Exploring Human-Automation Interactions in Next-Generation Airport Surface Movement Operations

an Agent-based Modelling Approach

Malte von der Burg & Alexei Sharpanskykh
Air Transport Operations, Faculty of Aerospace Engineering
Delft University of Technology, The Netherlands
M.F.vonderBurg@tudelft.nl, O.A.Sharpanskykh@tudelft.nl

***Abstract*—Both EASA and SESAR JU have outlined roadmaps towards an autonomous air traffic management system. Their long-term vision aligns with our previous work in which we studied the operational consequences of fully automated airport surface movement operations (ASM Ops) using a multi-agent system model that conducts multi-agent motion planning. However, such an automated system needs to be placed in a sociotechnical context, in which human operators and system users would be able to interact with automation in a seamless way. In the context of the SESAR JU project "ASTAIR - Auto-Steer Taxi at Airport", workshops and interviews were conducted with human operators to identify operational requirements and preferred interactions with automated systems. In this paper, we explore how the requirements can be modelled and implemented in the multi-agent system for automated planning of ASM Ops to enable interactions with human operators. To this end, we present a conceptual agent-based framework for human-automation teaming in ASM Ops, introduce interactive tools for ATCOs to engage with the system, and showcase these interactions through a series of use cases derived from the workshops as well as historic data. Illustrated through the interactions of a fictitious Air Traffic Controller at Amsterdam Airport Schiphol, we demonstrate their practical implementation and broader aspects of such interactions in next-generation ASM Ops.**

***Keywords*—human-automation teaming; multi-agent system; agent-based modelling; airport surface movement operations; air traffic control**

I. Introduction

The air transport industry faces the dual challenge of rising passenger demand and the requirement to reach net-zero emissions by 2050 [1]. At large airports that play a vital role in the air transport system, these challenges are amplified as infrastructural expansions are insufficient to facilitate the predicted growth [2]. Rising airport congestion may reduce the predictability of an aircraft's taxi time, i.e. the duration an aircraft spends travelling over the airport surface between the runway and its stand. This may impact not only the individual flight but may also cause knock-on effects on the network [3]. Moreover, the workload for Air Traffic Control Officers (ATCOs) may increase [4], which can in turn reduce the operational efficiency.

Increasing the level of automation is seen as one viable solution: towards 2050 and beyond, both EASA and SESAR JU define a vision and roadmap to eventually reach autonomous air traffic management, i.e. level 3 in EASA's AI roadmap [5], or level 4-5 in the roadmap of SESAR JU [6]. Previous SESAR projects investigated how to increase the efficiency and predictability of taxiing operations through automation [7], [8]. Furthermore, aircraft manufacturers are developing technology for allowing airplanes to taxi autonomously [9].

In previous work, we created a detailed, computational multi-agent system (MAS) model for autonomous airport surface movement operations (ASM Ops). We implemented and tested this model and its underlying path planning algorithm in Python [10]. Furthermore, we assessed the operational consequences for different concepts of fully automated ASM Ops [11], [12]. However, in real-world implementations of automated systems, humans will nonetheless play a crucial role: the automation needs to be maintained and extended by developers, and human operators must oversee that it continues to function as expected and intervene if necessary. Thus, human operators and automated systems must interact with each other seamlessly and effectively, forming a human-automation team (HAT).

In comparison to current ASM Ops, the tasks of human operators will change: the horizon of decision-making moves towards the tactical and strategic levels, as the operational level is managed foremost by automation, like sending 4D-trajectories and clearances automatically to (auto-)pilots. The interaction must be designed to reduce the cognitive load on human operators by automating routine tasks and providing decision support. However, it must also ensure that operators remain engaged and informed to prevent skill degradation and ensure quick, effective intervention when necessary [13].

The SESAR JU project "ASTAIR - Auto-Steer Taxi at Airport" aims to develop interactive tools and adaptive AI algorithms to enhance the efficiency, safety, and sustainability of ASM Ops [14]. In the project's context, workshops and interviews with human operators from three different airports (IATA: AMS, CDG, FRA; ICAO: EHAM, LFPG, EDDF) were carried out to determine requirements and desirable interactions between humans and automated systems within ASM Ops. Despite the different characteristics of the taxiway networks of the respective airports, the experts recognized the benefits that higher levels of automation and human-automation teaming would bring for their operations, some of which being:

- to increase the shared situational awareness, e.g. by providing better estimates of the engine-start times, or by keeping track of as well as managing clearances,

- to optimize e.g. the selection of a suitable holding point, the time point to cross an active runway, or the pushback procedures,
- to improve safety, e.g. by prohibiting accidents between service vehicles and aircraft at intersections.

Based on the discussions, regular occurring and characteristic situations were synthesized into use cases that are further studied in the project [15].

In this paper, we explore how the findings from these stakeholder workshops and interviews can be embedded into the automated multi-agent system and its path planning algorithms, with a focus on its technical implementation. To this end, we first outline the conceptual agent-based framework for the human-automation teaming, and describe the related interactive tools with which human operators can examine and manipulate the planned paths (Section II). The MAS model from our previous work forms the automation-side of the teaming. In Section III, we then present different exemplary interactions between ATCOs and the MAS, motivated by the workshop results. After introducing each case, we describe how the interaction was modelled and implemented in the multi-agent system context, and discuss its broader application within ASM Ops. We end by summarizing the directions for future work and concluding remarks.

II. Agent-based Model of Human-Automation Teaming

In this section, we outline the schematic agent-based framework of human-automation teaming for ASM Ops, which is visualized in Fig. 1. As main operator roles within ASM Ops, the **ATCO** and the **Pilot** are shown on the left-hand side in the figure. These human operators communicate with others, e.g. ground personnel and airline staff, to maintain shared situation awareness and to make time-critical decisions. While the ATCOs act as supervisors in the system, the pilots carry out the commands provided by the **MAS** from our previous work on fully automated operations [10]. To this end, the MAS automatically receives the position and speed of all surface movements, and sends the planned, conflict-free 4D-trajectories to the flight decks. To carry out the path planning, the MAS requires data input such as the flight schedule (FS) and runway configuration (RMO) from the ATCOs, and provides informative data like the estimated taxi time to them. To separate the MAS-internal functionalities from the interactive elements, we introduce the **Interaction Agent** that handles the interaction between human operators and the MAS. We further outline the roles and tasks as well as the interactive tools in the following.

A. Role of the Human Operators

In general, *human operator* refers to any person acting within the ASM Ops, such as air traffic control officers (ATCOs), pilots, airline staff, or ground personnel. While human operators interact among each other, they can obtain relevant information from the MAS model, and interact with the Interaction Agent

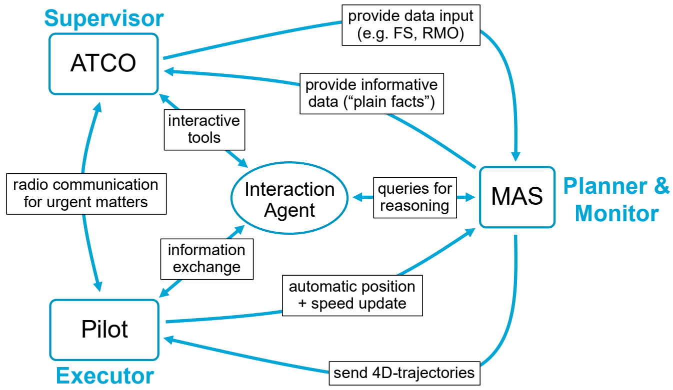

Figure 1. Schematic agent-based framework of human-automation teaming for airport surface movement operations. FS: flight schedule, RMO: runway mode of operation

to get an overview of the upcoming traffic situation, update data inputs and predictions, or seek decision-support. The ATCOs interact most extensively and with varying degree of complexity with the Interaction Agent to fulfil their responsibility of ensuring safe and efficient ASM Ops. In this paper, we thus focus mainly on their interactions with the automated side of the HAT.

In contrast, the (auto-)pilots primarily carry out the commands given by the MAS. While pilots will likely remain in control of the aircraft during taxiing in the mid-term, auto-pilot system may eventually take over the nominal control of the aircraft to follow the 4D-trajectories, which the MAS sends via an appropriate datalink system to the flight decks. In general, the pilots can obtain any relevant information by interacting with the Interaction Agent. Likewise, they can update aircraft-related data like the required engine-start time, which the MAS then accounts for when planning the aircraft's trajectory.

The pilots shall contact the ATCOs via radio for urgent matters, and can also pose non-critical requests. Allowing for such direct communication between ATCOs and pilots within the HAT does not only ensure quick response times in urgent situations, but was also seen as providing job satisfaction in interviews with operational experts. In the long term, the non-critical human-to-human communication could shift further towards human-automation interaction, given that the automation-side of the teaming will have gained the required additional capabilities to process such context-specific requests.

B. Role of the MAS

Based on the A-CDM milestones as well as the data provided by the ATCOs, the MAS calculates conflict-free 4D-trajectories for all surface movements using the Routing Algorithm outlined in [10]. The multi-agent motion planning algorithm sets priorities between the aircraft to deconflict their concurrent routes, and accounts for their shapes and kinematic properties. The 4D-trajectories are sent to the flight decks, and are automatically updated if necessary. Based on the automatic position and speed updates of all surface movements, the MAS also monitors that the planned trajectories are executed as instructed. When deviations occur, the MAS adjusts the routes where possible to minimize the impact, and alternatively replans the routes of the affected aircraft.

In the non-nominal case that it cannot find a viable solution, it requests the ATCOs to step in and resolve the situation, for which the MAS provides any decision-support needed. Besides the regular path planning according to the flight schedule, the MAS can also provide decision-support by calculating paths between specific points in the taxiway layout of a single or multiple aircraft. As such, it can be queried to obtain alternative routing options for one aircraft, or to test the effects that pending decisions would have on the overall traffic.

C. Interaction Agent and its Interaction Tools

The Interaction Agent acts as intermediary between the human operators and the MAS. To this end, it provides interfaces as well as interactive tools. We assume that the Interaction Agent as well as the MAS are working reliably so that human operators have sufficient trust in them. Furthermore, we assume that the human operators have sufficient training to use the interactive tools effectively. The relevant tools used in the interaction examples in Section III are outlined in the following.

1) Updating Flight Schedule: The start and goal locations as well as the start times of aircraft are obtained from the flight schedule (FS) provided through the A-CDM milestones and inputs from human operators. When updated schedule information is available, the ATCOs can update the FS entries accordingly.

2) Adjusting Activity Sequence: Per aircraft, the MAS creates an activity sequence to account for the different operations, e.g. following a specific path for pushback, travelling to and holding at a holding point, etc. The ATCOs can adjust this sequence and its elements: they can for instance change the pushback path that must be followed, or the required holding duration. Likewise, data such as the estimated engine start-up duration is stored within the activity sequence and may be altered by human operators.

3) Setting High-Level Parameters and Constraints: The ATCOs can set and adjust various high-level parameters. These affect the path planning of all, a group of, or certain aircraft, and are valid for an extended period of time. Examples are:

- adjusting the conformance level to ATC procedures, further discussed in Section III-A,
- blocking certain taxiway segments for maintenance work,
- adjusting the speed limits in general or in certain areas such as bay areas e.g. due to adverse weather conditions,
- setting general priority levels between aircraft groups e.g. arriving and departing aircraft.

Moreover, ATCOs can set constraints for specific aircraft that have a direct impact on their route. While some are valid for the entire taxiing, others are timed, i.e. issued for a specific duration. Examples are:

- constraining the start or goal location in case an aircraft shall leave before or arrive after a certain time point,
- selecting a certain location to be passed during taxiing, affecting the aircraft's activity sequence,
- assigning a specific priority relation between two aircraft, i.e. one has right of way over the other.

4) Fast-forward simulation with/without Change-overlay: The MAS plans the 4D-trajectories ahead of time, and deconflicts all routes within a pre-defined planning window. Thus, the ATCOs can inspect how the traffic will evolve in the upcoming period through a fast-forward simulation. The aircraft are categorized into one of six ICAO-types with coloured circles indicating the associated type sizes as shown in Fig. 2 a). A cyan-coloured tug symbol indicates that a pushback-truck is coupled to the aircraft. A colour scheme is used to indicate the predicted engine status, with the aircraft colour being green when engines are switched off, orange during engine-start, and red when they are running. An exemplary visualization of these symbols and colours is provided in Fig. 2 b). As additional functionality, any aircraft path can be plotted into the visualization with the executed part shown by a darker colour and the remaining path in a lighter colour w.r.t. the colour of the circle around the aircraft whose route is displayed. As example, Fig. 2 c) shows the blue and orange trajectories. Furthermore, the differences between two routing plans can be visualized with a change-overlay: one or multiple previous/inferior route alternatives are depicted by "shadow-aircraft" that are connected with a grey line to the new/superior solution.

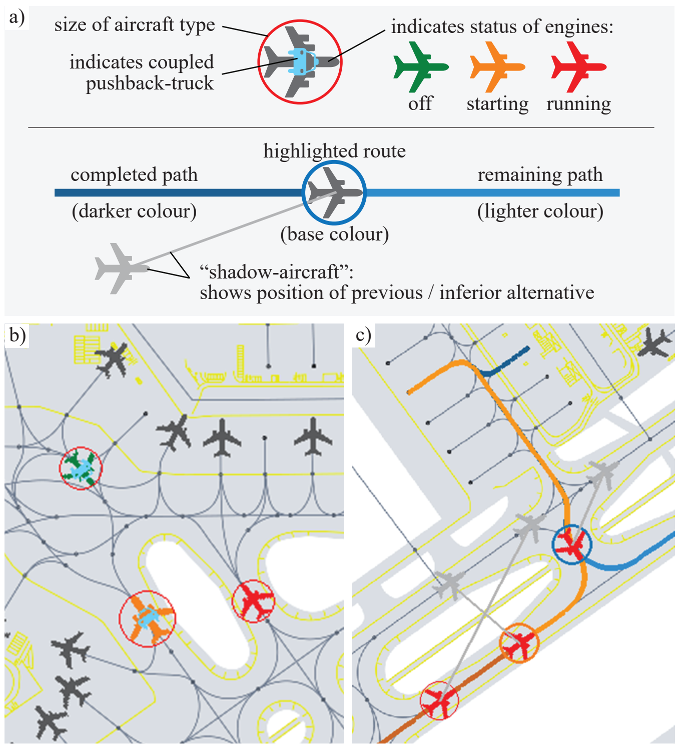

Figure 2. Illustration of fast-forward simulation with a) providing an overview of the symbols and colours, b) example without change-overlay, and c) example with change-overlay

5) Calculation of SAMP Routes: The MAS can calculate single-agent motion planning (SAMP) routes to quickly assess different routing options. These routes show the fastest possible paths that account for all relevant constraints, but are not yet coordinated with other traffic.

6) Calculation of MAMP Routes: Like the regular path planning, the MAS can also calculate alternative multi-agent motion planning (MAMP) routes based on altered input data, updated activity sequences, or changed high-level parameters or constraints like those mentioned in Section II-C3. However, to generate conflict-free trajectories for the selected set of vehicles, more computational time is needed than for calculating alternative SAMP routes.

7) Adjusting Support Level: The ATCOs can adjust the level of the notifications and decision-support provided by the Interaction Agent. With a high support level, the Interaction Agent carries out various analyses (see example in Section III-C5) and prompts the ATCOs to take action.

III. Modelling and Implementation of Human-MAS Interactions

In the following, we showcase the abilities of the MAS to engage in interactions with humans using its interaction tools. We provide a general motivation, exemplary use cases, and a technical elaboration how such interactions could be modelled and implemented in the multi-agent system context. To make the cases more illustrative, we follow the fictitious Air Traffic Controller Anne van der Waal through her shifts at Amsterdam Airport Schiphol. We assume Anne to represent an experienced ATCO who has received sufficient training to interact with the multi-agent system in an efficient way. Besides the technical elaboration, we also discuss the broader applicability of the presented interaction mechanisms.

A. Case "Conformance to Taxiway Procedures of ATC"

1) Motivation: In today's operations, ATCOs use procedures and rules to create aircraft flows through the taxiway system. At Amsterdam Airport Schiphol, the main taxiways TWY-A, TWY-B, TWY-C, and TWY-D are used primarily in one direction as visualized in Fig. 3. However, ATCOs may deviate from these taxiway procedures on their own discretion. In the historic data, most aircraft indeed follow the standard direction, with a few exceptions to e.g. go around an aircraft that is holding on a taxiway to start its engines. For fully automated operations, adhering to such procedures is not necessary, and paths could be optimized without such constraints, increasing the operational efficiency. However, such free-flow routes may not be comprehensible for human operators, especially when the traffic is dense. Furthermore, should the automation fail, the human operators may not be able to resolve and continue the operations. Therefore, for human-automation teaming, the ATCOs must be able to adjust the conformance level to taxiway procedures of the path planning, which we explore in the following.

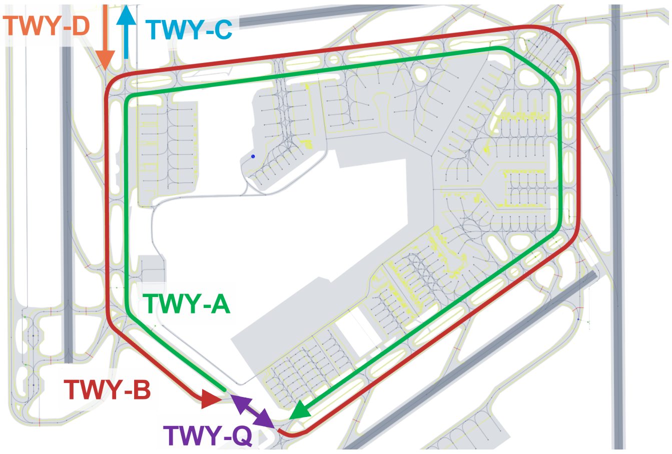

Figure 3. Standard taxiway directions as part of ATC procedures at Amsterdam Airport Schiphol

2) ATCO-MAS Interaction: Anne starts her shift in low traffic levels, and feels comfortable overseeing the operations with path planning having a low conformance to ATC procedures, and thus sets a low conformance level. A couple of hours later, with traffic increasing, she raises the conformance to a high level, and the MAS takes the new setting into account in the subsequent planning rounds.

3) Technical Elaboration: The standard taxiway directions visualized in Fig. 3 must be known to the MAS. The conformance to these taxiway procedures of ATC can then be adjusted by changing the associated cost of traversing a taxiway segment c_{seg} during path planning, by multiplying it with a cost factor c_{TWY}:

$$c_{seg} = (t_{taxi} + c_d * d_{taxi}) * c_{TWY} \tag{1}$$

with the taxi time t_{taxi} and taxi distance d_{taxi} along that segment. $c_d = 0.1\,\mathrm{s/m}$ to convert the distance to unit time. For example, when $c_{TWY} = 5$, it is five-times more expensive to traverse that taxiway segment in comparison to one with $c_{TWY} = 1$ for identical taxi time and distance. Since the sum-of-cost is minimized during path planning, non-standard taxiway directions that are assigned higher values for c_{TWY} are less likely chosen.

To analyse the impact of c_{TWY} on the conformance level, we conduct a sensitivity analysis by varying c_{TWY} between 1 to 5 with the two days of operational data used in [11], [12]. As shown by the red lines in Fig. 4, the amount of traversing taxiway segments in the non-standard direction sharply decreases for higher values of c_{TWY}. With $c_{TWY} > 1.2$, the path planning of the MAS adheres more to the standard taxiway directions than the paths taken in the historic operations (black lines). However, higher values of c_{TWY} cause the path planning to be more constrained, leading to a slight increase in taxi times as shown by the box-and-whisker plots per runway in Fig. 5. Nonetheless, the delay patterns that we studied in our previous work [11] remain unaffected by changes in the adherence to the taxiway procedures.

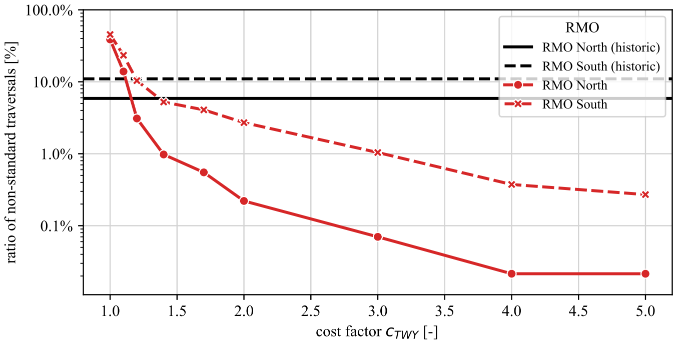

Figure 4. Ratio of non-standard taxiway traversals w.r.t. all traversals for two operational days with varying runway mode of operations (RMOs). Black lines indicate ratios of the historic operations, red lines those of the simulated operations for varying values of the cost factor c_{TWY}.

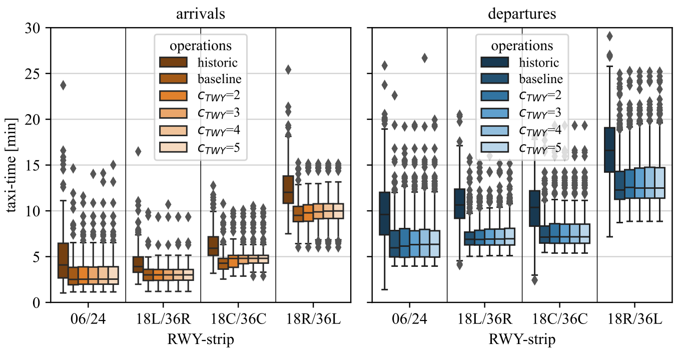

Figure 5. Comparison of taxi times per runway strip between historic operations and simulations with varying values of the cost factor c_{TWY}

4) Discussion: In the future, instead of letting the ATCOs set the conformance level manually, they could opt for letting the MAS do so dynamically, e.g. based on the number of flights to be routed in each planning round and the learned preferences of ATCOs. Moreover, other high-level routing parameters could be adjusted as well, for example:

- Setting default priorities, e.g. between arrivals and departures or between aircraft and ground vehicles: giving higher priorities to a certain group will tend to decrease their taxi times. Nonetheless, the ATCOs can deviate from these default priorities by assigning a specific priority-value to any vehicle.
- Setting the general speed limit for the entire taxiway system or certain areas like the bay areas, foremost dependent on the weather situation. Potentially, a data-driven AI subsystem could pose recommendations based on historic weather patterns.
- Setting the minimal or maximal duration that an aircraft holds at a remote holding location. This will impact the decision-support provided by the MAS for e.g. arriving aircraft whose stand is still unavailable for a certain duration.

B. Case "Departing Aircraft is Delayed"

1) Motivation: In real-world airport operations, delays frequently arise out of various reasons, and must be dealt with. When planning the taxi routes of aircraft, delays occurring prior to the predicted start of the route can be counteracted by updating the prediction and replanning the route accordingly. Many delays may remain unknown to automated systems, and updating the predictions may require coordination among human operators as well as their expertise and problem-solving skills. Therefore, to achieve effective human-automation teaming, such changes and prediction updates must be steadily supplied to the MAS. In the following, we explore one such example.

2) ATCO-MAS Interaction: 10 min prior to the planned pushback time of AC-1, ATC is informed by the ground handler that multiple passengers have not boarded yet, delaying the ready-time by at least 6 min. Anne updates the **flight schedule** of AC-1 accordingly, and the MAS automatically replans all affected routes.

3) Technical Elaboration: All routes of aircraft are based on the flight schedule: per flight, the respective start point and time as well as the goal location are extracted from it when forming its activity sequence. The corresponding activities are updated when flight schedule entries change. The updates are then accounted for when the MAS replans the routes of all flights that are or will be taxiing within the planning window. If necessary, the replanning can also be triggered rule-based or manually. In the example, the target off-block time as start time of the route is updated.

4) Discussion: In a similar way, any adaptations to the flight schedule are handled, e.g.:

- Assigning a new stand to an arriving aircraft: the goal location in the activity sequence is adapted.
- Allocating deicing to a departing flight in winter conditions: intermediate activities are inserted into the sequence that demand the aircraft to taxi to one of the deicing locations at which it has to hold for a specific time to receive the deicing.
- Changes to the takeoff slot assigned by Eurocontrol (i.e. CTOT-slot): the constraining time to be at the runway is adjusted, potentially also affecting the holding time at the stand and/or remote holding location.

In general, any flight schedule change may lead to knock-on effects, e.g. trigger a stand conflict (i.e. the delayed departing flight is blocking the stand that an arriving aircraft is assigned to) that have to be resolved as well. This may entail further interactions between the MAS and the human operators. However, as such interdependent effects occur mostly after the original cause, there is likely more time for the human operators to request and act upon recommendations from the MAS or make informed decisions to resolve them.

C. Case "Alternative Pushback Path"

1) Motivation: At Schiphol, the standard pushback and push-pull paths dependent on the aircraft types are defined in the airport manuals for all stands (see [16]). We integrated these paths in the airport layout and use them as part of the activity sequence for outbound flights. However, based on the historic data as well as on interviews with operational experts, it can be concluded that the ATCOs deviate from these standard pushback procedures in around 20 % to 30 % of times to further optimize the flows and taxi times of the involved aircraft. Thus, the MAS must accommodate to receive and process such informed changes based on the experience of the ATCOs. In the following, we explore two examples that showcase how such ad-hoc changes are dealt with in the MAS model.

2) ATCO-MAS Interaction - Part 1: The amount of traffic is low, and Anne decided to lower the **level of support** provided by the Interaction Agent. Using the **fast-forward simulation**, she checks how the traffic will likely evolve around bay area D. She notices that in around 5 min the departing aircraft AC-2 will start its engines in a location that blocks the stand entry for the arriving aircraft AC-3 (Fig. 6a). From her experience, she knows that a longer, non-standard pushback further into the bay area is unproblematic, and manually adapts the pushback path of the departing aircraft. She then queries the MAS to compute **alternative MAMP routes** for all involved aircraft. After reviewing the pending changes to the aircraft trajectories using the **fast-forward simulation** (Fig. 6b), she accepts the proposed new routing plans.

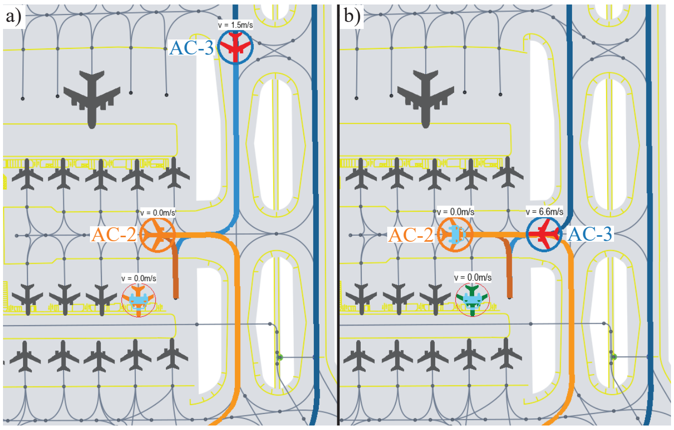

Figure 6. Fast-forward animations of a) the original trajectories and b) the adapted routes when AC-2 uses a longer pushback path

3) Technical Elaboration - Part 1: By specifying a longer pushback path, the activity sequence is automatically updated: during the pushback-activity, the aircraft must follow the new path. While the changes are still pending the acceptance by the ATCO, the Interaction Agent keeps both the original as well as the new activity sequences in cache. Likewise, it keeps a copy of the original MAMP solution. When queried by the ATCO, it replans the potential routes of all agents using the new activity sequence of the departing aircraft AC-2. Once the new solution is accepted by the ATCO, the MAS automatically sends the new routes to the affected aircraft.

4) ATCO-MAS Interaction - Part 2: With higher traffic load, Anne has increased the **level of support** provided by the Interaction Agent. She receives a notification from it to review upcoming routing plans for aircraft departing from bay area D/E as visualized in Fig. 7: by following the standard pushback path, AC-4 (blue route) will be trailing AC-5 (green route), followed by AC-6 (orange route). This will cause delays for both AC-4 and AC-6. Since other traffic is not involved as validated with the **fast-forward simulation**, Anne changes the pushback path of AC-4 so that the aircraft can leave the bay area directly after engine start. She lets the MAS calculate the adapted **MAMP routes** (Fig. 7b). The Interaction Agent notifies Anne that the alternative routes result in an improved runway sequence (Fig. 7d), reducing the taxi times of AC-4 by 2:40 and AC-6 by 1 min. Anne accepts the changed routes.

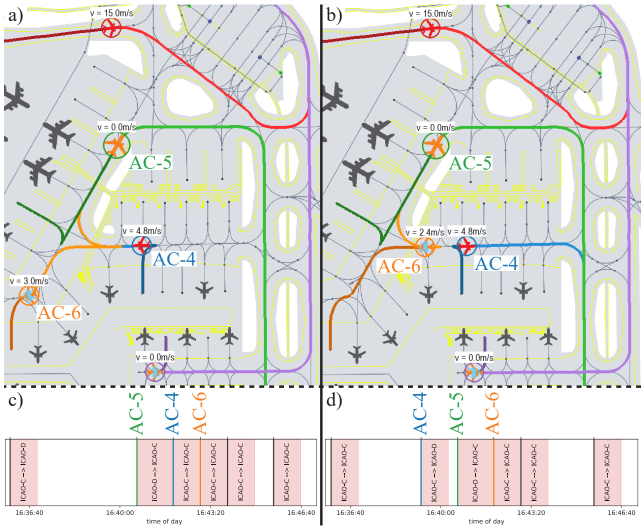

Figure 7. Fast-forward animations of a) the original trajectories and b) the adapted routes when AC-4 uses an alternative pushback path, with the respective runway sequences shown in c) and d)

5) Technical Elaboration - Part 2: By increasing the provided level of support, the Interaction Agent automatically carries out various analyses of the concurrent routes to find those that likely profit from manual changes. In the case above, the two aircraft AC-4 and AC-6 are delayed in the same bay area due to having lower priority than AC-5, which hints at a possibility to improve the flow by manual changes. Like in Section III-C3, the input of the ATCO leads first to an update of the activity sequence and second to a recalculation of the trajectories. The Interaction Agent then identifies the core changes as described above and lets the ATCO validate the new routes by animating the changes.

6) Discussion: The MAS stores the non-standard pushback paths per stand. Over time, and dependent on the chosen level of decision-support, the MAS can then automatically suggest one or multiple of these paths, after having analysed that the impact on other traffic is minimal. Still, the ATCO will need to validate and accept the suggested change. Eventually, when a high level of automation is chosen, the MAS can incorporate these alternative pushback paths directly in the route optimization, reducing the workload of the ATCO to review the suggestions.

D. *Case "Unavailability of Stand"*

1) Motivation: The stands form one of the bottlenecks in the capacity of airports, and the aircraft stand allocation is a manifold problem of its own. Due to delays, arriving ahead

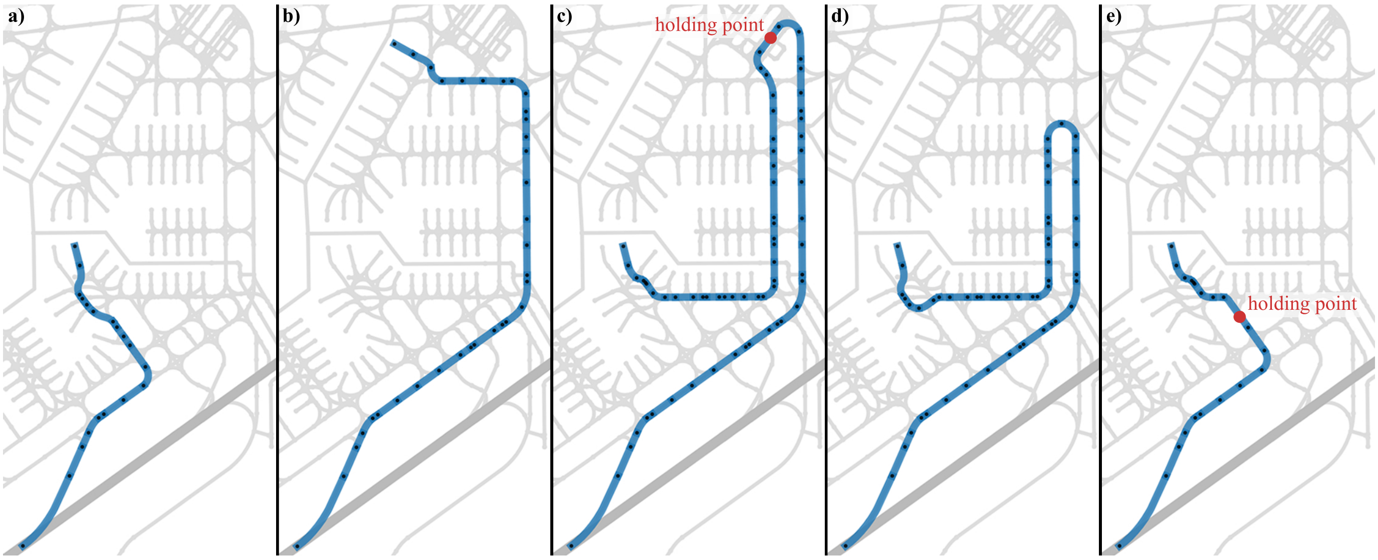

Figure 8. The stand of an arriving aircraft is unavailable: SAMP paths of a) original route, b) rerouting to a different stand, c) holding at a remote holding point, d) detour with minimal speed, and e) holding at a non-standard holding point

of time, or allocation constraints among others, a stand may still be occupied by another aircraft at the time that an arriving aircraft could enter it. The route of the arriving aircraft must thus be adapted. Such cases are visible in the historic data, but were also raised in interviews with experts as both a common operational challenge and interesting use case for airside automation. In the following, we provide an interactive example how the unavailability of a stand can be resolved using the agent-based framework of human-automation teaming.

2) ATCO-MAS Interaction: Shortly before the flight AC-7 is landing at Schiphol's runway 06, Anne is informed by the ground handler that the chosen stand D04 is blocked / unavailable for another 20 min. Since the aircraft is landing soon, the MAS has already planned a conflict-free route. The ATCO displays the planned route (path displayed in Fig. 8a), and notices that the aircraft would arrive approximately 15 min too early at the stand. First, Anne places a pending **goal-constraint** at the stand D04 for the blocked time period, and requests the Interaction Agent to compute initial alternatives without accounting for other traffic, i.e. **SAMP routes**. The Interaction Agent displays three initial solution strategies: (1) assign an alternative stand to the aircraft (path b), (2) let it hold at one of the remote holding points (path c), and (3) let the aircraft take a detour along TWY-B and TWY-A with minimal taxi speed (path d). With these options in mind, Anne uses the **fast-forward simulation** to get an impression of the traffic situation in the upcoming 30 min. In turn, she disregards all three options: the alternative stand (1) is in another bay area and will likely not be acceptable for the airline, all remote holding points (2) are already occupied and would require extensive changes to the routes of other aircraft, and she deems (3) to be impractical given the amount of traffic in the upcoming period. From her experience, Anne knows that the aircraft could alternatively hold in bay area C/D[1]. She selects the chosen point as non-standard holding location at which AC-7 shall hold for 15 min, and lets the MAS compute the **SAMP route** (path e). After checking her resolution option using again the **fast-forward simulation**, she accepts the changes to the **activity sequence** of aircraft AC-7. The MAS carries out partial replanning of the **MAMP routes** and notifies Anne that the changes have taken effect.

3) Technical Elaboration: Through the goal-constraint, AC-7 is not allowed to arrive at the stand prior to the end time, which is taken into account as strict requirement during path planning by the MAS. The Interaction Agent uses a set of options such as rerouting to another stand, or holding at a remote holding location among others to create appropriate activity sequences and let the MAS determine the corresponding SAMP-paths. As none of these route alternatives appear suitable to the ATCO, another activity sequence is created from the inputs provided by the ATCO: the chosen holding location is added as intermediate goal at which the aircraft must wait for the selected duration. The MAS carries out a partial replanning by deconflicting the routes of all aircraft that are affected by the new route of AC-7.

4) Discussion: In a similar manner, the ATCOs can add further constraints (see Section II-C3) as well as additions to the activity sequence such as passing specific waypoints or waiting durations to any aircraft route. This enables a co-design of the routes that are then deconflicted by the MAS using alternative, unconstrained paths and/or speed adjustments. When such non-standard routing and resolution strategies are used more frequently, the Interaction Agent can learn to provide these alternatives automatically, and eventually use them to already adapt the original planning where possible.

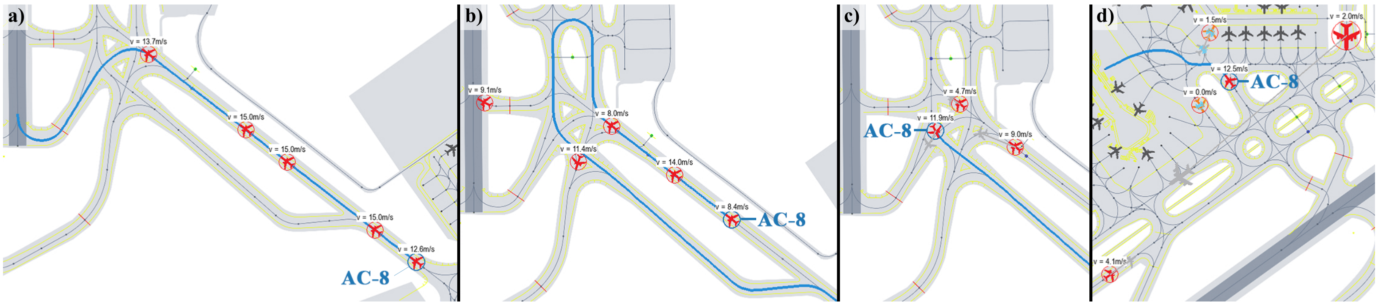

Figure 9. The blue-circled aircraft declares an emergency: a) shows the original route to the runway, b)-d) its route back to the stand with the effects on other traffic shown by shadow-aircraft

E. Case "Emergency demands Aircraft to Return to Stand"

1) Motivation: Non-nominal situations may occur infrequently, but often require non-standard resolution strategies. Especially in such situations, the automated side of the human-automation teaming must provide a flexible interface for effective decision-support. In the following, we assess a scenario in which a departing aircraft must return to its stand due to an emergency. To challenge the path planning algorithm, we chose an example with much surrounding traffic close to the bottleneck of TWY-Q (Fig. 3).

2) ATCO-MAS Interaction: Aircraft AC-8 is following TWY-Q on its way to runway 36C, when the pilots declare an emergency via radio call to ATC. They request to be guided back to the stand as fast as possible. Anne quickly assesses the situation. She updates the **activity sequence** of the aircraft and adjusts the **priority** assigned to the aircraft. The MAS recalculates the **MAMP routes** of AC-8 and all affected aircraft accordingly.

3) Technical Elaboration: The activity sequence must be redefined so that the aircraft is not routed further to the runway, but instead back to the original or alternative stand. Dependent on the emergency and general traffic situation, the ATCO can adjust the priority of the aircraft in comparison to other traffic: a high priority yields a fast return route while potentially more traffic is affected, whereas a low priority potentially lengthens its taxi time but creates minimal nuisance to the routes of other aircraft. For various exemplary emergency situations, changing the priority of the emergency-declaring aircraft did not have a significant impact on its own as well as the trajectories of affected aircraft. This suggests that the Routing Algorithm is able to recover to the desired level from such situations. For the default priority, Fig. 9 visualizes the original route of AC-8 (a) as well as its route back to the stand. The effects on other traffic with respect to the original planning are visualized by shadow-aircraft.

4) Discussion: In urgent situations, ATCOs could also deviate more extensively from the regular procedures. For instance, they could let the aircraft turn immediately from TWY-A back to TWY-B which is usually not allowed since no taxiway centreline exists for this turn. Moreover, they could let the aircraft stop immediately, order emergency response teams to it, and reroute all surrounding traffic. We leave the further elaboration of such non-nominal scenarios to future work.

[1] In the analysed historic data, aircraft frequently hold for multiple minutes in certain locations in the taxiway system.

IV. Future Work

In the context of the ASTAIR project, we plan to conduct human-in-the-loop experiments with ATCOs and pilots to validate and refine the interactions presented above. We expect that in further discussions with the operational experts more reoccurring cases will emerge and possibly new interaction tools will be identified. These as well as other non-nominal situations can then further improve the decision-support provided by the Interaction Agent and the MAS. Moreover, we plan to incorporate different degradation modes into the MAS model to handle failures of the automation and conflict-free routing. A thorough analysis will be necessary to assess the system's resilience, to explore how to relay critical information to ATCOs to handle such failures effectively, and to ensure safe operations in any situation.

We acknowledge that human factors like workload, stress, experience levels, and potential biases can significantly influence the decision-making in real-world settings. Thus, future research should also focus on creating a cognitive and behavioural model of the human operators within the Interaction Agent so that it can better predict and respond to the mental state of the human operators interacting with it. With such a model, the Interaction Agent could also aid in keeping the human operators better engaged and support in maintaining and honing their skills.

V. Conclusion

In this paper, we explored how human-automation interactions in airport surface movement operations (ASM Ops) can be modelled and integrated algorithmically into the MAS model on autonomous ASM Ops from our previous work. To this end, we specified a schematic agent-based framework of human-automation teaming for ASM Ops. The interactions between human operators and the automated side are facilitated by the Interaction Agent providing interactive interfaces and tools. By examining use cases derived from workshops and interviews with stakeholders, we have demonstrated how ATCOs can oversee, influence, and co-design the conflict-free 4D-trajectories computed by the MAS model, ensuring operational safety and efficiency. The proposed agent-based model has proven to be adaptive and flexible as well as to perform at the desired level, advancing the realization of successful human-automation teaming in ASM Ops.

Acknowledgment

This work has received funding from SESAR Joint Undertaking under grant agreement No 101114684 under European Union's Horizon 2020 research and innovation programme. The authors would like to thank our project partners in the ASTAIR project for their helpful feedback on this paper, as well as the valuable insights shared during workshops and interviews. We also wish to thank the reviewers for their constructive comments, which helped us to improve our work.

References

[1] IATA, *Net zero 2050: Operational and infrastructure improvements*, 2022.

[2] Eurocontrol, *European aviation in 2040 - challenges of growth*, Eurocontrol Statistics and Forecast Service, 2018.

[3] Eurocontrol, *Data snapshot #22 on lower summer taxi-out times*, 2021.

[4] Z. Chua, M. Cousy, M. Causse, and F. Lancelot, "Initial assessment of the impact of modern taxiing techniques on airport ground control," in *Proceedings of HCI-Aero 2016*, ACM, 2017.

[5] *AI roadmap 2.0: Human-centric approach to AI in aviation*, EASA, 2023.

[6] *Automation in air traffic management: Long-term vision and initial research roadmap*, SESAR JU, 2020.

[7] "SESAR Joint Undertaking — MOTO - the Embodied Remote Tower." (2016), [Online]. Available: https://www.sesarju.eu/projects/moto (visited on 09/25/2023).

[8] "SESAR Joint Undertaking — TaCo - Take Control." (2017), [Online]. Available: https://www.sesarju.eu/projects/taco (visited on 09/25/2023).

[9] L. Blain. "Autonomous Airbus aces autopilot taxi, takeoff and landing tests," New Atlas. (2020), [Online]. Available: https://newatlas.com/aircraft/airbus-attol-autonomous-airliner/ (visited on 05/07/2024).

[10] M. von der Burg, J. Kamphof, J. Soomers, and A. Sharpanskykh. "Towards Autonomous Airport Surface Movement Operations Using Hierarchical Multi-Agent Planning." (2024), [Online]. Available: https://papers.ssrn.com/abstract=4916874 (visited on 08/12/2024), pre-published.

[11] M. von der Burg and A. Sharpanskykh, "Multi-Agent Planning for Autonomous Airport Surface Movement Operations," in *SESAR Innovation Days 2023*, 2023.

[12] M. von der Burg and A. Sharpanskykh, "Studying the Operational Consequences of Automated Engine-Off Taxiing using Multi-Agent Planning," in *ICRAT-24*, 2024.

[13] M. Rieth and V. Hagemann, "Automation as an equal team player for humans? A view into the field and implications for research and practice," *Applied Ergonomics*, vol. 98, no. 4, 2022.

[14] "SESAR Joint Undertaking — ASTAIR - Auto-Steer Taxi at Airport." (2023), [Online]. Available: https://www.sesarju.eu/projects/astair (visited on 09/25/2023).

[15] ASTAIR, "D1.2 - Workshops Report," 2024.

[16] "Schiphol - standaard pushback per positie," Schiphol. (), [Online]. Available: https://www.schiphol.nl/en/operations/page/sleep-en-pushbackbewegingen/ (visited on 08/30/2022).