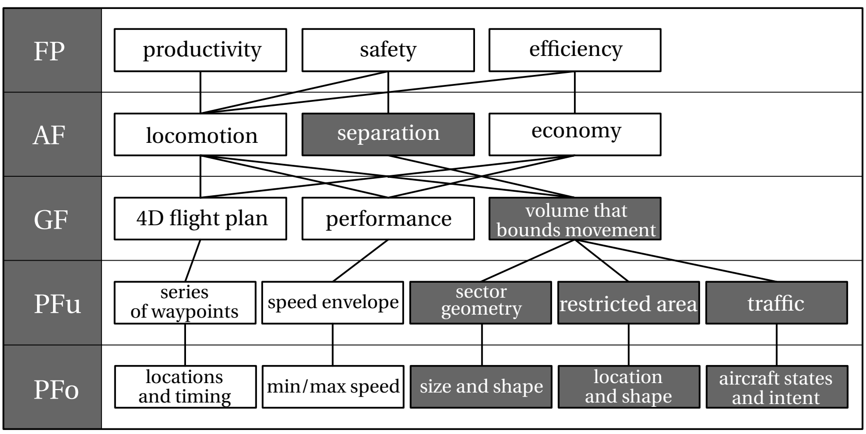

Figure 3.8: Elements in the AH related to the *external* constraints acting on a single flight plan.

In many control systems, a set of rules, operating limits, and other provisions are introduced to assure overall safety and efficiency of operations. These act as *intentional* buffers and safe-guards to cope with contingencies caused by system variability [28]. Typically, in tactically controlled ATC sectors, the lateral safety provision for positive separation between aircraft is set to 5 nautical miles. Any breach of this space is seen as an unsafe 'loss of separation' event.

In the scenario of fully mature trajectory-based air traffic control, the intent of all aircraft is broadcast, updated, and distributed to ATC centers by digital data-link. This allows automation to work with an accurate and up-to-date prediction of the (future) traffic movements. Using this information, avoidance zones can be mapped and visualized on top of horizontal and time-based Solution Spaces.

Horizontal constraints

Each point inside the horizontal control space represents a position where an intermediate waypoint can be inserted into the original trajectory, resulting in a new (feasible) trajectory consisting of two constant speed segments. This is shown in Fig. 3.9(a) where waypoint $\vec{p}$ is inserted into the path of the observed aircraft (A_{obs}). By neglecting accelerations, decelerations and turn dynamics, the resulting predicted locomotion of A_{obs} can be described by linear interpolation between the three time-metered points ($\vec{a}$, $\vec{p}$, $\vec{b}$).

Similarly, the predicted positions of all other traffic can be described by linear interpolation between their flight plan waypoints. This allows to compute the closest point of approach (CPA) of all other traffic (over time) with respect to the new trajectory. If the CPA is smaller than the separation minima, the new trajectory is predicted to lead to a loss of separation, deeming that trajectory as unsafe. When the same computation is performed for all points inside the horizontal control space, avoidance zones can be determined and superimposed as shown in red in Fig. 3.10(a).

In essence, the placement of an intermediate waypoint at any point inside an avoidance zone will lead to a trajectory that violates the separation constraints at some point in the future. In the case illustrated in the figure, the current trajectory of the aircraft passes through an avoidance zone, indicating that a loss of separation is predicted to occur if no action is taken.

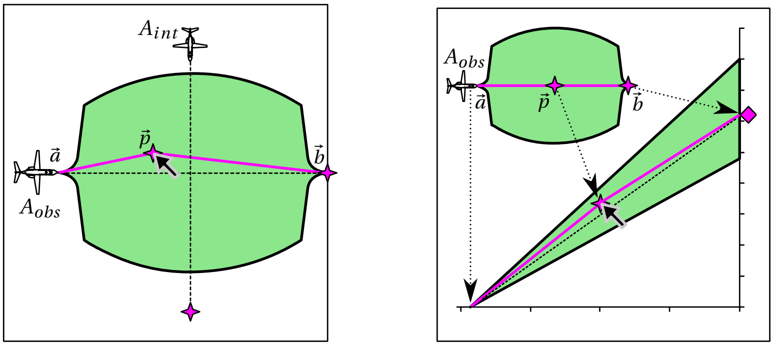

(a) Horizontal control space (b) Time-based control space

Figure 3.9: The relationship of intermediate positions in the horizontal and time-based control spaces to trajectory modifications.

Time-based constraints

Each point in the time-based control space represents a physical position along the *current lateral trajectory* at a given future time. Fig. 3.9(b) illustrates how an intermediate waypoint $\vec{p}$ can be inserted into the flight plan on the TSD, changing the metering time at that point. Similar as for the horizontal control space, the CPA of all other traffic with respect to the new trajectory can be computed. By performing this computation for all points inside the time-based control space, avoidance zones can be generated and superimposed as shown in red in Fig. 3.10(b).

Note that when manipulating the trajectory in the TSD, the lateral path of the observed aircraft remains unchanged. Assuming the trajectories of the other traffic are fixed, the size, shape, and locations of the avoidance zones caused by the traffic also remain unchanged, regardless of any timing modifications to the observed aircraft. This means that avoidance zones outside of the time-based control space can also be visualized.

A loss of separation is predicted when the time-space trajectory of the observed aircraft passes through an avoidance zone, as shown in the figure. A breach of the avoidance zones in the horizontal and time-based control spaces is always *mutually inclusive*. That is, if there is a predicted loss of separation, both the time-space line and horizontal path will pass through an avoidance zone. However, avoidance zones that do not immediately lead to losses of separation do not always have a clear counterpart in the other control dimension.

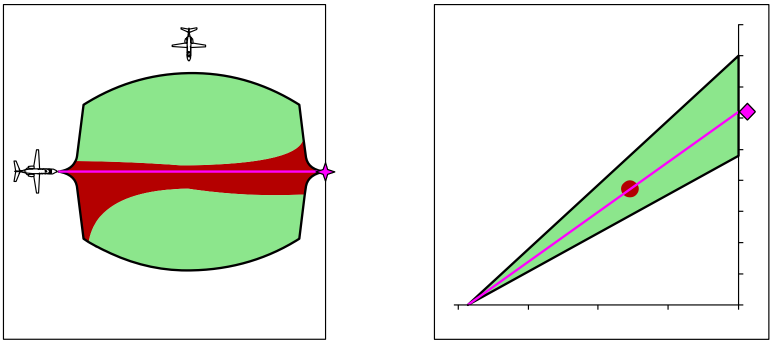

(a) Horizontal control space (b) Time-based control space

Figure 3.10: Avoidance zones caused by other traffic mapped on the horizontal and time-based control spaces.

Intentional constraints

Although the Solution Spaces illustrated in Fig. 3.10 visualize the control space in the form of feasible and infeasible areas (i.e., go and no-go zones), not every control action is equally preferable in terms of safety, efficiency, and flexibility. In order to steer the controller away from implementing such strategies, additional *intentional constraints* can be imposed and mapped onto the control spaces [28]. In this study we do not aim to provide absolute guidelines for the size and shape of these buffers, but merely provide an illustration of how these can be integrated in the Solution Space concept.

Placing waypoints close to the border of avoidance zones could rapidly lead to unsafe situations in case of slight deviations from the planned path. To account for this, additional (intentional) separation buffers can be computed and visualized in the control spaces as shown in amber in Fig. 3.11. The intentional separation zones then represent the areas in which 'control' will lead to other traffic passing in close vicinity to the separation minimum. The shape and size of these zones also hints at how additional separation provisions propagate throughout the control spaces. Control actions that require large deviations from the original path are uneconomical in terms of fuel-burn (inefficient) and leave little control space for further follow-up manipulations (reduce flexibility). Also, flying close to the limits of the speed envelope will leave little room to compensate for any deviations from the original plan by using further speed control. These intentional cautionary zones can be computed and visualized in the control spaces as shown in dark green in Fig. 3.11.

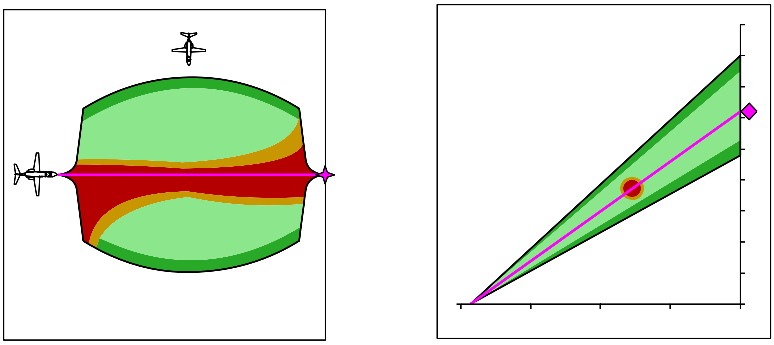

(a) Horizontal control space (b) Time-based control space

Figure 3.11: Intentional constraints added to the horizontal and time-based control spaces.

3.4.3 Integrated display concept

Fig. 3.12 shows a sketch of the lay-out of the integrated Solution Space display concept. The plan view display (PVD) is located on the left-hand side, and the TSD is located on the right-hand side. An aircraft is currently selected, and its horizontal and time-based Solution Spaces are visualized. The selected aircraft has a predicted loss of separation with respect to the traffic, which can be visually confirmed by the horizontal and time-based trajectories passing through an avoidance zone. Additionally, the conflicting aircraft is highlighted in red on the electronic plan view display. Next, the typical work-flow of monitoring and interacting with the Solution Space interfaces will be discussed.

Monitoring

The work-flow of monitoring and interacting with the integrated display concept is shown in Fig. 3.13. When no aircraft is selected, the PVD acts as a traditional ATC radar screen, showing a top-down view of the traffic, sector lay-out, entry and exit waypoints and restricted areas. In this case, no Solution Spaces are visualized on the PVD, and the TSD is empty. In this *monitoring* display state, the presented information allows the controller to monitor conformance to the overall system goals at a glance, and directly identify if control actions are necessary.

The aircraft labels show the callsign, the current speed, and the predicted delay with respect to the fix RTA. The trajectories of all aircraft are shown by thin lines, without the explicit visualization of trajectory waypoints or additional in-

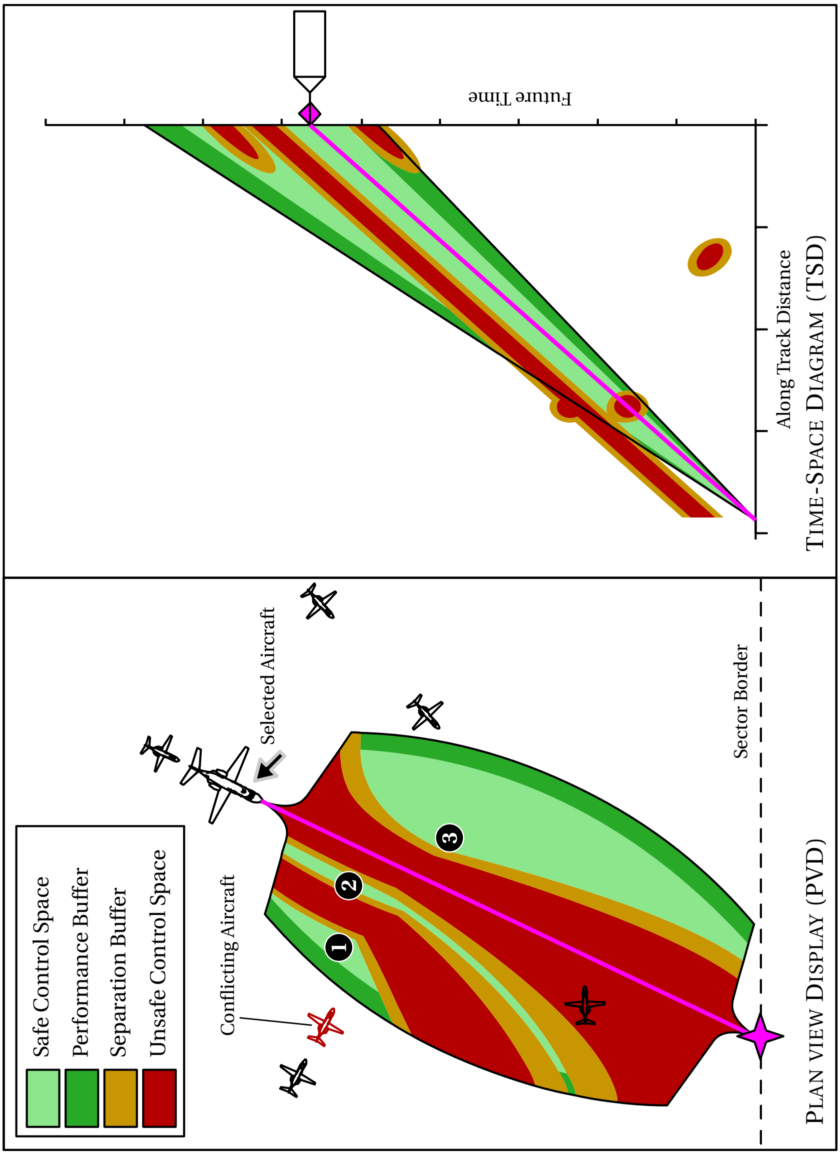

Figure 3.12: Overview of the integrated Solution Space display concept.

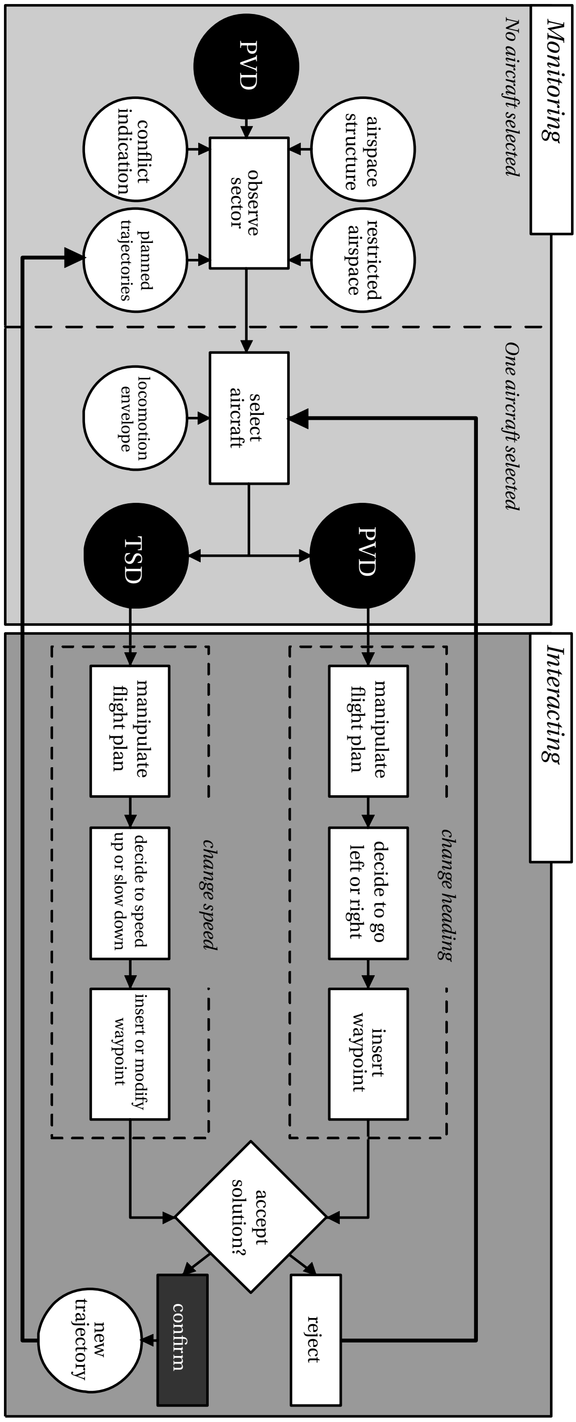

Figure 3.13: Workflow of monitoring and interacting with the integrated Solution Space concept.

formation. A slider along the time-axis in the TSD can be used to make a *ghost projection* of the predicted traffic movements in time, showing how the traffic will dynamically evolve. Aircraft that are predicted to have a loss of separation are highlighted in red. By hovering over a conflicted aircraft with the mouse cursor, the other aircraft involved the conflict (one or more) are also highlighted. This can be used to gain more insight about the nature of a conflict without directly zooming in on individual aircraft level.

At all times, the controller can select an aircraft of her choice on the PVD. Upon clicking the chosen aircraft, the horizontal and time-space trajectories and Solution Spaces are visualized in real-time *on both* sub-displays. In this display state, hovering with the mouse cursor over an avoidance zone in the PVD or TSD will highlight the aircraft that causes that zone on the PVD. Conversely, hovering with the mouse cursor over an aircraft on the PVD will highlight its associated avoidance zones (if any) on the horizontal and time-based Solution Spaces. This enables a controller to mentally 'connect' the information shown on the two displays.

Interacting

If the controller has identified that a trajectory manipulation is required, she can select that aircraft on the PVD. Manipulating the trajectory of the selected aircraft can be done *in each of* the two displays, both in isolation and in combination. These manipulations loosely correspond to current-day ATC vectoring, namely changing the flight direction (PVD) and the speed (TSD). The following paragraphs illustrate how the controller can reason about-, and manipulate the trajectory of an aircraft.

Horizontal manipulations

Fig. 3.14(a) shows the Solution Space on the PVD for a selected aircraft in a crossing conflict. Here, the controller has the option to either reroute the selected aircraft in front or behind the other traffic. Fig. 3.14(b) shows how the controller can insert an intermediate waypoint by clicking on the Solution Space with the mouse cursor. Fig. 3.14(c) illustrates how the aircraft trajectory is split in two equal speed segments, and a new Solution Spaces are constructed for both path segments. If the controller is satisfied with the modification, she can send the new trajectory to the aircraft for execution.

The horizontal Solution Spaces do not limit the allowed control actions by the controller, however, when a location is selected in the avoidance zone, this will lead to a new trajectory that is still in conflict with the traffic. Similarly,

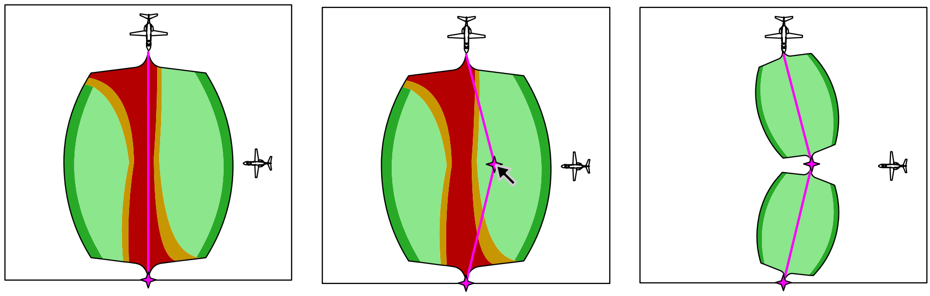

(a) Plan View Display (b) Horizontal modification (c) PVD after modification

Figure 3.14: Horizontal manipulations on the PVD.

the controller can choose to place a waypoint outside of the Solution Spaces all together, but that will lead to a delay at the sector exit point.

When observing the more complex scenario of Fig. 3.12, the controller could choose to either steer the selected aircraft behind both crossing aircraft (❶), steer the aircraft to fly between the crossing aircraft (❷), or steer the aircraft in front of the crossing traffic (❸). It can be deducted from the Solution Spaces that, in this case, a preferred solution in terms of safety and efficiency would be to steer in front.

Time-Space manipulations

Fig. 3.15(a) shows the time-based Solution Space of the same aircraft selected in Fig. 3.14(a). The conflict can be resolved by expediting or delaying the aircraft to the conflict location, resulting in it passing either in front or behind the traffic. Fig. 3.15(b) shows how the controller can insert an intermediate waypoint by clicking on the Solution Space with the mouse cursor. The time-space trajectory is then split into two segments as shown in Fig. 3.15(c), for which new Solution Spaces are generated. Here, the first segment is flown at a slower speed (sloped steeper), and the second at a faster speed in order to meet the original time constraint. The selected aircraft will pass behind the traffic, indicated by the avoidance zone that is located under the time-space trajectory.

The controller could also choose to resolve the conflict by delaying or expediting the aircraft along its entire trajectory by dragging the aircraft label on the time-axis up or down. However, that would violate the time constraint at the metering fix. It is not possible to place a waypoint outside of the time-based Solution Spaces because the selected aircraft cannot fly faster or slower than its maximum and minimum speed. Note that modifying the speed along the tra-

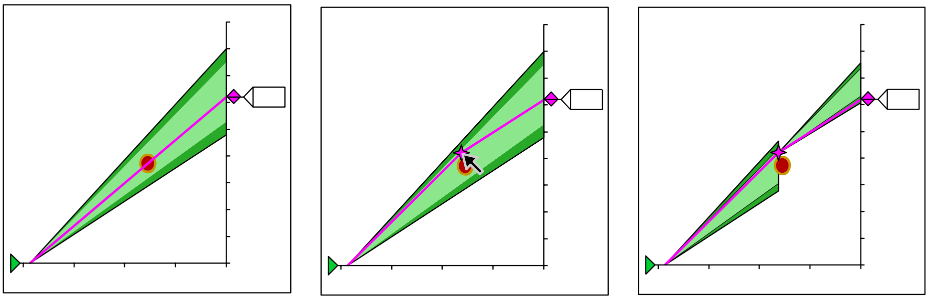

(a) Time-Space Diagram (b) Time-based modification (c) TSD after modification

Figure 3.15: Time-based manipulations on the TSD.

jectory widen the horizontal Solution Space for the slower segment, and narrow it for the faster segment.

When observing the time-based Solution Space in the more complex scenario of Fig. 3.12, a number of things can be observed. The trailing aircraft causes a stretched avoidance zone above the time-space trajectory of the selected aircraft. Thus, delaying the selected aircraft is likely to cause an 'overtake' conflict. The two avoidance zones at the end of the time-space trajectory (attached to the time axis) are caused by the two aircraft that are planned to exit the sector at the same metering fix. The two crossing aircraft cause avoidance zones at the same point along the trajectory, but at different times. From the TSD it can be concluded that the only feasible control option that respects all constraints is to steer the selected aircraft *between* the crossing aircraft, albeit inside the intentional buffer.

3.4.4 Real-time Implementation

A prototype of the integrated Solution Space concept has been implemented in a computer-based ATC simulator, and proved that Solution Space representations can be visualized in real-time (Fig.3.16). No significant performance issues were found when running at 60Hz on a current consumer-grade desktop with an NVIDIA GTX 970 (4GB VRAM) GPU, and using a 30", 2560 x 1600 pixel display. The horizontal and time-based Solution Space visualizations were computed and generated by utilizing pixel-based *shader* computations on the Graphics Processing Unit (GPU).

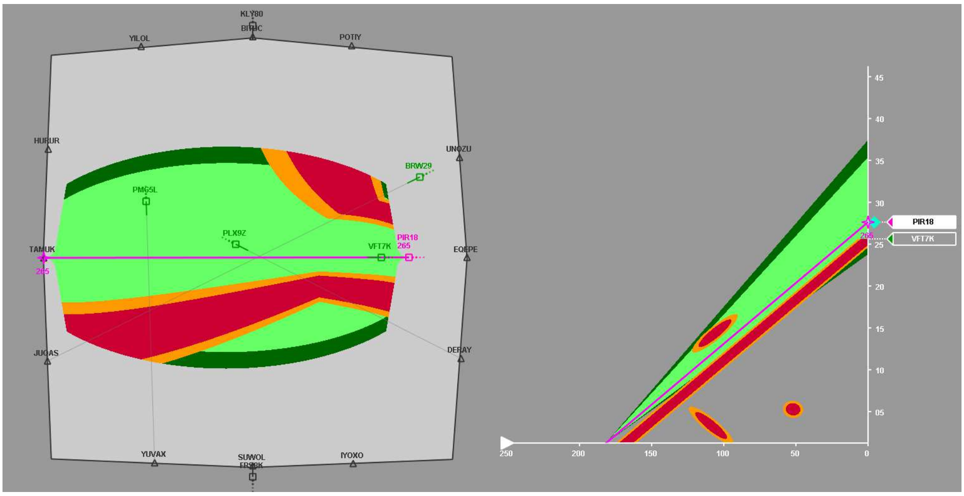

Figure 3.16: Screenshot of the computer-based implementation of the Integrated Solution Space concept, showing the horizontal and time-based Solution Spaces for a selected aircraft. The colors have been adapted for print.

3.5 Skills, Rules, and Knowledge

Although the workflow of monitoring and interacting with the interface is shown in Fig.3.13, it does not encompass the cognitive process of decision-making behind the choice for a certain control action. To determine the control strategies that can be applied with the interface, an analysis of the worker competencies has been performed by using Rasmussen's classification of skill-based, rule-based, or knowledge-based behavior (SRK taxonomy,[30], [31]). The decision ladder (3.17), typically used in the Activity Analysis dimension of CWA, has been used as a tool to investigate various modes of control. The ladder is entered at the left bottom, when a situation triggers the need for a control action. The left hand side shows the various steps in the *analysis* of the problem, and the right hand side shows the shows the steps in *planning* an acceptable solution.

A solution is formulated by following the ladder, whether by going all the way up to the knowledge-based domain, or by taking rule-based shortcuts. These shortcuts are typically based on stored rules or procedures that have been derived empirically during previous occasions, either learned during training, or stored from successful solutions in similar situations. A distinction can be made between lower and higher level rule-based behavior. Lower levels are close to skilled-based behavior and are typically control strategies that always work for a given situation without further thought. Higher level rule-based coordination is based on familiarity and/or explicit know-how that requires a higher level of cognition.

For the task of tactical, trajectory-based conflict resolution with the Solution Space interface, three rule-based shortcuts have been identified that are shown in the decision ladder (numbered ❶, ❷ and ❸). The decision flow map in Fig.3.18 accompanies the decision ladder, and shows the cognitive steps preceding each shortcut in more detail. The decision flow map is entered when the interface detects a conflict between two or more aircraft. The first step for the controller is to observe and assess the situation using the available information. The involved aircraft can be identified by their red labels. Further, the planned trajectories of all other traffic are shown as thin lines, aiding in creating an initial mental model of the situation.

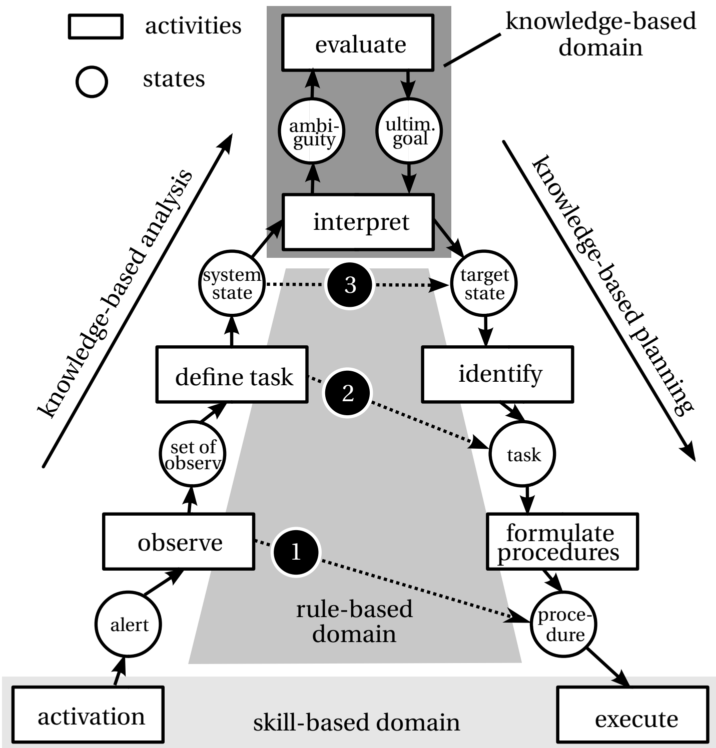

Figure 3.17: The decision ladder as a method to classify skill-, rule- and knowledge-based control. The three rule-based shortcuts are shown and numbered 1, 2 and 3.

Shortcut 1. Using the observations, the controller can determine whether immediate action is required or whether there is time to formulate a more elaborate strategy. When immediate action is required, the controller can select one of the conflicting aircraft and observe its Solution Space to check for the availability of safe control space (both horizontally or time-based). If not, the controller can select the other conflicting aircraft to rapidly observe which one has (the most) available control space. In taking this shortcut, typically any available solution will be acceptable, and control is mostly reactive.

Shortcut 2. When no immediate action is required, the controller can test solutions against best practices and/or if they conform to currently applied strategies. A typical best practice in ATC is for instance, letting slower aircraft pass behind faster aircraft. A current strategy can be, for instance, solving all identical conflicts (or traffic patterns) in a similar geometric way. By selecting the corresponding aircraft and observing its available control space, the controller can confirm the validity of that strategy.

Shortcut 3. If shortcut 2 cannot be taken, the controller can fall back on strategies learned by training, or strategies that have been successfully applied for similar scenarios in previous experiences. The availability of this shortcut mainly depends on the level of training, experience, and expertise of the controller to recognize familiar patterns and couple them to a suitable strategy.

Knowledge-based behavior. When no rule-based shortcut can be made, knowledge-based strategies have to be formulated. Knowledge-based behavior is typically applied in unfamiliar or complex situations, and relies on the operators' mental model of the work domain. Here, one or more strategies can be formulated and tested towards the overarching goal. This type of control can be classified as 'creative problem solving' and requires the highest level of cognition.

In practice, multiple iterations can be made through the decision flow map. The controller can view the Solution Spaces of multiple aircraft before selecting a suitable control strategy. Typically, lower level rule-based shortcuts, if available, will be preferred over higher level solutions. The number- and availability of such shortcuts is highly dependent on the level of training, experience and expertise of the controller.

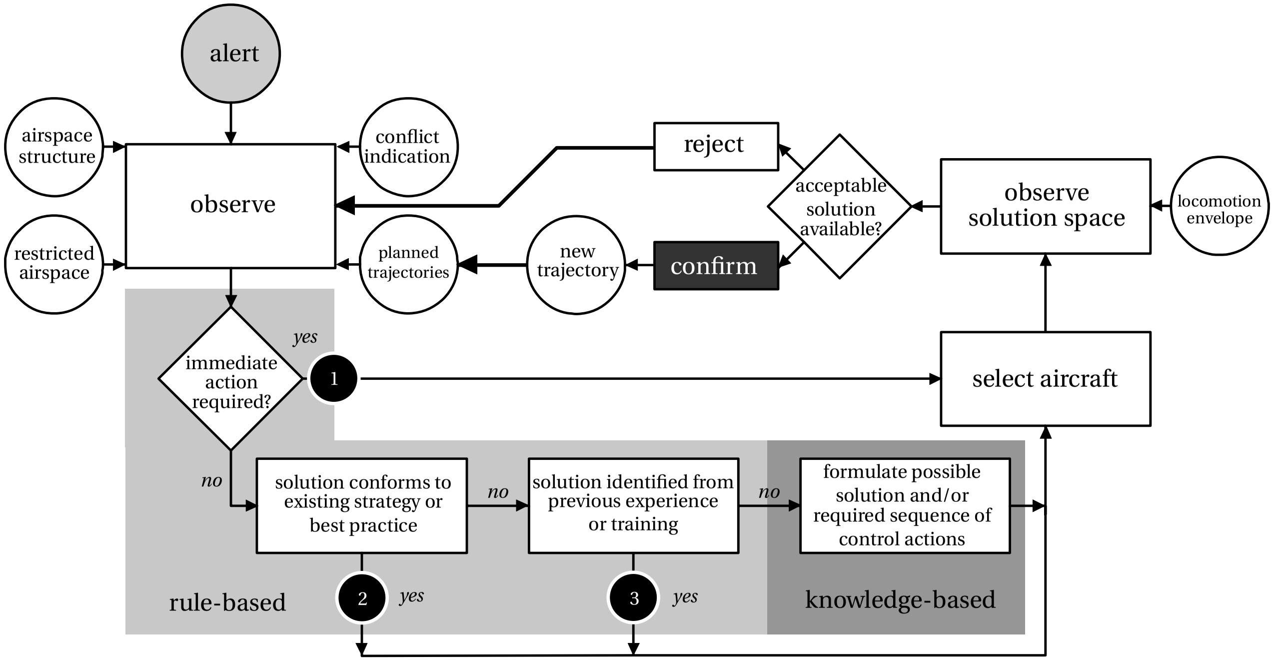

Figure 3.18: Decision flow map for formulating a strategy to resolve trajectory-based conflicts. The three rule-based shortcuts are shown and numbered 1, 2 and 3.

3.6 Conclusion

The ultimate goal of this work has been to design a prototype decision-support tool for real-time, tactical, 4D trajectory-based air traffic control. The design of the Solution Space concept presented in this chapter is inspired by ecological interface design, and focuses on visualizing action-relevant control spaces. These Solution Spaces show the available control space for horizontally rerouting the trajectory of an aircraft and modifying the speed along the trajectory, overlaid with avoidance areas caused by external constraints such as restricted areas, weather, and other traffic. The concept supports air traffic controllers to observe and modify 4D-trajectories in a way that allows them to exercise their own preferred—individual-sensitive—control strategies.

3.7 Recommendations for Future Work

The manipulation of the vertical profile of a 4D-trajectory has not been taken into account in this work. This is an area that firstly *must* be addressed in future work in order to gain more *traction* in the operational community. The addition of vertical Solution Spaces, however, will greatly increase the complexity of the control task, and bring new challenges to the display design; aircraft performance varies widely at different altitudes, which has a direct impact on the overall available control space. Next, the concept presented in this paper focuses on 4D-trajectory management in ideal—deterministic—conditions. *External uncertainties* caused by non-linear atmospheric conditions (i.e., wind, pressure, and temperature variations), variance in pilot response time and datalink delays, varying levels of hardware capabilities, varying aircraft navigation precision, etc., have not been taken into account. These factors must be addressed in order to gain more *acceptance* in the operational community. The main challenge will be to keep a careful balance between visualizing operational complexity versus the concept's usability in future developments.

Bibliography

[1] "Concept of Operations for the Next Generation Air Transportation System, Version 3.2," Joint Planning and Development Office, Washington, DC, USA, Report, 2011.

[2] "SESAR 2020 Concept of Operations Edition 2017," EUROCONTROL, Brussels, Belgium, Report, 2017.

[3] "European ATM Master Plan 2020 Edition: Executive View," SESAR Joint Undertaking, Brussels, Belgium, Report, 2019.

[4] C. A. L. Westin, C. Borst, and B. G. Hilburn, "Strategic Conformance: Overcoming Acceptance Issues of Decision Aiding Automation?" *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 41–52, 2016.

[5] M. Bekier, B. Molesworth, and A. M. Williamson, "Tipping point: The narrow path between automation acceptance and rejection in Air Traffic Management," *Safety Science*, vol. 50, no. 2, pp. 259–265, 2012.

[6] M. Bekier and A. M. Williamson, "Why Air Traffic Controllers Accept or Refuse Automated Technology," in *16th International Symposium on Aviation Psychology*, 2011, pp. 615–620.

[7] M. Mulder, R. Winterberg, M. M. Van Paassen, and M. Mulder, "Direct Manipulation Interfaces for In-Flight Four-Dimensional Navigation," *International Journal of Aviation Psychology*, vol. 20, no. 3, pp. 249–268, 2010.

[8] B. J. A. Van Marwijk, C. Borst, M. Mulder, M. Mulder, and M. M. Van Paassen, "Supporting 4D Trajectory Revisions on the Flight Deck: Design of a Human-Machine Interface," *The International Journal of Aviation Psychology*, vol. 21, no. 1, pp. 35–61, 2011.

[9] S. B. J. Van Dam, M. Mulder, and M. M. Van Paassen, "Ecological Interface Design of a Tactical Airborne Separation Assistance Tool," *IEEE Transactions on Systems, Man & Cybernetics, Part A*, vol. 38, no. 6, pp. 1221–1233, 2008.

[10] P. Hermes, M. Mulder, M. M. Van Paassen, J. H. L. Boering, and H. Huisman, "Solution-Space-Based Analysis of the Difficulty of Aircraft Merging Tasks," *Journal of Aircraft*, vol. 46, no. 6, pp. 1995–2015, 2009.

[11] M. S. Eby, "A Self-organizational Approach for Resolving Air Traffic Conflicts," *The Lincoln Laboratory Journal - Special issue on air traffic control*, vol. 7, pp. 239–254, 2 1994.

[12] H. R. Idris, R. Vivona, J.-L. Garcia-Chico, and D. Wing, "Distributed Traffic Complexity Management by Preserving Trajectory Flexibility," in *Proceedings of the 26th IEEE/AIAA Digital Avionics Systems Conference*, 2007, pp. 1–13.

[13] H. R. Idris, T. El-wakil, and D. J. Wing, "Trajectory Planning by Preserving Flexibility: Metrics and Analysis," in *Proceedings of the AIAA Guidance, Navigation, and Control (GNC) Conference and Exhibit*, 2008, pp. 1–14.

[14] M. J. Whiteley and I. Wilson, "PHARE Advanced Tools Problem Solver," EUROCONTROL, Brussels, Belgium, Final Report, 1999.

[15] T. Prevot, P. U. Lee, N. Smith, and E. A. Palmer, "ATC Technologies for Controller-Managed and Autonomous Flight Operations," in *Proceedings of the AIAA Guidance, Navigation, and Control (GNC) Conference and Exhibit*, 2005, pp. 1–43.

[16] J.-P. Nicolaon, G. Dean, and W. Miller, "An Overview of ARC 2000 Version 3 from the Operational Point of View," EUROCONTROL experimental center, Bretigny-sur-Orge, France, Tech. Rep. 286, 1995.

[17] L. Goldmuntz, J. T. Kefaliotis, L. A. Kleiman, R. A. Rucker, S. L., and D. Weathers, "The AERA Concept," DOT/FAA, Washington D.C., USA, Tech. Rep. FAA-EM-81-3, 1981.

[18] W. P. Niedringhaus, "A Mathematical Formulation for Planning Automated Aircraft Separations for AERA 3," DOT/FAA, Washington D.C., USA, Tech. Rep. DOT/FAA/DS-89/20, 1989.

[19] H. Erzberger, T. A. Lauderdale, and Y. C. Chu, "Automated Conflict Resolution, Arrival Management, and Weather Avoidance for Air Traffic Management," in *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, Institution of Mechanical Engineers, 2012, pp. 930–949.

[20] G. A. Mercado-Velasco, M. Mulder, and M. M. Van Paassen, "Analysis of Air Traffic Controller Workload Reduction Based on the Solution Space for the Merging Task," in *Proceedings of the AIAA Guidance, Navigation, and Control (GNC) Conference and Exhibit*, 2010, pp. 1–18.

[21] M. van Gool and H. Schroter, "PHARE Final Report," EUROCONTROL, Brussels, Belgium, Tech. Rep. DOC 99-70-09, 1999.

[22] T. Degré and X. Lefèvre, "A Collision Avoidance System," *The Journal of Navigation*, vol. 34, pp. 294–302, 2 1981.

[23] K. J. Vicente and J. Rasmussen, "Ecological Interface Design: Theoretical Foundations.," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 22, no. 4, pp. 589–606, 1992.

[24] C. M. Burns and J. Hajdukiewicz, *Ecological Interface Design*. Boca Raton, FL, USA: CRC Press, 2004.

[25] B. G. Hilburn, C. A. L. Westin, and C. Borst, "Will Controllers Accept a Machine That Thinks Like They Think? The Role of Strategic Conformance in Decision Aiding Automation," *Air Traffic Control Quarterly*, vol. 22, no. 2, pp. 115–136, 2014.

[26] C. Borst, “Ecological Approach to Pilot Terrain Awareness,” Ph.D. Dissertation, Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands, 2009.

[27] J. Ellerbroek, “Airborne Conflict Resolution In Three Dimensions,” Ph.D. Dissertation, Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands, 2013.

[28] J. Comans, “Visualizing Rules, Regulations, and Procedures in Ecological Information Systems,” Ph.D. Dissertation, Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands, 2017.

[29] M. Tielrooij, A. C. In ’t Veld, M. M. Van Paassen, and M. Mulder, “Development of a Time-Space Diagram to Assist Air Traffic Controllers in Monitoring Continuous Descent Approaches,” in *Air Traffic Control*, M. Mulder, Ed., Sciyo, 2010, ch. 7, pp. 135–147.

[30] J. Rasmussen, “Skills, Rules, and Knowledge; Signals, Signs, and Symbols, and other Distinctions in Human Performance Models,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 13, no. 3, pp. 257–266, 1983.

[31] ——, *Information Processing and Human-machine Interaction: An Approach to Cognitive Engineering*, A. P. Sage, Ed., ser. North-Holland series in system science and engineering. New York, NY, USA: Elsevier Science Publishers, 1986, vol. 12.

CHAPTER 4

EXPERTISE, STRATEGIES AND ROBUSTNESS

One machine can do the work of fifty ordinary men. No machine can do the work of one extraordinary man. —*Elbert Hubbard (Writer, Philosopher, 1856 – 1915)*

This chapter presents the first human-in-the-loop experiment with a partial implementation of the Solution Space concept. First, the visualization of horizontal Solution Spaces—named Travel Space in this chapter—is briefly introduced, followed by an investigation of the various control strategies by means of a decision ladder and an information flow map. A new metric is presented to quantify how robust control actions are to further system disruptions. The set-up of the human-in-the-loop experiment is presented, and the results of the experiment are presented and discussed. The chapter is concluded with a discussion of the results.

Paper Title	Expertise Level, Control Strategies, and Robustness in Future Air Traffic Control Decision Aiding
Authors	R.E. Klomp, C. Borst, M.M. van Paassen, and M. Mulder
Published in	IEEE Transactions on Human-Machine Systems, Vol. 46, Issue. 2, 2015, pp. 255-266.

ABSTRACT

The introduction of 4D trajectory-based operations will require the development of new and more advanced 'human-centered' decision support tools for future air traffic controllers. One approach to the design of human-centered decision aids is Ecological Interface Design, that focuses on visualizing the boundaries of safe system performance rather than prescribing predetermined strategies or discrete solutions. Previous studies with ecological interfaces in the aviation domain revealed that humans sometimes opted for control actions close to these boundaries, giving rise to a general concern about the robustness of control actions. The goal of this study has been to empirically investigate how effectively an ecological interface for 4D trajectory management, as developed in a previous study, supports the preservation of airspace robustness. For this purpose, a metric has been developed to evaluate both minimum and average sector-based and control-based robustness. Special attention was paid to quantifying and measuring the effect of expertise level on the robustness of human-generated control actions. Results of a human-in-the-loop experiment indicate that expert participants were most robust in their control actions, as compared to either skilled or novice participants. This result suggests that boundary-seeking control actions with ecological interfaces are mainly dependent on the level of expertise and the control strategies of the end user.

4.1 Introduction

FUTURE Air Traffic Management (ATM) is expected to undergo a paradigm shift in the way in which aircraft are controlled [1], [2]. The addition of *time* as an explicit control variable within the concept of Trajectory-Based Operations (TBO) will allow the controller to plan aircraft movements farther in advance, and will allow for the behavior of the system and its components to become more predictable. As a result, TBO will enable the human controllers to step away from their current 'hands-on' form of tactical control, and to become more strategic airspace managers.

The increased dimensionality and abstraction of this new task will require sophisticated computerized tools to separate, organize, and expedite flows of traffic safely and efficiently. There is a general consensus by operational communities that the human controller should remain actively involved in the control loop, and should remain responsible for the safety of operations. Therefore, a so-called 'human-centric' approach to automation is explicitly embraced in order to mitigate known human performance-related issues, such as impaired system understanding, skill degradation, transient workload peaks, and complacency [3]–[6]. However, a definite division between the role of the human controller and automation, and the full extent of these new control tasks are still unclear. As a result, it also remains unclear along what control strategies those computerized tools should guide controllers to meet all system and human-performance requirements.

One possible approach to the development of human-centered decision support tools is the constraint-based 'ecological' approach [7]–[9]. The rationale behind this approach is to let the decision support tool provide the human operator with a set of functional constraints that follow directly from the work domain (or ecology), rather than to provide him/her with explicit solutions. When visualized, these functional constraints will provide the controller with a 'map' of the available control space that does not directly govern a pre-determined control strategy. The benefits of such an approach would be that the operator remains the active decision maker, and that all control actions are safe and good enough as long as they do not violate the functional constraints. In other words, a constraint-based approach would support a myriad of different control strategies, ranging from local short-term fixes to organizing the airspace-wide flow of traffic.

Previous work in aviation has shown, however, that when supported by constraint-based interfaces, operators sometimes opt for solutions that are close to the border between safe and unsafe actions [10]–[13]. That research was primarily focused on egocentric flight deck applications where pilots needed to separate themselves from external hazards, such as terrain or other traffic. When such 'tight' solutions in narrow control spaces are applied in Air Traffic Control (ATC), this could potentially reduce the overall robustness of the airspace to cope with unforeseen events, and therefore might degrade the long-term system stability.

Although it can be argued that tight control actions are promoted—or at least enabled—by constraint-based representations, it can also be argued that the underlying goals, strategies and expertise of the operator play important roles in how such an interface is used. For example, pilots are usually unaware of, or not directly concerned with, streamlining operations within a sector, but are primarily focused on maximizing their own gain. As a result, such form of control can be labeled as more opportunistic. However, in a tightly-coupled air traffic control environment, the effect of one control action on the overall robustness of the system is often not immediately salient, and depends on many interrelated factors (e.g., other traffic, congested areas, preferred routing).

Experienced air traffic controllers have been frequently shown to perform risk aversive control strategies such as formulating backup plans, or by maintaining additional separation buffers between aircraft [14]. Such strategies are focused on mitigating the risk for safety-critical events to arise, and are learned both by formal training and through work experience. Therefore, given that ecological representations allow for a wide variety of control strategies, the level of training, expertise, and experience of the controller is expected to be an important factor in how such an interface is used.

In this chapter we investigate how three user groups with different levels of expertise (i.e., novice, skilled and expert groups) use a constraint-based interface that aims to support them in a future air traffic control task. The decision support interface used in the human-in-the-loop experiment—the previously developed Travel Space Representation (TSR)—is primarily designed for local trajectory revisions of individual aircraft [15]. As such, the goal of this research is to empirically investigate how different expertise groups implicitly take *global* system goals into account when they are working with a constraint-based interface designed to resolve *local* system perturbations. To capture, quantify, and compare the robustness of control actions and between the user groups, this chapter introduces a metric that reflects higher-order and long-term system stability goals in a centralized control setting. Additionally, it is investigated how control strategies shift under the influence of varying levels of perturbation (i.e., from few local to many airspace-wide perturbations) and varying initial traffic structures (i.e., from initially structured corridors to unstructured traffic).

The structure of this chapter is as follows. First, the practical use of the TSR by human controllers, and various classifications of control strategies is discussed. This is followed by an analysis of the robustness metric that has been developed for trajectory-based operations. Next, the experimental design is presented, followed by the results, discussion, and conclusions.

4.2 Travel Space Representation

Inspired by the principles of Ecological Interface Design (EID)[16], [17], the Travel Space Representation is a constraint-based decision support tool that visualizes the boundaries of safe control for the task of short-term trajectory-based air traffic control [15]. Rather than providing one or more discrete optimal trajectory advisories, the TSR visualizes a set of constraints that bound safe and feasible control actions to reroute a selected aircraft.

The general shape of the TSR is determined by the internal aircraft performance constraints. More specifically, the TSR represents the space in which the selected aircraft can be rerouted without exceeding its speed envelope or bank angle limits, but can still realize its planned time at the next waypoint. The additional constraints resulting from external factors (e.g., other traffic and restricted areas) are mapped on top of this shape in the form of no-go areas. Only the horizontal plane has been supported in this work.

The interface focuses on supporting the controller with the task of resolving local perturbations within a single sector that has been de-conflicted a-priori. That is, all aircraft are assumed to follow a predefined 4D path that is initially conflict free. However, as a result of unforeseen events such as delays in other sectors, or the presence of adverse weather, the controller will be required to realign them in order to ensure safe operations. More details on the design of the TSR can be found in previous work [15]; in this article the user interaction with the TSR is central.

4.2.1 Practical use of the Travel Space Representation

The TSR is a Direct Manipulation Interface (DMI) that allows the Air Traffic Controller (ATCo) to select and modify the trajectory of an aircraft by means of click and drag operations with a mouse input device. To illustrate how the TSR can support the controller in a manual trajectory revision task, Fig. 4.1 shows three subsequent images of its use in a hypothetical traffic scenario. The task considered here consists of de-conflicting a selected aircraft (A_{obs}), and rerouting it around a restricted airspace (RA) while meeting the planned sector exit time at waypoint FIX. The initial situation is illustrated in Fig. 4.1(a) in which the observed aircraft, the conflicting aircraft (A_{int}), and the restricted airspace are shown.

When the observed aircraft is selected by clicking on it with the mouse cursor, its TSR is visualized as illustrated in Fig. 4.1(b) (the shaded area extending from the aircraft to point FIX). In this scenario, the TSR consists of two *safe fields of travel* and one *restricted field of travel*. The safe fields of travel indicate the areas in which, when rerouting through any point, the resulting trajectory will both be feasible (i.e., a trajectory that adheres to internal performance limits) and safe (i.e., a trajectory that will resolve the conflict).

Fig. 4.1(b) shows how the controller can select and accept a position within the safe field of travel using the mouse to reroute the aircraft (point WP, indicated by the star symbol). Supported control actions with the TSR include deleting waypoints, modifying the location of a waypoint by dragging it to a

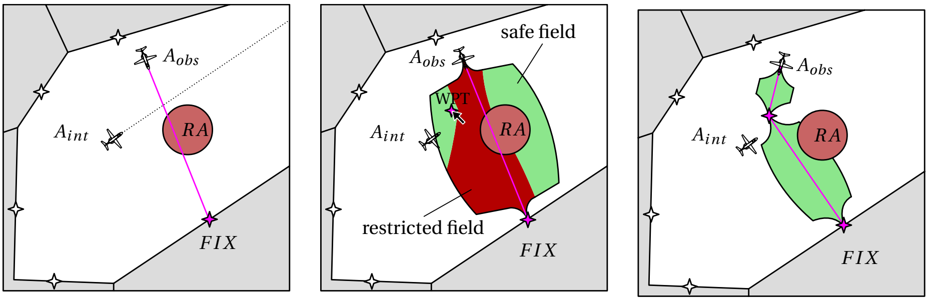

(a) En-route scenario with conflicting traffic and a RA

(b) The TSR and placement of an intermediate waypoint

(c) Resulting trajectory for the observed aircraft

Figure 4.1: Travel Space Representation (TSR) support for the task of manual trajectory revision of an observed aircraft by air traffic control.

new position, and modifying the target time of arrival at a waypoint by using the mouse scroll wheel. By right-clicking on the display, a drop-down menu appears that allows the controller to either reject the modifications, or to accept and send the new trajectory to the aircraft.

Fig. 4.1(c) shows the resulting valid resolution for the selected aircraft by introducing the new waypoint. The original straight trajectory is divided into two equal-speed segments, calculated such that the arrival time at the sector exit fix remains unchanged. Subsequently, a new TSR is visualized for both trajectory segments. It must be emphasized that the TSR is solely a visualization of the task-related constraints and does not physically limit the resolutions that the controller can choose. The controller could still place a waypoint at a location inside the restricted field of travel or outside the TSR. Those solutions will either trigger a conflict or will not guarantee meeting the timing constraints at the sector exit point.

4.2.2 Control strategies and level of expertise

The constraints visualized by the TSR provide the controller with an insight into the set of valid control options for a single selected aircraft. However, when a perturbation involves more than one aircraft, these constraints do not directly impose a predefined set of tasks or control strategies. For example, when two (or more) aircraft are in conflict, the controller must decide whether to resolve the conflict by manipulating one aircraft, or by performing a cooperative resolution. Also, the sequence in which the controller manipulates the aircraft, and the chosen rerouting geometry will both affect the overall dynamics within the airspace. Especially in more complex scenarios that involve the manipulation of multiple aircraft (i.e., control actions are more tightly-coupled), or in situations in which the work domain itself is less predictable (e.g., inherent uncertainties in the execution of trajectories, in unstructured airspace), the ability and expertise of the controller to perform goal-oriented control on a higher conceptual level will become increasingly important to assure the stability of operations.

To investigate how the level of expertise of the controller relates to the decision making process when resolving a perturbation with the constraint-based TSR, six control strategies have been examined. These control strategies have been identified by means of expert opinions and the observed usage of the TSR in a previous validation experiment. For this purpose, the Skills, Rules and Knowledge (SRK) framework developed by Rasmussen and the decision ladder, shown in Fig. 4.2(a), have been used to determine the rule-based shortcuts associated with each of these strategies, labeled 1 to 6 [18], [19]. An information flow map of these control strategies is depicted in Fig. 4.2(b), and shows the unique cognitive chains through the decision making process. The same figure also shows how these strategies have been mapped on the decision ladder.

The various control strategies and their coupling to the level of expertise of the controller can be described as follows:

Novice strategies are mainly expressed by 'if-then'—reactive—types of control that only employ simple rule-based shortcuts. For example, when using the TSR, a novice strategy could be to simply select a perturbed aircraft and to reroute it through a safe field of travel without any specific rationale (strategies 1 and 2). Such strategies will mainly focus on a single aircraft, and will not shy away from resolutions in narrow control spaces. When novice controllers attempt to apply more elaborate strategies, extensive knowledge-based reasoning will be required (strategies 3, 4, 5 and 6), and the effectiveness will be limited by the lack of a comprehensive understanding of the control task.

Skilled strategies rely on rule-based shortcuts that are built upon a basic understanding of the system and its dynamics. To a limited extent, a sequence and/or prioritization of control actions is made. For example, a controller could determine the sequence of aircraft in which to solve a conflict (strategies 3 and 4), or prioritize the order of perturbations to revolve (strategy 5). Such strategies can be seen as satisfactory local, short-term solutions, but will often fail to integrate the long-term airspace stability goals.

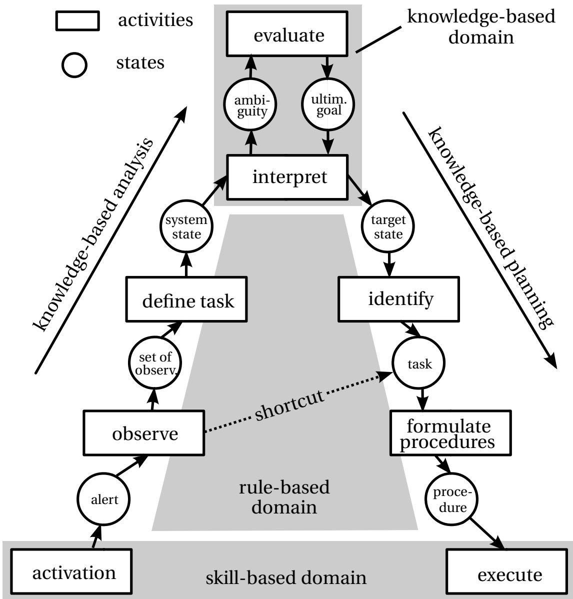

(a) The decision ladder as a tool for identifying skill-, rule- and knowledge-based control strategy

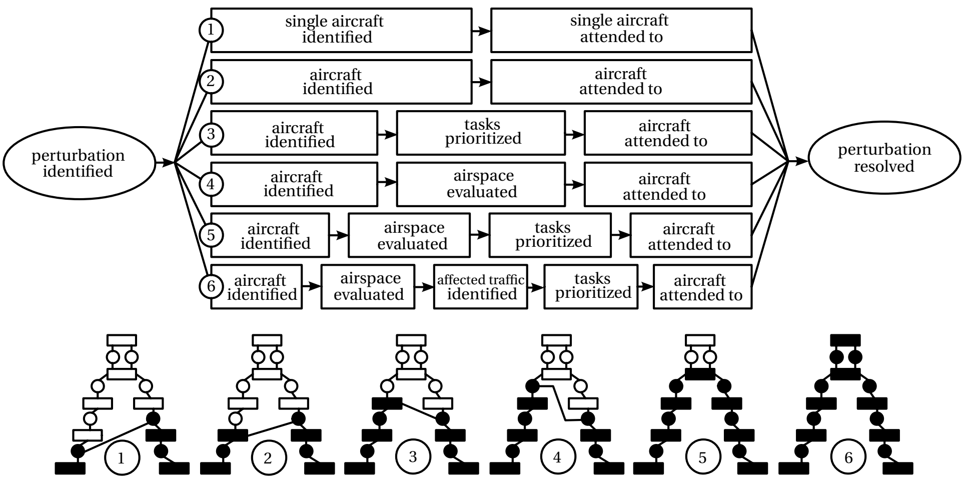

(b) The information flow map illustrating six control strategies to resolving a perturbation with the Travel Space Representation (Adapted from [20]).

Figure 4.2: The decision ladder and information flow map illustrating the cognitive steps of various control strategies when using the Travel Space Representation

Expert strategies typically foster all elements of the most elaborate—pro-active—strategy in the information flow map (strategy 6). However, this does not imply that, in order to perform expert strategies, one must always traverse through the complete decision ladder. Through training, previous experience and a deeper understanding of the work domain, experts will be able formulate high quality rule-based shortcuts that expedite the decision making process. Expert behavior in the current air traffic control system has been found to involve, amongst others, planning multiple steps ahead, formulating backup plans, and maintaining additional buffers in terms of separation to cope with uncertainties [14]. Such strategies focus heavily on finding long-term *robust* resolutions rather than applying short-term fixes.

Applying novice control strategies will typically result in solutions that are less robust to cope with airspace-wide uncertainties than expert strategies. Although the TSR visualizes the explicit constraints that support simple rule-based shortcuts (i.e., the go and no-go areas), the interpretation thereof, and the level of expertise of the controller are expected to largely determine the overall quality of the solutions. For example, simply clicking somewhere in a safe field of travel (reactive novice control) could result in a situation in which the maneuvering space of other traffic is reduced, whereas when using the TSR as a validity check for a thought-through sequence of control actions (pro-active expert control), this could have been foreseen and prevented. The robustness of a control action itself could then be used as an indication of the level of expert behavior of the controller.

4.3 Quantifying Robustness

In order to investigate the various strategies performed by controllers when using the TSR, a metric has been developed with the purpose of quantifying the robustness of control. This metric is not intended to reflect measures that address all aspects of robustness in air traffic operations, but instead has been developed to enable a quantitative post-hoc, between-subject comparison in this study. The metric enables an evaluation of the robustness of all airspace users in a given sector, and moreover, the difference in their robustness as a result of a control action with the TSR. In essence, this allows to determine whether a trajectory revision for one flight has a positive or detrimental effect to the overall stability of the system. This metric finds its origin in the flexibility preservation metric proposed by Idris [21], [22], and is based upon quantifying the probability of a trajectory to remain feasible despite probabilistic disturbances in its execution.

4.3.1 Robustness

Robustness is a quantitative measure of *trajectory flexibility* that has been defined as *"the ability of a flight to adhere to planned trajectory and imposed constraints, despite probabilistic random state deviations from that trajectory"* [21]. Such a trajectory will for instance remain feasible (i.e., no conflicts or restricted airspace violations materialize) despite a deviation in speed and/or heading at a certain point. In its original implementation, this metric has been used as a factor to quantify and select the 'best' trajectory from a set of recursively generated trajectories [21]. For this study the metric has been modified and discretized to allow for a post-hoc investigation of the robustness of control actions.

4.3.2 Point-based robustness

In this metric, robustness can be seen as a point-based attribute at each position along a given trajectory. This point-based robustness acts as the basic building block with which trajectory-based, sector-based and control-based robustness can be derived. Consider an observed aircraft with a given time-based intent. At each point in time that aircraft is assumed to be at a predicted state (t, x, y, V, ψ). Taking the predicted state as a starting point, the robustness at that point, $RBT(t)$, can be quantified by the probability of feasibility $P_f(t)$ of the aircraft to successfully reach a set of next states at time $t + \Delta t$.

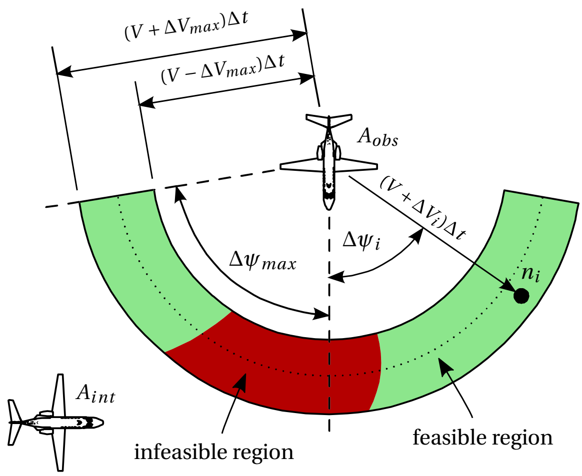

Figure 4.3: Sketch of the point-based robustness geometry and the resulting area of probabilistic feasibility at a discrete point.

To illustrate this, Fig. 4.3 shows the predicted position at time t of an observed aircraft (A_{obs}) along its trajectory. At that point, a set of disturbances acting on the heading and speed can be modeled using a probabilistic disturbance model. In Fig. 4.3 these disturbances are represented by a heading offset ($-\Delta\psi_{max} \leq \Delta\psi \leq \Delta\psi_{max}$), and a disturbance in the velocity of the aircraft ($-\Delta V_{max} \leq \Delta V \leq \Delta V_{max}$). When considering the propagation of a single heading and speed disturbance (respectively, $\Delta\psi_i$ and ΔV_i) for a time interval Δt, a new probe segment (n_i) can be constructed. Neglecting wind, this results in a disc-shaped area within which the aircraft is predicted to be at time $t + \Delta t$. When the disturbances are discretized, this results in $N(t)$ probe segments, each with probability $P_i(t)$, such that the sum of probabilities is equal to one. For this study, the probability of each probe segment is assumed to be equal, resulting in:

$$P_i(t) = \frac{1}{N(t)}, \text{ with } \sum_{i=1}^{N} P_i(t) = 1 \tag{4.1}$$

When checking the feasibility of each segment instance with respect to the system constraints (i.e., other aircraft predicted positions and restricted area intrusions), the set of $N(t)$ probe segments can be divided into two mutually exclusive sets—the set of feasible $N_f(t)$ and unfeasible $N_i(t)$ segments. In Fig. 4.3, the green area in the disc indicates the set of feasible segments, and the red area indicates the set of infeasible segments due to a predicted loss of separation with the other aircraft, A_{int}. Per definition, the robustness at the considered point of the trajectory is then given by:

$$RBT(t) = \frac{N_f(t)}{N(t)} \tag{4.2}$$

4.3.3 Trajectory-based robustness

When point-based robustness is evaluated along a given trajectory (trajectory itself again discretized in N points), the robustness of the trajectory can be evaluated over time. Moreover, two distinct trajectory-based robustness measures can be derived as follows:

- *Minimum trajectory robustness:* the point of least robustness along a trajectory, given by:

$$RBT_{T_{min}} = \min\{RBT(t)\}$$

- *Average trajectory robustness:* the average robustness along a trajectory, given by:

$$RBT_{T_{avg}} = \frac{1}{N}\sum_{i=1}^{N} RBT(t)$$

These measures allow for a quantitative comparison between two or more trajectory instances. Such a comparison is, for example, the difference in robustness between two valid conflict resolution geometries by using the TSR. Consider the Travel Space of an observed aircraft (A_{obs}) that is in conflict with a second aircraft (A_{int}) as shown in Fig. 4.4. The figure shows two valid resolution strategies, one by resolving the conflict by passing the other aircraft in front (T_f) and one by passing behind (T_b). By observing the TSR alone, the implications on the resulting robustness of any of the two resolutions is not immediately apparent (i.e., both trajectories require approximately the same amount of added track length and increase in speed). Here, when applying a novice control strategy (strategies 1 and 2 from Fig. 4.2(b)), no distinct preference would be given to either of the resolutions.

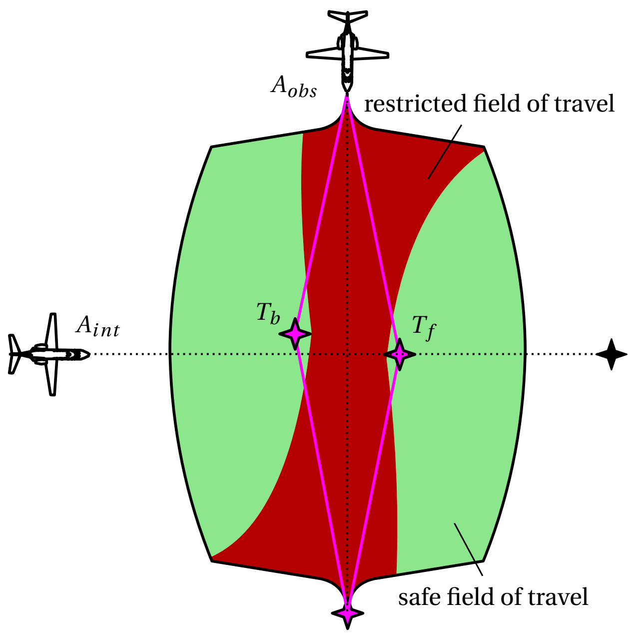

Figure 4.4: Trajectory-based conflict resolution for a crossing pair of aircraft using the Travel Space Representation.

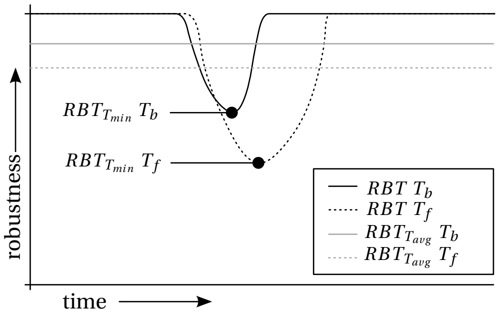

Figure 4.5: Robustness versus time for two valid trajectory resolutions.

In Fig. 4.5 a sketch is given of the robustness over time of the *observed aircraft* for both maneuvers. This sketch shows that there is a significant difference when looking at the resulting robustness; the passing behind maneuver (standard ATC practice) 'preserves' the trajectory-based robustness better than passing in front. The minimum robustness is highest for the passing behind resolution. This indicates that at the most 'critical' point along the trajectory (i.e., the point with least tolerance to disturbances), resolution T_b is more probable to remain feasible than T_f. Further, by comparing $RBT_{T_{avg}}$, it can be seen that the average robustness against disturbances along the trajectory is also higher for the passing behind maneuver. In essence, when only considering the observed aircraft, passing behind would reflect a more expert resolution strategy.

4.3.4 Sector-based robustness

In the previous example, the metric was used to evaluate the difference in robustness for only the observed aircraft. Here, the influence of the resolution to the robustness of aircraft A_{int} was not taken into account. To quantify the robustness of controller strategies, however, it is of more interest to measure the difference in robustness on a sector-wide level.

Similar to the trajectory-based robustness, a minimum and average sector-based robustness can be derived. The minimum sector robustness, $RBT_{S_{min}}$, is defined to be equal to the lowest value of trajectory-based robustness of all aircraft in a given sector. This measure then relates to the point in time at which a controlled aircraft is least robust to disturbances. For instance, in case of a predicted loss of separation in the sector, $RBT_{S_{min}}$ will be zero. The average sector robustness, $RBT_{S_{avg}}$, is defined as the average of $RBT_{T_{avg}}$ of all aircraft, and can be used as a more general indication of the current state of robustness of the sector.

4.3.5 Control-based robustness

To investigate how a given control action influences the robustness of the system, a per-aircraft difference can be calculated of the trajectory-based robustness at a moment in time just before, and just after a control action is issued. These two measures then become $\Delta RBT_{T_{min}}$ and $\Delta RBT_{T_{avg}}$, respectively the difference in minimum and average robustness of the trajectory as a result of the control action. A positive value of $\Delta RBT_{T_{min}}$ for a given trajectory indicates that the tolerance to disturbances at the point of least robustness (i.e., the bottleneck of the trajectory) has increased. Conversely, when the value is negative, a state disturbance at the bottleneck is more likely to violate one or more constraints. Note that the location of the bottleneck itself can also change due to the control action. Similarly, a positive value of $\Delta RBT_{T_{avg}}$ indicates that the average robustness to disturbances has increased along a given trajectory, and a negative value indicates a decrease. For both measures, a zero value indicates that the robustness of the observed aircraft is not affected by the control action. These measures allow for a quantitative investigation of how a certain control action influences the robustness of all other aircraft.

4.3.6 Control strategies and robustness

It is expected that the difference between the novice, skilled and expert control strategies with the TSR will be reflected by the robustness metrics. The solutions in narrow control spaces, and lack of planning ahead in novice strategies will likely result in both the lowest sector-based and control-based robustness. Skilled strategies will be more robust on a local level (i.e., high trajectory-based robustness of the controlled aircraft), but will likely fail to score high on overall sector-based robustness. Finally, expert strategies such as maintaining additional separation buffers and anticipating the progress of the system-state when manipulating the traffic will likely result in both the highest values of control-based and sector-based robustness.

4.4 Human-in-the-loop Experiment

To evaluate how the level of expertise of controllers affects robustness when using the constraint-based TSR, a human-in-the-loop experiment was performed. Three groups of participants with an increasing level of expertise were asked to manage various scenarios of trajectory-based air traffic without the aid of any automated advisories (i.e., by using the TSR alone). The goal of the experiment was twofold: (1) to investigate what type of control strategies the three groups of controllers apply when using the TSR and (2) to identify whether their control strategies shift under varying traffic and airspace settings (from low to high complexity).

4.4.1 Participants

The experiment was performed with a total of twelve participants divided into three groups with increasing operational ATC experience. The novice group consisted of four Ph.D. students who perform flight-deck and/or ATM-automation related research (3 male, 1 female, average age of 30). None of the novice participants received any prior training in operational air traffic control. The skilled group consisted of four domain experts who are currently working as professionals in ATM research and development (4 male, average age of 54). Finally, the expert group consisted of four operational Area Control Center (ACC) air traffic controllers, two of which were fully certified and the other two in on-the-job training (2 male, 2 female, average age of 27).

4.4.2 Procedure

First, the participants were given an initial briefing of the general concept of TBO and an outline of the experiment. This was followed by a 20 minute training session during which the participants were asked to follow an interactive script to become familiarized with the TSR and its functionality. The training ended when the script was completed, and when the participants indicated that they had a good understanding of how to use the TSR to manage the traffic.

In the main experiment, the participants were asked to manage traffic within a fictional two-dimensional sector under various initial conditions. During each run the overall goal was to plan and guide the traffic through the controlled sector safely (i.e., without losses of separation or restricted area intrusions) and efficiently (i.e., adhere to timing constraints at the sector exit points). After the initialization of an experiment scenario, the participants controlled the traffic by issuing changes to the 4D-trajectories of each individual aircraft by manipulating waypoints using the TSR. The resulting trajectories were automatically calculated, and were executed by the aircraft on acceptance by the participant. No additional input was required by the experimenter during each run. Each experiment ended with a short debrief of approximately 10 minutes. During the debrief the participants were asked to provide feedback on the positive and negative aspects of the TSR as a means to control air traffic.

4.4.3 Apparatus

The evaluation was performed on a dedicated software-based ATM platform, running on a single computer. The TSR was integrated in a traditional plan view display (PVD), providing a top-down view of airspace and air traffic. The TSR was presented on a 30 inch screen (60 Hz LED, 2560 x 1600 pixels) placed in front of the participant. Input was given by a standard mouse input device, and control options could be selected by on-screen drop-down menus.

4.4.4 Independent variables

The experiment followed a mixed design with the three levels of expertise (novice, skilled and expert) as a between-group variable, and two within-subject independent variables, that were:

- *Orderliness:* the initial traffic orderliness, with two levels: structured traffic (T_S) and unstructured traffic (T_U), and
- *Perturbation:* the number of aircraft in the traffic sample that were required to be rerouted in order to prevent losses of separation or restricted airspace intrusions, with three levels: small perturbation (P_S), medium perturbation (P_M), and large perturbation (P_L).

The *orderliness* variable defined the initial traffic set-up of the scenario. The rationale for choosing this variable has been that the form in which TBO in general will be implemented is not yet definite. Therefore both structured (fixed route-structures) and unstructured (free-routing) traffic conditions have been considered in this research. In structured traffic (T_S), all aircraft traversed the sector through a number of structured—predictable—streams as shown in Fig. 4.6(a). This implied that aircraft initially traversed the sector in-trail on a limited set of fixed routes. In unstructured traffic (T_U), all aircraft would enter and exit the sector by a unique combination (entry/exit point) of the eight fixed waypoints as shown in Fig. 4.6(b). In general, unstructured traffic is less predictable, and will require more knowledge-based analysis and planning by the controller to find robust resolutions to perturbations. Only one baseline structured and one baseline unstructured initial traffic scenario were used in all six scenarios.

The perturbation variable was defined by the minimum number of aircraft that the participant initially had to realign in order to resolve all conflicts and restricted area violations. The initial conflicts were purely geometrical; all aircraft entered the sector with the same speed, thus catch-up and/or overtake scenarios were not considered. In the small perturbation (P_S) condition three aircraft pairs were deliberately put into a conflict and had to be rerouted (Fig. 4.6(a)). In the medium perturbation (P_M) condition a restricted area (circular area with a radius of 10 NM) was added in the sector at a location that required the additional rerouting of five aircraft. In the large perturbation (P_L) condition the restricted area was placed at a location that required a total of seven additional aircraft to be rerouted (Fig. 4.6(b)). A higher level of perturbations will reduce the available time for the controller to analyze control strategies, and will reduce the overall control space in which an aircraft can be rerouted. This is foreseen to invoke lower level rule-based strategies.

The control variables (i.e., the variables that remained fixed) in the experiment were: the sector area, size and shape, the availability of the Travel Space decision-support tool, the size and shape of the restricted area, and the initial traffic sample in structured and unstructured conditions.

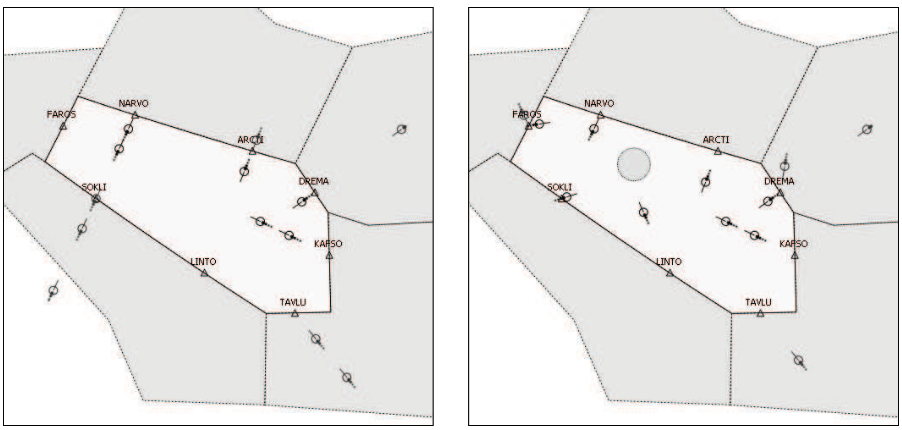

(a) T_S-P_S condition (b) T_U-P_L condition

Figure 4.6: Screen shots of two distinct experiment conditions.

4.4.5 Dependent measures

The following dependent measures were used to investigate the effect of the level of expertise, the traffic orderliness and the perturbation scale on the effectiveness of the TSR:

- *Safety:* The number of losses of separation and restricted area intrusions per condition,
- *Control performance:* The number of control actions the participants performed to resolve a scenario,
- *Sector-based robustness:* The overall minimum and average robustness of the aircraft trajectories that is an indication of the expertise level of control strategies,
- *Control-based robustness:* The effect of the individual control actions on the trajectory-based robustness of all aircraft within the sector, and
- *Performance:* defined by the additional track length flown by the rerouted aircraft

4.4.6 Scenarios

The participants were asked to manage traffic in a hypothetical en-route sector ($\approx$ 40.000 KM^2) under the six different control conditions (Table 4.1). The rotation of the sector varied uniquely between scenarios consisting of the same (baseline) traffic structure to avoid a control bias due to scenario recognition. The names of the waypoints and aircraft were also varied in each run to prevent the aforementioned bias.

Each scenario presented approximately 15 aircraft and eight sector entry/exit points and lasted 24 minutes in scenario-time. The simulation ran at four times the normal speed, such that each scenario lasted six minutes in real-time. The average traffic density was set to approximately 8 controlled aircraft at any given point, with exception of the first and last minute of the scenarios (real-time), in which the traffic either built up or reduced to compensate for the absence of hand-overs in between sectors, and the lack of verbal communication.

All aircraft entered the controlled sector at FL300 through one of eight fixed waypoints on the sector border and were given an initial (straight) 4D trajectory leading towards one of the other waypoints. Aircraft could only be controlled laterally (i.e., vertical manipulation of the trajectories was not possible), and

Table 4.1: Definition of the six experiment conditions.

condition	orderliness	perturbation
T_S-P_S	structured	small
T_S-P_M	structured	medium
T_S-P_L	structured	large
T_U-P_S	unstructured	small
T_U-P_M	unstructured	medium
T_U-P_L	unstructured	large

only if they were physically inside the sector. Nevertheless, aircraft inbound to the sector were shown in gray when approaching the sector, such that the participants had ten minutes (scenario-time) to prepare for future traffic situations. Further, the performance characteristics of all aircraft have been simulated by a single generic aircraft type.

The initial conditions of each scenario were set such that the controller had to resolve a fixed number of perturbations (conflicting pairs of aircraft and avoiding restricted areas) by manipulating the trajectories of individual aircraft. However, the control actions themselves could create new conflicts and restricted area intrusions further ahead in time.

4.4.7 Hypotheses

It was hypothesized that the TSR would allow all participants to safely control the traffic and resolve all perturbations, but also that the level of expertise of the controller would have an influence on the overall robustness of the solutions. That is, a higher level of expertise would result in control strategies that preserve a higher level of sector-based and control-based robustness. It was also hypothesized that the control strategies of participants would shift under the various experiment conditions. Unstructured traffic and an increasing level of perturbations were foreseen to invoke strategies that use lower level rule-based shortcuts for all participants.

4.4.8 Data analysis

To test for between-group and within-group effects, Kruskal-Wallis and Friedman tests have been performed, respectively. Here, the significance level (α) has been set to 0.05. Post-hoc, pair-wise comparison tests between experimental conditions featured either three or four planned pair-wise comparisons to investigate the effect of the traffic orderliness and the scale of the perturbation respectively.

For the post-hoc calculation of sector-based and control-based robustness, the maximum heading disturbance ($\Delta\psi_{max}$) has been set to 80 degrees, and the maximum speed disturbance (ΔV_{max}) to 20 kts IAS ($\approx$ M.05, or 30 kts ground speed). For the calculation of point-based robustness, the heading range was discretized in steps of 5 degrees, and the speed range in steps of 2.5 kts IAS, resulting in a total of 561 probe segments (n_i) per point. Point-based robustness was sampled at every second along the trajectory. The duration of a probabilistic disturbance (Δt) was set to 120 seconds (a typical ATC Short Term Conflict Alert time window). Although the magnitude of the disturbances might seem high compared to what could reasonably be expected in TBO, such 'extreme' values are expected to magnify the between-group and within-subject variation of robustness (i.e., only relative differences are compared).

4.5 Results

4.5.1 Safety

Out of the 1,116 controlled flights, three safety-critical events occurred, two losses of separation and one restricted area intrusion. The first loss of separation (skilled participant in the T_S-P_M condition) occurred as a result of a previous control action, after which a resulting short-term multi-aircraft conflict could not be resolved in time. The second loss of separation (novice participant in the T_U-P_L condition) occurred when the participant attempted to resolve a two-aircraft conflict, but forgot to 'send' the modified trajectory the the aircraft. The participant only noticed his/her mistake after the actual loss of separation had occurred. The restricted area intrusion occurred for a skilled participant in the T_U-P_M condition, and was only actively attended to when the aircraft was close to the center of the circular restricted area. Further, minimum sector-based robustness (i.e., the minimum control space in a given run) is also an implicit measure of safety, and will be discussed in a next paragraph.

4.5.2 Control performance

Fig. 4.7 shows a box plot of the number of control actions per condition, grouped per participants group. The figure shows that the expert participant group, on average, performed a higher number control actions in any given condition. This difference is especially pronounced in the low perturbation (T_S-P_S, T_U-P_S) and unstructured traffic conditions (T_U). Closer inspection of the data revealed that the expert group was frequently adopting a *nursing* strategy. That is, they implemented a larger number of small trajectory revisions rather than fewer, but more coarse control actions.

The experiment condition was found to significantly affect the number of control actions (Friedman: $\chi^2(5) = 44.28, p < 0.01$). A post-hoc Wilcoxon test showed that more control actions were issued in scenarios with a higher initial perturbation level. This was to be expected, considering that a larger perturbation level required more control actions by design.

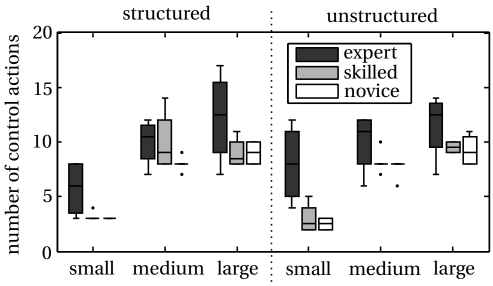

Figure 4.7: Box plot of the number of control actions per condition.

4.5.3 Sector-based robustness

Fig. 4.8(a) shows a box plot of the *minimum* sector-based robustness per condition, and grouped per participant group. The figure represents the lowest value of robustness for any aircraft during a run. On average the minimum robustness is highest for the expert participant group. This could indicate that experts prefer to avoid steering in narrow control spaces when using the TSR. The difference in minimum robustness between the three groups is even more pronounced in the unstructured traffic conditions, which could indicate that in more complex traffic patterns, the experts increase the spatial separation buffer.

Further, the experimental condition was found to significantly affect the minimum robustness (Friedman $\chi^2(5) = 14.18, p = 0.02$). The minimum robustness is significantly lower for unstructured traffic in the small (Bonferroni correction: $\alpha = 0.05/3 = 0.01$, Wilcoxon: $z = -2.67, p < 0.01$) and large perturbation (Wilcoxon: $z = -2.67, p < 0.01$) cases. This can be an effect of the higher number of trajectory crossings, that inherently reduces control space in unstructured traffic conditions.

Fig. 4.8(b) shows a box plot of the *average* sector-based robustness per condition, grouped per participant group. This metric reflects the average trajectory-based robustness of all aircraft during a run, and relates to the average robustness of the airspace. The average sector-based robustness was found to differ significantly between participant groups (Kruskal-Wallis $H(2) = 10.75, p < 0.01$). A post-hoc Mann-Whitney U-test showed that the average robustness was significantly higher for the expert participant group compared to skilled participants ($U = 1, z = -2.72, p < 0.01$) and novices ($U = 0, z = -2.88, p < 0.01$). The higher average robustness indicates that expert strategies indeed preserve robustness better than skilled and novice strategies.

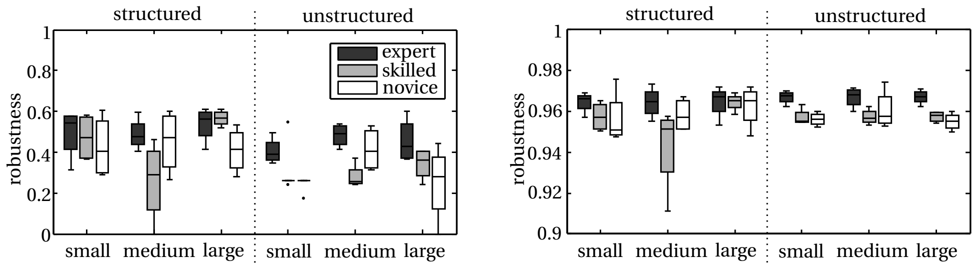

(a) $RBT_{S_{min}}$ (b) $RBT_{S_{avg}}$

Figure 4.8: Box plot of the minimum and average sector robustness per condition.

4.5.4 Control-based robustness

Fig. 4.9(a) and Fig. 4.9(b) show box plots of the total number of aircraft whose *minimum* trajectory-based robustness had been affected due to operator control actions. Fig. 4.9(a) shows the number of flights that have been positively affected, and 4.9(b) shows the number of flights that have been negatively affected. In this measure, two or more control actions that affect the minimum trajectory-based robustness of a given flight count multiple times.

Both figures show that, on average, control actions by the expert group affected the minimum trajectory-based robustness of more aircraft compared to the skilled and novice participants. This trend is most likely caused by the fact that the expert participants perform a larger number of minor trajectory revisions to resolve any given scenario; even for small corrections, any control action is likely to (albeit slightly) affect the minimum trajectory-based robustness of other aircraft.

Figures 4.9(c) and 4.9(d) show box plots of the *average* (i.e., per control action) and *cumulative* (i.e., summed per experiment condition) change in $\Delta RBT_{T_{min}}$ of all affected flights. Fig. 4.9(c) shows that the magnitude of $\Delta RBT_{T_{min}}$ per control action is, on average, smaller for the expert participant group compared to the other groups. Fig. 4.9(d), however, shows that the cumulative sum does not vary as much. This again strengthens the finding that participants in the expert group perform more, but finer, trajectory revisions.

Fig. 4.9(d) shows a reverse in the between-group trend between structured and unstructured traffic conditions. That is, in structured traffic conditions that are closer to the current form of ATC operations, experts add *less*, or equal total robustness compared to the other groups. In unstructured traffic, however, experts add *more* total robustness. This indicates that the expert participants maintain additional safety-buffers in less familiar traffic conditions.

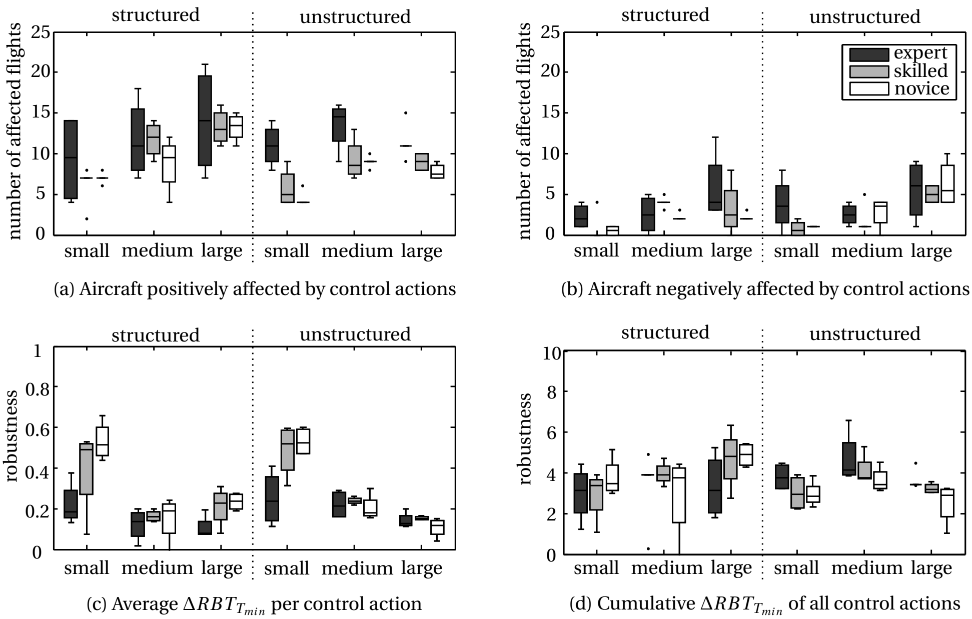

(a) Aircraft positively affected by control actions

(b) Aircraft negatively affected by control actions

(c) Average $\Delta RBT_{T_{min}}$ per control action

(d) Cumulative $\Delta RBT_{T_{min}}$ of all control actions

Figure 4.9: Box plots of the number of aircraft affected by control actions, and the average and cumulative contributions to $\Delta RBT_{T_{min}}$ by these control actions, per condition.

4.5.5 Performance

Fig. 4.10 shows a box plot of the number of added track miles per condition, grouped per participant group. This measure shows the total deviation in nautical miles with respect to the original straight aircraft trajectories. The figure shows that, on average, the number of added track miles is higher for the expert group. This trend is especially pronounced in the unstructured traffic conditions, and is in line with the results of minimum and average scenario-based robustness. Though not statistically significant, visual inspection of the data showed that the participants in the expert group often opted for more conservative resolution strategies rather than minimizing track deviation by steering close to the constraint boundaries.

Further, the experiment condition was found to have a significant effect on the added track miles (Friedman: $\chi^2(5) = 38.86, p < 0.01$). A post-hoc Wilcoxon rank test showed that the perturbation scale significantly influenced the added track miles in the unstructured traffic conditions. Here, the added track miles are significantly higher in the medium perturbation case compared to the small perturbation scale (Bonferroni correction: $\alpha = 0.05/4 = 0.01$, Wilcoxon: $z = -3.06, p < 0.01$), and significantly lower in the large perturbation condition compared to the medium perturbation case (Wilcoxon $z = -2.90, p < 0.01$). Closer inspection of the data showed that this unexpected result is likely due to scenario design. Although the medium perturbation scale condition in unstructured traffic required less aircraft to be rerouted, the aircraft had to deviate farther from their path.

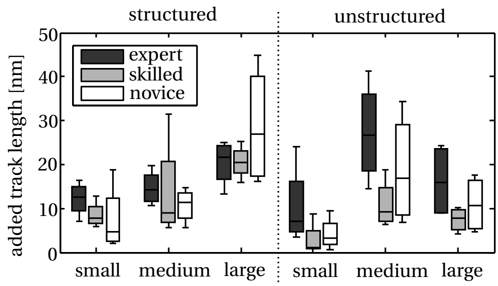

Figure 4.10: Box plot of added track miles per condition.

4.5.6 Debrief

During the post experiment debrief participants reported that the Travel Space Representation is fairly intuitive to use and provides a clear overview of the rerouting possibilities for a selected aircraft. Many participants indicated that they liked the fact that they could plan ahead due to the time-based nature of operations. A common remark given by participants from all groups was that they disliked the fact that constraints imposed by aircraft outside of the controlled sector were not shown in the TSR of a selected flight. This could result in an unanticipated conflict occurring when a new aircraft entered the sector.

Participants from all groups indicated that they sometimes could not link a restricted field of travel in the Travel Space Representation to a specific aircraft. Here, one novice participant commented that this reduced situation awareness, and one participant from the expert group commented not trusting the constraint zones. A second novice participant remarked that sometimes the green zones (e.g., the safe field of travel) form very narrow solution areas. An expert participant commented that trusting the tool felt doubtful as there were many 'close calls' between aircraft. A second expert participant reported not feeling comfortable working with this system. Finally, a general consensus was that more training would be necessary to get a better understanding of how the TSR responds under varying conditions.

4.6 Discussion

Ecological, constraint-based approaches to the design of decision support interfaces have shown to be promising in supporting more egocentric control tasks. However, there are concerns that such an approach would not be equally applicable to the development of decision support tools for the more centralized ATM domain. One concern is that a decentralized, distributed form of control could have a detrimental effect on the flexibility within the tightly-coupled airspace system to cope with inherent variability (i.e., reducing the control space within a sector). Secondly, there is a concern that constraint-based representations could promote boundary-seeking control by operators, resulting in less robust operations as they are close to the border between what is safe and unsafe.

The results from the human-in-the-loop experiment showed that the level of expertise of the operator mainly determined the nature of the chosen control actions when using the Travel Space Representation. The more experienced participants opted for more conservative solutions, and thereby maintaining a higher level of robustness, whereas domain experts and students with no operational experience often opted for tighter control spaces. This difference in control behavior can be explained by the following three factors:

1. The experts indicated that they had relatively low trust in the visualizations, making them more cautious about implementing 'tight' solutions. Apparently, the experts took constraints associated with uncertainties in the displayed information into account, which is something that the current Travel Space Representation does not make explicit. Such strategies rely heavily on knowledge-based analysis and planning.

2. The experts followed control strategies that were inherently more robust than the strategies adopted by skilled participants and novices (e.g., opting for 'passing-behind' control actions instead of 'passing-in-front' op-

tions). As some of the skilled participants and novices were not aware of this preferred strategy—that also appears to be more common practice in the current ATM domain—they were less robust in their control actions.

3. The experts were much more pro-actively controlling the traffic (as observed by the higher number of control actions), whereas the skilled participants and novices seemed to apply a more reactive form of control.

Based on these observations, in order to bring skilled and novice controllers closer to an expert performance level, it is recommended to make constraints associated with current practice (such as uncertainties and preferred strategies) more explicit in the representation. Presenting such information will provide additional information that underlies the trade-offs between robustness and efficiency, and will facilitate knowledge-based reasoning. Finally, although the de-centralized nature of control did not appear to negatively influence robustness in this study, it does not mean that this type of control is desirable in more dense traffic situations and/or in the presence of larger scale perturbations. In such cases a higher-level, flow-based control strategy might prove more efficient. That is, a representation supporting a strategy where constraints associated with aircraft flows and streams can be defined instead of manipulating constraints directly tied to the individual particles (i.e., aircraft). For future research these considerations will be further investigated.

4.7 Conclusion

The goal of the study has been to empirically investigate how effective an ecological interface for 4D trajectory management (i.e., the Travel Space Representation) supports the preservation of airspace robustness. In an experiment the initial traffic orderliness and the scale of trajectory perturbations have been been varied. The human-generated control actions were evaluated in terms of robustness, safety, and efficiency among three user groups with a varying level of expertise. Results show that robustness was mainly influenced by the expertise level of the participants. This indicates that ecological interfaces would be most effective for *experienced* domain experts, as they can make smarter, knowledge-based, decisions with regard to the tradeoff between safety and efficiency.

Bibliography

[1] "SESAR CONCEPT OF OPERATIONS," EUROCONTROL, Brussels, Belgium, Report, 2007.

[2] "Concept of Operations for the Next Generation Air Transportation System, Version 3.2," Joint Planning and Development Office, Washington, DC, USA, Report, 2011.

[3] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

[4] M. R. Endsley and D. B. Kaber, "Level of Automation Effects on Performance, Situation Awareness and Workload in a Dynamic Control Task," *Ergonomics*, vol. 42, no. 3, pp. 462–492, 1999.

[5] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation.," *IEEE Transactions on Systems, Man, and Cybernetics. Part A, Systems and Humans*, vol. 30, no. 3, pp. 286–297, 2000.

[6] ——, "Situation awareness, mental workload, and trust in automation: Viable, empirically supported cognitive engineering constructs," *Journal of Cognitive Engineering and Decision Making*, vol. 2, no. 2, pp. 140–160, 2008.

[7] K. J. Vicente and J. Rasmussen, "The Ecology of Human-Machine Systems II: Mediating Direct-Perception in Complex Work Domains," *Ecological Psychology*, vol. 2, no. 3, pp. 207–249, 1990.

[8] ——, "Ecological Interface Design: Theoretical Foundations.," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 22, no. 4, pp. 589–606, 1992.

[9] J. Rasmussen, "Ecological Interface Design for Reliable Human-Machine Systems," *The International Journal of Aviation Psychology*, vol. 9, no. 3, pp. 203–223, 1999.

[10] C. Borst, H. C. H. Suijkerbuijk, M. Mulder, and M. M. Van Paassen, "Ecological Interface Design for terrain awareness," *The International Journal of Aviation Psychology*, vol. 16, no. 4, pp. 375–400, 2006.

[11] S. B. J. van Dam, M. Mulder, and M. M. Van Paassen, "Ecological Interface Design of a tactical airborne separation assistance tool," *IEEE Transactions on Systems, Man and Cybernetics, Part A*, vol. 38, no. 6, pp. 1221–1233, 2008.

[12] C. Borst, M. Mulder, and M. M. Van Paassen, "Design and Simulator Evaluation of an Ecological Synthetic Vision Display," *Journal of Guidance, Control, and Dynamics*, vol. 33, no. 5, pp. 1577–1591, 2010.

[13] J. Ellerbroek, K. C. R. Brantegem, M. M. Van Paassen, N. de Gelder, and M. Mulder, "Experimental Evaluation of a Coplanar Airborne Separation Display," *Human-Machine Systems, IEEE Transactions on*, vol. 43, no. 3, pp. 290–301, 2013.

[14] J.-F. D'Arcy and P. S. Della, "Air Traffic Control Specialist Decision Making and Strategic Planning - A Field Survey," Federal Aviation Administration, Atlantic City, NJ, USA, Report, 2001.

[15] M. M. Van Paassen, C. Borst, R. E. Klomp, M. Mulder, P. van Leeuwen, and M. Mooij, "Designing for Shared Cognition in Air Traffic Management," *Journal of Aerospace Operations*, vol. 2, pp. 39–51, 2013.

[16] C. M. Burns and J. Hajdukiewicz, *Ecological Interface Design*. Boca Raton, FL, USA: CRC Press, 2004.

[17] J. M. Flach, K. J. Vicente, F. Tanabe, K. Monta, and J. Rasmussen, "An Ecological Approach to Interface Design," in *Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting*, 1998, pp. 295–299.

[18] J. Rasmussen, "Skills, Rules, and Knowledge; Signals, Signs, and Symbols, and other Distinctions in Human Performance Models," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 13, no. 3, pp. 257–266, 1983.

[19] ——, *Information Processing and Human-machine Interaction: An Approach to Cognitive Engineering*, A. P. Sage, Ed., ser. North-Holland series in system science and engineering. New York, NY, USA: Elsevier Science Publishers, 1986, vol. 12.

[20] E. M. Roth and A. M. Bisantz, "Cognitive Work Analysis," in *The Oxford Handbook of Cognitive Engineering*, Oxford University Press, 2013, ch. 14, pp. 240–260.

[21] H. R. Idris, R. Vivona, J.-L. Garcia-Chico, and D. Wing, "Distributed Traffic Complexity Management by Preserving Trajectory Flexibility," in *Proceedings of the 26th IEEE/AIAA Digital Avionics Systems Conference*, 2007, pp. 1–13.

[22] H. R. Idris, T. El-wakil, and D. J. Wing, "Trajectory Planning by Preserving Flexibility: Metrics and Analysis," in *Proceedings of the AIAA Guidance, Navigation, and Control (GNC) Conference and Exhibit*, 2008, pp. 1–14.

CHAPTER 5

INTENTIONAL CONSTRAINTS

An expert is someone who knows some of the worst mistakes that can be made in his subject, and how to avoid them. —*Werner Heisenberg (Theoretical Physicist, 1901–1976)*

In this chapter, the human-in-the-loop experiment of the previous chapter is repeated with the addition of intentional constraints to the original horizontal Solution Space. As in the previous chapter, the horizontal Solution Space is named Travel Space in this chapter. Results of a fast-time simulation are presented that quantify the effect of the conflict angle and additional separation on the robustness of control actions. Further, the design choices of visual mapping of the intentional constraints to the original horizontal Solution Space display are discussed. The set-up of the second human-in-the-loop experiment is presented, together with the experiment results. The chapter concludes with a discussion of the results.

ABSTRACT

Higher levels of automation and the introduction of novel decision-support tools will play a key role in enabling future trajectory-based operations in air traffic control. The ultimate goal is to increase the capacity, efficiency and flexibility of airspace-use, whilst retaining the current high standards of safety. In previous research, following the principles of Ecological Interface Design, a constraint-based decision-support tool has been developed for the task of tactical trajectory manipulation. Rather than presenting discrete optimized solutions to the controller, this Travel Space Representation visualizes the constraints for safe control in the form of a set of go and no-go areas. A human-in-the-loop experiment showed that although the interface enabled safe control by all participant groups, the quality of control in terms of robustness mainly depended on the level of expertise and prior experience of the controller. Less skilled controllers sometimes opted for solutions close to the boundary of safe control and/or for less preferable conflict resolution geometries. These results gave rise to a concern that such a representation could actually work against the flexibility of the system to cope with inherent system variability. Following these findings, the goal of the present study has been twofold; (1) to investigate the effect of conflict resolution geometry on resulting system robustness, and (2) to investigate how the minimum separation distance and the variance in aircraft speed affect the available control space in the Travel Space Representation. As a result, a set of intentional constraints has been identified that highlights the sub-set of control actions that are undesirable in terms of airspace robustness and trajectory adaptability. Results of a human-in-the-loop experiment show that the quality of control of novice and skilled participants better matches expert control when these intentional constraints are explicitly visualized on the interface.

5.1 Introduction

BOTH in Europe and the US, the concept of Trajectory-based Operations (TBO) is currently being investigated as a means to cope with the increasing demand placed upon the Air Traffic Control (ATC) system [1], [2]. A key aspect of TBO is to increase the airspace efficiency and capacity by stepping away from the current fixed airspace route-structure, allowing flights to fly user-preferred (i.e., direct) routes between their departure and arrival airport. Such trajectories are optimized, de-conflicted, and agreed upon beforehand, and fully define the gate-to-gate route in 4 dimensions (i.e., ground track, altitude and *time*). Although this concept reduces the need for tactical intervention by human Air Traffic Controllers (ATCos) by design, the myriad of unforeseen disturbances and the complexity of the new task require higher levels of automation and novel ATCo decision-support tools.

In previous work, the so-called *Travel Space Representation* (TSR [3]) was developed following the Ecological Interface Design framework (EID [4]–[6]). This framework is based upon identifying the functions, constraints and interdependencies within the work domain, and directly visualizing them to the human controller in the form of meaningful representations. The TSR represents the *causal* constraints for manipulating an aircraft trajectory by visualizing the physical boundaries of safe and unsafe operations as a set of *go* and *no-go* areas. Here, the *go* areas are defined as the *areas of safe travel* through which an aircraft can be rerouted safely whilst adhering to its internal (i.e., aircraft performance envelope) and external (i.e., other traffic, restricted areas) constraints. The *no-go* areas represent the space in which an aircraft rerouting will either cause a time-constraint violation or a separation violation with other traffic.

A human-in-the-loop experiment with participant groups of increasing levels of ATC experience and expertise showed that the TSR allowed all participants to manage the traffic safely [7]. However, post-hoc analyses of the robustness of individual control actions also revealed that controllers with a higher level of expertise often performed more *robust* control strategies compared to novice controllers. Such strategies can be attributed to, amongst others, preferring more conservative over tight control actions (i.e., not re-routing trajectories close to the border between the safe and unsafe areas), choosing inherently better conflict resolution geometries through knowledge-based analysis and planning, and more pro-actively controlling the traffic by resolving conflicts before they had been detected by the system.

In this study, the factors that contribute to robust control strategies are investigated by quantifying the robustness of control within the safe fields of travel in the TSR. A batch analysis is performed to analyze how varying conflict resolution geometries affect the resulting system robustness. Here, the robustness of various passing angles, from head-on to in-trail, are compared. A second analysis is made of the effect of varying the minimum separation distance and required speed increase on the available safe area of travel in the TSR. The ultimate goal of this research has been to design a set of constraints underlying robust control strategies that can be superimposed on the safe area in the travel space, in order to promote more robust control by novice controllers.

To test the effectiveness of the addition of the constraints in the TSR, the human-in-the loop experiment of the previous study has been repeated. For that purpose, the robustness of control of the novice (no prior ATC experience) and skilled (no operational ATC experience) participant groups with the inclusion of the constraints is compared to that of the expert (operational ATC experience) participant group without the additional visualization.

This chapter is structured as follows. First, the investigation and quantification of the factors that contribute to the robustness of control is discussed. This is followed by an analysis of how those factors propagate in the safe field of travel in the TSR, and how they can be transformed into a set of simplified intentional constraints that can be added to the TSR. Next, the experimental design is presented, followed by the results, discussion and conclusions.

5.2 Baseline Travel Space Representation

The current practice of tactical air traffic control is based on controllers providing discrete heading, altitude, and speed instructions to the flight deck by means of voice communication. However, in a fully transitioned trajectory-based ATC environment this form of *tactical vectoring* will no longer be possible. In order to—as much as possible—support current ATC practices, the baseline Travel Space Representation has been designed as a decision-support tool to enable *strategic trajectory-based vectoring*. With this tool, rather than providing discrete instructions, the controller can insert an intermediate waypoint into the trajectory of an aircraft resulting in a set of time-based strategic instructions.

To illustrate how the TSR can be used to manipulate aircraft trajectories, the hypothetical baseline scenario shown in Fig. 5.1 is discussed. Fig. 5.1(a) shows a selected (observed) aircraft in the controlled sector that is navigating to the sector exit waypoint (FIX) and is planned to leave the sector at an agreed time. The travel space of A_{obs} is shaded in green; its boundaries delimit the zone where the introduction of an intermediate waypoint into the flight segment will result in a new trajectory that is feasible with respect to the aircraft performance limits (i.e., turn radius and maximum speed). Placing a waypoint outside of the travel space will result in a trajectory that violates the fixed time constraint at the exit waypoint due to the excess added track-miles (i.e., violates performance constraints).

A conflicting aircraft A_{int} is introduced in Fig. 5.1(b). As a result, part of the travel space of the selected aircraft becomes restricted. This restricted field of travel, shown in red, indicates the area in which the placement of an intermediate waypoint will result in a loss of separation with the other traffic at a given point along the trajectory, rendering this control action as unsafe (i.e., violates separation constraints).

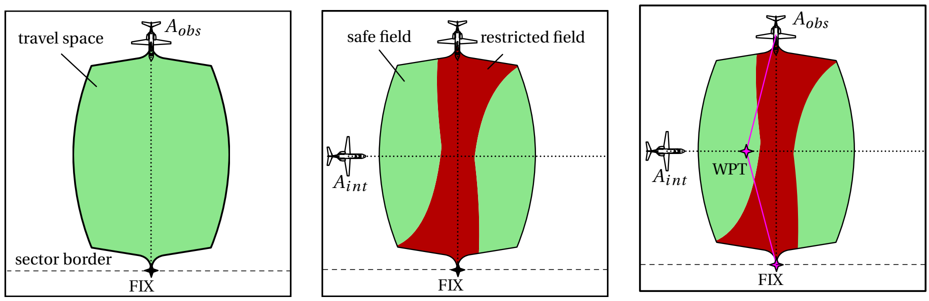

(a) The travel space of an observed aircraft

(b) Overlay of a restricted field of travel due to traffic

(c) Possible conflict resolution strategy

Figure 5.1: Hypothetical baseline scenario with and without other traffic.

Fig. 5.1(c) shows how the controller can identify and select any location in the safe field of travel to place an intermediate waypoint. The illustrated control action results in the original flight segment to be split in two new flight segments, with equal speed, such that the conflict is avoided and the planned time at the exit point remains unchanged. Although all control actions within the safe field of travel will result in a safe and feasible trajectory, the travel space in its baseline form does not provide information about the *quality* of the control action.

5.3 Geometric Conflict Resolution Robustness

Ideally, a human-machine system should promote 'anti-entropic' behavior [8], and thereby help to reduce complexity—or chaos—within the system as a whole. Although the TSR has shown to support safe control, the human-in-the-loop experiment performed in the previous study also showed that expert controllers used control strategies that were inherently more robust than those of the less skilled participants [7]. In that study *robustness* (RBT) was identified as a hidden motivation that experts use to perform adequate control decisions.

In ATC, common practices for separating aircraft implicitly increase the system robustness and decrease controller workload. Here, typically, a distinction can be made between *anticipated* and *assured* separation [9]–[11]. Anticipated separation are situations in which the trajectories of two or more aircraft conflict, but due to the geometrical lay-out and dynamics (e.g., speed difference, varying climb and descent rate), are anticipated to not have a loss of separation. Common practice in ATC is to—as much as possible and practical—avoid relying on anticipated separation as such situations require continuous monitoring and are more prone to lead to unsafe situations [9].

Assured separation, on the other hand, is achieved by removing—or greatly mitigating—the possibility for a loss of separation to occur in the first place. This can be achieved by setting-up the traffic patterns such that crossing points are eliminated, traffic complexity is reduced and minimal monitoring is required; for instance, one frequently used control strategy to achieve assured lateral separation for two crossing aircraft is to vector one aircraft directly towards the current position of the other aircraft [10]. This assures that both aircraft will have passed each other at the intersection point.

Although fully assured separation is always preferable in terms of safety, in practice, and especially in more complex traffic scenarios, this is not always possible and is most often not very efficient. For example, when observing the traffic situation shown in Fig. 5.1(c), placing the intermediate waypoint on the far-left boundary of the travel space will result in the most safe traffic situation, but is the least efficient one in terms of added track miles and required speed change. Vice versa, placing the intermediate waypoint on the boundary between the safe and restricted field of travel will result in the most efficient trajectory, but leans heavily on anticipated separation.

In order to investigate and *quantify* how this trade-off between safety and efficiency in conflict resolution strategies affects the robustness of the system to cope with disturbances, a numerical batch-analysis has been performed for a *two-aircraft crossing scenario*. Here, the effects of conflict geometry, through varying both the *trajectory crossing angle* and *Closest Point of Approach* (CPA) between the aircraft on the resulting robustness are analyzed. For this purpose, a representative en-route setting has been simulated in which both aircraft execute their trajectory at FL300, and at a constant speed of 250kts IAS ($\approx$ M.67, or 400kts ground speed assuming no wind). Similar to current ATC operations, the minimum allowable lateral separation between the aircraft has been defined as 5NM.

5.3.1 Geometry and metric set-up

Fig. 5.2 shows a schematic overview of the geometrical set-up of the simulated scenario and the variables that have been used to compute robustness. The figure shows an observed (A_{obs}) and intruding aircraft (A_{int}) that cross each other's trajectory with crossing angle θ. Numerical simulations of the movements of the passing aircraft have been performed for various crossing angles, various values of the CPA, and by letting the observed aircraft both pass in front

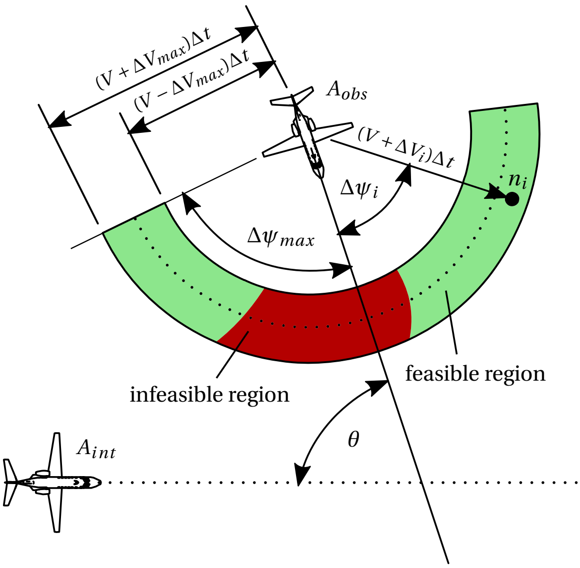

Figure 5.2: Sketch of the batch analysis geometry and the point-based robustness computation control variables.

and behind of the intruding aircraft. The crossing angle was varied between 20 and 160 degrees in steps of 10 degrees (i.e., from almost parallel to almost head-on), and the CPA was set to either 5NM (i.e., the minimum allowable separation), 6NM and 7NM. For all combinations of θ and CPA, the initial conditions were set such that the aircraft movements are simulated from exactly 10 minutes before the CPA, and continue for 10 minutes after the CPA to allow for a fair comparison of robustness between the runs.

To quantify robustness, the *point-based* and *trajectory-based* robustness metrics from the previous study have been used as described in [7]. To compute point-based robustness, a set of probabilistic disturbances are superimposed on the locomotion of the observed aircraft for a fixed duration (Δt) at a given point along its trajectory. As illustrated in Fig.5.2 the disturbances are represented by a heading offset ($-\Delta\psi_{max} \leq \Delta\psi \leq \Delta\psi_{max}$) and a speed offset ($-\Delta V_{max} \leq \Delta V \leq \Delta V_{max}$). For this analysis, the maximum heading disturbance ($\Delta\psi_{max}$) was set to 80 degrees, and the maximum speed disturbance (ΔV_{max}) to 20kts IAS ($\approx$ M.05, or 30kts ground speed). Propagating all combinations of these disturbances for the fixed duration Δt leads to the set of all possible future positions visualized by the disk-shaped form as shown in Fig.5.2.

For the computation of point-based robustness, the heading range has been discretized in steps of 5 degrees, and the speed range in steps of 2.5kts IAS, resulting in a total of 561 probe states (n_i) per point. Subsequently, each probe state can be evaluated whether it will result in a conflict with the other aircraft at time $t + \Delta t$. The point-based robustness is then given by the fraction of probe segments that will not lead to a conflict (i.e., feasible region) compared to the total number of probabilistic disturbances (i.e., the feasible *and* infeasible region) as shown in the figure.

To compute trajectory-based robustness, the point-based robustness was sampled at every second along the trajectory of the observed aircraft. From these samples, the average robustness (RBT_{avg}) and minimum robustness (RBT_{min}) were computed for each combination of θ and CPA. Finally, to investigate the effect of the duration of a probabilistic disturbance (Δt) on robustness, two sets of runs have been performed, one set with Δt set to 30 seconds (a typical TCAS Resolution Advisory time window), and one with Δt set to 120 seconds (a typical ATC Short Term Conflict Alert time window).

5.3.2 Batch-analysis results

Fig. 5.3 shows the six resulting plots of the trajectory-based robustness. Here, Fig. 5.3(a) and Fig. 5.3(b) show the minimum and average robustness of the observed aircraft passing behind the intruding aircraft. Similarly, the minimum and average robustness of the observed aircraft passing in front are shown in Fig. 5.3(c) and Fig. 5.3(d). Fig. 5.3(e) and Fig. 5.3(f) show the summed average of RBT_{min} and RBT_{avg} for passing in front and passing behind combined. In each graph, the crossing angle θ is indicated on the horizontal axis, and robustness on the vertical axis. Separate lines are plotted for the two durations of the model disturbance (Δt), and for the cases with a CPA of 5, 6, and 7NM. In the legend, the first number relates to Δt, and the second to the CPA (i.e., 30-5 relates to a Δt of 30 seconds and a CPA of 5NM).

Minimum robustness

The minimum robustness is a reflection of the probability of a loss of separation due to a disturbance at the most critical point along the paths of the crossing aircraft. Figures 5.3(a), 5.3(c), and Fig. 5.3(e) show that the CPA has the most profound effect on the resulting robustness for both the passing-in-front and passing-behind scenarios. This emphasizes that controlling close to the border of the restricted field of travel in the TSR will result in situations that are the least robust to absorb disturbances. That is, if one of the aircraft would deviate from its trajectory, there is a high probability that a subsequent control action is necessary to de-conflict the aircraft again, and thus requires continuous monitoring. Further, tight conflict resolution geometries will decrease the available control space of the involved aircraft over time, resulting in less options in case they need a further rerouting.

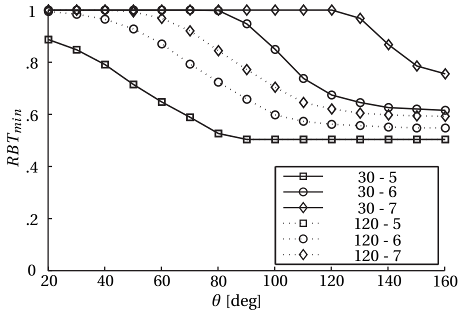

(a) RBT_{min} passing behind

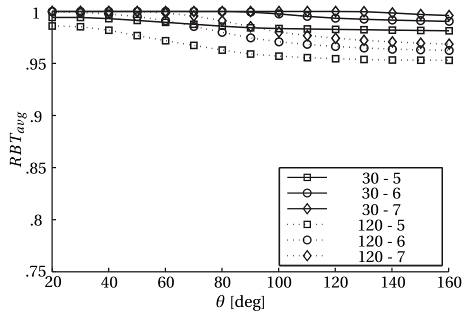

(b) RBT_{avg} passing behind

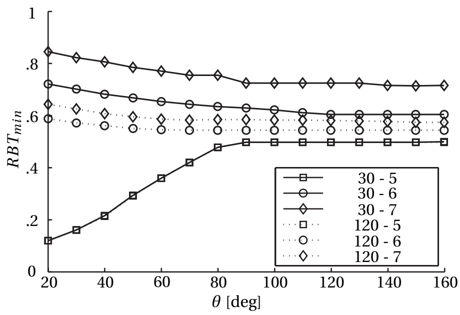

(c) RBT_{min} passing in front

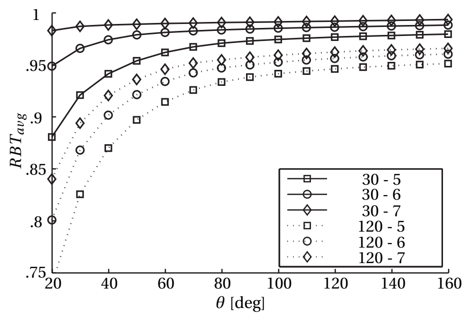

(d) RBT_{avg} passing in front

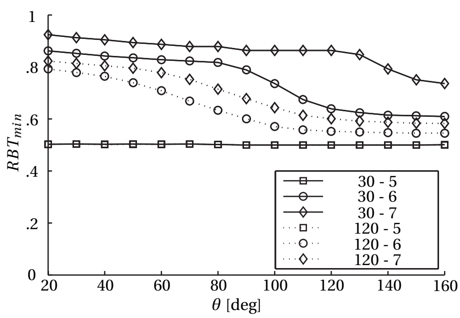

(e) RBT_{min} average

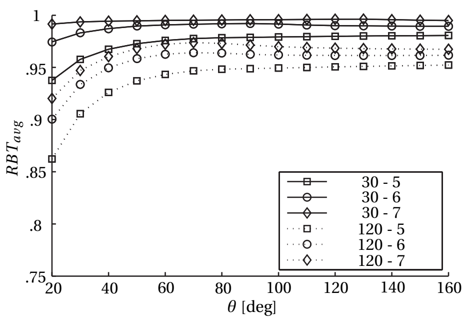

(f) RBT_{avg} average

Figure 5.3: Robustness Batch Analysis Results.

The figures show that the minimum robustness is higher when passing behind (Fig. 5.3(a)) compared to passing in front (Fig. 5.3(c)) for all angles and conditions; this indicates that a heading or speed disturbance of an aircraft passing in front of other traffic is more likely to result in a loss of separation than if it would pass behind.

In general, shallower (more in-trail) crossing angles result in a higher (better) minimum robustness. This can be attributed to the lower relative speed between the aircraft that reduces the impact of disturbances. For instance, when flying in-trail with 1NM additional separation, a heading or speed change is unlikely to cause a loss of separation for the relatively low look-forward times of 30 and 120 seconds. However, Fig. 5.3(c) shows one exception when passing in front with the minimum allowable separation distance (i.e., the 30-5 and 120-5 cases). In these cases, when at the exact point of minimum separation, any instantaneous deviation of aircraft flying in front is likely to—momentarily and marginally—violate the separation minimum resulting in a low minimum robustness.

Average robustness

The average robustness represents the time-averaged, point-based robustness along the trajectory of the simulated aircraft. Similar to the minimum robustness, the largest effector on average robustness is the CPA. Figs. 5.3(b) and 5.3(d) show opposite trends compared to average robustness with respect to the crossing angle. Here, the average robustness of the passing behind aircraft increases when passing at more shallow angles, but significantly decreases when passing in front. This effect is caused by the fact that when the aircraft are flying near parallel, for a relatively long period of time, disturbances to the flightpath of the leading aircraft could result in a loss of separation. It should, however, be taken into account that at small angle crossings the closure rate of both aircraft is much smaller than for near head-on crossings; when flying in-trail, if a loss of separation would threaten to occur, the controller would have ample time to intervene.

5.4 Robustness and Control Space

The batch analysis presented in the previous section shows that increasing the CPA between aircraft has a large positive effect on the resulting robustness. However, increasing the CPA also typically requires larger deviations from the original aircraft trajectories, thereby reducing the efficiency of the individual trajectories and overall airspace use. Such control actions are also less preferable in terms of resulting control space. That is, when the velocity of an aircraft is increased, the remaining control space for a subsequent control action decreases and its trajectory becomes less *adaptable*.

In this section the effect of increasing the minimum CPA to both the available control space in the TSR and the performance (required speed increase) of the controlled aircraft is investigated. For this purpose, a hypothetical *two aircraft perpendicular crossing scenario* has been simulated, with a geometry that is representative for the scale and size of the sector used in the human-in-the-loop experiment that will follow. Both aircraft execute their trajectory at FL300, and at a constant speed of 250kts IAS (≈ M.67, or 400kts ground speed assuming no wind). The maximum aircraft speed has been set to 300kts IAS, and the crossing point has been set to 50NM ahead the aircraft (≈ 7.5 minutes).

5.4.1 Separation

Fig. 5.4(a) shows the travel space of the observed aircraft and includes contour lines for restricted field of travel computed to achieve a minimum separation of 5NM through 10NM in steps of 1NM. As expected, the available safe field of travel reduces with increased separation and requires larger deviations from the initial direct trajectory. Further, in this symmetric traffic scenario, the restricted field of travel is also diagonally symmetrical; the shape and size of the available control space for passing in front or behind the intruding aircraft are equal.

5.4.2 Performance

Fig. 5.4(b) shows the travel space of the observed aircraft and includes contour lines for the required speed increase to achieve the fixed sector exit time from 250kts IAS through 300kts IAS in steps of 10kts IAS. 10kts IAS steps are chosen because this is the typical resolution in which ATC provides speed commands. Note that the travel space also contains the restricted field of travel caused by the 5NM minimum separation as this is seen as a ‘hard’ constraint. The figure shows that larger deviations from the current trajectory require a larger speed increase, and that the rate of this increase is higher when steering closer to the boundaries of the available control space.

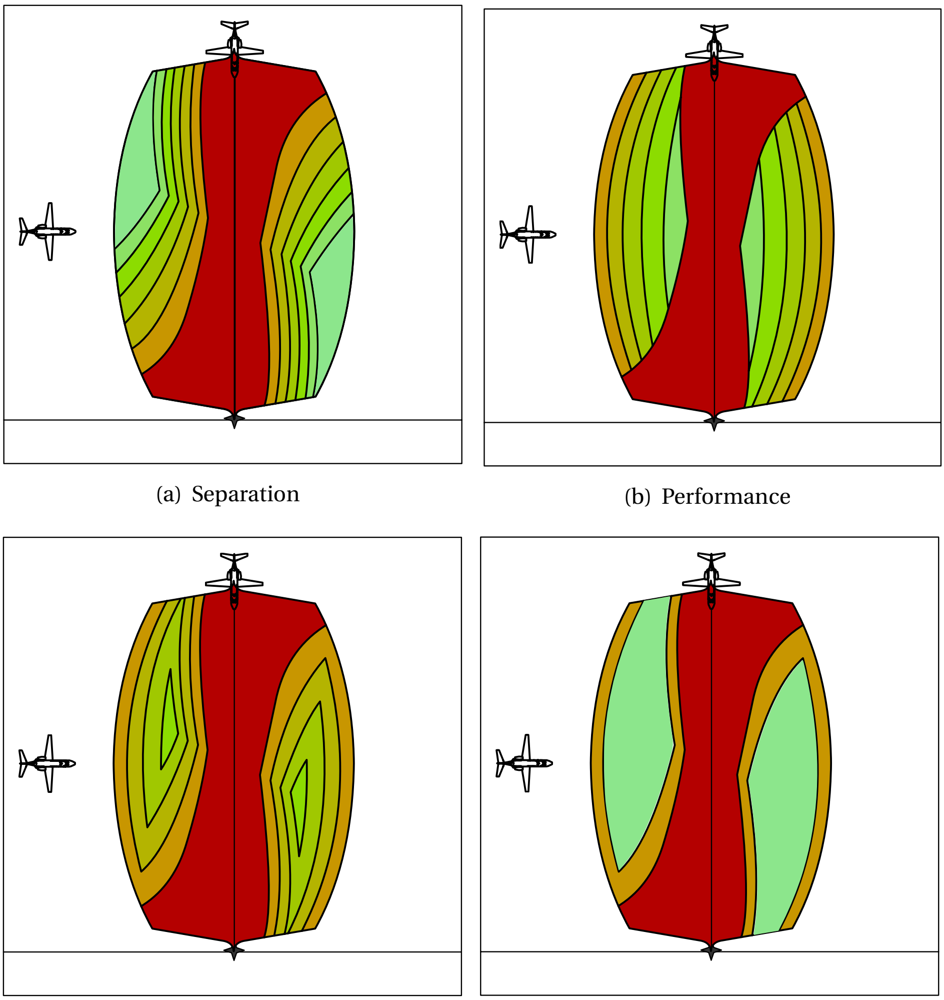

(a) Separation (b) Performance (c) Combined (d) Intentional

Figure 5.4: Performance and separation constraints overlaid on the TSR.

5.4.3 Practical use: Intentional Constraints

Fig. 5.4(c) shows the travel space with the separation *and* performance contour lines combined. The darker shaded areas thus indicate the regions in which the resulting trajectory will have less additional separation and/or will require a large speed increase. In practice, however, a quantitative trade-off between added separation and speed increase is not possible. Depending on the dynamic air traffic environment, additional separation could be preferred over a lower speed increase or vice versa. Therefore, presenting the information to the controller in the form of Fig. 5.4(c) could give misleading cues.

Perhaps a more meaningful representation is given in Fig. 5.4(d). Here, only the least robust and least adaptable control areas are highlighted; respectively, less than 1NM additional separation and less than 10kts IAS from the maximum aircraft speed. In this form, these zones act as cautionary zones, or *intentional constraints* [12], that indicate the areas in which control actions are less preferable. Controlling within the intentional constraints is possible and will not violate the separation and/or timing constraints (i..e, the total available control space is not reduced), but is likely to either compromise safety or limit the remaining control space.

To illustrate how the TSR including intentional constraints can promote more robust and preferred control actions, consider the asymmetric conflict scenario shown in Fig. 5.5. In this scenario, two aircraft flying at different speeds are in conflict and will have a loss of separation halfway along their trajectory. Fig. 5.5(a) shows the travel space representation when selecting the slower aircraft, and includes the intentional constraints. When observing the available control space, the safe area of travel is larger for passing-behind maneuvers compared to passing-in-front, and the minimum required deviation for passing-behind is smaller than when passing-in-front. Conversely, Fig. 5.5(b) shows the travel space representation when selecting the *faster* aircraft.

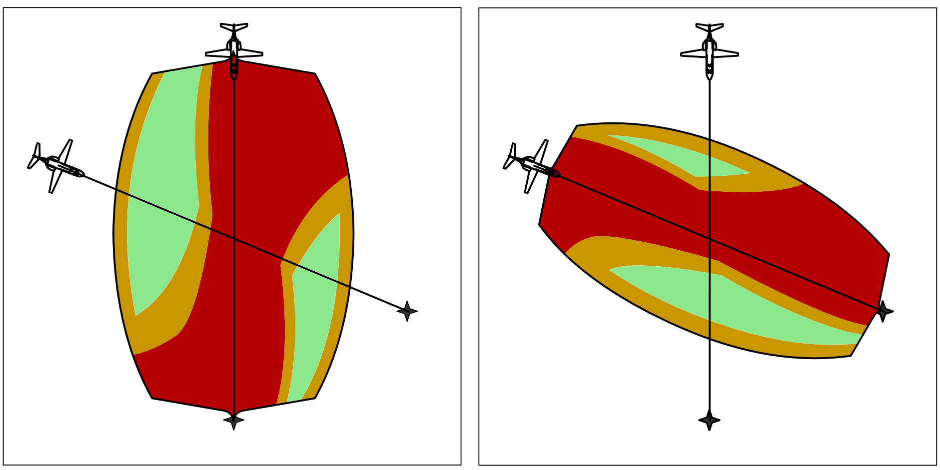

(a) TSR of the slower aircraft

(b) TSR of the faster aircraft

Figure 5.5: The TSR of a slower and a faster aircraft in an asymmetric conflict.

For that aircraft the available safe area of travel is larger for passing-in-front maneuvers as compared to passing-behind, and the minimum required deviation for passing in front is smaller than when passing behind. Based upon this information, and adhering to the intentional constraints, the preferred control strategy would lean towards the established ATC practice of steering the slower aircraft behind the faster aircraft with a safety margin for minimum separation.

5.5 Human-in-the-loop Experiment

To evaluate the effects of the addition of the intentional constraints in the TSR on control robustness, the human-in-the-loop experiment of the previous study has been repeated. In that experiment, three participant groups with an increasing level of expertise (novice, skilled and expert) were asked to manage various scenarios of trajectory-based air traffic with the TSR interface. A novice and a skilled participant group will re-run the exact same scenarios, but with the inclusion of intentional constraints in the TSR. The goal of the experiment has been to investigate whether including intentional constraints in the interface improves the robustness of control of the novice and skilled groups up to the level of the expert participant group. For that purpose, the results of both experiments have been combined into a new data-set.

5.5.1 Participants

Including the participants of the previous research, this study was performed with a total number of 20 participants, divided into three groups of operational experience. The novice group consisted of eight aerospace engineering students with either a BSc. or MSc. degree who perform flight-deck and/or ATM-automation related research (7 male, 1 female, average age of 28 years). None of the novice participants received any prior training in operational air traffic control. The skilled group consisted of eight domain experts who are currently working as professionals in ATM research and development (8 male, average age of 48). Finally, the expert group consisted of four operational Area Control Center (ACC) air traffic controllers, two of which were fully certified, and two who were involved in on-the-job training (2 male, 2 female, average age of 27).

5.5.2 Procedure

First, the participants were asked to complete a 20 minute interactive training session with the interface in which the participants followed an interactive script to become familiarized with the TSR and its functionality. The training ended when the script was completed, and when the participants indicated that they had a good understanding of how to use the TSR to manage the traffic.

In the main experiment, the participants were asked to manage traffic within a fictional two-dimensional sector under various initial conditions. During each run the overall goal was to plan and guide the traffic through the controlled sector safely (i.e., without losses of separation or restricted area intrusions) and efficiently (i.e., adhere to timing constraints at the sector exit points). After the initialization of an experiment scenario, the participants controlled the traffic by issuing changes to the trajectories of each individual aircraft by manipulating waypoints using the TSR including the intentional constraints. The resulting trajectories were automatically computed, sent to, and executed by, the aircraft on acceptance by the participant (via simulated data-link, no voice communication required). No additional input was required by the experimenter during each run.

5.5.3 Apparatus

The evaluation was performed on a dedicated software-based ATM platform, running on a single computer. The TSR was integrated in a traditional plan view display (PVD), providing a top-down view of airspace and air traffic. The TSR was presented on a 30 inch screen (60 Hz LED, 2560 x 1600 pixels) placed in front of the participant. Input was given by a standard mouse input device, and control options could be selected by on-screen drop-down menus.

5.5.4 Independent variables

The results of both the previous experiment and this experiment have been combined, resulting in a set of five participant groups given in Table 5.1. The study followed a mixed design with the three levels of expertise (novice, skilled and expert) and the two levels of the TSR lay-out (including and not including the intentional constraint visualization) as between-group variables. Similar to the previous experiment, the two within-participant independent variables were traffic orderliness (structured and unstructured traffic) and perturbation level (small, medium and large number of perturbations).

The control variables (i.e., the variables that remained fixed) in the experiment were: the sector area, size and shape, the availability of the Travel Space decision-support tool, the size and shape of the restricted area, and the initial traffic sample in structured and unstructured conditions.

Table 5.1: Definition of the five participant groups.

group	skill level	intentional constraints
$N1$	novice	no
$S1$	skilled	no
$E1$	expert	no
$N2$	novice	yes
$S2$	skilled	yes

5.5.5 Dependent measures

The following dependent measures were used to investigate the effects of adding the intentional constraints, the level of expertise, the traffic orderliness and the perturbation scale on the effectiveness of the TSR:

- *Safety:* The number of losses of separation and restricted area intrusions per condition,
- *Sector-based robustness:* The overall minimum and average robustness of the aircraft trajectories that is an indication of the expertise level of control strategies, and
- *Performance:* defined by the additional track length flown by the rerouted aircraft

5.5.6 Scenarios

The participants were asked to manage traffic in a hypothetical en-route sector ($\approx$ 40.000 KM^2) under the six different control conditions (Table 5.2). The rotation of the sector varied uniquely between scenarios consisting of the same (baseline) traffic structure to avoid a control bias due to scenario recognition. The names of the waypoints and aircraft were also varied in each run to prevent the aforementioned bias.

Each scenario presented approximately 15 aircraft and eight sector entry/exit points and lasted 24 minutes in scenario-time. The simulation ran at four times the normal speed, such that each scenario lasted six minutes in real-time. The average traffic density was set to approximately 8 controlled aircraft at any given point, with the exemption of the first and last minute of the scenarios (real-time), in which the traffic either built up or reduced to compensate

Table 5.2: Definition of the six experiment conditions.

condition	**orderliness**	**perturbation**
T_S-P_S	structured	small
T_S-P_M	structured	medium
T_S-P_L	structured	large
T_U-P_S	unstructured	small
T_U-P_M	unstructured	medium
T_U-P_L	unstructured	large

for the absence of hand-overs in between sectors, and the lack of verbal communication. All aircraft entered the controlled sector at FL300 through one of eight fixed waypoints on the sector border and were given an initial (straight) 4D trajectory leading towards one of the other waypoints.

Aircraft could only be controlled laterally (i.e., vertical manipulation of the trajectories was not possible), and only if they were physically inside the sector. Nevertheless, aircraft inbound to the sector were shown in gray when approaching the sector, such that the participants had ten minutes (scenario-time) to prepare for future traffic situations. Further, the performance characteristics of all aircraft have been simulated by a single generic commercial aircraft type. The initial conditions of each scenario were set such that the controller had to resolve a fixed number of perturbations (conflicting pairs of aircraft and avoiding restricted areas) by manipulating the trajectories of individual aircraft. However, the control actions themselves could create new conflicts and restricted area intrusions further ahead in time.

5.5.7 Hypotheses

It was hypothesized that the addition of the intentional constraints would improve the control robustness for the novice and skilled participant groups, and bring it closer to the level of expert control without intentional constraints. It was also hypothesized that this effect would be most significant in the scenarios with unstructured traffic and larger levels of perturbations as here the participants were expected to rely more the information provided by the TSR in order to make rapid rule-based shortcuts.

5.5.8 Data analysis

To test for between-group and within-group effects, Kruskal-Wallis and Friedman tests have been performed, respectively. Here, the significance level (α) has been set to 0.05. Post-hoc, pair-wise comparison tests between experimental conditions featured either three or four planned pair-wise comparisons to investigate the effect of the traffic orderliness and the scale of the perturbation respectively.

For the post-hoc calculation of sector-based and control-based robustness, the maximum heading disturbance ($\Delta\psi_{max}$) has been set to 80 degrees, and the maximum speed disturbance (ΔV_{max}) to 20 kts IAS ($\approx$ M.05, or 30 kts ground speed). For the calculation of point-based robustness, the heading range was discretized in steps of 5 degrees, and the speed range in steps of 2.5 kts IAS, resulting in a total of 561 probe states per point. Point-based robustness was sampled at every second along the trajectory. The duration of a probabilistic disturbance (Δt) was set to 120 seconds (a typical ATC Short Term Conflict Alert time window). Although the magnitude of the disturbances might seem high as compared to what could reasonably be expected in TBO, such 'extreme' values are expected to magnify the between-group variation of robustness (i.e., only relative differences are compared).

5.6 Results

5.6.1 Safety

In the baseline experiment without the visualization of the intentional constraints, out of the 1,116 controlled flights, three safety-critical events occurred. These events were two losses of separation (one novice in the T_U-P_L condition and one skilled participant in the T_U-P_L condition) and one restricted area intrusion (one skilled participant in the T_U-P_M). No losses of separation or restricted area intrusions were recorded for the expert participants.

In the experiment with intentional constraints, no losses of separation or restricted area intrusions were recorded for all 744 controlled flights. This is an indication that the addition of the intentional constraints to the TSR indeed promotes *safer* control strategies for less skilled participants. However, the loss of separation count in itself is a binary measure (i.e., loss of separation or no loss of separation), and therefore does not directly reflect the robustness of the overall air traffic. For that purpose, the minimum and average sector-based robustness are analyzed and discussed in the next paragraphs.

5.6.2 Sector-based robustness

Fig. 5.6 shows box plots of the sector-based robustness for the experiment with and without intentional constraints. Note that in the plots for the experiment with intentional constraints, the data of the expert participant group *without* intentional constraints have been added to allow for a direct comparison.

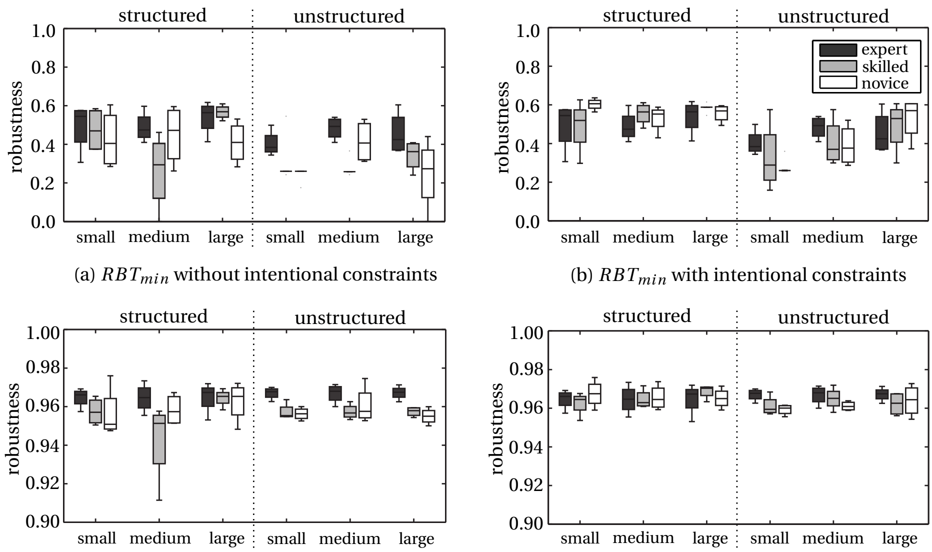

(a) RBT_{min} without intentional constraints

(b) RBT_{min} with intentional constraints

(c) RBT_{avg} without intentional constraints

(d) RBT_{avg} with intentional constraints

Figure 5.6: Box plots of the minimum and average sector-based robustness with and without intentional constraints.

Minimum robustness

Figs. 5.6(a) and 5.6(b) show box plots of the minimum robustness for each experiment condition without and with intentional constraints respectively. The minimum robustness measure reflects the least robust point-in-time for any aircraft during a run (i.e., the 'tightest' conflict resolution maneuver). In both figures the results of the novice and skilled participant groups are plotted compared to the expert group without intentional constraints.

The minimum robustness was found to be significantly higher with the intentional constraints visualized compared to without for both the novice (Mann-Whitney U-test: $U = 155$, $z = -2.75$, $p < 0.01$) and skilled ($U = 176.5$, $z = -2.30$, $p < 0.05$) participant groups. Further, in the experiment without intentional constraints, the minimum robustness was found to be significantly higher for experts compared to that of the novice and skilled participants. With the intentional constraints included, however, this significant difference was no longer measured; novices with intentional constraints compared to experts without: $U = 258$, $z = -0.608$, $p > 0.05$, and skilled participants with intentional constraints compared to experts without: $U = 264$, $z = -0.485$, $p > 0.05$.

Average robustness

Figs. 5.6(c) and 5.6(d) show box plots of the average robustness for each experiment condition without and with intentional constraints. This measure reflects the average robustness of all aircraft combined for the entirety of each experiment run. Again, in both figures the results of the novice and skilled participant groups are plotted compared to the expert group without intentional constraints.

Similar to the minimum robustness measure, the average robustness was found to be significantly higher with the intentional constraints visualized compared to without for both the novice ($U = 160$, $z = -2.64$, $p < 0.01$) and skilled ($U = 118$, $z = -3.505$, $p < 0.01$) participant groups. Again, the significantly higher average robustness for experts compared to that of the novice and skilled participants in the experiment without intentional constraints was no longer measured in the experiment with intentional constraints; novices with intentional constraints compared to experts without: $U = 208$, $z = -1.650$, $p > 0.05$, and skilled participants with intentional constraints compared to experts without: $U = 232$, $z = -1.155$, $p > 0.05$.

The figures also show that the spread of the average robustness with intentional constraints is smaller than without. This indicates that the performance and quality of control of the novice and skilled participants is more consistent when the intentional constraints are present.

5.6.3 Performance

Fig. 5.7 shows box plots for the added track miles due to control actions for the experiment without and with intentional constraints. As expected, the more complex experiment conditions require more and larger deviations from the original aircraft trajectories. However, no significant difference was found in added track miles with the visualization of intentional constraints compared to without for both the novice ($U = 251$, $z = -0.763$, $p > 0.05$) and skilled ($U = 214$, $z = -1.526$, $p > 0.05$) participant groups. These results indicate that using the additional separation buffer as visualized by the intentional constraints does–overall–not lead to less efficient aircraft trajectories.

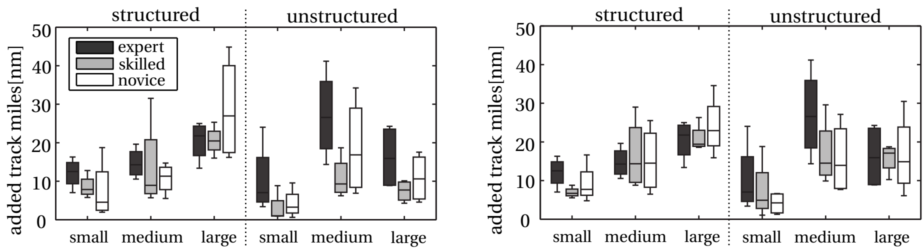

(a) added track miles without intentional constraints

(b) added track miles with intentional constraints

Figure 5.7: Box plots of the added track miles with and without intentional constraints.

5.7 Discussion

In the previous human-in-the-loop experiment with the TSR, a significant difference was observed in control performance between participant groups with varying levels of expertise. Expert controllers were found to frequently follow control strategies that were inherently more robust than those of less skilled participants. This gave a rise to concerns that in its original form, the TSR could actually work against maintaining robust traffic patterns and could reduce the overall flexibility of the airspace.

Although the baseline TSR visualizes the safe and unsafe control areas for re-routing aircraft, no additional information is provided on the quality of the resulting trajectories in terms of safety and efficiency. That is, the TSR only visualizes areas in which re-routings will or will not lead to a loss of separation, but it does not show the remaining separation buffers or efficiency resulting from a control action. Based on these observations, and to improve the performance of all operators, it was recommended to make the *soft constraints* related to robustness and efficiency more explicit in the representation.

The results from the batch analysis in this study show that robustness is mainly dependent on the size of the minimum separation buffers between aircraft. However, increasing separation is usually less preferable in terms of efficiency. In order to investigate how robustness and efficiency propagate in the TSR, contour lines for increments of respectively 1NM additional separation, and speed increments of 10kts, were added to the safe field of travel.

The multiple contour-line representation in itself was not found to be very meaningful or usable by the controller; on the one hand an objective trade-off between additional separation distance and speed increase cannot be made, and on the other hand, especially in more complex traffic patterns with multiple aircraft, the visualization becomes too cluttered for practical use. Therefore, the choice was made to only highlight the least robust and least efficient areas in the safe field of travel in the form of intentional constraints. These cautionary zones then aid the controller to avoid the least preferable control areas that will lead to 'tight' or inefficient solutions.

Further, the batch data showed that the crossing geometry of a pair of aircraft, especially at smaller separation distances, has various effects on the resulting robustness. Shallow, more in-trail crossing angles result in a higher instantaneous robustness (i.e., safer), but in a lower average robustness (i.e., requires more monitoring) compared to more head-on angles, and vice versa. In practice, however, conflict resolution geometry is seldom as symmetrical as the conditions investigated in the batch analysis. The speed increase that accompanies a certain control action adds a bias to the benefits of any specific crossing geometry. Typically, and as illustrated in the asymmetrical conflict example, the preferred conflict resolution geometry is evident from the shape, size and location of the available control space in the TSR.

The results from the second human-in-the loop experiment show that the robustness—and thereby the quality—of control of the less skilled and less experienced participant groups significantly improved when the intentional constraints were visualized; both the minimum and average sector-based robustness were significantly higher than in the previous experiment, and were much closer to that of the expert participants without intentional constraints. By comparing the robustness boxplots, the spread in minimum and average robustness also decreased when the intentional constraints were added, indicating more consistent control behavior. In addition, no losses of separation were recorded, indicating that tight solutions were most likely avoided.

Surprisingly, there is no significant difference, nor a clear trend, in the amount of added track miles with and without the addition of intentional constraints in the TSR. This shows that although the intentional constraints allow the less skilled operators to perform more robust control, the efficiency of their control strategy was not negatively impacted.

5.8 Conclusion

The goal of this study has been to enhance the ecological Travel Space Representation decision-support tool for 4D trajectory management by explicitly visualizing areas in which control actions will lead to inefficient solutions, or solutions with a low robustness. By investigating how various conflict resolution geometries affect robustness, increasing the minimum separation buffer was found to have the most positive effect. The 'best' and most efficient de-confliction geometry was found to heavily depend on the specific traffic situation. From these findings, intentional constraints that highlight the areas of up to 1NM additional separation and areas that require the aircraft to fly close to its maximum speed were added to the TSR. Results from the human-in-the-loop experiment in this study show that the intentional constraints significantly improved the robustness of control by less skilled operators, without reducing efficiency. Their performance was close to that of expert participants in a previous experiment, who worked without the intentional constraints, but are accustomed to apply these mentally because of their experience and honed workflow.

Bibliography

[1] "SESAR CONCEPT OF OPERATIONS," EUROCONTROL, Brussels, Belgium, Report, 2007.

[2] "Concept of Operations for the Next Generation Air Transportation System, Version 3.2," Joint Planning and Development Office, Washington, DC, USA, Report, 2011.

[3] M. M. Van Paassen, C. Borst, R. E. Klomp, M. Mulder, P. van Leeuwen, and M. Mooij, "Designing for Shared Cognition in Air Traffic Management," *Journal of Aerospace Operations*, vol. 2, pp. 39–51, 2013.

[4] K. J. Vicente and J. Rasmussen, "The Ecology of Human-Machine Systems II: Mediating Direct-Perception in Complex Work Domains," *Ecological Psychology*, vol. 2, no. 3, pp. 207–249, 1990.

[5] ——, "Ecological Interface Design: Theoretical Foundations.," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 22, no. 4, pp. 589–606, 1992.

[6] J. Rasmussen, "Ecological Interface Design for Reliable Human-Machine Systems," *The International Journal of Aviation Psychology*, vol. 9, no. 3, pp. 203–223, 1999.

[7] R. E. Klomp, C. Borst, M. M. Van Paassen, and M. Mulder, "Expertise Level, Control Strategies, and Robustness in Future Air Traffic Control Decision Aiding," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 2, pp. 255–266, 2015.

[8] E. Hollnagel and D. D. Woods, *Joint Cognitive Systems*, E. Hollnagel and D. D. Woods, Eds. Taylor & Francis Group, LLC, 2005, ch. 6, pp. 113–133.

[9] J.-F. D'Arcy and P. S. Della, "Air Traffic Control Specialist Decision Making and Strategic Planning - A Field Survey," Federal Aviation Administration, Atlantic City, NJ, USA, Report, 2001.

[10] S. Fothergill and A. Neal, "Conflict-Resolution Heuristics for En Route Air Traffic Management," in *Proceedings of the Human Factors and Ergonomics Society 57th Annual Meeting*, 2013, pp. 71–75.

[11] E. M. Rantanen and A. Nunes, "Hierarchical Conflict Detection in Air Traffic Control," *The International Journal of Aviation Psychology*, vol. 15, no. 4, pp. 339–362, 2005.

[12] J. Comans, "Visualizing Rules, Regulations, and Procedures in Ecological Information Systems," Ph.D. Dissertation, Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands, 2017.

Chapter 6

Control Strategies

Strategy without tactics is the slowest route to victory. Tactics without strategy is the noise before defeat. *—Sun Tzu (Military General, Strategist, Philosopher, ca. 544BCE – 496BCE)*

This chapter presents the final human-in-the-loop experiment with the enhanced *Solution Space concept that featured both the horizontal and time-based Solution Space representations. The chapter and experiment are focused on the various control strategies that can be applied with the interface. The chapter begins with an analysis of the effect of both horizontal and time-based conflict resolution strategies on the efficiency of resulting trajectories. Next, an overview of the experiment interface and its elements is provided. The experiment set-up is described in detail, followed by a presentation of the results. The chapter concludes with a discussion of the results.*

ABSTRACT

The expected shift towards a concept of trajectory-based air traffic control using four-dimensional trajectories, stored in automated support tools, and exchanged through digital data-link, will inevitably bring a large shift in the work of the air traffic controller. The introduction of higher levels of automation and advanced human-automation decision-support tools is pivotal to the success of such operations. In previous work, the Solution Space concept was introduced as a decision-support interface for tactical air traffic control that visualizes action-relevant affordance zones for real-time trajectory management. Previous human-in-the-loop experiments with the interface showed that controllers with varying levels of expertise and experience use the interface in different ways. In those experiments, only lateral modifications to the trajectories were supported. In this chapter, the interface has been extended with support to also modify the speeds and timings along trajectories. Adding time and speed as explicit control variables increases the overall control space, and thus allows for executing new control strategies. This work presents the results of a human-in-the-loop experiment that focused on the use of the interface and applied control strategies of the participants under various airspace routing structures. Results show that all participants could manage the traffic safely, however, the use of the interface varied greatly between participants.

6.1 Introduction

THE work of the Air Traffic Controller (ATCo) is foreseen to fundamentally change with the introduction of 4D trajectory-based operations (TBO). Although the ultimate goal of the Air Traffic Management (ATM) system will remain unchanged (i.e., supporting the production of flights in a safe and economical manner), the means to control traffic, and the available *control strategies* will be different. TBO steps away from control by *discrete* commands by voice (i.e., vectoring: speed, altitude, and heading commands) towards modifying and communicating (a segment of) planned aircraft trajectories via digital data-link. Due to the added complexity and increased dimensionality of this task, controllers will no longer be able to perform their work without advanced decision-support tools and higher levels of supporting automation [1], [2].

In the operational concept of TBO, the role of the controller will lean more towards higher level airspace management and decision making, whilst routine tasks and data integration will be taken over by automation [2]. However, the means to modify trajectories, the scale and extent of automation support, and the division of roles between the human and automation are not yet fully defined. In the past, replacing highly manual tasks by more automation has shown to cause numerous problems such as skill degradation, reduced situation awareness, and over-reliance in automation [3]–[5]. This emphasizes the importance to address and overcome these problems in an early design phase.

In ongoing research, the design of a new *ecological* decision-support interface for 4D trajectory manipulation has been proposed that visualizes *Solution Spaces* on the electronic Plan View Display (PVD) [6]–[9]. Following an approach based upon the Ecological Interface Design framework (EID, [10], [11]), the interface concept aims to visualize the *action-relevant* control spaces that bound the set of feasible, safe, and *goal-oriented* control actions. That is, the display allows the controller to visually observe the available control space for re-routing aircraft safely (i.e., avoid losses of separation), whilst adhering to the planning constraints imposed on the trajectory. In essence, this allows the controller to remain fully-in-the loop (i.e., actively in control), and is foreseen to mitigate human-automation mismatches.

A previous human-in-the-loop experiment showed the viability of the Solution Space concept, but also showed that *how* controllers with various levels of expertise and training used the interface differed [9]. In that experiment, trajectories could be manipulated laterally to de-conflict and re-route aircraft (speed and altitude control was not supported). It was observed that, typically, novice participants with little to no prior ATC experience performed more *boundary-seeking control*, steering aircraft close to the border of safe and unsafe actions. Such strategies focused on minimizing path deviation by opting for tight solutions, thereby, in essence, using the visualization of the control spaces (go/no-go areas), and with that, the control space boundaries, as a basis for their strategy. In contrast, expert participants frequently structured traffic, executed *current* best practices, and maintained additional separation buffers.

Expert strategies make the work of the air traffic controller more structured and predictable, but might be less preferable and economical from a pilot and airline perspective. A justification for the use of such strategies can be that the current method of vectoring does not take uncertainty into account. That is, controllers gain situation awareness about the current state of the traffic by voice communication, observation, and by mental integration of "raw" data on the radar screen. Control strategies are then formulated by taking the myriad of *foreseen* variability and external factors such as wind, atmospheric conditions, airspace status, etc., into account to find a solution that is robust against these disturbances [12]. This can be best described as a form of *workload management*, or mitigation, where control strategies follow from training and experience, and are aimed at reducing *future* workload [12], [13].

Although the exact form and extent in which trajectory-based operations will be introduced is not yet fully specified, the underlying concept will inherently make the system more predictable by partially closing the loop in the execution phase of flight [2]. The *same* definition of the trajectory is shared by ATC and the flight-deck, and will act as an active navigation target for the aircraft. This mitigates most current open-loop uncertainties, and thereby also the need to explicitly manage workload by formulating more structured, but less optimal solutions. Perhaps then, more efficient, tighter, and less structured solutions will be preferred; that raises the question for *whom* and for *what task* the automation should ultimately provide support?

In this chapter, the interface has been extended with a time-space diagram (TSD, [14]) that enables the controller to adjust timings (and thereby speed) at waypoints along the trajectory as an additional way to de-conflict aircraft. Although the means to manipulate the speed along a trajectory increases the overall *action-relevant* control space, the lateral (i.e., horizontal) and time-based Solution Spaces are highly coupled, adding an extra degree of *complexity*. That is, manipulating the aircraft speed along the trajectory directly affects the available horizontal control space, and vice versa [15]. This raises the question of how well controllers are able to connect the information between the horizontal and time-based domains, and how well they are able to divide their attention between the two representations.

A following question is *how* controllers use the interface, and *what strategies* they could or should apply to de-conflict and re-route traffic. Rather than solely manipulating the lateral path of the aircraft, controllers can now opt for resolving conflicts by speeding up or slowing down. The effectiveness and availability of different control strategies is highly dependent on the specific task and airspace environment. In current operations, the frequency in which controllers use various separation techniques greatly differs from one type of ATC facility to another [12]. Speed control is typically only used extensively in approach control where the main task is to merge traffic and create in-trail arrival-streams, but is less desirable for departure and upper area control. In upper area control, separation is mainly performed by altitude separation, or by lateral path modifications [12].

To investigate the effect of controller expertise and the type of *airspace structure* on the control strategies applied with the interface, this work presents a human-in-the-loop experiment with the extended Solution Space interface. Six expert and six novice participants were asked to manage en-route traffic in both converging (i.e., merging), and diverging traffic structures. Participants were trained how to manipulate and de-conflict trajectories laterally and in time, but were *not* instructed on the (the benefits of) any specific strategies. After each run, participants were asked to formulate their control strategies, and to comment on their preference of time-based over lateral control in that scenario. After the experiment, participants were asked to rate and provide comments on the level of difficulty of the experiment, the ease of use of the interface, and the interpretation and understanding of the Solution Spaces. Further, an analysis of the logged data has been performed to investigate the participants' on-screen-activity, and to analyze the resulting traffic patterns.

The chapter is structured as follows. First, an analysis will be given of the efficiency of lateral and time-based control strategies in the en-route setting. This is followed by a detailed description of the interface and its various elements as used in the human-in-the-loop experiment. The next two sections describe the experiment set-up, and the results of the experiment, respectively. The chapter is concluded with a discussion on the results, recommendations for future work, and a conclusion.

6.2 Control Strategy Efficiency

The addition of the time-space diagram to the Solution Space interface enables controllers to manipulate time targets at waypoints along aircraft trajectories. In essence, the controller can thereby manipulate the aircraft speed along different segments of the trajectory. As a result, aircraft can now not only be separated laterally, but also by using speed commands (i.e., speeding up or slowing down). To investigate what strategy is more efficient for the *controlled aircraft*, the effect of both strategies on total *fuel-burn* as well as the *required speeds* has been analyzed.

The baseline scenario for the analysis is a two-aircraft crossing conflict, and is illustrated in Fig. 6.1. Here, both aircraft fly along a direct trajectory from one side of the sector to the other, crossing each other at the center. The two conflict resolution strategies for the controlled aircraft (A_{con}) that have been investigated are also illustrated in the figure; ❶, conflict resolution by flying a constant speed *dog-leg* to maneuver behind the conflicting aircraft (A_{int}), and ❷, flying at a slower speed initially to pass behind the other aircraft, and then speeding up again. In both cases, the required speeds have been computed such that the sector exit time of the controlled aircraft remain unchanged.

The aircraft used for this investigation represent a popular medium-range passenger aircraft, and were simulated using the BADA aircraft performance model [16]. Similar to the human-in-the-loop experiment in this work, the aircraft initially fly at 30,000ft (FL300) at an indicated airspeed (IAS) of 250kts ($\approx$

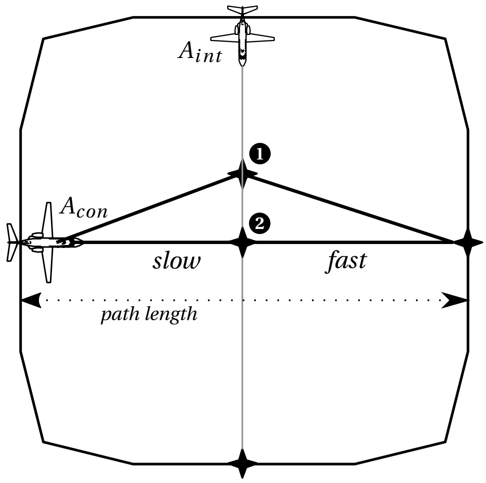

Figure 6.1: The baseline scenario used for analyzing control strategy efficiency.

395kts ground speed). That speed is close to the optimum long range cruise speed (LRC, minimum fuel-burn per NM) at that altitude. At FL300, the maximum and minimum speeds of the aircraft are 315kts IAS (≈ 490kts ground speed) and 195kts IAS (≈ 310kts ground speed), respectively.

The length of the initial trajectory was varied; the crossing point occurs at *half* that distance. With these parameters, the time to the crossing point (TTCP) is approximately 15 minutes for an initial trajectory of 200NM. Further, the minimum separation distance between the aircraft was set to 5NM, and accelerations and decelerations were not taken into account (i.e., the trajectories were represented by two constant speed segments).

Fig. 6.2 shows the effect of the two strategies on the additional fuel burn and required speeds. The figures firstly show that the scenario can be resolved by using a dog-leg for conflicts 15NM out and farther (initial length 30NM, TTCP ≈ 2:20). The speed strategy is only feasible for conflicts of 45NM out and onwards (initial length 90NM, TTCP ≈ 6:50). Using these strategies to resolve closer conflicts will result in either exceeding the speed envelope, or violating the exit time constraint.

Fig. 6.2(a) shows that a considerable speed increase is required to resolve close-by conflicts using a dog-leg, but that this rapidly becomes less for larger distances; a dog-leg flown 52kts faster is required to resolve a conflict 15NM ahead, at 50NM this is already down to 3kts, and at 100NM it is less than 1kts. This directly relates to the smaller required relative path deviation. The speed-only strategy requires relatively large speed changes between the *slow* and the

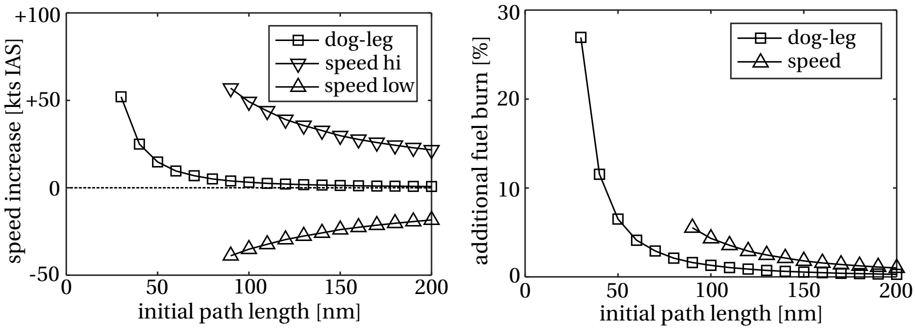

(a) required speed increase

(b) additional fuel burn

Figure 6.2: Additional fuel burn and speed increase required to resolve a conflict by using either a dog-leg, or by speed only.

fast leg; for an initial path length of 200NM, the two trajectory segments still require a speed difference of -18kts and +22kts, respectively.

Fig. 6.2(b) shows the additional fuel burn in percentage to resolve the conflict with respect to the fuel burn required to fly the original path (also computed using the BADA performance model). The figure shows that—as to be expected—resolving conflicts farther out results in a considerably smaller penalty to the fuel burn as compared to resolving conflicts closer by. Further, conflicts resolved by using a dog-leg are *always* more fuel efficient than slowing down and speeding up.

The analysis shows that for the two aircraft crossing scenario, flying a dog-leg is—in every case—more efficient in terms of fuel burn, and has less impact on the speed profile compared to a pure speed-based strategy. One could then argue that, in this case, controllers would always prefer to issue lateral commands over speed commands when reasoning about the the consequences for the affected aircraft. However, issuing lateral deviations to all aircraft in an initially structured sector (i.e., streams of traffic) will reduce the overall structure of the airspace. Less structure is, in itself, not good or bad from a system perspective, but could increase controller workload in terms of monitoring. The question then is, in what cases would controllers choose which strategy?

To investigate when and why controllers will provide lateral and/or speed commands using the extended Solution Space concept, the human-in-the-loop experiment in this chapter aimed at identifying (motivations for) different control strategies under different initial traffic structures. The next section describes the experiment interface in more detail.

6.3 The Interface

Figure 6.3 shows a sketch of the software-based implementation of the extended Solution Space concept as used in the human-in-the-loop experiment. The sector lay-out, traffic structure, and color scheme shown in the figure are representative of one of the experiment scenarios. The main parts and sub-elements in the interface are numbered, and are described below.

The interface is divided into two main parts positioned side-by-side. The electronic plan view display (PVD, ❶) is located on the left hand side, and shows a top-down representation of the controlled sector and aircraft. When no aircraft is selected, only the PVD section of the interface is displayed. When the controller selects an aircraft on the PVD, a time-based representation of its trajectory is shown in the time-space diagram (TSD, ❷) on the right hand side. Inspired by the Ecological Interface Design framework (EID [10], [11]), constraints for manipulating the selected trajectory are directly mapped on the interface in the form of *action-relevant* control spaces.

6.3.1 Plan view display

The PVD shows the border of the controlled sector, the sector entry and exit waypoints, and, in this scenario, includes a circular restricted area around which traffic has to be diverted ❸. The controlled sector size is 200x200NM, and the horizontal range of the PVD is set to 350NM. Although the size of this sector is larger than most current upper area sectors, it reflects the TBO vision of Functional Airspace Blocks (FABs) in which single larger sectors are allocated in a *flight-centered* manner, better facilitating direct-routing than smaller geographically-bounded sectors [2]. Aircraft are represented by small square symbols, supported by labels showing the aircraft callsign and current speed. The symbols are trailed by 'history dots' showing the aircraft previous position, and include a speed vector that shows the projected aircraft position in 90 seconds. The symbols are-color coded to indicate the state of the aircraft; uncontrolled ❹, controlled but not yet manipulated ❺, controlled and previously manipulated ❻, in conflict with other traffic ❼, highlighted ❽, and selected ❾. Further, the trajectories of all non-selected aircraft are indicated by thin lines.

Aircraft can be selected on the PVD by left-clicking on their symbol which prompts the visualization of its lateral and time-based Solution Space. The color coding of the Solution Spaces is as follows: safe control space ❿, intentional separation or performance constraint ⓫, and unsafe control space ⓬. Hovering over an aircraft symbol or an unsafe control space with the mouse cursor highlights the associated aircraft and all unsafe zones caused by that aircraft on both displays (8) (13). Waypoints along the selected trajectory are shown as magenta stars, and show the commanded speed of the previous segment (14).

6.3.2 Time-space diagram

The time-based representation of the trajectory of the selected aircraft is shown on the TSD (15). Here, the horizontal axis represents the along track distance to go (DTG) to the sector exit point (16). The vertical axis represents the time that the aircraft is planned to be at that DTG (17); the currently selected aircraft has approximately 165NM to go to the exit fix, and will arrive there in approximately 25 minutes. The original sector exit time is indicated by a cyan diamond along the time axis. Labels of all other controlled aircraft that are flying towards the *same* exit fix are also shown along the time axis, and are color coded by aircraft state using the same rules as on the plan view display (18). These aircraft can be selected by left-clicking the mouse on their label. Further, the TSD features a so-called 'time-slider' (19). By dragging this slider up and down along the time-axis, the projected traffic state at that time is shown on the PVD.

6.3.3 Trajectory manipulation

The interface allows for direct manipulation of the trajectories, and assumes that modified trajectories can be up-linked to aircraft using digital data-link. By holding down the 'control' key (ctrl) on the keyboard while working on the PVD, a new waypoint is attached to the mouse cursor, and can be inserted into the trajectory by left-clicking on a position inside the sector (typically inside a safe control space). Similarly, by holding down the control key while working on the TSD, a waypoint is attached to the time-space trajectory at the DTG where the mouse cursor is located, and can also be inserted by left-clicking the mouse. Waypoints can be dragged up and down within the available control space in the time-space diagram to change the planned arrival time at that point. Deletion of waypoints on either display is performed by holding down control and right-clicking on that waypoint. Concurrent modifications can be performed on both displays, and the trajectory modifications are reflected in both Solution Spaces in real-time. When the controller is satisfied with the new trajectory, one can either press the 'enter' key to up-link the updated trajectory to the aircraft, or deselect the aircraft to discard the changes.

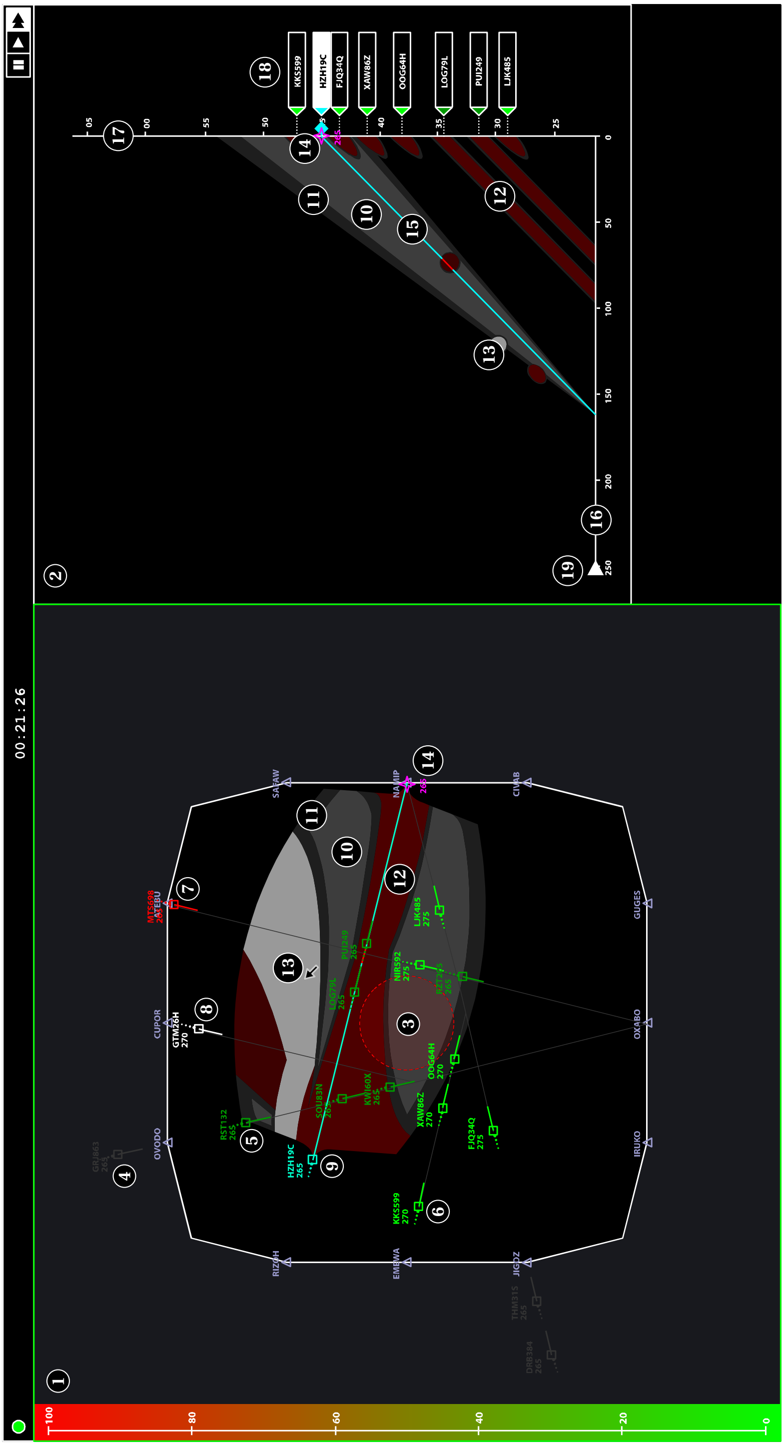

Figure 6.3: The interface and color scheme as used during the experiment.

6.4 Experiment

To evaluate how human operators use the integrated Solution Space concept to formulate and implement their control strategies, twelve participants with varying levels of expertise and experience were asked to control en-route air traffic in four hypothetical scenarios with varying initial route structures and sector complexity.

6.4.1 Apparatus and simulation platform

The software-based implementation of the Solution Space concept ran on a desktop computer with a single screen (30 inch, 60Hz LED, 2560 x 1600 pixels). The lateral and time-based Solution Space visualizations were computed and generated in real-time (60Hz) using pixel-based *shader* computations on the on-board NVIDIA GTX 970 (4GB VRAM) Graphics Processing Unit (GPU). A schematic of the experiment set-up is shown in Fig. 6.4. All input such as aircraft selection, trajectory manipulation and up-link was performed by using standard mouse and keyboard devices.

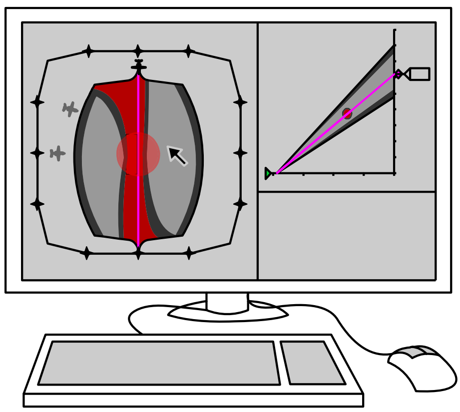

Figure 6.4: Schematic overview of the experiment set-up.

6.4.2 Participants

A total of twelve participants, divided into two groups, participated in the experiment. The first *novice* group consisted of six M.Sc. students who perform research related to ATM and flight-deck automation, but did not have prior formal training in air traffic control (5 males, 1 female, average age μ=27 years). One student, however, did have prior experience controlling traffic with the lateral and time-based Solution Spaces. The second *expert* group consisted of six academic researchers that had followed a formal air traffic control training course at the Netherlands Aerospace Laboratory (6 males, average age μ=46 years). Of these experts, four had participated in previous experiments with the lateral Solution Space interface, and two hold an active private pilot license.

6.4.3 Independent variables

The experiment followed a mixed design with both between-group and within-participant variables. The between-group variable was the *expertise* of the participants. The within-participant independent variables were the *structure* of the traffic and the presence (or not) of a *restricted area* (RA) in the sector.

Traffic structure. The two traffic structures featured in the experiment were *converging* and *diverging* patterns as illustrated in Fig. 6.5. Fig. 6.5(a) shows the converging structure in which crossing aircraft enter the sector from one of six waypoints, and exit through a single waypoint on the opposite side. This structure requires the aircraft to be merged at the sector exit point. Fig. 6.5(b) shows the *diverging* traffic structure in which crossing aircraft enter the sector from one of two waypoints, and exit through one of three waypoints on the opposite side.

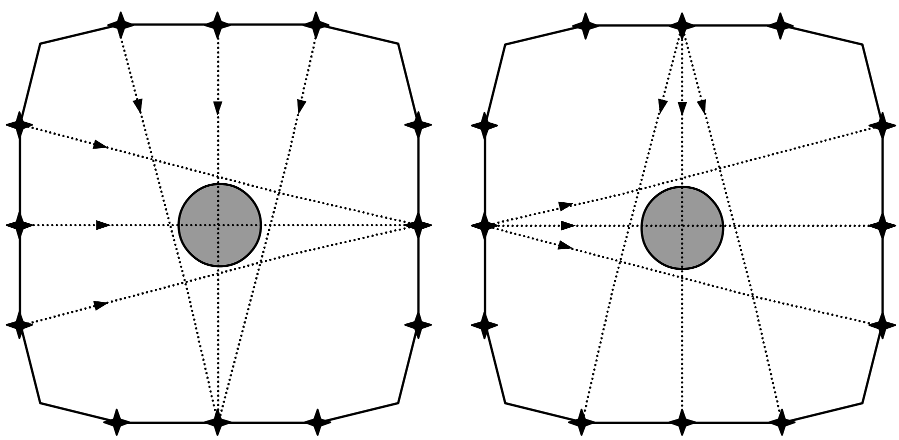

(a) Converging traffic condition with a restricted area

(b) Diverging traffic condition with a restricted area

Figure 6.5: The controlled sector lay-out showing the two different traffic structures and location and relative size of the restricted area.

Restricted Area. The second variable was the presence of a generic *restricted area* (RA) inside the sector. In practice, such avoidance zones can result from numerous unforeseen events, such as hazardous weather or temporary closed airspace. A circular zone in the middle of the sector represented the restricted area, affecting two of the six traffic streams in both traffic structures. In scenarios without a RA, participants could opt to solve all conflicts that occurred purely by manipulating the aircraft speed or purely by lateral control. In scenarios with a RA, controllers were forced to laterally re-route aircraft coming from two streams around the zone, but could still further de-conflict aircraft by using speed.

The combinations of the two traffic structure and restricted area conditions formed the four experiment conditions as given in Table 6.1.

Table 6.1: Definition of the four experiment conditions.

id	condition	traffic	restricted area
1	T_C	converging	none
2	T_C-RA	converging	present
3	T_D	diverging	none
4	T_D-RA	diverging	present

6.4.4 Procedure

The experiment consisted of two parts. First, a training session of approximately one hour in which the participants were made familiar with the interface. After that, the main experiment was conducted, and lasted approximately one and a half hours (including breaks). No preparation was required beforehand by the participants.

Training. Training was performed by completing an interactive *paper-based script* that led the participants through a total of ten practice scenarios using the interface. During the training session, participants were encouraged to talk out loud and ask questions when something was unclear. Participants received no instructions on specific control strategies or efficiency. In the first scenario, participants were made familiar with the basic visual representations of the interface, and the methods for manipulating a single aircraft trajectory. The second scenario focused on conflict detection, Solution Space interpretation, and conflict resolution in a two-aircraft crossing scenario. In the third scenario, the interpretation of the Solution Spaces for more complex, multi-aircraft situations was practiced. This scenario was followed by a paper-based quiz in which participants had to link aircraft on a sketch of the plan view display to their respective time-space representation in the TSD. Completing this test ensured that the participants had a good understanding of how the Solution Spaces on the time-space diagrams relate to the position and orientation of aircraft on the plan view display. In scenario four, a restricted area was added in the sector, and focus lay on how to manage the overall traffic flow. The last six scenarios were *open* training runs in which participants could practice applying various control strategies at their own discretion.

Main Experiment. In the main experiment, the participants were asked to manage the traffic in the four experiment scenarios using the tools and skills that they had learned during the training. The overall objective was to resolve conflicts between aircraft, re-route aircraft around a restricted area if present, and to adhere as close as possible to the original planned sector exit times. After initialization, the participants were free to implement their own control strategy and to use both lateral (PVD) and/or time-based (TSD) trajectory manipulations to complete the scenarios. During the main experiment, no additional input was required from the experimenter.

6.4.5 Dependent measures

The dependent measures consisted of both objective and subjective measurements. The primary goal of *safety* was measured in terms of the number of losses of separation and restricted area intrusions. To investigate *how* participants used the interface and what *control strategies* they applied, all lateral- and time-based trajectory modifications were logged, together with the resulting traffic movements. This allowed for a post-hoc replay of the scenarios to observe the emerging traffic patterns and strategies. *Efficiency* was measured by evaluating the additional fuel burn and flown track miles due to control actions.

After each scenario, participants were asked to formulate their control strategy and their preference for lateral over time-based control in an intermediate questionnaire. Together with the logged data, this provided an insight into *what tasks* the participants prioritized; did they retain or introduce structure into the traffic to reduce their own workload, or did they prioritize minimizing deviations for individual aircraft. After the experiment, the participants were asked to fill in a questionnaire with both multiple-choice and open questions regarding various aspects of the interface. These questions focused on preference for lateral or time-based control, interpretation of the Solution Spaces, and the overall control strategies that they applied.

6.4.6 Scenarios

Participants were asked to control traffic in a symmetrical 200NM by 200NM semi-square en-route sector as shown in Fig. 6.5. The circular restricted area in the sector center had a radius of 20NM. The *same* sector was used during both the training phase and main experiment so that the participants were familiar with the scale and dynamics of the traffic. The simulation ran at four times speed; in a set of pre-experiment validation runs, it was found that participants had ample time to formulate and implement their control strategies with this setting. Each run lasted exactly 15 minutes in *real-time*, representing a traffic sample of one hour *scenario-time*.

The scenarios were automatically generated using an iterative computer-based algorithm such that; (1) an aircraft entered the sector every 90 seconds scenario-time from 0 to 55 minutes (37 aircraft total); (2) 10 losses-of separation would occur without control; (3) in scenarios including a restricted area, a total of 15 aircraft needed to be laterally re-routed, and (4) no merging conflicts existed at the sector exit points. This resulted in two unique traffic samples, one converging and one diverging, in which the traffic was more or less equally divided over the six different streams. For the converging and diverging cases, the *same* traffic sample was used in the scenarios with and without a restricted area. The average time for an aircraft to cross the sector was approximately 30 minutes scenario-time. At the busiest point in a scenario, participants would have control over approximately 20 aircraft simultaneously.

Aircraft trajectories could be modified laterally on the PVD (i.e., re-routing), and time-based control could be performed on the TSD (i.e., modifying waypoint timings or speed). Aircraft could only be controlled when they were physically inside the sector, but inbound aircraft were visible outside the sector 10 minutes (scenario-time) before entering. Finally, the scenarios were rotated and mirrored between experiment runs, and all callsigns and waypoint names were randomly generated to mitigate the chance of traffic sample recognition.

6.4.7 Hypotheses

The following hypotheses are formulated regarding the strategies and use of the interface under the four experiment conditions:

H.1. More experienced participants are hypothesized to prefer lateral, or a combination of lateral and speed control over pure speed control to minimize the penalty in efficiency for the aircraft. It is also expected that they will formulate the tradeoffs in their applied strategies more explicitly in the subjective questionnaires.

H.2. The structure of the airspace is hypothesized to have an effect on the *type* of control used to manage the traffic. It is expected that participants will prefer lateral control in converging traffic, and time-based control in diverging traffic in order to retain the structure of the traffic flows.

H.3. The presence of a restricted area is hypothesized to further shift the overall control preference towards lateral control, as compared to without the restricted area.

H.4. Control strategies are hypothesized to vary greatly *between* individual participants, but it is expected that individual participants will rely on the *same* set of motivations to solve all scenarios.

6.5 Results

6.5.1 Safety

Out of the 1,776 controlled flights in the experiment, no losses of separation occurred. One restricted area intrusion was recorded for a novice participant in the diverging traffic condition. Replay of the logged data showed that the participant had selected and observed the aircraft, but did not re-route it. As a result, the aircraft flew through the center of the restricted area.

6.5.2 Overall interface use

Fig. 6.6 shows two box plots representing *how* the participants used the interface during the experiment runs. The results are presented irrespective of the strategies applied by the individual participants.

Fig. 6.6(a) shows a box plot of the total number of trajectory modifications per scenario. Fig. 6.6(b) shows a box plot of the percentage of lateral commands compared to all given commands. Here, 100% indicates that only lateral modifications were performed, and 0% indicates that only speed modifications were performed. These modifications were either adding waypoints on the plan view display, or adding waypoints and modifying timings on the time-space diagram. The counts and percentages shown in both plots are also irrespective of whether the modifications were actually sent to the aircraft, and thus also account for *solution probing* by the participants.

There was no significant difference in the use of lateral and speed commands between participants in the novice and expert groups. However, the figures show a large variation in how *individual* participants used the interface. A clear example of the individual variance can be seen in Fig. 6.6(b) in

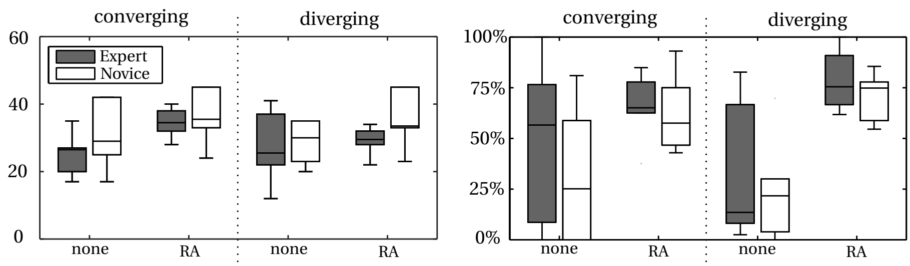

(a) Number of trajectory modifications (b) Percentage of lateral modifications

Figure 6.6: Box plots related to the use of the interface.

the converging condition with no restricted area; one expert *exclusively* used time-based manipulations stating his strategy as: *"I kept the initial route structure the same"*, and one other expert exclusively used lateral control stating: *"I merged streams of traffic to simplify the scenario"*.

The participant expertise did not have a significant effect on the *number* of performed trajectory modifications (Fig. 6.6(a)), however, the spread was larger within the novice group ($\mu = 44.9$, $\sigma = 35.3$), as compared to the expert group ($\mu = 29.0$, $\sigma = 7.2$). This was mainly the result of an outlier in the novice group who performed more than 60 trajectory modifications in *all* scenarios (outside the range of Fig. 6.6(a)); all other participants performed less than 50 modifications in any scenario.

The scenario condition significantly affected the number of modifications (Friedman: $\chi^2(3) = 15.15, p < 0.01$) and the type of modifications (Friedman: $\chi^2(3) = 20.33, p < 0.01$). Post-hoc Wilcoxon tests showed that participants performed significantly more trajectory modifications when a restricted area was present. Also, significantly more lateral commands were given compared to the scenarios without a restricted area. This was expected because 15 additional aircraft needed to be laterally re-routed to avoid the restricted area. The traffic structure did significantly affect the number and type of modifications.

6.5.3 Individual interface use

Fig. 6.7 provides a more detailed overview of the interface use per participant to investigate the individual control strategies. The height of the bars indicates the total number of trajectory *updates* performed by participants in each scenario. Each trajectory update (up-linked to the aircraft by pressing the 'enter' key) consisted of one or more concurrent trajectory modifications.

The total number of trajectory *modifications* in each scenario is shown beneath the bars, and is equal to or higher than the number of updates by definition. The count is again irrespective of whether the modifications were actually sent to the aircraft, and directly relate to the counts in Fig. 6.6(a). Further, the colors inside the bars indicate the fraction of lateral and speed modifications with respect to all modifications.

The figure emphasizes the difference in control strategies between the individual participants. Where some participants solved all scenarios by mainly using lateral control (e1, e2, n3), others chose more exclusively for speed-based solutions (e4, e6, n1, n4, n5, n6). Also the number of trajectory manipulations varied greatly between participants. Some solved scenarios with a small num-

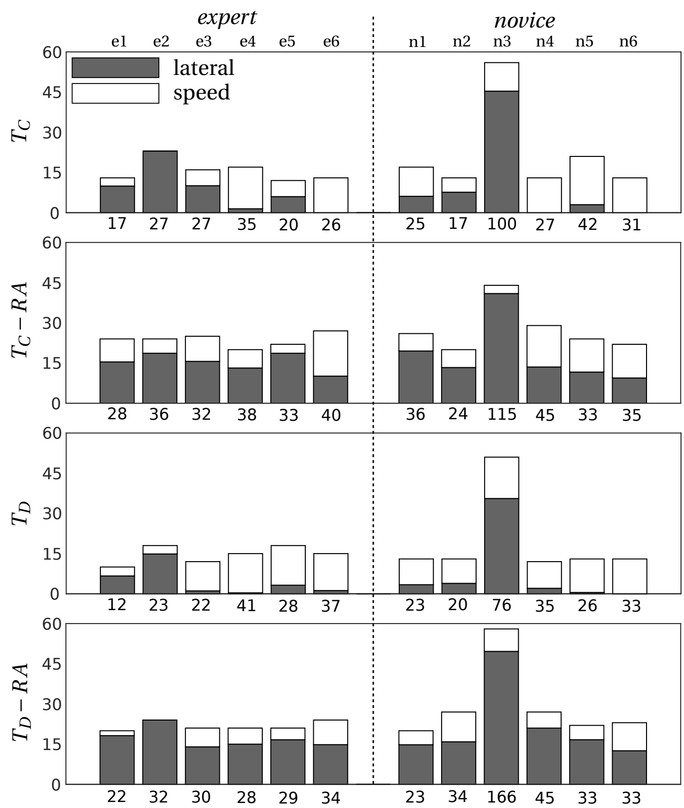

Figure 6.7: Number of trajectory updates per participant, sorted by group, and given for each scenario. Each update consisted of one or more modifications. The values beneath the bars indicate the total number of modifications. The fraction of lateral and speed modifications is indicated by shaded areas inside the bars.

ber of modifications (e1, maximum of 28 modifications in a scenario), and others frequently chose for more complex solutions, issuing multiple lateral and speed changes to aircraft (n3, maximum of 166 modifications in a scenario).

The variance in preference for speed and/or lateral control was also reflected in the post experiment questionnaire, where four participants indicated they preferred speed control over lateral control (e4, e6, n4, n6), four participants indicated they preferred lateral control over speed control (e1, e2, e3, n3), and four participants indicated that they did not have a clear preference (e5, n1, n2, n5).

6.5.4 Control strategies

Although there was a large variation in the use of the interface between participants, from observations during the experiment, feedback on the applied control strategies in the questionnaires, and by viewing the scenario replays, three distinctive strategies could be identified:

Preserve route structure. (e4, e6, n1, n4, n5, n6) Especially in scenarios without a restricted area, some participants opted to resolve conflicts primarily by using *speed commands* to preserve the original route structure and to minimize lateral deviations. Relatively large speed changes were required to de-conflict all traffic for this strategy to succeed. One participant (n6) described this strategy as: *“use lateral commands to steer around restricted areas, otherwise, use speed to de-conflict aircraft”*.

Merge traffic streams. (e1, e2, n3) Other participants mainly used *lateral commands* to simplify the streams of traffic and to minimize the number of crossing points between routes. Frequently, these participants would also modify the trajectory of aircraft that were not in need of re-routing to adhere to the new traffic pattern. Two examples of this strategy are illustrated in Fig. 6.8. Fig.6.8(a) shows how adjacent traffic streams can be merged to avoid the restricted area. Fig.6.8(b) shows a route-merging strategy for converging traffic without a restricted area. This strategy was described by a participant (e2) as *“use lateral control where possible, only change speed where necessary”*.

Minimize control impact. (e3, e5, n2) Some participants stated that their strategy focused on minimizing path deviations and speed changes for all aircraft. They frequently chose *co-operative* resolutions over larger single aircraft modifications. Other participants indicated that they respected a maximum speed change when controlling the speed of aircraft. When it was not possible

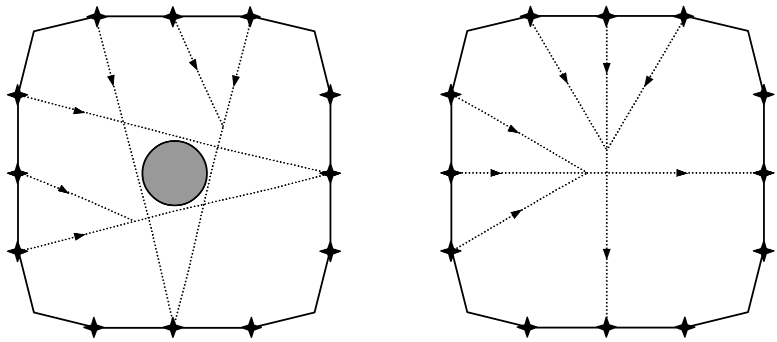

(a) Merging routes to create a predictable traffic pattern

(b) Merging routes to minimize crossing points

Figure 6.8: Examples of two observed strategies for creating structure in the original traffic patterns.

to resolve a conflict within this limit, they would revert to lateral control. Minimizing performance penalties and/or performing co-operative resolutions are not explicitly supported nor promoted by the interface. Participants that performed this type of control typically had a real-world operational background in aviation. One participant (e5) described his strategy as: *"I tried to build trains of aircraft. Also I tried to minimize any change in time and speed. This resulted in mostly lateral solutions first, and small changes to speed second"*.

6.5.5 Efficiency

No significant effect on added track miles or additional fuel burn was found between the expert and novice group, however, the variation between the individual participants and strategies was again reflected. Fig. 6.9 shows the added track length and additional fuel burn resulting from the participants' trajectory manipulations, summed over all runs, and grouped by control strategy; preserve route structure (*preserve*, speed control), merge traffic streams (*merge*, lateral control), and minimize control impact (*minimize*, minimal control).

In the control efficiency study in this chapter, lateral control was found to always be more efficient in terms of fuel-burn than a speed-based solution for a crossing conflict scenario. The figure, however, does not show a clear and definitive correlation between the three strategies and the additional fuel burn and added track length. By reviewing the scenario replays, the strategies themselves were not found to be clear-cut in all cases, and were frequently combined, or implemented in different ways.

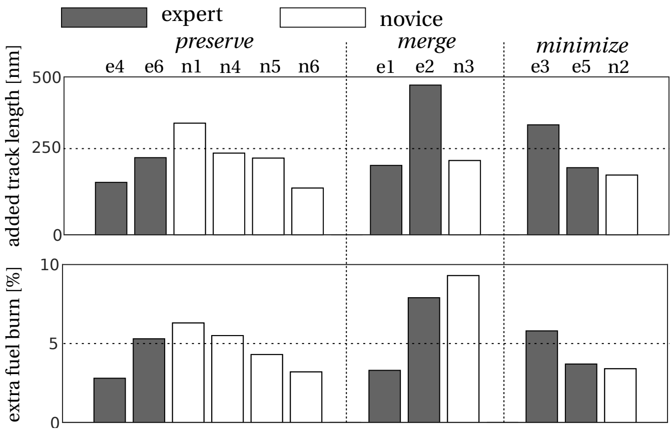

Figure 6.9: Added track length and extra fuel burn per identified control strategy.

For instance, a participant who used the *preserve* strategy (e4), in which speed changes are preferred over lateral control, also frequently applied the *minimize* strategy, making smaller speed changes to multiple aircraft instead of one large speed change. As a result, the additional fuel burn was relatively low. Another participant who applied the *merge* strategy (e2), which favors lateral changes over speed control, often re-routed aircraft that were not in need of control to better structure the sector, and resulted in a relatively high additional fuel burn. The different strategies can, therefore, not be conclusively ranked in terms of efficiency due to the high variation in their individual implementation.

6.6 Discussion

Results show that the extended Solution Space interface enabled both novice and expert participants to control the traffic without any loss of separation. However, as hypothesized, the applied control strategies were found to vary greatly between individual participants (**H.4**). On the one hand this indicates that the interface is *individual sensitive*, and allows enough freedom for controllers to formulate and implement their own unique strategies (i.e., non-normative). It highlights a trait which is frequently attributed to ecological interfaces; that they step away from optimality in favor of robustness, and allow the operator to shift strategies when necessary [17].

On the other hand, because ecological interfaces are in favor of showing the full control space, and thereby not occluding less optimal solutions, it also emphasizes that extensive *training* and *experience* will still remain *es-*

sential in order to create *expert* controllers. Although the Solution Space interface shows the physical and intentional constraints that bound the feasible *spatially-oriented* control space for aircraft (i.e., constraint-based), the current implementation, however, does not make salient the implications of control actions on resulting efficiency. In today's operations, with little integrated information readily available, controllers deal with this inherent complexity through training, best practices, and extensive experience. In essence—and as shown in the experiment—although ecological interfaces can enable novice participants to perform *adequate* control, they are ultimately intended for *expert* operators.

In a previous experiment with the lateral Solution Space, novice controllers were found to more frequently steer to the border between safe and unsafe actions (boundary-seeking control), where experts typically incorporated a larger safety margin (more separation) [9]. In that experiment, the lateral Solution Space only indicated safe and unsafe zones. In a follow-up experiment, the interface was modified to include *intentional* constraints, visualizing an additional separation and aircraft performance buffer. Results of that experiment showed that with the intentional constraints, the control strategies of novice participants better matched those of the experts (Chapter 5).

One solution could then be to promote efficient strategies by explicitly visualizing efficiency inside the Solution Spaces as intentional constraints. However, this could make the interface overly complex and cluttered. Ultimately, the sweet spot of what information can and should be presented needs to be investigated further, and will be a trade-off between interface complexity and usability, whilst not trivializing the underlying work domain complexity.

6.6.1 Controller expertise

It was hypothesized that the expert participants would prefer lateral, or a combination of lateral and speed control over pure speed commands to minimize the penalty in efficiency for aircraft (**H.1**), however, this was not reflected in the experiment results. On the one hand, an explanation could be that the participants in the expert group were not active air traffic controllers, and did not have any on-the-job experience in operational air traffic control. On the other hand, it could also indicate that the new form of trajectory-based air traffic control will require new control strategies, different training, and perhaps even new or other controller *competences*.

Another reason for the variance in strategies between participants could be that the experiment featured a simplified representation of its real-world counterpart. External factors that normally influence decision making in air traffic control such as prediction (in)accuracy, environmental factors, communication delays, etc., were not present. That, combined with the abstract nature of the Solution Space representations, likely led to a certain extent of *gamification*; participants sometimes viewed and solved scenarios as abstract puzzles rather than real-world operations.

Although some participants provided a motivation for their strategies based upon consequences for the affected aircraft, others would often greatly be steered by the information provided by the interface, and would not avoid choosing less efficient solutions. Also, because of the absence of any explicit briefing, training, and indication of the effect on efficiency of different control strategies, the optimization thereof typically did not play a factor in the strategies and types of control that participants applied.

6.6.2 Individual strategies

It was hypothesized that the airspace structure (i.e., converging and diverging traffic) would have an effect on the *type* of control used to manage the traffic (**H.2**). Lateral control was expected to be used more often in converging traffic, and time-based control in diverging traffic to better retain the structure of the traffic flows. Results of the experiment showed that this was not the case, and that lateral control was only used significantly more with the presence of a restricted area, what was to be expected because 15 aircraft had to be re-routed laterally around that zone (**H.3**).

Individual controller preference was found to be the most important factor in the choice for mainly lateral or speed-based strategies. The results also show that the individual strategies did not vary much between the different scenarios. That is, controllers that indicated that they preferred lateral control over speed control showed the same bias in all scenarios. Vice versa, controllers that preferred speed control over lateral control frequently only used lateral control to steer aircraft around restricted areas. Individual participants typically did not apply a great variety of strategies, nor did they frequently switch strategies under different airspace settings.

Further, the question was raised for whom and for what task the interface should ultimately provide support. Three overarching control strategies were identified in the human-in-the-loop experiment; (1), preserving the route structure, and (2), merging traffic streams, and (3) minimizing control impact. The first two strategies are focused on alleviating the work of the *air traffic controller* (i.e., centralized reasoning). Preserving the route structure and merging traffic streams create more predictable flows of traffic and subsequently require less resources to monitor (non-)conformance. The third strategy is more focused on minimizing the impact of control for the individual aircraft (i.e., decentralized reasoning). This indicates that for whom the interface provides support is directly coupled to for what task the user intends to use the interface. The Solution Space interface itself thus does not impose either centralized or decentralized control.

6.6.3 Limitations and recommendations for future work

One limitation of this work was that no active air traffic controllers with real-world operational experience participated in the experiment. In a previous study on en-route air traffic control strategies by Redding *et al.* in the 1990's, it was found that expert controllers tended to use a greater variety of strategies, and that their strategy usage varied with airspace context [18]. Those observations were not reflected in the results of the human-in-the-loop experiment. It is expected that a more profound difference in control strategies would be observed between novices and active air traffic controllers.

The limited duration and scope of the instructions and training was likely also insufficient for—especially novice—participants to familiarize with both the full complexity of the interface and the underlying work domain. The training itself mainly focused on how participants could interact with the interface to fulfill the task of de-conflicting and re-routing aircraft, but did not elaborate on the many secondary tasks and considerations that air traffic controllers take into account and actively balance in their work.

As a result, only a small number of participants indicated that trade-offs on efficiency played any role on the selection and execution of their control strategies. To promote more efficient strategies for all controllers, efficiency could be explicitly visualized inside the Solution Spaces as intentional constraints (perhaps even only as a training aid). Adding intentional constraints has shown to increase awareness and improve the control behavior of, typically, more novice participants (Chapter 5, [19]).

Finally, the Solution Space concept currently allows the controller to perform lateral and time-based modifications to aircraft trajectories, however, modification of the vertical flight path are not yet supported. This, in essence, reduces the available control space of 4D trajectory management to 3 dimensions; lateral (2D) and time (1D). To gain more traction within the operational community, support for vertical modifications must be included. Especially in en-route sectors, crossing conflicts are typically be resolved by altitude separation in current ATC operations.

6.7 Conclusions

A human-in-the-loop experiment was performed with the integrated Solution Space concept in which twelve participants, grouped by expertise (novice and expert), were asked to control en-route air traffic under different airspace conditions. Results show that all participants could safely control the traffic; no losses of separation were recorded. No significant differences in control strategies or preference for lateral or speed control were found between the expert and novice participant groups. However, the applied control strategies varied greatly between *individual* participants.

Although strategies varied between participants, individuals rarely changed their *own* strategy under different airspace conditions. Three overarching control strategies have been identified; (1), preserving the route structure, and (2), merging traffic streams, and (3) minimizing control impact. The first two strategies are focused on creating structure, and thereby reducing controller workload, the latter strategy focuses on minimizing the impact on efficiency for individual aircraft. This indicates that the Solution Space is individual-sensitive and allows the controller to manage traffic in different ways.

Bibliography

[1] "Concept of Operations for the Next Generation Air Transportation System, Version 3.2," Joint Planning and Development Office, Washington, DC, USA, Report, 2011.

[2] "SESAR 2020 Concept of Operations Edition 2017," EUROCONTROL, Brussels, Belgium, Report, 2017.

[3] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

[4] M. R. Endsley and E. O. Kiris, "The Out-of-the-loop Performance Problem and Level of Control in Automation," *Human Factors*, vol. 37, no. 2, pp. 381–394, 1995.

[5] C. E. Billings, "Human-Centered Aircraft Automation: A Concept and Guidelines," NASA Ames research Center, Moffet Field, CA, USA, Report, 1991.

[6] M. M. Van Paassen, C. Borst, R. E. Klomp, M. Mulder, P. van Leeuwen, and M. Mooij, "Designing for Shared Cognition in Air Traffic Management," *Journal of Aerospace Operations*, vol. 2, pp. 39–51, 2013.

[7] R. E. Klomp, M. M. Van Paassen, C. Borst, and M. Mulder, “Joint Human-Automation Cognition through a Shared Representation of 4D Trajectory Management,” in *Proceedings of the 2nd SESAR Innovation Days*, D. Schaefer, Ed., Braunschweig, Germany: EUROCONTROL, 2012, pp. 1–7.

[8] R. E. Klomp, C. Borst, M. M. Van Paassen, M. Mulder, D. Nieuwenhuisen, A. Maij, M. Mooij, and A. van Drunen, “Designing for Joint Human-Automation Cognition through a Shared Representation of 4D Trajectory Management,” in *Proceedings of 17th International Symposium of Aviation Psychology*, Dayton, OH, USA, 2013, pp. 518–523.

[9] R. E. Klomp, C. Borst, M. M. Van Paassen, and M. Mulder, “Expertise Level, Control Strategies, and Robustness in Future Air Traffic Control Decision Aiding,” *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 2, pp. 255–266, 2015.

[10] K. J. Vicente and J. Rasmussen, “Ecological Interface Design: Theoretical Foundations.,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 22, no. 4, pp. 589–606, 1992.

[11] C. M. Burns and J. Hajdukiewicz, *Ecological Interface Design.* Boca Raton, FL, USA: CRC Press, 2004.

[12] J.-F. D’Arcy and P. S. Della, “Air Traffic Control Specialist Decision Making and Strategic Planning - A Field Survey,” Federal Aviation Administration, Atlantic City, NJ, USA, Report, 2001.

[13] C. Borst, R. M. Visser, M. M. Van Paassen, and M. Mulder, “Exploring Short-Term Training Effects of Ecological Interfaces: A Case Study in Air Traffic Control,” *IEEE Transactions on Human-Machine Systems*, vol. 49, no. 6, pp. 623–632, 2019.

[14] M. Tielrooij, A. C. In ’t Veld, M. M. Van Paassen, and M. Mulder, “Development of a Time-Space Diagram to Assist Air Traffic Controllers in Monitoring Continuous Descent Approaches,” in *Air Traffic Control*, M. Mulder, Ed., Sciyo, 2010, ch. 7, pp. 135–147.

[15] R. E. Klomp, R. Riegman, C. Borst, M. Mulder, and M. M. Van Paassen, “Solution Space Concept: Human-machine Interface for 4D Trajectory Management,” in *Proceedings of the 13th USA/Europe Air Traffic Management Research and Development Seminar*, Vienna, Austria, 2019.

[16] P. Baulleret, “User Manual for the Base of Aircraft Data (BADA) - Revision 3.1,” EUROCONTROL experimental center, Bretigny-sur-Orge, France, Tech. Rep. EEC Technical/Scientific Report No. 1998-034, 1998.

[17] C. Borst, J. M. Flach, and J. Ellerbroek, “Beyond Ecological Interface Design: Lessons From Concerns and Misconceptions,” *IEEE Transactions on Human-Machine Systems*, vol. 45, no. 2, pp. 164–175, 2015.

[18] R. Redding, J. Ryder, T. Seamster, J. Purcell, and J. Cannon, “Cognitive Task Analysis of En Route Air Traffic Control: Model Extension and Validation.,” Human Technology, Inc., McLean, VA, USA, Report, 1991.

[19] J. Comans, “Visualizing Rules, Regulations, and Procedures in Ecological Information Systems,” Ph.D. Dissertation, Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands, 2017.

CHAPTER 7

DISCUSSION AND CONCLUSIONS

In this chapter, the discussions and conclusions of previous chapters are summarized and combined to provide an overarching discussion on the design and employment of ecological air traffic controller decision-support tools for 4D trajectory-based operations. Limitations of the work in this thesis, together with real-world implications and recommendations for future work are provided. The chapter concludes with a summary of the most important insights.

7.1 Retrospective

THE transition towards Trajectory-based Operations (TBO) and the introduction of higher levels of automation in Air Traffic Management (ATM) has a real potential to enhance the capacity and efficiency of current operations. In a fully mature state, traffic movements will be more predictable, airspace use will be more flexible, and trajectories will be more optimized. Although many of the organizational and technological challenges are being addressed—mainly in isolation—the precise task of the air traffic controller, and the form and extent of automation support, still remain unclear.

The work in this thesis took a Cognitive Work Analysis (CWA, [1], [2]) approach to investigate the implications on human-automation task allocation in TBO (Chapters 2, and 3). Here, the Work Domain Analysis (WDA), as a first step in CWA, focused on identifying the set of *universal goals and functions* inherent to the trajectory-based *ecology*, initially independent of any specific division of roles or automation solutions. This resulted in the conceptualization of "*Solution Spaces*" as a shared representation for 4D trajectory-based air traffic control. By using the *same* definition of the Solution Spaces as a basis for both the work of the human controller and (any level of) automation support, *mismatches in cognition* between human controllers (i.e., their mental model), and automation (i.e., decision-support tools and algorithms) are sought to be mitigated by design.

The work then progressed to taking an Ecological Interface Design (EID [3], [4]) approach to develop a decision-support interface concept for tactical, trajectory-based air traffic control (Chapter 3). By taking the WDA as a basis, elements of the ecology have been translated to *goal-relevant Solution Spaces* that directly visualize the available control space to human controllers. Three human-in-the-loop experiments with (partial) implementations of the interface were presented, both to refine the concept (Chapter 5), and to evaluate how controllers use the Solution Space representations to formulate their control behavior and strategies (Chapters 4, 5, and 6).

The aim of this final chapter is to reflect on the research questions and associated research challenges, and to provide a broader discussion in the context of the combined work rather than the individual studies. Recommendations for future work, and a discussion on the implications thereof is provided. Further, limitations following from both the scope and assumptions, and limitations of the human-in-the-loop evaluations are discussed. The chapter concludes with a summary of the most important insights.

7.1.1 Finding a Common Ground (RQ1)

The first part of this research focused on the Work Domain Analysis (WDA) for trajectory-based operations, and the implications thereof on the design of decision-support interfaces. In Chapter 2, a (partial) WDA was performed, resulting in an Abstraction Hierarchy (AH, [5]) decomposition of the work domain for the three step-wise phases of refinement of 4D trajectories as laid out in the SESAR Concept of Operations (SESAR ConOps, [6]). In Chapter 3, a more detailed AH was constructed for the tactical control phase, providing the underlying foundations for the design of the Solution Space decision-support concepts. These parts of the work aimed to answer the first research question:

> **Research question 1**
>
> What are the underlying system goals, functions, relationships, and constraints inherent to the work domain of tactical 4D trajectory-based air traffic control, and what are the implications on interface design?

Common ground and system purpose. To answer the first research question, three distinct approaches to automation can be overlaid on the Abstraction Hierarchy as illustrated in Fig. 7.1. Fig. 7.1(a) shows the AH for a system with little to no automation support. In such a system, information about the state of its lower level components is typically made salient to the operator in a single sensor/single indicator (SSSI) fashion. The operator must use her cognition to integrate this information into a comprehensive mental model of the work domain (stippled box) in order to align the individual states to the overarching system goal. Conversely, Fig. 7.1(b) shows the AH for a fully automated system. In such systems, the operator sets the desired outcome, after which

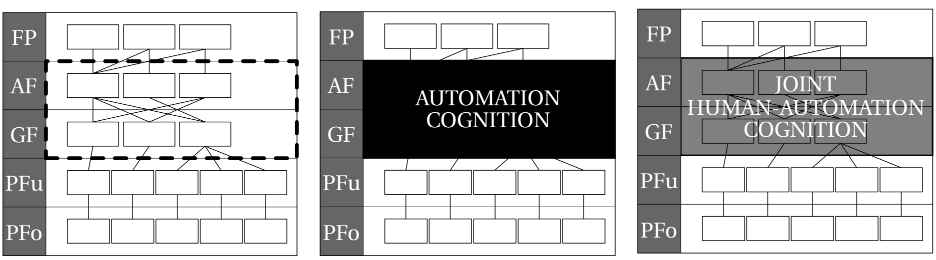

(a) No Automation (b) Full Automation (c) Joint System

Figure 7.1: Three distinct approaches to automation overlaid on the AH.

automation acts fully automatically and opaque according to an internal set of rules and strategies to achieve that state. Finally, Fig. 7.1(c) shows the AH for a joint human-machine (i.e., joint cognitive) system.

The common ground that then somehow needs to be made salient to both operators and automation hinges on the *integration of task-relevant information* on the higher levels of abstraction. A first challenge therefore was to identify the elements on the Abstract Function (AF) level that describe the basic principles and processes (i.e., mechanisms) in the work domain, and the elements on the Generalized Function (GF) level that further specify this in terms of system solutions.

The proposed Abstraction Hierarchy for tactical trajectory-based air traffic control (Chapter 3) is shown in Fig. 7.2. Using this decomposition as a basis, *Solution Spaces* were conceptualized as the common ground for operators and automation, representing the *comprehensive set of feasible control space*, constrained by both the *internal constraints* of the individual aircraft, and by *external constraints* within the air traffic control environment. Many elements, components, and functions in the AH—typically at lower levels of abstractions—are intrinsic to the (ATM) work domain itself and are often immutable. However, as Rasmussen identified early on, models at low levels of abstraction are related to a specific physical world that can serve *several purposes*, but models on the higher levels of abstraction are closely related to a *specific purpose* that can be met by several physical arrangements [5]. This implies that the purpose of the system itself affects the useful physical arrangements of the underlying func-

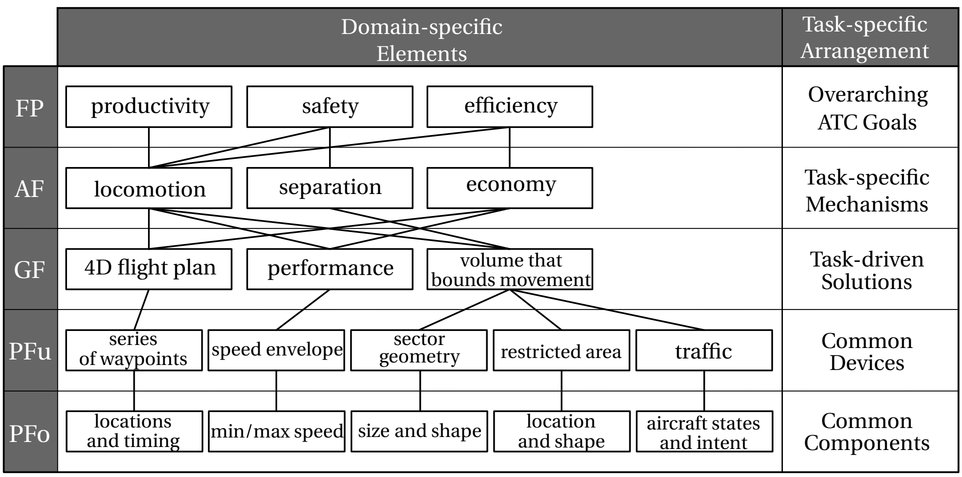

Figure 7.2: The categorized Abstraction Hierarchy for tactical 4D TBO.

tions. This will be particularly true for the different environments within the ATC domain. The work of an en-route controller is very different from that of an *area* controller or *terminal* controller.

When taking the AH in Fig. 7.2 as a basis, any given *arrangement* of elements on the different levels of abstraction for 4D trajectory-based ATC can (generally) be categorized as *common*, *task-driven*, or *task-specific*. Elements on the Physical Form (PFo) and Physical Function (PFu) levels make up the basic building blocks of trajectory-based operations, irrespective of specific tasks, priorities, or strategies; 4D flight plans are always defined by a series of 4D waypoints, aircraft should never exceed their performance envelope, sector geometry and restricted areas (e.g., terrain, hazardous weather, closed airspace) will always constrain the control space, separation violations between aircraft must be avoided. Such elements can be considered as *truly common* in all stages of TBO.

The arrangement, form, and importance of these building blocks on the task-driven *solutions* on the Generalized Function (GF) abstraction, however, are highly dependent on the specific *control task* and *time scale*. For instance, the vertical flight path and terrain avoidance are more relevant in terminal control than in en-route control, the margins for (safe) control within the performance envelope are very different at high and low altitudes, area and terminal control are all typically more dynamic and time-critical than en-route control.

Finally, the *mechanisms* to reach the overarching system goals are highly dependent on the task-specific *intentional constraints* (rules and regulations) and *strategies*. Separation rules are typically time-based in terminal control, and lateral and vertical in area and en-route control. The *economy of operations* is defined in terms of throughput in terminal and area control, where it is specified as minimizing delays and trajectory deviation in en-route control. Although the underlying conceptualization of Solution Spaces is relevant in all phases of 4D trajectory-based operations, the applicability and usefulness of any *specific implementation* is highly dependent on the *specific scope* of the work domain.

Implications for interface design. While the concept of Solution Space in itself is fully four-dimensional and represents the comprehensive set of feasible control space for trajectory-based air traffic control, it would be impractical (and likely impossible) to represent the full multidimensional set of Solution Spaces in a single interface. Any (partial) representations of Solution Spaces should match the task-specific solutions of the work domain, and any interface should be tailored to support the task-specific mechanisms. For instance, the work of the area- and terminal controller will be more focused on the time-based Solution Spaces, and the work of the en-route controller will be more focused on horizontal and vertical views of Solution Spaces.

Even within the same ATC domain, Solution Space-based interfaces will likely be different for the individual controller roles; the executive and planner controller might prefer having different 'views' of the same work domain. Although the common underlying elements will be the same for any interface, how these elements are combined and integrated will be different for each specific task. Finally, it is important that any representations of the Solution Spaces should match the controllers' mental model of work domain, and support their control strategies. Previous studies have found this to be critical in order to gain operational *acceptance* [7]–[9].

7.1.2 Visualizing the Solution Space (RQ2)

Although the WDA provided guidance in identifying what functions and constraints affect the Solution Spaces, it did not provide a ready-made recipe for the shape and form in which that information should be presented. The second part of this research therefore looked at how the Solution Space concept can be translated into functional visual representations for en-route control decision-support interfaces (Chapters 3 and 5). These parts of the work aimed to answer the second research question:

> Research question 2
>
> In what visual form can the constraints and functions that govern the tactical 4D trajectory-based control task best be mapped to a decision-support interface such that it supports the work and strategies of air traffic controllers as best as possible?

Framework, design templates, and tooling. The Solution Space visualizations presented in this work are inspired by EID, which aims to make work domain constraints, both causal and intentional, salient on an interface in such a way that operators can directly perceive and act within the action-relevant control space. The main challenge was to translate the Solution Spaces into visual forms that effectively show the feasible control space of a selected aircraft within the constraints of the overall ATC environment.

The work in this thesis took existing (ATC) interface templates as a basis, and progressed to map the 4D trajectory-based Solution Spaces within them. Although one could argue that this evolutionary approach could limit the creativity of visualizing the Solution Space, however, using existing templates provides controllers with a familiar basis, and limited the required amount of training in the human-in-the-loop experiments [10]. The horizontal Solution Space was mapped on a traditional radar Plan View Display (PVD), and the time-based Solution Space was mapped on a Time-Space Diagram (TSD). The computer-generated visualizations of the Solution Spaces in the interface concepts took advantage of relatively recent developments in Graphics Processing Units (GPU) technology that allowed for per-pixel brute force computations that could be presented and updated in real-time.

Although a Solution Space itself is four-dimensional by definition, the separate visualizations were represented on co-planar two-dimensional displays. The decoupling of the control dimension visualizations (i.e., horizontal path, speed/time, and altitude) on separate displays will require controllers to divide their attention between the displays to explore all possible control strategies. This also posed a second challenge: although presented in separate screen regions, manipulating the Solution Spaces on one plane directly affects the Solution Space on the other. To overcome this, the concept of *visual momentum* [11] was applied to the Solution Space displays, attempting to make the means-end links, and interconnectivity between the visualizations more salient. For instance, conflicting aircraft were highlighted on the displays providing a *long shot* of system conformance towards the overarching goal of safety. *Perceptual landmarks* were added to link the solution space of one aircraft to the other traffic; when hovering with the mouse over an unsafe control space within the Solution Space of one aircraft, the conflicting aircraft would be highlighted on all other displays. However, this coupling was not always found to be immediately intuitive due to the abstraction of the Solution Spaces, and required extensive training in the human-in-the-loop experiments.

Visual form and constraints. The visual form of the Solution Spaces, bound by the internal aircraft locomotion constraints (i.e., basic control space), primarily followed from the task-specific mechanics and design choices made within the overall scope assumptions. For the *horizontal* Solution Spaces, the main assumption was that the sector exit times of the aircraft are fixed. This assumption followed from the SESAR vision of trajectory-based operations in which all trajectories will be fixed—gate to gate—in time and space before execution [6], [12]. The main task for the tactical air traffic controller will then be to resolve conflicting or unfeasible trajectories, but to minimize trajectory deviations and delays to avoid snowball effects in other sectors.

En-route air traffic controllers in current operations, however, very rarely command changes to the (cost optimized) speed of aircraft, but prefer heading and/or altitude changes in favor of efficiency [13]. To align more to current operations, a different design choice for the construction of the horizontal Solution Spaces could have been to assume a *variable exit time* and *fixed speed* along the trajectory. The basic form of the horizontal Solution Spaces would then remain the same, but instead of representing control actions requiring a higher speed, the control space would represent control actions causing a sector exit delay. This would provide the controller a view of a different subset of task-specific control space within the same overarching multidimensional Solution Space.

In this thesis, the TSD was primarily used as a tool to *de-conflict* aircraft by modifying the speed along segments of a given horizontal trajectory. In the visualization of the time-based Solution Space, however, the fixed sector exit time was not explicitly taken into account. This contradicts the previous assumption of a fixed sector exit time that underpins the horizontal Solution Space. As a result, the *perceived* control space in the TSD was larger than the *actual* control space when strictly enforcing the exit constraints. Explicitly fixing the exit time would align more closely to the mechanization of the horizontal Solution Space concept, but would also greatly reduce the available control space.

Previous research has shown how the TSD can be used as a useful tool for *planning* and *sequencing* streams of inbound traffic in area and terminal control [14], [15]. For this task, to modify the sector exit times of trajectories using the TSD and time-based Solution Spaces was the primary means of control. This again illustrates that use and usefulness of different subsets of the Solution Spaces depend on the specific task and time scale of the work domain scope.

Next, the shape and size of the unsafe control space and intentional constraints (i.e., separation and performance buffers) overlaid on the overall control space directly follows from the assumptions about rules, regulations, and further design choices. As shown in Chapter 5, increasing or decreasing separation minima or intentional performance buffers will have a large impact on the shape and size of 'safe' control space. All the above again emphasizes that any specific implementation and visualization of the Solution Space concept is highly dependent on the task-specific assumptions on the system constraints (i.e., fixed exit time, fixed speed, etc.), and is further subject to assumptions of the intentional constraints and design choices.

Finally, in this work, the high computational intensity of computing and visualizing the en-route Solution Space representations in real-time warranted a number of simplifications. For the *computation* of the Solution Spaces, aircraft were treated as discrete point-masses (no turn radius) and path segments were discretized (no acceleration and deceleration). On the physical scale and time frame of en-route control, however, these simplifications did not have a significant impact on the usability of the concept.

Interface mechanization. Adding additional data and visualizations to existing controller interfaces will inevitably lead to a higher level of interface complexity and clutter. When made visible on the PVD, the horizontal Solution Space of one aircraft could obscure the view of other traffic. When working on the TSD screen region, awareness of the physical locations of aircraft could be lost. Further, air traffic control is a *centralized* task, but the Solution Spaces present the control space for individual aircraft. There could be a danger that controllers over focus on individual aircraft and could lose focus of the overall traffic state. These are crucial aspects that must be addressed in any specific implementation of Solution Space-based interfaces. To overcome this, in this work, the Solution Spaces of an aircraft were only visualized when the controller explicitly selected that aircraft on the interface. When no aircraft was selected, the controller will have to observe the traffic and reason about their actions and control strategy before implementing it.

The interface mechanization should further aim to support the work and strategies of air traffic controllers as best as possible. Both the interactions with the interface and the visual representation of the Solution Spaces should be tailored for the specific controller tasks. This could also entail providing different views of the Solution Spaces for different control strategies to *the same* controller; the controller could select to view a subset of the Solution Spaces bound by constraints that match a specific strategy. The visualization and mechanization of the horizontal Solution Spaces aimed to support a means of integrated *4D-vectoring* that resembles the current method of discrete vectoring (i.e., discrete speed and heading commands). The visualization and mechanization of the time-based Solution Spaces aimed to support the task of speed control as a means of separation.

7.1.3 Experimental Evaluation (RQ3)

The third part of this research focused on human-automation interaction with partial implementations of the Solution Space decision-support concept. Chapter 4 investigated the effect of the level of *expertise* of controllers on the use and *robustness* of control using the horizontal Solution Space interface. Chapter 5 repeated the previous experiment, but with the addition of intentional separation and performance buffers to the Solution Space visualizations. Finally, Chapter 6 focused on controller *strategies* with a combined horizontal and time-based Solution Space interface. These parts of the work aimed to answer the third research question:

> Research question 3
>
> How do human controllers with different levels of expertise use (partial) implementations of the Solution Space concept for tactical 4D trajectory-based control, what strategies do they apply, and what lessons can be learned?

Expert interfaces for expert controllers. Results of the three human-in-the-loop experiments in this thesis showed that all participants, from novice to experts, could use the partial Solution Space decision-support interfaces to control traffic in a 4D trajectory-based environment. Their prior experience with ecological interfaces and/or ATC ranged from none at all to being active air traffic controllers that had participated in previous experiments. This, however, should not be misinterpreted as that either the Solution Space representations are 'simple' or intuitive, nor that the task of tactical trajectory-based air traffic control is in any way a trivial one; the ATC domain is an inherently complex environment, and ecological interfaces neither strive to (or would be able to) simplify the complexities of the underlying work domain, nor trivialize the control task itself. This is underlied by Ashby's Law of Requisite Variety, that in the context of the ecological approach states that for any representation to be effective for control, it must be as complex as the problem to be solved [10], [16]. Although one could argue that part of the complexity has been removed by assumptions that narrowed the scope (see recommendations), the aforementioned is articulated by the fact that significant differences were observed in *how* individual controllers interpret and use the interfaces to formulate and execute their strategies.

In Chapter 4, three participant groups with different levels of expertise (novice, skilled, and expert) were asked to control en-route traffic using the horizontal Solution Space under various traffic and sector conditions. Expert controllers were found to follow inherently more *robust* control strategies than the skilled and novice participants. They were also found to more pro-actively control the traffic, whereas novices applied more reactionary control. However, experts also indicated that they had relatively *low trust* in the Solution Space representation, as it marked 'tight' solutions—although strictly feasible—as safe control space. In Chapter 5, the experiment of Chapter 4 was repeated for the novice and skilled participant groups, but here with additional separation and performance buffers added to the Solution Space visualizations. Results showed that the robustness of control strategies and consistency of control significantly improved for all participant groups.

These results can largely be explained when observing the decision flow map for the en-route Solution Space concept presented in Chapter 3 (Fig. 7.3). Assuming that a controller has only received a basic explanation of the interface, and has no prior experience in ATC, their control strategies will be fully guided by the presented information (shortcut ❶); the controller is (visually) alerted that there is a conflicting trajectory, and then without much further reasoning selects that aircraft and chooses a(ny) 'safe' location inside the Solution Space that resolves the conflict (i.e., rule-based control). By explicitly visualizing additional intentional safety buffers to such more novice controllers, their options for trajectory modifications are effectively limited to more robust solutions (as observed in Chapter 5).

Controllers with higher levels of expertise will be able to better reason about the Solution Space visualizations and the overall traffic state (shortcuts ❷ and ❸). Although the Solution Space interfaces themselves present the *egocentric* control space of individual aircraft, air traffic controllers perform their work from a *centralized* perspective. The results from the experiment show that the expert controllers were able to 'see beyond the information presented', and to integrate that information into a more comprehensive mental model of the overall control task. This allowed them to progress further into the rule-based domain, and ultimately allowed them to perform *critical thinking* and formulate knowledge-based strategies. This likely explains their relatively high number of control actions and lower level of trust in the presented information.

Further, as shown in Chapter 3, the Solution Space visualizations themselves do not only show *go* and *no-go* areas, but also implicitly contain additional information on the traffic structure. That is, for instance, the *location* of restricted zones due to other traffic implicitly indicates the crossing sequence (e.g., which aircraft crosses the other one in front or behind) and its shape implicitly indicates the crossing geometry (e.g., more head-on or more overtaking). With training and experience, controllers will be able to 'see' more information and interrelationships in the *same* Solution Space visualizations. This emphasizes that extensive training and experience with ecological interfaces remains crucial to achieve and maintain a high level of controller expertise.

Control Strategies. In Chapter 6, two participant groups with a different level of expertise were asked to control en-route traffic under varying traffic structures using both the horizontal and time-based Solution Space representations. Although the Solution Spaces show the causal and intentional con-

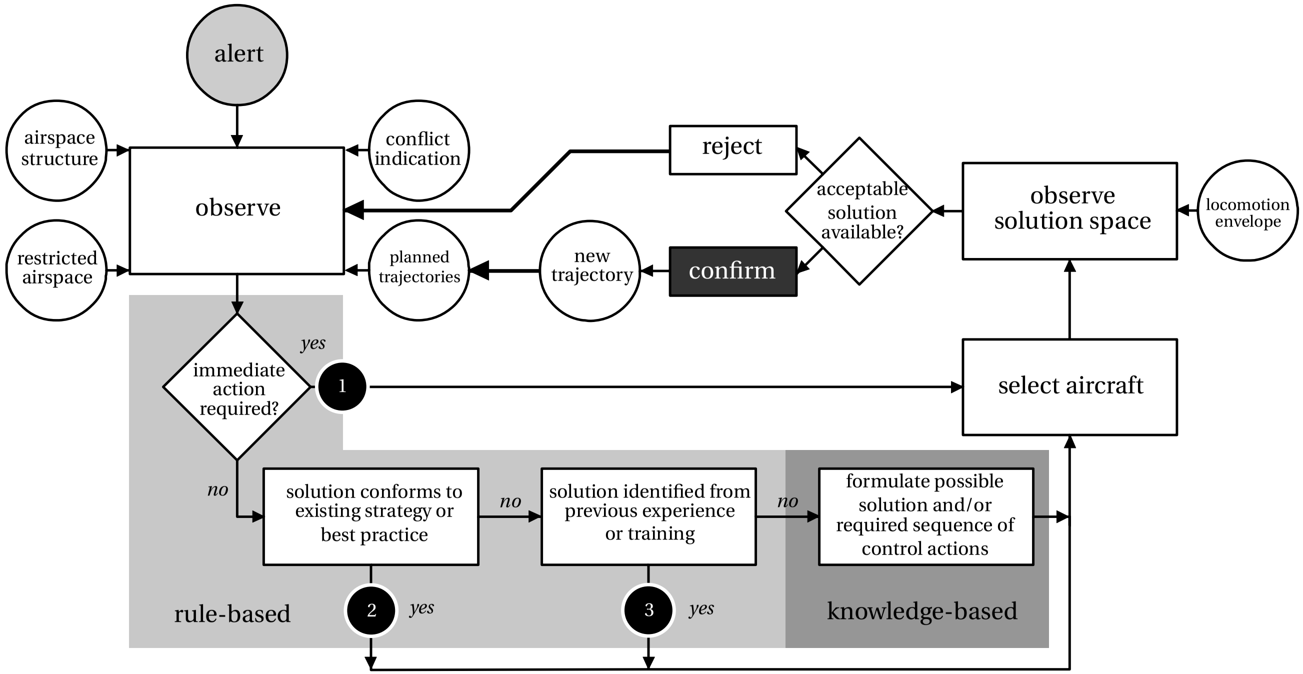

Figure 7.3: Decision flow map for formulating a strategy to resolve en-route trajectory-based conflicts. The three rule-based shortcuts are shown and numbered 1, 2 and 3 (taken from Chapter 3).

straints that bound the control space of individual aircraft, they do not infer a specific strategy or sequence of control actions. Results from the experiment showed that the observed control strategies were found to vary greatly between participants. However, the individual participants themselves typically preferred a certain strategy, and applied that strategy in all experiment runs. This indicates that the Solution Space concepts are *individual sensitive*, and allow freedom for controllers to formulate and implement their own unique—non-normative—strategies.

Although some control strategies will be more 'optimal' than others when seen from a strict perspective, ecological interfaces trade optimality in favor of flexibility. On the one hand, this can viewed as a strong point because operators will be able to adapt and shift strategies when necessary [17]. On the other hand, however, allowing the full freedom of control (especially without extensive training) can lead to unnecessarily inefficient operations. For instance, results from the experiment in Chapter 6 showed that some participants would merge streams of traffic only to create a more (optically) structured flow, but as a result increased trajectory deviations, thereby reducing efficiency.

The question then is, how, and to what extent should or could optimality better be promoted without decreasing the full flexibility of ecological interfaces? Controllers could be guided to more optimal strategies by adding or modifying intentional constraints. Automation could propose efficient control strategies within the Solution Spaces that the controller can then accept or reject (control by consent). Training could more explicitly be focused on optimality. Answering this question will become increasingly important in future work with the Solution Space concepts in order to gain acceptance by all involved stakeholders (e.g., air traffic control organizations, airlines and aircraft operators, and other airspace users).

7.2 Recommendations

The work in this research did not come without challenges and limitations. In the search for a common ground for tactical, trajectory-based air traffic control, the scope and and assumptions directly affected the complexity of the work domain and control task itself. The visualization and mechanization of the partial Solution Space interfaces followed directly from design decisions and trade-offs. Each human-in-the-loop experiment posed its own unique challenges. This section aims to address the main challenges and limitations, to place them in a broader context, and to provide recommendations for future research.

7.2.1 On Finding a Common Ground

Work domain complexity. The scope and assumptions in this research reduced the inherent complexity of the work domain when compared to a real-world trajectory-based air traffic control environment. Although this research showed the potential of the conceptualization of Solution Space as a viable underpinning for the design of decision-support interfaces, the following two limitations must be addressed to gain acceptance by the operational community.

First, the research considered a *fully deterministic* work domain. That is, in the computation of Solution Spaces, only deterministic internal constraints (performance envelope), and external constraints (intent of other traffic) were taken into account. As a result, control actions in the human-in-the-loop experiments were always executed immediately and absolutely (i.e., without uncertainty). *Fuzzy* constraints such as *uncertainties* in the prediction and execution of trajectories, *variability* in atmospheric conditions (i.e., wind, temperature, pressure), and *delays* in communication, for instance, were beyond the scope. Adding fuzzy constraints will be challenging, and will require to revisit and expand upon the work domain analysis. Where the boundaries of feasible and safe actions were well defined and deterministic in this thesis, fuzzy constraints will add *probabilistic attributes* to the concept of Solution Space. Rather than defining strict safe and unsafe areas, Solution Spaces will progress more towards representing probability zones (i.e., high or low likelihood of conformance towards the overall system goals). Further, taking varying atmospheric conditions into account (such as 4D wind fields) will distort the current symmetricity of the Solution Spaces. Although such fuzzy constraints can be included in the computational model of Solution Spaces, the biggest challenge will be to present this information in a way that resonates well with the mental model of the (human) controller.

Second, the partial implementations of the Solution Space concept did not encompass the full four-dimensionality of trajectory-based operations. The concepts allowed for horizontal and/or time-based control, but not for vertical control. To better support the work and strategies of operational air traffic controllers, the full—highly-coupled—control space will have to be integrated into the ecological decision-support interfaces. A main challenge here will be to unify the mechanization of trajectory modifications with the partial Solution Space visualizations in such a way that the implications of control in one dimension to the shape and size of the control space in the other dimensions is clear. That is, the definition of a 4D trajectory is *ground referenced*, whereas the aircraft performance envelope is referenced to its aerodynamic and instantaneous properties (e.g., max thrust, lift-drag polar, gross-weight, etc.) with respect to *local atmospheric conditions*; airspeed, vertical path, ground speed, climb and descent gradients are all highly coupled.

Location-dependent fuzzy (atmospheric) constraints, in combination with the full four-dimensionality of trajectory-based operations can rapidly let the complexity of the Solution Spaces explode. Addressing this will not only be a computationally intensive challenge, but will inherently and inevitably add greatly to the work domain and task complexity. Managing interface complexity and balancing workload, whilst upholding situation awareness will be crucial to address in future work when progressing to a more operational concept.

One recommendation to keep complexity manageable is to add *operational constraints* to the control task, effectively reducing the degrees of freedom of control. Typically, in terminal control, arriving and departing aircraft follow fixed—altitude constrained—routes that ensure separation by design. Crossing airways in en-route sectors are typically separated using different flight levels. Although limiting the degrees of freedom of control will decrease the flexibility of airspace use, structuring the airspace in a convenient way *by design* could support a more standardized form of operations—with standard strategies and solutions—that can significantly reduce the overall control task complexity.

Levels of automation. The Solution Space decision-support interfaces presented in this thesis primarily used automation as a means for the functions of *information acquisition* and *information analysis*. Information about the state and intent of traffic was integrated and presented in a way that action-relevant control spaces could directly be perceived by controllers. The subsequent functions of *decision selection* and *action implementation* were primarily manual tasks performed by human controllers.

The level of automation thus mainly occurred prior to any point of decision (i.e., information automation). Parasuraman et al. proposed that in such systems, automation can be applied on a continuous scale—from none to full—for all of these four functions [18]. In Fig. 7.4, an approximation of the levels of automation in this work is shown by ❶. Lowering the level of *information automation* in trajectory-based operations is practically impossible because the complexity of the control task all but excludes a manual alternative. However, increasing the level of *action automation* will not be without risks and direct consequences to the task of the human controller.

Ecological interfaces strive to make the underlying constraints of the work domain salient to operators (*domain transparency*). The decomposition in the domain-specific elements in the Abstraction Hierarchy is—by definition—

independent of both the specific task and the executive actor(s). That is, the functions and constraints hold for all specific tasks, and must be adhered to by all actors, either human, automated, or in a joint effort. The Solution Space interfaces in their current form, therefore, do not rule out adding more automation to support the controller in workload balancing. However, increasing automated decision selection and action implementation implies that automation will plan within, and act upon a specific subset of the constraint space.

Increasing the levels of decision making automation will shift the task of the air traffic controller more from active control towards supervisory control. One example of supervisory control is in the form of control by consent (or exception). Fig. 7.5 illustrates how automation could propose one or more positions within the horizontal Solution Space for a conflicted aircraft. In this scenario, the automation proposes two control actions based upon preselected strategies (e.g., minimize deviations, optimize robustness). The controller could then either select one of the two proposed solutions, or reject the advisory and opt for manual control (default: rejection after 10 seconds inactivity). The increase of action automation in this example on the overall levels of automation is indicated by ❷ in Fig. 7.4.

Such a shift is not good or bad in itself, but it will be crucial to address the following. First, the processes behind information analysis and decision selection by automated agents should be made transparent to the operator (*agent transparency* [19], [20]). Without agent transparency, the controller could lose *situation awareness* (what is it doing?) or lose *trust* in automation (why is it

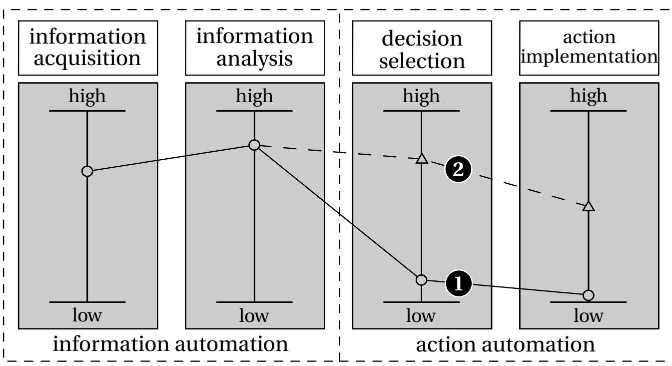

Figure 7.4: The four main automation functions applied on a continuous scale as proposed by Parasuraman et al. [18]. ❶ indicates the level of automation of the Solution Space concept proposed in this work. ❷ indicates the level of automation with a higher level of action automation (control by consent).

doing that?). This will be especially important when automation proposes solutions that are non-conformal to standard ATC practices. However, there is a risk that too much transparency might over-saturate the controller to be practically usable [21]. Second, even if the automation is transparent and perfectly reliable in almost all cases, *complacency* and reduced *vigilance* could emerge (the automation can't be wrong) [22]–[24]. This could lead to situations in which the controller accepts risky or dangerous solutions without detecting failures in the underlying automation. Third, automation is particularly useful in nominal situations as a means for reducing or balancing workload. However, automation will typically fail more frequently in non-nominal situations. There could then be a risk for *task saturation* caused by the unevenly distributed workload between the automation-supported task and hands-on manual control [18], [25].

Increasing the levels of action automation, but simultaneously expecting humans to remain ultimately responsible and take over when automation fails might seem paradoxical. Unless society is willing to assign full control to automation, and accept automation responsibility in case of (catastrophic) failures, any form of action automation must be transparent and unambiguous enough for the operator to understand and intervene. That implies that the operator herself must remain educated and skilled in the tasks and complexities taken over by automation.

In the control by consent example in Fig. 7.5, automation is responsible (in part) for decision selection, but the human executes the action implementation. This form of interaction can be seen as a *serial* work flow in which both humans and automation perform part of the task towards the final control de-

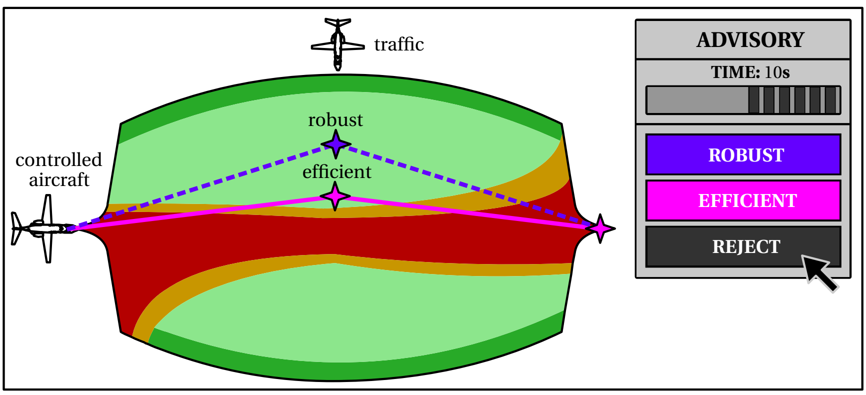

Figure 7.5: An example of action automation in the form of control by consent in the horizontal Solution Space decision-support concept.

cision. Previous studies have raised concerns about this approach, and showed that situation awareness and *out-of-the-loop* performance problems can in fact be worse compared to fully manual control [26]–[28]. Therefore, it is questionable that this type of automation will enhance operator performance and operational safety.

Perhaps this step should be skipped in favor of higher, and more *collaborative* forms of *human-automation teaming* [29]. Automation and humans could work in *parallel* and share the *same* tasks. One project in which this type of collaboration is currently being explored is project ARGOS by Maastricht Upper Area Control (MUAC) [30]. Here, the envisioned scenario is that human controllers and automation both control aircraft inside the same sector; automation is in full control of basic flights whilst the controller focuses on more complex flights. In this case, automated agents can be seen as a digital co-workers, performing the same tasks. For this collaboration to succeed, it is crucial that the human controllers gain a high level of trust and acceptance in the automation. Solution Space interfaces could (initially) play a role in gaining trust and acceptance by providing transparency to the controller about the reasoning and strategies of the automated agents. Simultaneously, the controller could use the same Solution Space interfaces as decision-support tools in performing the manual control task for complex flights.

Another division could be to design individual-sensitive automation by letting the controllers directly set and/or manipulate the underlying agent logic. Humans then steer automation on higher—more abstract—representations of the control task, but remain fully in-the-loop. One example of this type of collaboration using Solution Space interfaces was explored by Ten Brink et al.[31]. Here, controllers could add geometrical exclusion zones inside the horizontal—airway-based—Solution Space to steer the strategies of an underlying path-planning algorithm. Results showed that the Solution Space visualization, in combination with the exclusion zones, supported the controllers' understanding of the underlying algorithms.

7.2.2 On Visualizing the Solution Space

Centralized Solution Space? The ecological Solution Space Representations presented in this work visualize the egocentric control space of individual aircraft within the constraints posed by the air traffic control environment. Controllers must still observe, select, and manipulate the trajectories of individual aircraft. The task of an air traffic controller, however, is to perform control from a *centralized* standpoint. One could then argue that the current representational forms do not match the task of the controller, and that perhaps a more centralized view of the Solution Space would be beneficial. During this research, the concept of a multi-aircraft horizontal Solution Space representation, and flow-based Solution Spaces were also explored [32], [33].

Fig. 7.6 shows a schematic representation of the prototype multi-aircraft horizontal Solution Space concept in which two aircraft in a crossing conflict are de-conflicted. Fig. 7.6(a) shows the initial situation in which both aircraft are selected, and the resulting *combined* Solution Space for both aircraft is shown. In this interface concept, the controller can select a 'safe' location in the combined Solution Space through which both aircraft are re-routed, effectively controlling them simultaneously. Fig. 7.6(b) shows a control strategy that will minimize the path deviation, sharing the deviation equally between the aircraft. The usefulness and effectiveness of the (horizontal) multi-aircraft Solution Space concept, however, was found to be highly dependent on the specific traffic situation [32]. Multi-aircraft control showed benefits for controlling in-stream (trailing) aircraft and in simple crossing conflicts such as in Fig. 7.6. However, in more complex traffic scenarios, controllers preferred single aircraft control because it allowed them to better optimize individual trajectories. The latter is apparent when observing the four-aircraft Solution Space shown in Fig. 7.7. The available safe control space is relatively small, and a combined control action will lead to a situation in which deviations are unequally shared. Further, the resulting geometric situation in which all crossing aircraft merge at a single point in the sector would be viewed as unacceptable in real-world operations.

A second prototype focused on visualizing the Solution Space for selected airways (i.e., flow control), multiple aircraft control (grouping control), and critical point control [33]. Fig. 7.8 shows a schematic representation of how, by selecting an airway, the Solution Space is visualized for re-routing all (three) aircraft on that airway. Additionally, this prototype visualized the robustness of control actions inside the safe field of travel, indicated by a color gradient. As for the multi-aircraft horizontal Solution Space, the benefits of this concept over single aircraft control were highly dependent on the specifics of the traffic scenario. Further, the high computational load required for computing robustness in the Solution Spaces, even in scenarios with a limited amount of traffic, resulted in that the Solution Spaces could currently not be updated in real-time.

Based on the research and human-in-the-loop experiments with the multi-aircraft and flow-based Solution Space representations, these concepts should not immediately be abandoned, but rather be seen as a possible addition to the Solution Space-based decision-support toolkit. In situations that lend themselves well for multi-aircraft control, the controller could opt for this type of control rather than modifying multiple individual trajectories. The ability to modify individual aircraft trajectories, however, will remain a fundamental basic means of control, and cannot be replaced by multi-aircraft control only.

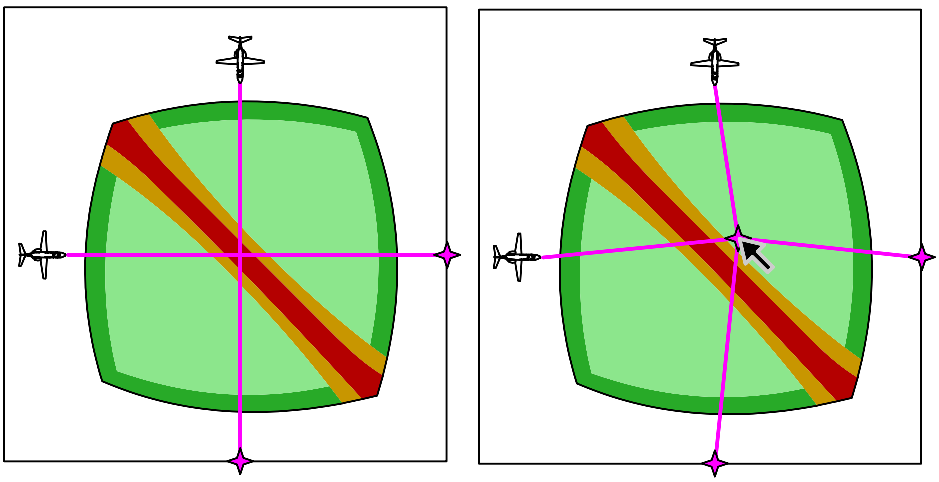

(a) Two aircraft Solution Space (b) Combined control action

Figure 7.6: Horizontal multi-aircraft Solution Space for a two aircraft crossing conflict.

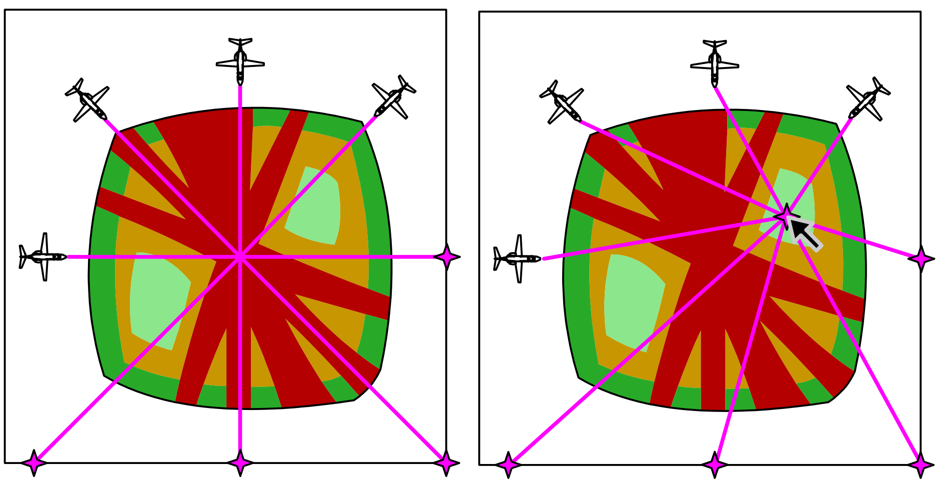

(a) Four aircraft Solution Space (b) Combined control action

Figure 7.7: Horizontal multi-aircraft Solution Space for a four aircraft crossing conflict.

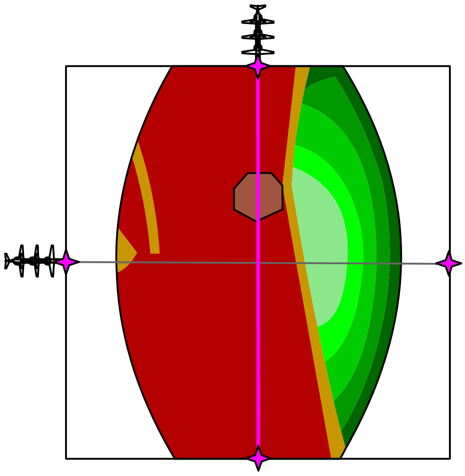

(a) Airway-based Solution Space (b) Airway-based control action

Figure 7.8: Horizontal airway-based Solution Space.

Centralized versus de-centralized control. The Solution Space concept was inspired by previous work that explored the visualization of affordance zones as a means for trajectory re-planning on the flight deck [34]. This concept focused on *egocentric* trajectory manipulation for (ownship) weather and hazard avoidance. For the Solution Spaces presented in this thesis, the intent of other traffic was added to the visualization of the available control space, and the concept was used in a centralized environment. A question that arises is: could the Solution Space concept be used in a de-centralized environment as a means of self separation by pilots? One of the major differences between a centralized and de-centralized environment is that in the centralized environment, a single controller organizes the traffic from a top-down view. In a de-centralized environment, a form of *implicit coordination* between pilots will be required. A previous study with multi-actor—egocentric—self separation was performed that highlighted a number of issues with this concept [35], [36]. Although this study focused on state-based control (i.e., modifying heading and speed), the same issues will likely also arise in trajectory-based control. Similar to two people passing each other in a hallway, difficulties with implicit coordination can lead to ones conflict resolution strategy interfering with that of the other. Many times, conflicts were found to be resolved by one aircraft instead of cooperatively. In case of cooperative resolutions, path deviation was often found to be unequally shared between aircraft. Typically, each pilot will seek to minimize their own path deviation, which sometimes resulted in tight solutions and minor separation violations. Although de-centralized Solution Space representations could be used for small trajectory adjustments by pilots, a centralized form of organization has clear benefits in structuring and optimizing the overall traffic movements in the absence of explicit de-centralized coordination.

Visualizing complexity. One of the biggest challenges towards developing more operational Solution Space decision-support concepts that convey the full four-dimensionality of control, uncertainty, and fuzzy constrains, will be to transform the exponential increase in complexity into useful visual representations. That is, to convey all relevant information in a form that is understandable and interpretable by humans, without ignoring or simplifying the underlying work domain complexities. The air traffic control task is, in essence, an ever morphing, highly coupled four-dimensional puzzle. Rather than visualizing deterministic shapes, *probability visualizations* and (propagation in) uncertainty of control spaces will start to play a key role. Further, harmonizing the partial Solution Space visualizations in a way that the coupling between control actions in multiple dimensions is perceivable in an intuitive way is most certainly not a straightforward task.

One of the most evident approaches to reduce the necessary visual complexity of visual representations is to reduce the control task complexity itself. As mentioned previously in the recommendations on work domain complexity, the control task itself could be simplified by adding additional operational constraints that reduce the degrees of freedom of control. For instance, if aircraft fly along predetermined, vertically separated airways, the control task is reduced to one of keeping in-trail separation by speed control alone. Further, standardized strategies and solutions could also remove the necessity to display the full complexity and dimensionality of the Solution Spaces; only the parts of the Solution Spaces relevant to the specific situation and strategy need to be conveyed. The Solution Space visualizations could then be made *adaptable*. Controllers can choose what part(s) of the Solution Spaces they want to zoom in to or hide, and switch between different views in different situations. In conclusion, rather than ignoring any work domain complexities, the control task complexity itself can be simplified by design so that the required adaptable Solution Space visualizations remain perceivable and manageable for humans.

7.2.3 On Experimental Evaluation

Simulation realism. The human-in-the-loop experiments in this research were performed in a simulated environment that presented a simplified view of trajectory-based operations. First, the simulations did not provide the controller with the capability to control all degrees of freedom of trajectories. Especially the absence of vertical control in the simulated en-route environment limited the set of conflict resolution strategies. Conflicting trajectories that would normally be separated vertically now had to be separated horizontally (in-plane). Especially the expert participants marked this as an important limitation of the current interface design. Second, the simulated airspace environment did not reflect the full complexity of the real world control task. All experiments provided the controller with *full and immediate* authority. Secondary functions such as delayed communications, atmospheric disturbances, non-conformity of execution by aircraft, and multi-sector coordination were not simulated. This likely decreased the sense of urgency (responsibility), workload, and task-demand compared to real world scenarios.

Simulation realism could have been increased by adding the above functions. However, experiment design itself is always a fine balance between *operational realism* and *experiment control*. On the one hand, simplifying realism can present unrealistic scenarios that can be approached as a—far from real life—arcade game by the participants. On the other hand, especially in shorter duration experiments as presented in this thesis, adding more complexities can create unwanted noise that confounds—and thereby obscures—the experiment results.

One observation during all three experiments was that participants would typically only focus on an aircraft as soon as it entered the sector, re-plan its trajectory, and then not touch that aircraft for the remainder of the scenario. There was no need to monitor the conformance of flights once inside the sector, and thus, choosing 'tight' solutions did not have any negative implications. One way to encourage participants to perform a higher level of critical thinking and exert more caution in formulating their control strategies would be to introduce uncertainties between the trajectory forecast and its actual execution. This could, for instance, be provoked by introducing external factors such as variable wind fields causing aircraft to drift from their planned trajectories. The assumption of drift is in line with the capabilities of current aircraft; in most modern Flight Management Systems (FMS) it is possible to set a Required Time of Arrival (RTA) at a certain waypoint. The FMS then uses the pilot input of weights, cruise flight level, and expected wind to compute a speed schedule to arrive at the waypoint at the desired time. This is an open loop process; the FMS does not actively recompute the speed schedule, and the actual arrival time *will* drift from the RTA. Instead, the pilots will receive a message in case the deviation becomes larger than a pre-determined margin (typically 15 seconds).

Expertise and training. The limited time for training in the human-in-the-loop experiments limited the participants' potential to fully understand and effectively use all the information presented in the Solution Space visualizations. The interfaces presented in this work are highly information rich, and contain many intricacies and interrelationships that are not directly salient in the immediate shape and form of the visualizations. Solution Space is not a static entity, but will dynamically change as time progresses (i.e., narrow down or open up). Many times during the experiments, participants remarked that they came to new 'deeper' insights about the Solution Spaces and other aspects of the interfaces. This emphasizes that the spatial and temporal aspects of Solution Spaces in dynamic environments cannot be learned without extensive hands-on experience and training in *equally dynamic* environments.

Although participants were still actively learning about the Solution Space visualizations and the interfaces during the experiments, experts had more fundamental knowledge about the control task as a result of their real world experience. This likely explains the observed differences in control strategies compared to more novice participants; where novice participants were still actively gaining knowledge about the work domain and the control task itself, expert controllers could apply their current knowledge base to their strategy selection. Experts therefore used the Solution Space representations in a *different* way than novices, using it more as a means of strategy validation than decision-selection.

Longitudinal studies have shown that ecological interfaces can lead to more *consistent* and *superior performance* compared to traditional interface designs over time[37], [38]. The studies also showed that ecological interfaces can lead to a more *functionally organized knowledge base* about the control task when participants took a deep approach to learning [38]. Extending the experiments in this thesis to more longitudinal studies would therefore most likely improve the consistency, performance, and knowledge base of novice participants. However, the extent of the knowledge base would only go as far as the scope of the work domain and information provided in the interfaces. The Solution Spaces themselves do not convey information about best practices, preferred strategies, optimality, etc. Therefore, it is unlikely that novices would be able to reach the same level of operational expertise as expert controllers by self-learning alone.

Air traffic controllers undergo years of education and (on-the-job)training before they are are allowed to independently control traffic. A majority of training is not solely focused on the technicalities of *how* to control air traffic, but on gaining a deeper understanding of the work domain by practicing and evaluating simulated and real world scenarios. A previous study showed that ecological interfaces can play an important role in the early stages of deep knowledge development [39]. Results from that study suggest that training with ecological interfaces can change how participants thought about and approached their control problem, even after removing the ecological support. The Solution Space interfaces, therefore, could be useful as a training tool for novices and experts alike to gain more insight in the dynamics of trajectory-based operations. Even if the concept of Solution Space itself is not directly used in decision-support interfaces in a future operational setting. However, additional training will be required on the Solution Space interfaces themselves.

7.3 Conclusions

This thesis showed how the construct of *Solution Space* can act as a common basis for the design and development of human-centered ecological decision-support tools for tactical trajectory-based air traffic control. Although the work primarily focused on the en-route air traffic control environment, the underlying concept can be extended for use in other tactical air traffic control environments. However, the *usefulness* and *effectiveness* of any partial representations of Solution Spaces will be highly dependent on the specifics of the control task and control strategies for which it purposed.

In this work, the deployment of automation was limited to *information automation* only, but the presented concepts do not rule out the addition of higher levels of *action automation*. However, any form of action automation must be transparent and unambiguous enough for the operator to understand and intervene; the operator herself must remain educated and skilled in the tasks and complexities taken over by automation when automation fails. Here, Solution Space visualizations can play a role in both gaining insight into the strategies and actions of automation, and in decision-support when performing the fully manual control task.

The three human-in-the-loop experiments showed that participants with different levels of expertise could use the partial Solution Space decision-support interfaces successfully, albeit in a different way. Novices typically used the Solution Spaces as a means of *decision selection* where experts frequently used it as a means of *strategy validation*. Adding intentional constraints to the Solution Spaces showed to improve the control robustness of novice participants, however, experts will already know how to avoid wandering into this terrain. This emphasizes that ecological interfaces, by themselves, do not remove or replace the advantages of expertise and deep domain knowledge gained by extensive training and experience. Ecological interfaces could be beneficial in improving ones knowledge base about the work domain, however, the depth of knowledge will only go as far as the information presented on the interfaces.

The *deterministic* nature of the Solution Space visualizations and scenarios in the human-in-the-loop experiments in this work significantly reduced the complexity, urgency, and workload of the participants. To progress to a more operational concept of Solution Space, the work domain analysis should be revisited in future work to include *uncertainty* and the *full dimensionality* of the control task. This will inevitably increase the complexity of both the underlying work domain and any subsequent iterations of Solution Space interfaces. Trade-offs will need to be made between the degrees of freedom of control and the control task complexity in order to keep *workload* manageable, and upholding *situation awareness* practicable by human operators.

Bibliography

[1] J. Rasmussen, A. M. Pejtersen, and L. P. Goodstein, *Cognitive Systems Engineering*, ser. Wiley series in systems engineering October. Wiley-Interscience, 1994, p. 396.

[2] K. J. Vicente, *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. Mahwah, NJ: Lawrence Erlbaum Associates, 1999, pp. 1–415.

[3] K. J. Vicente and J. Rasmussen, "Ecological Interface Design: Theoretical Foundations.," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 22, no. 4, pp. 589–606, 1992.

[4] C. M. Burns and J. Hajdukiewicz, *Ecological Interface Design*. Boca Raton, FL, USA: CRC Press, 2004.

[5] J. Rasmussen, "The Role of Hierarchical Knowledge Representation in Decision Making and System Management," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 15, no. 2, pp. 234–243, 1985.

[6] "SESAR 2020 Concept of Operations Edition 2017," EUROCONTROL, Brussels, Belgium, Report, 2017.

[7] C. A. L. Westin, C. Borst, and B. G. Hilburn, "Strategic Conformance: Overcoming Acceptance Issues of Decision Aiding Automation?" *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 41–52, 2016.

[8] M. Bekier, B. Molesworth, and A. M. Williamson, "Tipping point: The narrow path between automation acceptance and rejection in Air Traffic Management," *Safety Science*, vol. 50, no. 2, pp. 259–265, 2012.

[9] M. Bekier and A. M. Williamson, "Why Air Traffic Controllers Accept or Refuse Automated Technology," in *16th International Symposium on Aviation Psychology*, 2011, pp. 615–620.

[10] M. M. Van Paassen, J. Ellerbroek, M. Mulder, and J. M. Flach, "Ecological Interface Design for Vehicle Locomotion Control," *IEEE Transactions on Human-Machine Systems*, vol. 48, no. 5, pp. 541–555, 2018.

[11] D. D. Woods, "Visual momentum: A concept to improve the cognitive coupling of person and computer," *International Journal of Man-Machine Studies*, vol. 21, no. 3, pp. 229–244, 1984.

[12] "European ATM Master Plan 2020 Edition: Executive View," SESAR Joint Undertaking, Brussels, Belgium, Report, 2019.

[13] J.-F. D'Arcy and P. S. Della, "Air Traffic Control Specialist Decision Making and Strategic Planning - A Field Survey," Federal Aviation Administration, Atlantic City, NJ, USA, Report, 2001.

[14] R. E. Klomp, M. Mulder, M. M. Van Paassen, and M. J. Roerdink, "Redesign of an Inbound Planning Interface for Air Traffic Control," in *Proceedings of the AIAA Guidance, Navigation, and Control (GNC) Conference and Exhibit*, American Institute of Aeronautics and Astronautics, Aug. 2011.

[15] R. E. Klomp, M. M. Van Paassen, M. M., and M. J. Roerdink, "Air Traffic Control Interface for Creating 4D Inbound Trajectories," in *Proceedings of the 16th International Symposium of Aviation Psychology*, Dayton, OH, USA, 2011, pp. 263–268.

[16] W. R. Ashby, *An Introduction to Cybernetics*. London, U.K.: Chapman and Hall, 1956.

[17] C. Borst, J. M. Flach, and J. Ellerbroek, "Beyond Ecological Interface Design: Lessons From Concerns and Misconceptions," *IEEE Transactions on Human-Machine Systems*, vol. 45, no. 2, pp. 164–175, 2015.

[18] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation.," *IEEE Transactions on Systems, Man, and Cybernetics. Part A, Systems and Humans*, vol. 30, no. 3, pp. 286–297, 2000.

[19] J. L. Wright, J. Y. C. Chen, and S. G. Lakhmani, "Agent Transparency and Reliability in Human–Robot Interaction: The Influence on User Confidence and Perceived Reliability," *IEEE Transactions on Human-Machine Systems*, vol. 50, no. 3, pp. 254–263, 2019.

[20] A. Bhaskara, M. Skinner, and S. Loft, "Agent Transparency: A Review of Current Theory and Evidence," *IEEE Transactions on Human-Machine Systems*, vol. 50, no. 3, pp. 215–224, 2020.

[21] R. R. Hoffman, S. T. Mueller, G. Klein, and J. Litman, "Metrics for explainable AI: Challenges and prospects," arXiv preprint, arXiv:1812.04608, 2018.

[22] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

[23] R. Parasuraman and V. Riley, "Humans and Automation: Use, Misuse, Disuse, Abuse," *Human Factors*, vol. 39, no. 2, pp. 230–253, 1997.

[24] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "Situation awareness, mental workload, and trust in automation: Viable, empirically supported cognitive engineering constructs," *Journal of Cognitive Engineering and Decision Making*, vol. 2, no. 2, pp. 140–160, 2008.

[25] N. B. Sarter, D. D. Woods, and C. E. Billings, "Automation Surprises," in *Handbook of Human Factors & Ergonomics, 2nd ed.* G. Salvendy, Ed., Wiley, 1997, pp. 1926–1943.

[26] M. R. Endsley and E. O. Kiris, "The Out-of-the-loop Performance Problem and Level of Control in Automation," *Human Factors*, vol. 37, no. 2, pp. 381–394, 1995.

[27] M. R. Endsley and D. B. Kaber, "Level of Automation Effects on Performance, Situation Awareness and Workload in a Dynamic Control Task," *Ergonomics*, vol. 42, no. 3, pp. 462–492, 1999.

[28] D. B. Kaber and M. R. Endsley, "The Effects of Level of Automation and Adaptive Automation on Human Performance, Situation Awareness and Workload in a Dynamic Control Task," *Theoretical Issues in Ergonomics Science*, vol. 5, no. 2, pp. 113–153, 2004.

[29] M. R. Endsley, "From here to Autonomy: Lessons Learned from Human–automation Research," *Human Factors*, vol. 59, no. 1, pp. 5–27, 2017.

[30] EUROCONTROL. (2019). "Blueprint for High-performance International ATM Services," [Online]. Available: `https://skybrary.aero/sites/default/files/bookshelf/5341.pdf` (visited on 10/27/2022).

[31] D. S. A. Ten Brink, R. E. Klomp, C. Borst, M. M. Van Paassen, and M. Mulder, "Flow-based Air Traffic Control: Human-machine Interface for Steering a Path-planning Algorithm," in *Proceedings of the 2019 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, Bari, Italy: IEEE, 2019, pp. 3186–3191.

[32] R. Nagaraj, R. E. Klomp, C. Borst, M. M. Van Paassen, and M. Mulder, "Exploring the Potential Benefits of Multi-aircraft Trajectory Manipulation in Future Air Traffic Control," in *Proceedings of the 2019 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, Bari, Italy: IEEE, 2019, pp. 3699–3704.

[33] J. Pinto, R. E. Klomp, C. Borst, M. M. Van Paassen, and M. Mulder, "Design of an Ecological Flow-based Interface for 4D Trajectory Management in Air Traffic Control," in *Proceedings of the 5th CEAS Air and Space Conference*, Delft, The Netherlands, 2015, pp. 1–10.

[34] B. J. A. Van Marwijk, C. Borst, M. Mulder, M. Mulder, and M. M. Van Paassen, "Supporting 4D Trajectory Revisions on the Flight Deck: Design of a Human-Machine Interface," *The International Journal of Aviation Psychology*, vol. 21, no. 1, pp. 35–61, 2011.

[35] J. Ellerbroek and M. M. Van Paassen, "Evaluation of a separation assistance display in a multi-actor experiment," in *Proceedings of the 16th International Symposium on Aviation Psychology*, 2011, pp. 1–6.

[36] J. Ellerbroek, "Airborne Conflict Resolution In Three Dimensions," Ph.D. Dissertation, Faculty of Aerospace Engineering, Delft University of Technology, Delft, The Netherlands, 2013.

[37] K. Christoffersen, C. N. Hunter, and K. J. Vicente, "A Longitudinal Study of the Effects of Ecological Interface Design on Skill Acquisition," *Human Factors*, vol. 38, no. 3, pp. 523–541, 1996.

[38] ——, "A Longitudinal Study of the Effects of Ecological Interface Design on Deep Knowledge," *International Journal of Human-Computer Studies*, vol. 48, no. 6, pp. 729–762, 1998.

[39] C. Borst, R. M. Visser, M. M. Van Paassen, and M. Mulder, "Exploring Short-Term Training Effects of Ecological Interfaces: A Case Study in Air Traffic Control," *IEEE Transactions on Human-Machine Systems*, vol. 49, no. 6, pp. 623–632, 2019.

Appendix A

GPU-based Solution Space Implementation

A description is provided in this appendix of how Solution Spaces were computed and visualized in the various prototype interfaces presented in this thesis. By utilizing recent technologies in Graphics Processing Unit (GPU) parallel programming, Solution Spaces were able to be generated in real-time using consumer-grade hardware without performance issues. Pseudocode is given for the computation of both the horizontal and time-based Solution Spaces.

A.1 Methodology

TRADITIONALLY, most algorithms and computations running on computers are performed by using the Central Processing Unit (CPU). The CPU has a high clock frequency and allows for fast serial computations in computation loops. The CPU has a flexible architecture with a large instruction set making it suitable to be used for everything from running an operating system to posing as an opponent in a game of digital chess. With the early introduction of computer generated imagery (CGI) in the scientific and military field, and more recently accelerated by the use of CGI in other professions such as the entertainment industry (i.e., movies, computer games), the CPU architecture increasingly became a bottleneck in handling the required amount of computations to power these applications.

This led to the introduction of the Graphics Processing Unit (GPU) as a separate processor next to the CPU. The main difference between the CPU and GPU architectures is that where the CPU has a small amount of cores (typically 2-12) with a high clock frequency, the GPU has a large amount of smaller cores (typically 2.000-10.000) with a lower clock frequency. The instruction set of the GPU is much more limited than that of the CPU, and is mainly optimized for matrix and vector computations. The power of the GPU, however, lies in that all cores can perform parallel computations within the same clock cycle; where the CPU can perform a few floating point operations each cycle, the GPU can perform thousands.

Computing and generating the horizontal and time-based Solution Spaces by the GPU cores in the prototype interfaces in this thesis allowed them to be visualized and updated in real-time. The underlying frameworks that were used are OpenGL 2.0 and GLSL (OpenGL Shader language) 1.2. Input to the GPU is an *intent buffer* containing details of the trajectory of the selected aircraft and all other traffic (i.e., waypoint locations and timings). The intent buffer used in this work was in the form of a 32-bit float RGB texture, and was filled by coding the aircraft intent (x,y,t) into the color slots. Output is a *frame buffer* containing the colored pixels of the Solution Space representations.

The experiment scenarios simulated up to 20 aircraft with a maximum of 20 waypoints for each flight. No significant performance issues were found when running at 60Hz on a consumer-grade desktop with a NVIDIA GTX 970 (4GB VRAM) GPU, and using a 30", 2560 x 1600 pixel display (> 200M pixel computations per second). The following paragraphs illustrate how the horizontal and time-based Solution Space visualizations were generated using the intent buffer in the *rasterization* process of the GPU.

A.2 Pseudocode

A.2.1 Generation of the horizontal Solution Space

Fig.A.1 illustrates the rasterization and computation process on the plan view display to generate the horizontal Solution Space for a selected aircraft (A_{sel}). Each grid-square in the figure represents one on-screen pixel for which the following computations are performed in sequential order to determine its color. In pseudocode, for each on-screen pixel p:

1. Determine the real-world spacial location corresponding to that pixel (p_x,p_y), and
2. Determine the closest trajectory segment of the selected aircraft to that position, and
3. Determine if the point is abeam that segment within a given cutoff angle (e.g., 80°), if not, discard, if so
4. Compute the distance of the track passing from the segment start point through p to the segment end point, and
5. Compute the required speed (V_{req}) to arrive at the segment end point at the original time, and
6. Determine if V_{req} is within the feasible locomotion envelope of the aircraft, if not, discard, if so,
7. Compute the closest point of approach (CPA) of other traffic to the two resulting trajectory segments using the intent buffer, and
8. Color the pixel in the following order of priority:

   a) If CPA $< 5NM$: unsafe control space color, else

   b) If CPA $< 5NM + \Delta D_{sep}$: separation buffer color, else

   c) If $V_{req} \leq V_{min} + \Delta V_{min}$, or $V_{req} \geq V_{max} - \Delta V_{max}$: performance buffer color, else

   d) Safe control space color

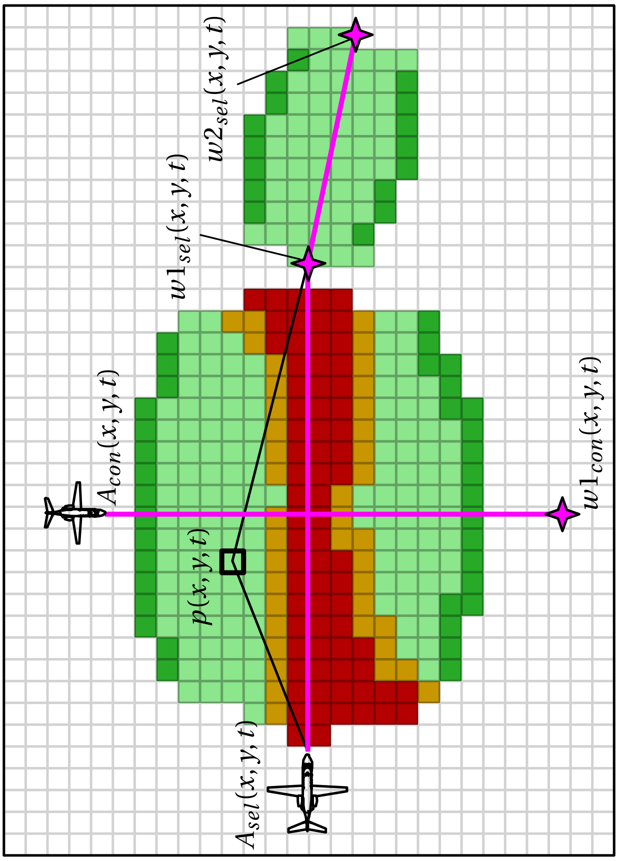

Figure A.1: Raterization for the lateral Solution Space

A.2.2 Generation of the time-based Solution Space

Similar to the previous paragraph, fig.A.2 illustrates the rasterization and computation process of *time-based* Solution Space for the selected aircraft. For each on-screen pixel p:

1. Determine the distance to go (DTG) to the exit fix and time at that pixel (p_d, p_t) corresponding to the graph axis, and
2. Determine to what trajectory segment of the selected aircraft that DTG belongs, discard if no valid segment, else
3. Determine the real-world spacial location of the pixel (p_x,p_y) at that DTG using the intent buffer, and
4. Compute the required speed (V_{req}) to arrive at that position at p_t from the segment start point, and
5. Compute the CPA of other traffic to (p_x,p_y) at time p_t, and
6. Color the pixel in the following order of priority:

    a) If CPA $< 5NM$: unsafe control space color, else
    b) If CPA $< 5NM + \Delta D_{sep}$: separation buffer color, else
    c) If $V_{req} < V_{min}$, or $V_{req} > V_{max}$: discard, else
    d) If $V_{req} \leq V_{min} + \Delta V_{min}$, or $V_{req} \geq V_{max} - \Delta V_{max}$: performance buffer color, else
    e) Safe control space color

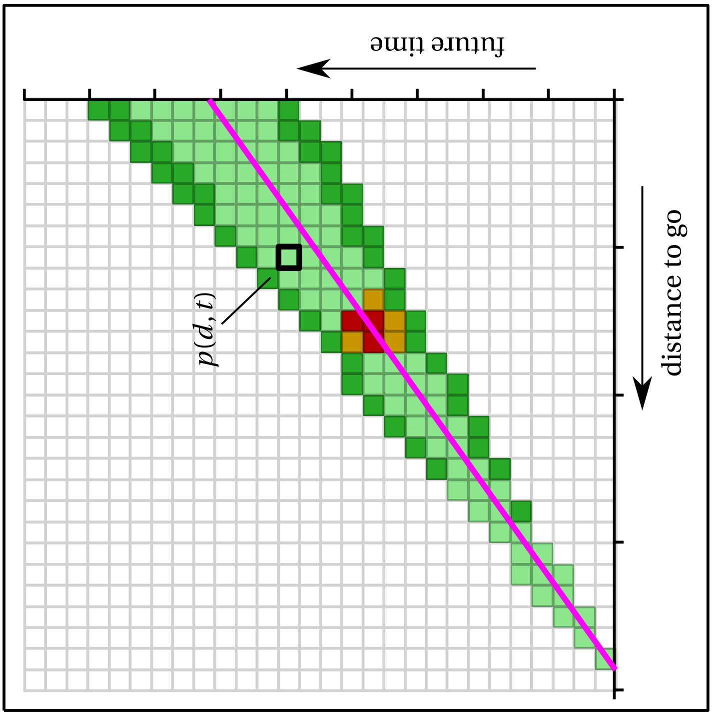

Figure A.2: Raterization for the time-based Solution Space

Appendix B

Interactive Training Script

This appendix contains the experiment briefing and interactive training script for the human-in-the-loop experiment performed for the strategies analysis in Chapter 6. The training consisted of a total of ten interactive scenarios. Participants were asked to closely follow the instructed steps throughout the first four scenarios, but were free to solve the last six scenarios in their own way. A screenshot of the initial setting of each of the first four scenarios is provided after the instructed steps.

B.1 Experiment Briefing

Purpose of the training

IN order to have a good understanding of how to perform your role as future air traffic controller in the main experiment, all tools and features that are available to you in the experiment simulator will be described in this training session. The training will be in the form of an interactive step-by-step script that will guide you through a number of scenarios. Each scenario will focus on a specific learning objective. At certain points during the scenario you may be required to answer one or more questions to test your understanding so far.

Your main task in the experiment will be to manage the traffic safely, and to try to adhere as much as possible to the initial traffic structure. That is, to keep any trajectory deviations and delays at the sector exit point as small as possible. This will be explained in more detail during the training.

Please try to talk out loud and try to motivate your reasoning for the decisions you make during the training scenarios. Read all the instructions carefully and don't hesitate to ask questions if something is unclear. During the training you are free to ask questions or ask for help, but in main experiment you will be asked to control the traffic without external interference.

Airspace and traffic

The controlled airspace used in the training scenarios and in the main experiment are artificial, en-route upper airspace sectors that are designed especially for this experiment. All aircraft resemble a generic type of medium-sized commercial airliner and have equal performance characteristics (equal speed envelope, acceleration, etc.). You will be able to manipulate the route and the speed of the aircraft, but vertical movements are not supported. All aircraft in this experiment fly at the same altitude, so vertical separation will not be possible.

B.1.1 Training Scenario 1

Part 1: System functionality and basic representations

The simulation is paused at this point, so please take the time to carefully read each following step:

1. The experiment simulator is built up by two separate screens:

- PVD (Plan View Display): The screen on the left hand side shows the top-down radar view of the sector, the entry and exit waypoints and all aircraft. The controlled sector in the training session has 12 unique entry and exit points, and in this scenario there is one controlled aircraft (callsign: BMS02N). You will use this screen to manipulate the route of the aircraft.
- TSD (Time-Space Diagram): The screen on the right hand side is a so-called Time-Space Diagram and will visualize information about the trajectory of a selected aircraft in terms of distance and time. A more in-depth explanation will follow later on.

Basic information on the PVD

2. The heading of BMS02N is indicated by a speed vector that is currently aligned with its route. The tip of the speed vector indicates the position at which the aircraft will be when following the current heading for 90 seconds at the current speed (indicated in the label; 265kts). A longer speed vector therefore indicates a faster flying aircraft. The aircraft is flying towards exit point TAMUK, shown by its route indicated by a thin line.

3. Highlight BMS02N and its route by hovering over it with the mouse in the PVD.

4. Left-click on the highlighted aircraft to select it.

5. The selected aircraft and its route will turn cyan. The waypoints along the route of a selected aircraft are visualized by magenta star symbols. BMS02N has one active waypoint that is located at the sector exit point (TAMUK). The planned speed towards this point is shown below the star symbol (also 265 KTS).

6. The shaded area that has appeared along the route of BMS02N is the so-called Travel Space of the aircraft. The Travel Space shows the area in which the aircraft can be rerouted and will still be able to arrive at the originally planned time at the sector exit point. Note that any deviation from the current direct route to the exit point will lead to a longer trajectory, and as a result, the aircraft will have to fly faster to reach the original exit time. The travel space is therefore bound by the speed envelope of the aircraft. That is, the travel space is bounded by the maximum speed that the aircraft can fly.

7. The darker shaded area at the edge of the Travel Space shows the region where the aircraft is required to fly close to its maximum speed. This is less efficient in terms of fuel usage, and therefore should, if possible, be avoided.

Basic information on the TSD

8. The TSD (right screen) shows the time-space representation of the trajectory. Here, the x-axis indicates the distance from the sector exit point along the current trajectory. The y-axis indicates future time. The cyan line represents the trajectory of the aircraft. Observe that at the current time (00:00), the aircraft has approximately 175 nautical miles to fly until reaching the exit point. The arrival time of the aircraft at the sector exit point is approximately at (00:26), and is indicated by the intersection of the line with the time-axis.

9. The position of the aircraft label along the time axis in the TSD indicates the current exit time of the aircraft. The cyan diamond along the time axis indicates the originally planned exit time of the aircraft. Note that these are now the same, but in case of a delay they will be different.

10. The speed envelope of the aircraft is also represented in the TSD by a shaded area. Similar to the Travel Space, the darker shaded area indicates the less efficient speeds. The intersection of the area with the time axis indicates the possible arrival times of the aircraft at the exit and is now approximately from (00:23) to (00:35). Note that if the aircraft would fly slower than currently (i.e., arrive at a later time), the time-space line will be steeper. Vice versa, a more shallow line indicates a faster flying aircraft.

11. Further, the white triangle at the left-bottom of the TSD is a slider that can be used to make a projection of the future aircraft movements. Drag the slider up to see the expected position of the aircraft in future time on the PVD.

12. Deselect the aircraft with a right mouse click on the TSD or PVD. The time-slider in the TSD will also reset to the initial position.

Part 2: Trajectory manipulation

Route manipulation on the PVD

1. The route of an aircraft can be modified in the PVD by adding or deleting

waypoints. Please select BMS02N in the PVD.

2. Hold CTRL to enter route manipulation mode. A waypoint symbol will be attached to the mouse cursor.

3. Hold CTRL and left click on a position inside the Travel Space to insert an intermediate waypoint into the trajectory of the selected aircraft.

4. You can see that the route has been split-up into two segments and the aircraft route passes through the newly created waypoint. The two new segments will have an equal speed (check that by the speed indication label under the waypoints).

5. Observe in the TSD that the sector exit time of the aircraft has not changed (the label and cyan star coincide).

6. Also notice that the new waypoint is visible in the TSD, and that, as for the Travel Space, the speed/time constraints have been split over the two segments.

7. Delete the waypoint by pressing CTRL and right clicking on it when it is highlighted.

Route manipulation on the TSD

8. Waypoints can also be added, manipulated and deleted on the TSD. Press and hold CTRL when the mouse cursor is in the TSD. A waypoint now appears attached to the time-space line of the aircraft.

9. Holding CTRL, left click somewhere on the time-space line of BMS02N to insert a waypoint into the trajectory of that aircraft.

10. Move the mouse over the new waypoint in the TSD to highlight it and left click and drag to change the planned arrival time at that waypoint.

11. Note that you can manipulate the arrival time of the aircraft at both the intermediate waypoint and at the sector exit point. Also notice how the Travel Space on the PVD is directly influenced by speeding up or slowing down the aircraft. In general, the area of the Travel Space will increase when the aircraft is delayed.

12. Delete the waypoint in the TSD by pressing CTRL and right clicking on it when it is highlighted.

13. So far you have only modified the 'probe' trajectory of the aircraft. Any changes made here have not been sent to the aircraft and the aircraft would continue to fly along its original trajectory if the simulator was running.

14. Manipulate the route of BMS02N (add a waypoint and/or change the timing at that waypoint) and press ENTER to send it to the aircraft. You will notice a message at the top left corner of the PVD that confirms that the trajectory of the selected aircraft has been updated. Manipulated aircraft are shown in a brighter shade of green. You can see this after the aircraft is deselected.

15. Deselect the aircraft by right-clicking on the PVD or TSD. Deselecting an aircraft will also cause any changes made to the probe trajectory to be reset. You can use this cancel the probe trajectory.

Part 3: Dynamic traffic

1. When the simulator is running, select the aircraft and observe how it maneuvers along the updated trajectory in the PVD. Also observe how the time-axis moves down as time progresses in the TSD. In accordance, you can see that the along-track distance of the aircraft to the exit point will decrease. Practice adding, manipulating, deleting and sending updated trajectories for BMS02N.

2. Every 2nd minute a workload rating scale will appear on the left side of the PVD. Please indicate your experienced workload at that time (0 to 100, low to high) by clicking in this scale.

3. Press the fast forward on the right top corner of the simulator above the TSD. The simulator will start running at 4x speed (fast-time). You may continue to the next scenario when you feel comfortable with manipulating the route of the aircraft.

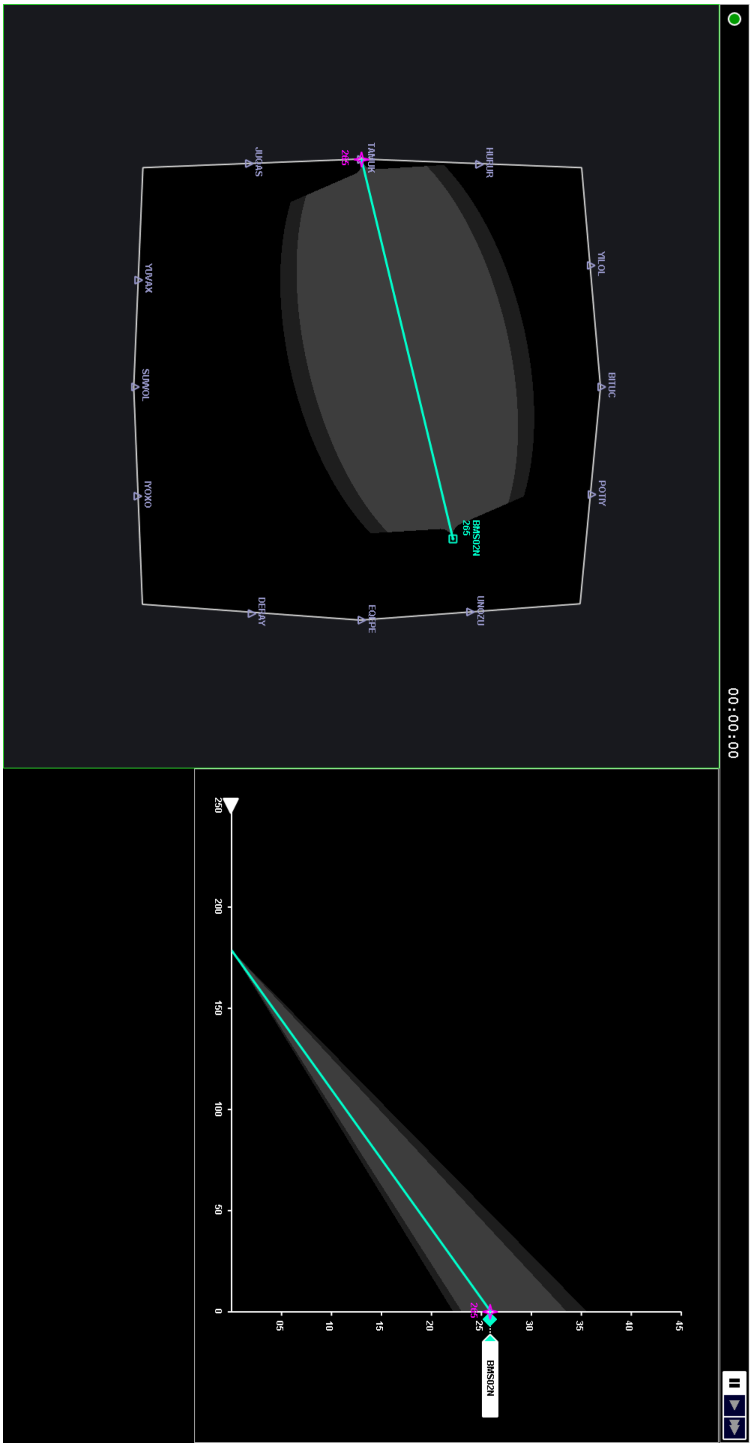

Figure B.1: Training scenario 1 at initialization.

B.1.2 Training Scenario 2

Part 1: Conflicts on the PVD

1. A red colored aircraft symbol indicates that that aircraft is expected to have a loss of separation at some point in the future. A loss of separation occurs when the lateral separation of the aircraft is less than 5 nautical miles with respect to other traffic.

2. Use the time slider in the TSD to investigate where and when the loss of separation will occur (do not yet select an aircraft). The circles around the projected aircraft positions have a radius of 2.5 nautical miles, hence, a loss of separation occurs when these circles overlap.

3. Select one of the aircraft on the PVD.

4. Notice the red and yellow part of the trajectory of the selected aircraft (not in the Travel Space, but along the trajectory line itself). The red section indicates the location of the projected loss of separation for that aircraft. The yellow portion of the line indicates where the aircraft will have a separation of between 5 and 6 nautical miles (close proximity to a loss of separation).

5. Also notice that a large red zone is present in the Travel Space of the aircraft. The red zone—or restricted field of travel shows all the locations that are unsafe to place a waypoint in. When a waypoint is placed somewhere in the restricted zone, the new trajectory will lead to a conflict with other traffic.

6. The dark gray boundary around the restricted field of travel indicates that if a waypoint is placed in that area, the new trajectory will be in close proximity to a loss of separation. Note that in the case that an aircraft has a close proximity to a loss of separation, its symbol will become yellow instead of red.

7. Hover over the restricted field of travel in the Travel Space on the PVD with the mouse to highlight the aircraft that causes this zone. Left click on the highlighted zone to select the other aircraft. You can see how the Travel Space of both aircraft is affected by the other aircraft.

8. Add a waypoint somewhere in the restricted field of travel for the selected aircraft and check with the time slider in the TSD that the conflict has indeed not been resolved.

9. If possible, add a new waypoint in the gray area of the travel space to resolve the conflict. If this is not possible, delete the other created waypoint first. Check the validity of this conflict resolution with the time slider in the TSD.

10. Please delete all newly created waypoints for both aircraft before continuing to the next part.

Part 2: Conflicts on the TSD

1. Select one of the aircraft on the PVD

2. Notice the restricted field of travel in the TSD. This restricted area represents the locations in time and distance to go for the selected aircraft that are occupied by other traffic. A conflict will occur if the time-space trajectory of the aircraft passes through such a zone.

3. Similar to the Travel Space, the dark gray boundary around the restricted field of travel indicates that if the time-space trajectory passes through this area, the trajectory will be in close proximity to a loss of separation.

4. Hover over the restricted field of travel in the TSD with the mouse to highlight the aircraft that causes this zone. Left click on the highlighted zone to select the other aircraft. You can also see here how the Travel Space of both aircraft is affected by the other aircraft.

5. Solve the conflict by changing the arrival time at the sector exit for one of the aircraft and check the validity of this solution by using the time slider in the TSD.

6. In this scenario it is possible to solve the conflict and to let both aircraft arrive at the sector exit point at their originally planned time by adding an intermediate waypoint in the TSD. Experiment with such a solution for a given aircraft and check the solution with the time slider.

7. Please delete all newly created waypoints for both aircraft before continuing to the next part.

Part 3: Dynamic conflict resolution

1. When the simulator is running, select an aircraft and observe how the restricted fields of travel evolve in the Travel Space and TSD. Also note that the available control space becomes smaller as the aircraft close in.

2. Practice conflict resolution with the simulator running. You could, for instance, try to perform a cooperative resolution in which the conflict is resolved by giving both aircraft a small path deviation (spatial or time), rather than manipulating only one aircraft. This will reduce the relative path deviation for each individual aircraft.

3. Press the fast forward on the right top corner of the simulator above the TSD. The simulator will start running at 4x speed (fast-time). You may continue to the next scenario when you feel comfortable with manipulating the route of the aircraft.

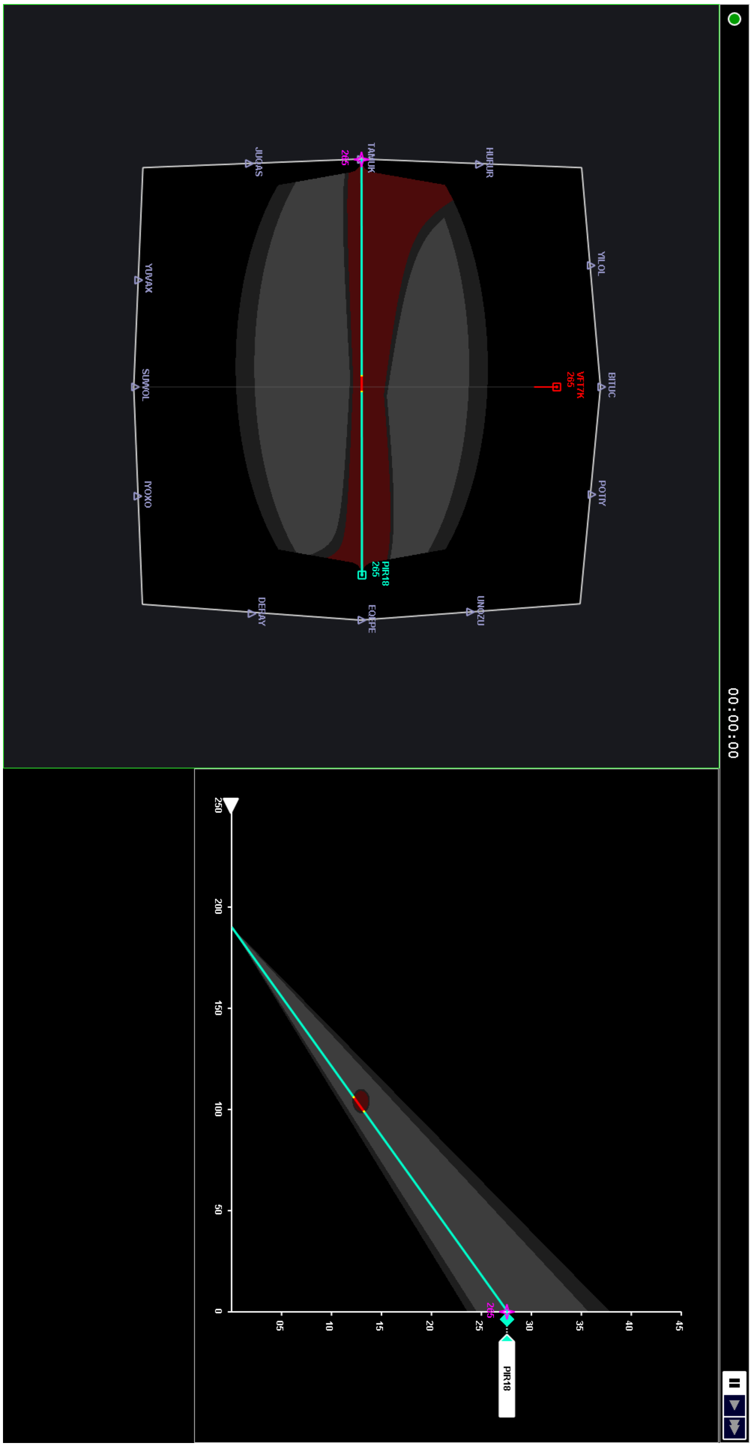

Figure B.2: Training scenario 2 at initialization.

B.1.3 Training Scenario 3

Part 1: Restricted field interpretation on the TSD

1. In this scenario there are seven controlled aircraft. Five are currently inside the sector and two aircraft (KLY80 and FRS8K) will enter from point BITUC and SUWOL respectively in the near future. None of the aircraft are in conflict with each other (no aircraft are red). You can check the predicted evolution of the traffic by using the time slider in the TSD.

2. Note that restricted fields of travel in the Travel Space and TSD are only shown for aircraft that are inside the controlled sector. The zones caused by KLY80 and FRS8K are therefore not represented yet. Aircraft that enter from outside the sector may have conflicts with other aircraft inside the sector, but will only be flagged once they enter the sector.

3. Please select the aircraft VFT7K.

4. The shape and location of the restricted fields of travel in the TSD provides additional information about the crossing geometry and relative movements of the traffic.

5. All restricted zones that lie under the time-space trajectory in the TSD represent aircraft that will pass in front of the selected aircraft. How many aircraft will pass in front? Check this by hovering over the restricted fields to find out which aircraft they belong to, and by using the time slider.

6. All restricted zones that lie above the time-space trajectory represent aircraft that will pass behind the selected aircraft. How many aircraft will pass behind? Check this by hovering over the restricted fields to find out which aircraft they belong to, and by using the time slider.

7. The location of a restricted zone along the x-axis of the TSD indicates where along the trajectory the other aircraft will pass. You can see that PLX9Z and BRW29 will cross at around the 100 nautical mile to go mark. PMG5L will pass at a further point along the trajectory at around the 50 mile mark.

8. An in-trail aircraft (PIR18) is indicated by a restricted field along the entire trajectory of the selected aircraft. In this case the restricted zone is above the time-space line indicating that the aircraft is in-trail and behind VFT7K. Note that delaying the selected aircraft at the sector exit point will cause an in-trail conflict (overtake) with PIR18.

9. Note that in the TSD the labels are also shown of all other aircraft that have the same sector exit point as the selected aircraft. In this case you can see the label of PIR18. Additionally, by clicking on the label you can select the other aircraft.

10. Please switch between selecting aircraft VFT7K and PIR18 on the TSD. Because these aircraft fly along exactly the same trajectory, the shape and location of the restricted zones caused by the other traffic remain the same for both aircraft.

11. Select one of the in-trail aircraft (VFT7K or PIR18).

12. The shape of the restricted zones in the TSD also provides information about the crossing angle of the other traffic. A pure 90 degree crossing (PMG5L) will show up as a circular restricted zone. Check this in the TSD.

13. A shallow crossing (BRW29) will result in a forward slanted ellipse-shaped restricted zone. Check this in the TSD.

14. A head-on crossing (PLX9Z) will look like a backward slanted ellipse-shaped restricted zone. Check this in the TSD.

15. As a direct result of the crossing geometry and the shape of the restricted zone, can you reason whether a head-on conflict or a shallow conflict is harder to resolve with a speed change alone? As a hint: imagine that the ellipses of BRW29 and PLX9Z are located on the time-space line of the selected aircraft.

16. Please select the aircraft BRW29.

17. By only looking at the conflict zones in the TSD, can you reason how many aircraft will pass in front of the selected aircraft, how many will pass behind, how many crossing points there are along the trajectory, and where and what the passing geometry looks like (shallow or head-on crossing)?

18. When KLY80 and FRS8K enter the sector a head-on conflict will occur, could you reason what the shape of a head-on restricted zone would look like in the TSD?

Part 2: Dynamic restricted fields

1. In this scenario you are free to manipulate the trajectories of the aircraft and see what the influence is on the conflict zones in the TSD.

2. Press the fast forward on the right top corner of the simulator above the TSD. The simulator will start running at 4x speed (fast-time). You may continue to the next scenario when you feel comfortable with manipulating the route of the aircraft.

Part 3: Exercise

The following figure shows a schematic representation of a sector and three TSDs. Each TSD shows the time-space line and restricted zones for a specific aircraft in the sector (1, 2, 3, or 4). Which TSD belongs to which aircraft?

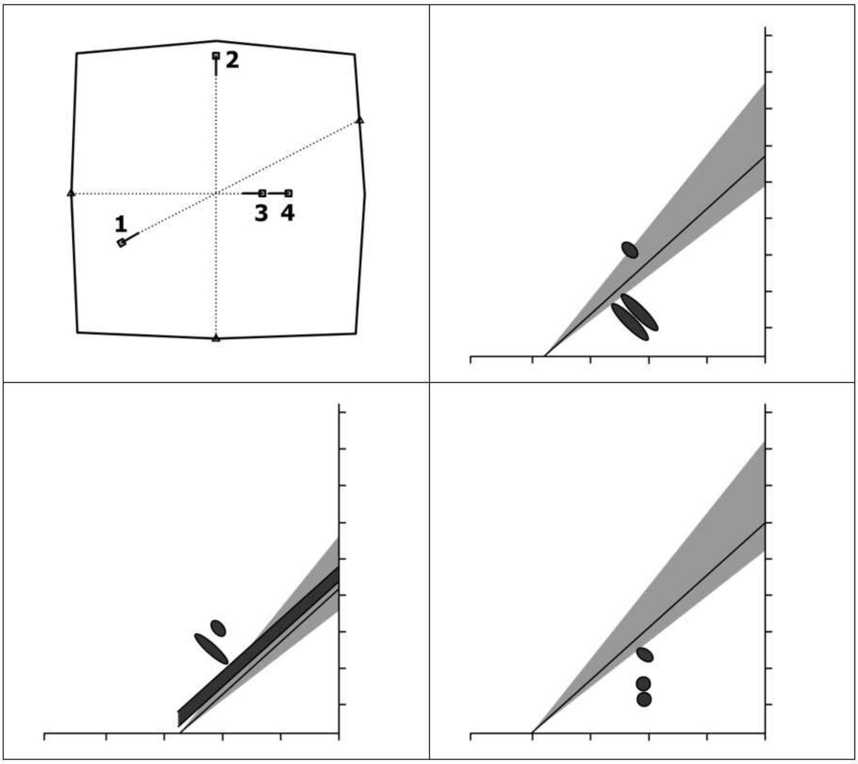

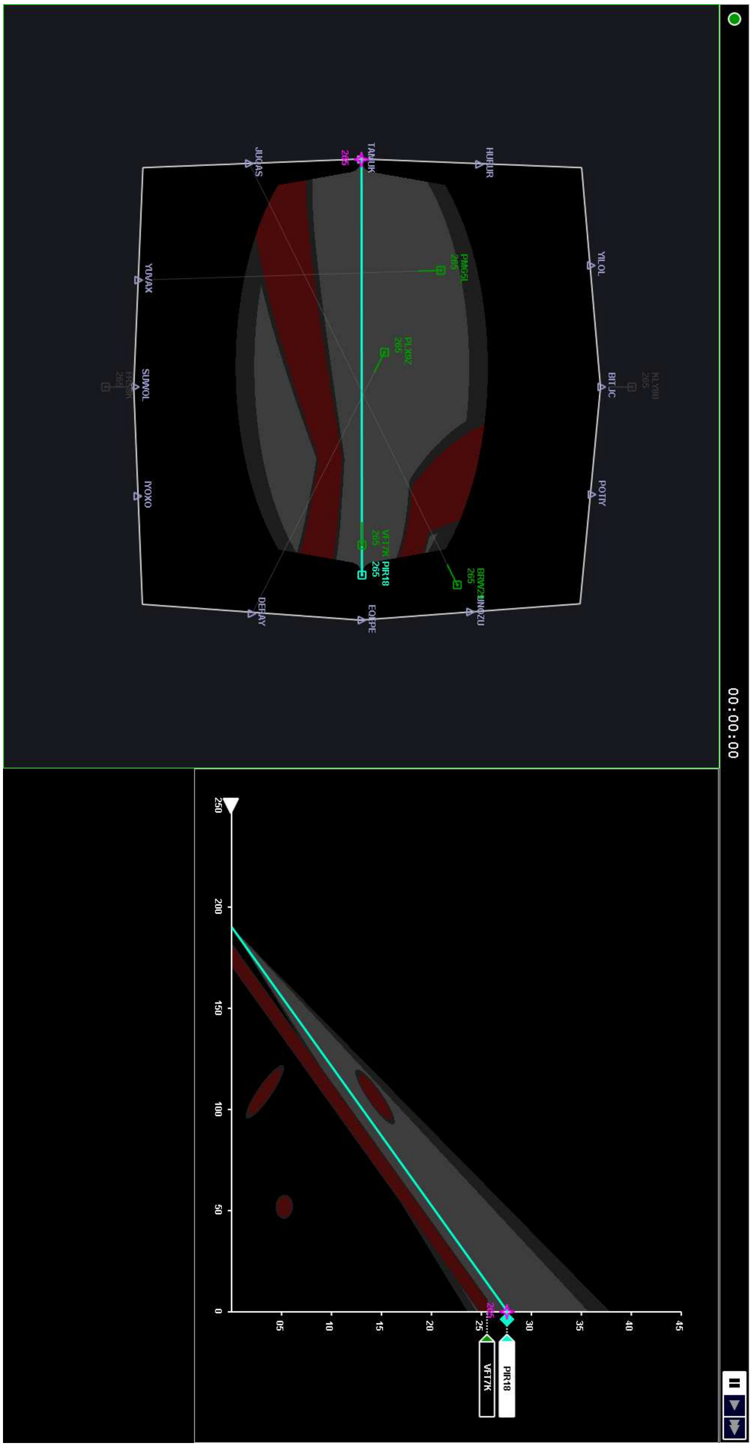

Figure B.3: Training scenario 3 at initialization.

B.1.4 Training Scenario 4

Part 1: Restricted airspace and traffic flows

1. This scenario features a restricted airspace in the middle of the controlled sector. The restricted airspace represents hazardous weather or a no-fly zone and should be avoided by all aircraft. There is no direct indication that an aircraft will fly through the zone (i.e., color change), you should monitor that by looking at the trajectories of the aircraft.

2. Currently none of the active aircraft is planned to fly through the restricted zone. However, there is one conflict between PIR18 and FTM6R.

3. Please select aircraft PIR18.

4. Add a waypoint in the Travel Space on the trajectory of LWB54 and HTZ78, and on a safe location (i.e., merge the routes of the selected aircraft and the two other southbound aircraft). Press enter to send the new trajectory to the aircraft.

5. Investigate the TSD of PIR18 to see how the restricted areas in the TSD of a trajectory merge look. It can be seen that from the added waypoint to the sector exit point, the two other aircraft will fly in-trail and behind.

6. Also investigate the TSD of HTZ78 and LWB54 to see how the restricted area of an aircraft looks that merges on the trajectory of a selected aircraft.

Part 2: Basic dynamic scenario

1. In this scenario you are free to manipulate the trajectories of the aircraft to resolve any further conflicts or restricted airspace crossings.

2. Press the fast forward on the right top corner of the simulator above the TSD. The simulator will start running at 4x speed (fast-time). You may continue to the next scenario when all incoming aircraft have entered the sector, are conflict free, and will fly around the restricted airspace.

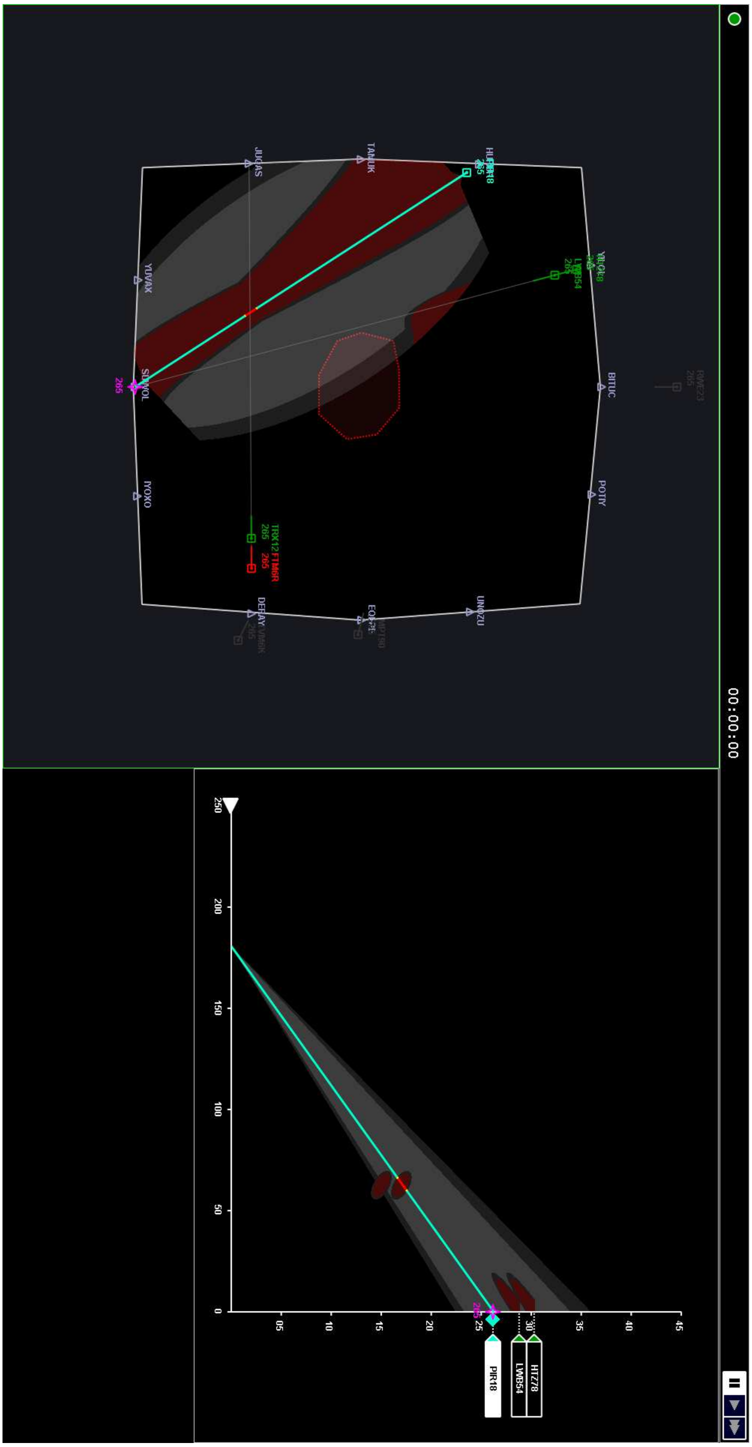

Figure B.4: Training scenario 4 at initialization.

B.1.5 Practice scenarios

1. In the previous scenarios you have been shown all the tools and features that are available to you to in the experiment simulator. The following training scenarios are intended as practice, to increase your experience, and to make you feel comfortable with performing your task in this experiment.

2. In each scenario you are free to manipulate the trajectories of the aircraft to resolve any further conflicts or restricted airspace crossings during the remainder of the scenario. Try to minimize any trajectory deviations and delays at the sector exit point.

3. The difficulty will slightly increase in each subsequent practice scenario.

4. You may continue to the next scenario when all incoming aircraft have entered the sector, are conflict free, and will fly around the restricted airspace.

5. The difficulty of the scenarios in the main experiment will be at more or less the same level as these practice scenarios.

6. At the start of each scenario, press the fast forward on the right top corner of the simulator above the TSD. The simulator will start running at 4x speed (fast-time).

ACKNOWLEDGEMENTS

THE trajectory of my PhD journey has finally reached its final destination after having stayed aloft for over more than a decade. This shows that uncertainty and variability does not only affect the execution of planned trajectories in the air traffic control domain alone. Although having been a long flight, it was an invaluable and memorable experience that I would not have wanted to miss. Never did I have the feeling of flying alone and therefore would like to start by thanking everyone who—even in the smallest ways possible—contributed to the successful completion of this thesis.

In particular I would like to thank my promotors, Max and Clark, without whose guidance and unrelenting patience I could not have achieved this result. Max, it has been an absolute pleasure to work with you. From the beginning of my MSc thesis in 2008 and throughout my PhD research you always managed to free up time, and never failed to provide valuable input and new insights. Clark, I really appreciated our—frequently daily—interactions and cooperation, especially during many of the most intensive parts of the process (conferences, paper deadlines, experiment preparations, etc.). Oftentimes you were the motor that kept me going. To both, your unmatched sense of humor and contagious laughter, be it about supersized burritos, attempts to visit the USAF museum, or just random everyday matters is something that I will never forget.

Many thanks to all current and former staff at the Control and Simulation (C&S) department that I was blessed to cross paths with (which are many). I am bound to forget to mention a name or two, my apologies in advance, but here goes nothing. Rene, thank you for your valuable guidance and direction in the initial phase of this journey. Maarten, my long-time room mate and fellow department dinosaur that beat me to the finish line only by a few months. My fellow PhD researchers (in no specific order): Liguo, Deniz, Dyah, Laurens, Wouter, Jazdi, Mariam, Gustavo, Sophie, Tommaso, Jaime, Jan Comans, Jan Smisek, Sherry, Tao, Kasper, Ivan, Kimberly, Kirk, Dirk, Sarah, Diana, Jelmer, Isabel, Ezgi, Julia, Matej, Ewoud, Annemarie, Jerom, Daniel, Carl, Ivan, Joao, Sjoerd, Emmanuel, Gijs, Rowenna, Wei, Ye Zhou, Ye Zhang. Many thanks to Gesa, who co-authored our 2014 paper that won the best student paper award, it was a pleasure to collaborate with you! The staff members: Daan, Erik-Jan, Coen, Jacco, Joost, Junzi, Olaf, Andries, Harold, Ferdinand, Menno, Alwin, Bart, Guido, Cristophe, Marilena and Bob. Hans and Xander, many thanks for allowing me participate in, and contribute to one of the most exiting practical classrooms that exists. Thanks to the MSc students that I was honored to supervise and assist: Jorge, Rick, Radesh, Douwe, Jason and Matthijs. Matthijs, many thanks for our great discussions and for your current efforts in progressing this research to the next dimension (pun intended). All the best for your own PhD journey! Summarizing, I am very grateful for the tight-knit community and friendly and open atmosphere at C&S that presented me with many amazing experiences. From all the interesting conversations (both formal and informal) during coffee and lunch breaks to attending conferences in far-away places and floating in zero-G, it was truly an unforgettable journey!

Next, I would like to extend my thanks to my current employer, MPS, who since 2016 has always been proactive and flexible in providing me with the time and resources needed to work on my dissertation. Additionally, working with professional flight simulators has given me a lot of new insights about the operational side of air traffic control as perceived by the pilot. Most importantly, I would like to thank all my direct colleagues for their support, curiosity, optimism, and for continuously enquiring how my PhD was progressing throughout the past years.

I cannot forgo to mention my closest friends. Ewout and Bastiaan, to be honest, I really cannot thank you for contributing to this thesis. On the contrary, you always provided me with ample excuses and distractions not to work on it. But that's perfectly fine! Our countless shenanigans since high school, both big and small, have kept me social and sane. To all my other friends (you know who you are), I count myself blessed to have such a diverse group of wonderful individuals around me. You make life more interesting, thank you!

It goes without saying that I owe many thanks to my parents, Ulfert and Claudia, and my brother Alfred who never stopped encouraging me to "just finish the bloody thing already!". All jokes aside, thank you for always believing in me, even at times when I lost belief in myself. I could not have achieved this without the solid foundation that you have always provided me with. This result is as much yours as it is mine.

Finally, I would profoundly like to thank Massi for entering into my life in the final phases of writing the thesis. This book would not have been completed without your unfading patience, support, and love that empowered me to keep pressing on, even during the most difficult moments. We are proof that two different cultures make for a strong team together!

Rolf Klomp **The Hague, March 18th 2023**

For those whom seek to decode the binary (ARINC-429 BNR-coded) data on the cover: MSB: 28, LSB: 9, scaling: $180 \cdot 0.5^{20}$.

Curriculum Vitae

Rolf Edwin KLOMP

14-01-1984	Born in Voorburg, The Netherlands

Education

1987-2002	Primary School: Ysgol Bryn Coch, Mold (UK) Primary School: 't Klankbord, Leiden (NL) Secondary School: Adelbert College, Wassenaar (NL)
2002-2006	Bachelor of Science in Aerospace Engineering Delft University of Technology, Delft (NL)
2006-2010	Master of Science in Aerospace Engineering Delft University of Technology, Delft (NL)
2011-2023	Doctor of Philosophy in Aerospace Engineering Delft University of Technology, Delft (NL) Thesis: Solution Space Concept for Trajectory-Based Air Traffic Control Promotor: Prof. Dr. Ir. M. Mulder Co-promotor: Dr. Ir. C. Borst

Work Experience

2010-2017	Freelance engineering software consultancy Avisys, Den Haag (NL)
2016-2023	Senior simulator software engineer Multi Pilot Simulations, Groenekan (NL)

Awards

2014	IEEE Systems, Man, and Cybernetics Society Conference 2014 Best Student Paper Award, San Diego, CA (USA)

Rolf Edwin Klomp was born in Voorburg (NL) on January 14th, 1984. He spent his early years growing up in Zoetermeer (NL), and Mold (UK). He obtained his Atheneum diploma at Het Adelbert College in Wassenaar in 2002 with electives focused on mathematics, physics, and biology. Because of his enthusiasm about everything related to aviation from an early age, enrolling at the faculty of Aerospace Engineering at the Delft University of Technology (DUT) was a natural choice.

He completed his B.Sc. degree in Aerospace engineering in 2006. His interest in computer science and simulation was a primary motivation to pursue an M.Sc. degree at the department of Control and Simulation (C&S). During this period, he completed a seven month internship at the EUROCONTROL headquarters in Brussels, Belgium, working on enhancing digital aircraft performance models for air traffic control simulations. He obtained his M.Sc. degree in 2010 after a ten month collaborative project between DUT and Air Traffic Control the Netherlands (LVNL), during which his thesis focused on the development of ecological air traffic controller decision-support interfaces for inbound air traffic control.

After his graduation, he started working as a freelance software consultant on projects in various fields of engineering. But after one and half years, the calling to come back and—in parallel—continue his academic career led to the start of his journey towards obtaining a Ph.D. degree. The initial part of his work was performed within SESAR Work Package E project C-SHARE, a collaboration between DUT, the Netherlands Aerospace Center (NLR), and Thales. The project focused on enabling effective joint human-automation collaboration in trajectory-based air traffic control. In 2014 he won the Best Student Paper Award at the IEEE Systems, Man, and Cybernetics Society Conference in San Diego, CA (USA). During this time he also served as a teaching assistant for various practicals. Most noteworthy, the flying classroom practical for which he developed Android-based apps for hand-held tablets that visualize live aircraft parameters to the students in flight.

In 2016, Rolf started working as a simulation software engineer for Multi Pilot Simulations (MPS), an Utrecht (NL) based company that develops and produces professional flight simulators for flight training organizations and airlines. He is a software systems architect and a subject matter expert in, amongst others, the Flight Management System and core simulation framework. In this role he frequently supports on-site hand-overs of new simulators to customers, and certifications of new simulators together with civil aviation authorities.

List of Publications

[1] R. E. Klomp, M. Mulder, M. M. Van Paassen, and M. J. Roerdink, "Redesign of an Inbound Planning Interface for Air Traffic Control," in *Proceedings of the AIAA Guidance, Navigation, and Control (GNC) Conference and Exhibit*, American Institute of Aeronautics and Astronautics, Aug. 2011.

[2] R. E. Klomp, M. M. Van Paassen, M. M., and M. J. Roerdink, "Air Traffic Control Interface for Creating 4D Inbound Trajectories," in *Proceedings of the 16th International Symposium of Aviation Psychology*, Dayton, OH, USA, 2011, pp. 263–268.

[3] R. E. Klomp, M. M. Van Paassen, C. Borst, and M. Mulder, "Joint Human-Automation Cognition through a Shared Representation of 4D Trajectory Management," in *Proceedings of the 2nd SESAR Innovation Days*, D. Schaefer, Ed., Braunschweig, Germany: EUROCONTROL, 2012, pp. 1–7.

[4] M. M. Van Paassen, C. Borst, R. E. Klomp, M. Mulder, P. van Leeuwen, and M. Mooij, "Designing for Shared Cognition in Air Traffic Management," *Journal of Aerospace Operations*, vol. 2, pp. 39–51, 2013.

[5] R. E. Klomp, C. Borst, M. Mulder, G. Praetorius, M. Mooij, and D. Nieuwenhuisen, "Experimental Evaluation of a Joint Cognitive System for 4D Trajectory Management," in *Proceedings of the 3rd SESAR Innovation Days*, D. Schaefer, Ed., Stockholm, Sweden: EUROCONTROL, 2013, pp. 1–7.

[6] R. E. Klomp, C. Borst, M. M. Van Paassen, M. Mulder, D. Nieuwenhuisen, A. Maij, M. Mooij, and A. van Drunen, "Designing for Joint Human-Automation Cognition through a Shared Representation of 4D Trajectory Management," in *Proceedings of 17th International Symposium of Aviation Psychology*, Dayton, OH, USA, 2013, pp. 518–523.

[7] A. Maj, A. Van Drunen, T. J. J. Bos, and R. E. Klomp, "Human Factors Testing Results of a Joint Cognitive System for 4D Trajectory Management," National Aerospace Laboratory NLR, division Air Transport, Amsterdam, The Netherlands, Tech. Rep. NLR-TP-2013-272, 2014.

[8] R. E. Klomp, C. Borst, M. Mulder, and G. Praetorius, "Ecological Interface Design: Control Space Robustness in Future Trajectory-based Air Traffic Control Decision Support," in *Proceedings of the 2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, San Diego, CA, USA: IEEE, 2014, pp. 329–334.

[9] R. E. Klomp, C. Borst, M. M. Van Paassen, and M. Mulder, "Expertise Level, Control Strategies, and Robustness in Future Air Traffic Control Decision Aiding," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 2, pp. 255–266, 2015.

[10] J. Pinto, R. E. Klomp, C. Borst, M. M. Van Paassen, and M. Mulder, "Design of an Ecological Flow-based Interface for 4D Trajectory Management in Air Traffic Control," in *Proceedings of the 5th CEAS Air and Space Conference*, Delft, The Netherlands, 2015, pp. 1–10.

[11] R. E. Klomp, R. Riegman, C. Borst, M. Mulder, and M. M. Van Paassen, "Solution Space Concept: Human-machine Interface for 4D Trajectory Management," in *Proceedings of the 13th USA/Europe Air Traffic Management Research and Development Seminar*, Vienna, Austria, 2019.

[12] R. Nagaraj, R. E. Klomp, C. Borst, M. M. Van Paassen, and M. Mulder, "Exploring the Potential Benefits of Multi-aircraft Trajectory Manipulation in Future Air Traffic Control," in *Proceedings of the 2019 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, Bari, Italy: IEEE, 2019, pp. 3699–3704.

[13] D. S. A. Ten Brink, R. E. Klomp, C. Borst, M. M. Van Paassen, and M. Mulder, "Flow-based Air Traffic Control: Human-machine Interface for Steering a Path-planning Algorithm," in *Proceedings of the 2019 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, Bari, Italy: IEEE, 2019, pp. 3186–3191.

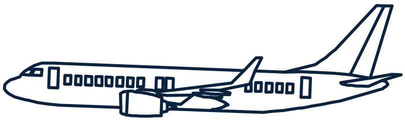

0110 0000 0110 0011
1001 0010 1100 1001

TUDelft

ISBN: 978-94-6473-071-5