Function Allocation for NextGen Airspace via Agents

Nathan Schurr
Aptima, Inc.
Woburn, MA 01801
nschurr@aptima.com

Paul Picciano
Aptima, Inc.
Woburn, MA 01801
ppicciano@aptima.com

Janusz Marecki
IBM T.J. Watson Research
Yorktown Heights, N.Y. 10598
marecki@us.ibm.com

ABSTRACT

Commercial aviation transportation is on the rise and has become a necessity in our increasingly global world. There is a societal demand for more options, more traffic, more efficiency, while still maintaining safety in the airspace. To meet these demands the Next Generation Air Transportation System (NextGen) concept from NASA calls for technologies and systems offering increasing support from automated decision-aiding and optimization tools. Such systems must coordinate with the human operator to take advantage of the functions each can best perform: The automated tools must be designed to support the optimal allocation of tasks (functions) between the system and the human operators using these systems. Preliminary function allocation methods must be developed (and evaluated) that focus on the NextGen Airportal challenges, given a flexible, changing Concept of Operations (ConOps).

We have begun making steps toward this by leveraging work in agents research (namely Adjustable Autonomy) in order to allow function allocation to become more dynamic and adjust to the goals, demands, and constraints of the current situation as it unfolds. In this paper we introduce Dynamic Function Allocation Strategies (DFAS) that are not static and singular, but rather are represented by allocation policies that vary over time and circumstances. The NextGen aviation domain is a natural fit for agent based systems because of its inherently distributed nature and the need for automated systems to coordinate on tasks maps well to the adjustable autonomy problem. While current adjustable autonomy methods are applicable in this context, crucial extensions are needed to push the existing models to larger numbers of human players, while maintaining critical timing. To this end, we have created an air traffic control system that includes: (1) A simulation environment, (2) a DFAS algorithm for providing adjustable autonomy strategies and (3) the agents for executing the strategies and measuring system efficiency. We believe that our system is the first step towards showing the efficacy of agent supported approach to driving the dynamic roles across human operators and automated systems in the NextGen environment. We present some initial results from a pilot study using this system.

Cite as: Function Allocation for NextGen Airspace via Agents, Nathan Schurr, Paul Picciano and Janusz Marecki, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1731-1738

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Applications and Expert Systems

General Terms

Human Factors, Algorithms

Keywords

Aviation Transportation, Human-agent Interaction, Adjustable Autonomy, Markov Decision Process

1. INTRODUCTION

This paper is describing an application of agents and agent research to an ongoing NASA research project named Airportal[1] Function Allocation Reasoning (AFAR). The main focus of this project is to understanding the roles and responsibilities of automation in the Next Generation Air Transportation System (NextGen). While we are currently in the first year of AFAR's 3 year effort, the project already shows potential for the commercialization of agent technologies.

The core components of the AFAR project include:

- AFAR Agent Framework - for monitoring and adjusting workflow;
- Dynamic Function Allocation Strategies (DFAS) - for determining the best role for automation over time;
- Comprehensive Metrics for Function Allocation (CMFA) - for quantifying the impact of these novel roles for automation and
- AFAR Simulation - for modeling and visualizing air traffic that is to be controlled by both human and automation.

All of these components are leveraged to conduct human in the loop experiments in order to explore the roles of automation in NexGen. Currently the AFAR project has just completed building the initial system which allows for human-agent interaction. The focus of the scenarios currently are on the East Tower of the Dallas Fort Worth Airport (DFW) and the air traffic controller position of ground controller. We have just completed our first pilot study experiment with actual air traffic controllers and are beginning to analyze

[1]Airportal refers to the area in and around an airport, which is the most congested and dangerous section of the National Airspace (NAS)

the results. Furthermore, the implementation of the system has led us to identifying extensions to recent work in Adjustable Autonomy in order to improve the performance of DFAS. The rest of the paper documents our approach, progress and insights. We begin by providing an overview of some of the aspects of AFAR industrial relevance.

- *Importance:* Commercial aviation transportation is on the rise and has become a necessity in our increasingly global world. There is a societal demand for more options, more traffic, more efficiency, while still maintaining safety in the airspace. This is not possible unless dramatic investments are made in automation and infrastructure. Automation will be playing a larger role to meet these needs, but automation designers need new tools to allow for greater and more flexible roles for the automation.

- *Rationale:* An agent based technology is a natural fit to the aviation transportation domains for the following reasons: (1) Airspace operations are inherently distributed; (2) Problems are often communications and coordination focused; (3) Agent based solutions provide more timely and integrated responses and (4) they increase the amount of humans and automated systems that can be simultaneously coordinated. In addition, existing research in agents (namely in Adjustable Autonomy) can be directly applied to aid in human-automation transfer of control.

- *Barriers*: The primary barrier to adoption of this technology is rigorous human in the loop evaluation of these methods to show feasibility and usefulness. In order to address this, AFAR is focused on not only development, but evaluation of these systems using actual air traffic controllers.

- *Financial:* Since development and evaluation is in the early stages, the exact value is difficult to quantify. However, DFAS holds great potential to industry designers of automated system in the flight deck and air traffic control towers and government customers wishing to evaluate said designs. In order to convey the economic need for NextGen, it is best to quote the FAA: "By 2022, FAA estimates that [the lack of NextGen] would cost the U.S. economy $22 billion annually in lost economic activity. That number grows to more than $40 billion by 2033 if the air transportation system is not transformed."

2. MOTIVATION

The responsibilities and privileges bestowed upon National Airspace (NAS) stakeholders are expected to dramatically evolve as the demand for the use of airspace continues to rise [5]. To meet these growing demands and provide the level of service, safety and security needed to sustain future air transport, the Next Generation Air Transportation System (NextGen) concept calls for technologies and systems offering increasing support from automated systems that provide decision-aiding and optimization capabilities. The automation must be designed to enable the human operators to access/process a myriad of information sources, understand heightened system complexity, and maximize capacity and throughput in the Airportal environment.

While NextGen promotes the introduction of advanced technologies to support these specific needs, fundamental to the Joint Planning and Development Office (JPDO) definition of the NextGen concept is the notion that such automated systems must coordinate with the human operator to "take advantage of the functions each can best perform" [5]. Under this specification, such automated tools must be designed to support the optimal allocation of tasks between the system and the human operators using these systems. That is, preliminary function allocation methods among operators and automated systems must be studied early in the design process, to ensure that the impact of these function allocations can be fully realized in an implemented system. This becomes even more important because NextGen will involve changing concepts of operations to meet future needs. Any function allocation methodology cannot be static, but must be capable of adjusting seamlessly and appropriately as required to meet technology developments as they arise.

There are three critical challenges that need to be addressed in supporting the early design of systems that allow for optimal, context-sensitive function allocation. First: preliminary function allocation methods must be developed that focus on the NextGen Airportal challenges, given a flexible, changing Concept of Operations (ConOps). Because the NextGen ConOps involves new classes of functions and responsibilities, with relatively unstudied consequences, methods for assessing the allocations must be adaptable to previously unplanned-for scenarios.

Second, the Human Systems Integration (HSI) community is in need of robust methods for evaluating the effectiveness of total system performance, where a system is defined as the humans and automated tools working collaboratively to meet productivity and performance goals. While traditional metrics (e.g., accuracy, completion time) can be used for assessing total system performance, these methods have generally been extended from human performance measurement approaches and often lack sensitivity for assessing the complex performance of the collaborative human/automation team.

Finally, such methods must leverage the extensive work that has been done in the last decade to evaluate the intricate relationship between automation and the human operator(s) [8, 9]. These solutions must capture the complexities of that empirical work in clear, concise, and intuitive design guidelines to support designers in developing systems and policies to exploit the capabilities of both human operators and automated systems. In addition to design guidelines that follow from empirical human-automation research, computational and formal models are also important tools that can be delivered to designers. Examples of these include formal models for automation verification [3], Bayesian modeling approaches (Inagaki, 1999), and mathematical models of automation reliance and compliance (Wickens & Dixon, 2006). A combination of both guidelines and more formal models is needed to support objectives as described by Safe and Efficient Surface Operations (SESO) and Coordinated Arrival/Departure Operations Management (CADOM) technical areas of Airportal.

3. AIRPORTAL FUNCTION ALLOCATION REASONING (AFAR)

To address these challenges we present the Airportal Function Allocation Reasoning (AFAR) project. AFAR's approach has three key components: (1) Leverages recent work on human-agent allocation to develop novel methods for function allocation, particularly the Dynamic Function Allocation Strategies (DFAS); (2) Provides Comprehensive Measurements for Function Allocation (CMFA), encompassing a multi-source measurement approach for evaluating function allocations among complex networks of distributed human-automation interactions, and (3) Establishes a clear path to engage NASA researchers throughout the process, from selecting critical research issues to use case generation, experimental design and execution, and finally transitioning to useful guidelines.

These guidelines are generated to help designers grasp the implications of allocating functions between humans and automation in the NextGen Airportal, whether those designers seek to develop design aid tools or augment existing ConOps. To accomplish this objective, we employ an iterative, recursive approach of development and evaluation. In addition, we leverage previous work from the human factors community in the understanding of human-automation interaction to drive both design and assessment. Thus, as CMFA is constructed to critically assess DFAS performance in human in the loop simulation, it concurrently provides an opportunity to evaluate the adequacy of the performance metrics and provides opportunities to improve CMFA for future evaluations of even other methods of function allocation. By spiraling this process up from lower fidelity environments (e.g., Distributed Dynamic Decision-making (DDD) simulator) to more sophisticated testbeds (e.g., NASA facilities, such as the Airspace Operations Lab (AOL) and Advanced Concepts Flight Simulator (ACFS)), we can make incremental improvements to both DFAS and CMFA, and manage risk as the simulations better approximate a real world environment. This provides the best opportunity to generating guidance that will be useful to NextGen stakeholders.

We are developing, evaluating and transition methods for human-automation function allocation to guide designers of NextGen technologies and procedures as they consider the critical issues of human-automation interaction. The AFAR project involves collaboration with NASA personnel throughout the 3-year effort. This approach leverages recent advances in artificial intelligence (namely Adjustable Autonomy) in order to allow function allocation to become more dynamic and adjust to the goals, demands, and constraints of the current situation as it unfolds. Consequently, AFAR introduces Dynamic Function Allocation Strategies (DFAS) that are not static and singular, but rather are represented by allocation policies that vary over time and circumstances. By casting function allocation as a problem of adjusting autonomy between human-automation (and even human-human) links, we can leverage the current techniques for optimizing both strategy selection and strategy timing (i.e., the how and when). The adjustable autonomy approach incorporates aspects governed by two vital factors: performance capability and deadline. In determining the optimal allocation between human and agent, adjustable autonomy considers not only the capability, availability, and appropriateness of an actor to perform a function, but the overall system state and temporal implications regarding the overall planning of higher level goals. Thus, an adjustable autonomy approach brings rigorous, testable methods for understanding how function allocation will need to adjust to the requirements of future concepts.

We expect multiple iterations of evaluation to occur to continually refine the Dynamic Function Allocation Strategies (DFAS) and the Comprehensive Measurements of Function Allocation (CMFA) methods based on the progression of the various simulations. The CMFA results from the combination of measures from three different sources: (a) observer-based data from trained individuals recorded in real time; (b) self-report data from operators collected during or after testing runs; and (c) system-based data derived directly from the test scenario.

Throughout this process, we focus on use cases to identify function allocation challenges relevant to the NextGen Airportal, to iteratively address both DFAS and CMFA issues. Indeed, design decisions which require strong human-automation integration are best validated by both empirical human performance testing combined with computational/mathematical model testing. Such evaluations are essential because prior research has shown that automated systems in many domains, including aviation, are not always used by human operators in the ways that designers intended [8]. In particular, human operators may not use automated aiding appropriately. While such automation disuse and misuse can be minimized through careful design and application of AFAR, human performance evaluation is necessary to ensure appropriate automation use [9]. This will enable AFAR's function allocations (DFAS) and evaluation methods (CMFA) to lead to more feasible and useful guidelines for automation and policy designers for the NextGen Airportal.

Critical to any program success is the assurance of a sound, disciplined systems engineering process that ties program elements and human-automation collaboration together. Our activities in the first year have centered on developing appropriate function allocations to foster NextGen Airportal initiatives. We are working with NASA to develop use cases that will highlight the challenges brought about by function allocation in the future NextGen Airportal. These use cases serve as scenarios to be carried throughout all three years of the AFAR effort, providing a common framework for our DFAS and the CMFA components.

4. DYNAMIC FUNCTION ALLOCATION STRATEGY (DFAS)

One of the most fundamental challenges of building a human-multiagent team is that of "adjustable autonomy" [11, 15, 12, 16, 17, 18]. AFAR builds upon this recent work on adjustable autonomy to create novel Dynamic Function Allocation Strategies (DFAS). In particular DFAS' human-agent automation allocation strategies are expanded to handle increased dynamics and agent workload.

Conventionally, adjustable autonomy refers to the ability of an agent to dynamically adjust its own autonomy, thereby altering the amount of control a human has over the agent at a particular time and context. Given the changing state of the environment and the team, it is beneficial for an agent to be flexible in its autonomy and not be forced to act either with full autonomy or with zero autonomy. The key to ad-

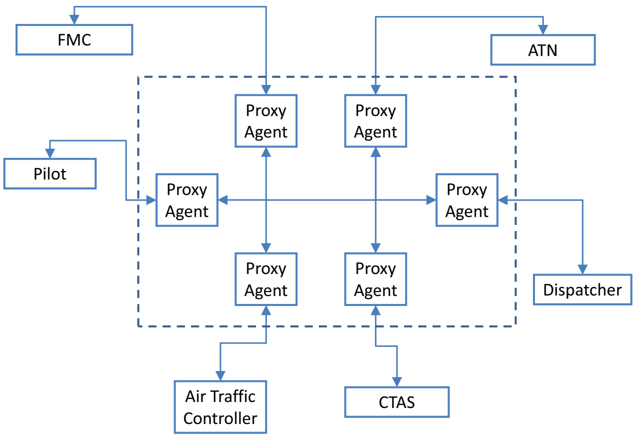

Figure 1: AFAR Multiagent Proxy Architecture

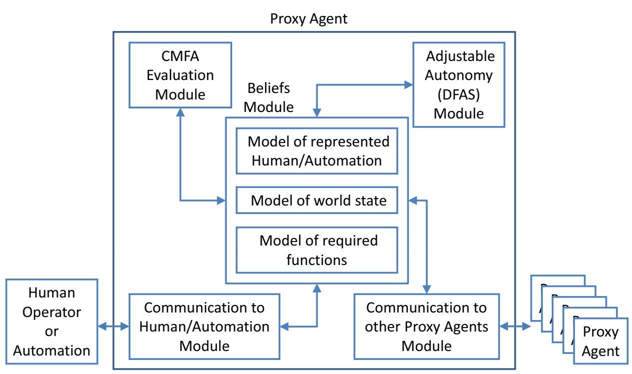

Figure 2: AFAR Proxy Agent Components

justable autonomy is to answer a question "when and where to transfer the control of a decision".

Figure 1 shows the AFAR multiagent proxy architecture applied to an example Airportal scenario that allows for dynamic, reactive function allocations (DFAS). This scenario has three human operators and three components of automation that all must coordinate in the context of the Airportal. The human operators are the Pilot, Air Traffic Controller and the Dispatcher, whereas the automation encompasses the Flight Management Computer (FMC), the Aeronautical Telecommunications Network (ATN), and the Center/TRACON Automation System (CTAS). Inside the dashed box are the proxy agents that represent the members (human operator or automation) of the overall Airportal team. The proxies each negotiate and react to allow for a function allocation to shift from a human operator to the automation, or even to another human operator. Figure 2 shows the internal structure of an individual proxy agent.

Each AFAR proxy agent has a number of modules that are guiding its reasoning over the functional allocation. In order to shape the functions and their respective desired allocations, a Beliefs Module keeps track of the state of the represented human or automation, the state of the external world and a list of required functions. The other modules read and write to the Beliefs Module as necessary. For example, the model of a represented human operator would include factors such as capabilities to perform given functions, estimates on how long it would take that human to perform a function, and the estimated quality of performance for each function. These models of human performance can be implemented in a variety of ways ranging from simple distributions to complex computational model of the human (eg., ACT-R or SOAR). The Adjustable Autonomy Module hence determines the Dynamic Function Allocation Strategy (DFAS) over time of which Airportal member (human operator or automation) should be allocated for a given function.

Recent work in adjustable autonomy [13], generalizes it to a process in which the control over a decision is dynamically transferred between a human and an agent in order to find the teammate (agent or human) that is best fit for making the decision. The choice of the best fit teammate depends on the state of the world, the locality of information among team members, the priority of the decision, and the availability of the teammate. At the same time, adjustable autonomy also factors in policy or doctrine to guide the allocation decision [2]. For example, if the nature of the decision involves lethality (e.g., firing a weapon) then only a human would be able to have autonomy. In addition, all of these factors change continuously over time and consequently, the best teammate to make a decision may change over time as well. In an example of a possible function allocation strategy for this context, the Traffic Flow Manager (Human Operator 1) determines that there is a need to go to super density operations for an arrival. This alters the allocation (autonomy) over the function of Landing Procedure over the Aircraft Flight Deck (Automation 1), which activates the autolanding system due to limited visibility. At that point, the allocation is switched to the Air Traffic Management Systems (Automation 2) that compute and transmit the taxi clearance to the aircraft before landing. After touchdown, the locus of control is transferred again, this time to the human operator piloting the aircraft (Human Operator 2). The pilot executes the necessary taxi to avoid runway incursions and arrive safely at the gate. In the above case, instead of having a single, predetermined function allocation that statically assigns roles and responsibilities to the automation or the human, there can be a dynamic function allocation strategy that suggests points in time to alter the allocation. In addition, there can be a reactive updates to this strategy as uncertainties decrease and events arise.

Consider the following example DFAS: $H_1T_1AT_2H_3T_3AT_4$. In the example, times at which the allocation should transfer to the next function performer (human/automation) are labeled as T_1,T_2,T_3 and T_4 respectively. Optimization algorithms for the DFAS determine not only the order of the allocation changes, but also the timings. Transitions of control between agents (human/automation) must be carefully orchestrated to avoid automation surprises. AFAR focuses on applying adjustable autonomy to the complex, dynamic domain of NextGen-Airportal operations where human automation function allocation must adapt due to changes in the aspects of the functions themselves.

5. DFAS SOLUTION

Early work on mixed-initiative and adjustable autonomy assumed that transfer of control was one-shot and permanent. Since such rigid transfer of control decisions are problematic, particularly in time-critical domains, researchers have since focused on multi-stage adjustable autonomy strategies [12, 18, 4, 16], effectively allowing for transfer of control to go back-and-forth between the agents and the humans. Due to their ability to handle the uncertainty about the availability of humans, Markov Decision Process (MDP) based methods for adjustable autonomy in particular have received a lot of attention [18, 15]. However, in order to ensure high-quality decisions, these works have used very fine-grained discretization of time, resulting in large MDP state spaces, and prohibitive algorithm runtimes. To remedy that, [13, 14] was the first to employ continuous time MDP solvers [1, 6, 7] to compute adjustable autonomy strategies. Yet, the strategies in [13, 14] assumed that only two actors, a human or an agent team, could at any point in time be in possession of decision making. Furthermore, these works assumed that human responsiveness (specified by a given probability distribution that does not change as the human accepts to make more and more decisions.

These two shortcomings of [13, 14], the inability to handle multiple human operators and the lack of implicit modeling of their changing responsiveness, need to remedied to address the requirements of the DFAS component. In the remainder of this section we show exactly that can be accomplished. We begin by recalling the formalism of Time Dependent Markov Decision Processes.

5.1 Time Dependent Planning

In many aviation situations, agents execute actions whose durations are uncertain and can only be characterized by continuous probability density function. A common approach to model such domains with continuous time characteristics has been to use the framework of semi-Markov Decision Process [10]. However, semi-MDP policies are not indexed by the current time and as such, might not have the desired expressivity when dealing with time-critical domains. For such domains, one can use the Time dependent MDP (TMDP) framework [1]. TMDP's approach to modeling continuous time is to create a finite, hybrid state-space where each discrete state has a corresponding continuous time dimension. This allows TMDP policies to be both accurate and conditioned on time limits, as required in our domains. To model its strategies, DFAS employs Interruptiple-TMDPs [14], a framework that extends TMDPs by allowing the actions to be interrupted.

The Interruptible TMDP (I-TMDP) model [1] is defined as a tuple $\langle S, A, P, D, R\rangle$ where S is a finite set of discrete states and A is a finite set of actions. P is the discrete transition function, i.e., $P(s, a, s')$ is the probability of transitioning to state $s' \in S$ if action $a \in A$ is executed in state $s' \in S$. Furthermore, for each s, a, s' there is a corresponding probability density function $d_{s,a,s'} \in D$ of action duration, i.e., $d_{s,a,s'}(t)$ is the probability that the execution of action a from state s to state s' takes time t. (The execution of actions can be interrupted in which point the agent reverts to the starting state s without incurring any cost.) The reward function, too, is time-dependent; $R(s, a, s', t)$ is the reward for transitioning to state s' from state s via action a completed at time t. Let Δ (referred to as the *deadline*) be the earliest time after which no reward can be earned for any action $a \in A$ and any states $s, s' \in S$. A policy π for a TMDP is defined as a mapping $S \times [0, \Delta] \to A$ and the expected utility of following a policy π from state s at time t is denoted as $U^{\pi}(s, t)$. The policy π^* which provides the highest expected utility for any state and time pair is then referred to as the optimal TMDP policy. Note, that π^* can be found efficiently using one of the existing TMDP solvers [1, 6, 7].

5.2 Adjustable Autonomy using I-TMDPs

We now show how to address the first shortcoming of [13, 14]: the inability to handle multiple human operators. DFAS requires that, before a role is performed, a role allocation about that role is made. For example, if the role is to "*land a plane within the next* 10 *minutes*", then the role allocation is the decision about who (either some human operator or the Automation[2].) will be piloting the plane to be landed. Note, that both role allocation and role execution tasks consume time and as such, they both need to be performed in a timely manner (within 10 minutes) for the role to be completed successfully.

For each new role that arrives, DFAS constructs the adjustable autonomy strategy using I-TMDPs as follows: The set of states is $S = \{s_{H_1}, s_{H_2}, \ldots, s_{H_N}, s_A, s_{end}\}$. If the system is in state s_A, the Automation has the control over the role; if the system is in state s_{H_n}, human operator n has the control over the role, for $1 \le n \le N$. The process starts in state s_A and stops if it transitions to state s_{end} or if the current time is greater than Δ. Whoever has the control over the role can either execute that role or transfer it to another entity. That is, $A = \{a_{execute}, a_{transfer(H_1)}, a_{transfer(H_2)}, ..., a_{transfer(H_N)}, a_{transfer(H_A)}\}$ where $a_{execute}$ results in the execution of the role, $a_{transfer(A)}$ results in the transfer of control to the Automation and $a_{tansfer(H_n)}$ results in the transfer of control to human operator n, for $1 \le n \le N$. The transition function is then defined as $P(s_A, a_{execute}, s_{end}) = P(s_{H_n}, a_{execute}, s_{end}) = 1$ for $1 \le n \le N$ and $P(s_x, a_{transfer(y)}, s_y) = 1$ where $x, y \in \{A, H_1, ..., H_N\}$ and $x \neq y$.

We assume that the transfer of control actions take less time (=has lower mean) than the role execution actions. Also, because the human operator can be busy executing other roles, we assume that the time it takes to execute the role is smaller for the Automation than for the human operator. Under these circumstances, what encourages the system to transfer the control to the human operators (who on average take longer to respond than the Automation) is the structure of the reward function. Specifically, we assume that the reward for executing a role (before the deadline) is higher if the role is executed by the human operator than when it is executed by the Automation. The adjustable autonomy strategy for the role is then identified by the I-TMDP policy π^* that prescribes the best action (=execute a role or transfer of control to some entity) for each state from the current time until the deadline Δ.

5.3 Human Operator Responsiveness

The second shortcoming of [13, 14] is that the (elicited from domain experts) responsiveness of the human operators used in determining the adjustable autonomy strategies remains unchanged across different strategies. In essence, even if according to an adjustable autonomy strategy for some earlier role there is a non-zero probability that human operator n will be in control of the role, the responsiveness

[2]For explanation purposes, we use in the following a single automation proxy as our formalism can be trivially extended to multiple independent automation proxies

of the human operator remains unchanged when the human operator is being considered to be assigned to some future role. For example, initial responsiveness of human operator n when asked to perform a role "*land a plane*" could be governed by a normal distribution $d_{s_{H_n}, a_{execute}, s_{end}} = \mathcal{N}(\mu = 4min, \sigma = 1min)$. However, if there is e.g. 20% chance that the human operator will have to land one plane (earlier role), the responsiveness of the human operator in future roles will no longer be governed by $\mathcal{N}(\mu = 4min, \sigma = 1min)$ but instead, with some distribution with a higher expected duration. In determining this updated distribution (new responsiveness of a human operator) DFAS uses transient analysis of the Markov processes [10] for the already determined adjustable autonomy strategies. The updated distribution is then fed into the I-TMDP model constructed to determine the optimal adjustable autonomy strategy for a newly arrived role.

6. IMPACT OF AFAR

Inherent to the NextGen concept of operations is the increasing collaboration between human operators and automated systems, driven by advanced models to support optimal function allocation of flight responsibilities. Optimization must be guided by the design of automated systems that exploit the unique strengths of both human and the systems they use to ensure productivity and safety. To fully realize the human-automation systems that will be central to the NextGen-Airportal, models that support function allocation must be used to drive design of these technologies early in the design process. These models must be flexible to the evolving concept of operations of the near future national airspace, and must provide innovative methods for defining function allocation that leverage existing research and tools.

To achieve this result, the AFAR approach is grounded in a comprehensive understanding of the work domain of the various NAS users and the proposed capabilities of the NextGen environment, resulting in novel methods for defining DFAS that support operationally-relevant function allocation and measurement capabilities. The AFAR approach leverages an analysis technique to generate a deep and formal understanding of the human operator's interaction with system technologies across a rich work domain. This deep understanding is used to define innovative methods for both defining and evaluating the functions that will be allocated between human pilots and automated system. The AFAR approach leverages existing work on dynamic allocation strategies for human-multiagent interaction. These adjustable autonomy strategies serve as a framework to adapt to both the needs of human operators and the requirements of NextGen-Airportal functional allocations.

In addition to providing innovative methods for identifying and developing strategies for function allocation (DFAS), the proposed approach uses iterative test cycles designed to increase in fidelity. This approach ensures that these strategies support human operators in meeting productivity and safety requirements, and that the DFAS is sensitive to changes in context due to shifts in operator characteristics, supporting technologies, flow management goals, or external conditions such as weather or other environmental disturbances. Central to these evaluations is the development of metrics, CMFA, specifically designed to assess human-automation interaction from a systems viewpoint, considering the unique and collaborative behaviors of the complex distributed network of humans and automated agents. The development of these metrics is expected to largely impact the NextGen engineering and research community, by providing valid and sensitive measures of human-automation interaction performance that extend beyond the application of traditional performance measures to the unique Airportal problem space. Finally, the developed models and methods for measurement are to be shared with NASA researchers and the NextGen research and engineering community. This applies data and conclusions drawn from knowledge elicitation sessions with NASA researchers and developers, as well as empirical evaluations conducted with simulated agents and human operators at the NASA Ames facility. The refined DFAS and CMFA will serve as a framework for design recommendations and research thrusts shared with the NASA community and research community at large.

In particular, NASA's and FAA's NextGen concept will benefit from dynamic allocations of tasks between human operators and automated systems in accordance with changing airport terminal situations. This research is expected to contribute to the NextGen-related goals by 1) developing system evaluation methods and performance metrics relevant to the Airportal environment, namely CMFA, 2) developing human-automation function allocation strategies and timings within an adjustable autonomy framework, namely DFAS, and 3) providing guidance regarding function allocation strategies and methods for assessing function allocations to designers of NextGen systems. The operational research questions to be addressed by this work, along with their associated Airportal and IIFD milestones, include:

- What roles and responsibilities allocated among humans (e.g., pilots, controllers, dispatchers) and automated systems are appropriate for the future airport environment? (milestones AP.2.A.3, AP.2.A.11, IIFD.3.1.3)

- What guidance can be provided to automation designers to support the design and evaluation of new operational concepts involving human-automation function allocations? (milestones AP.3.A.4, AP.2.A.10)

7. PILOT STUDY

We have recently completed implementation of the system as seen in Figures 3,4 and 5. The specifications of our latest pilot study implementation are as follows:

- Two participants
  - One current controller from BOS Logan Airport
- Each subject came in for a total of 3 sessions
  - First session was for training on the AFAR testbed
  - Second session involved 2 training and 2 pilot study scenarios
  - Third session involved 1 training and 2 pilot study scenarios
- Captured video
- Conducted surveys at the end of each scenario

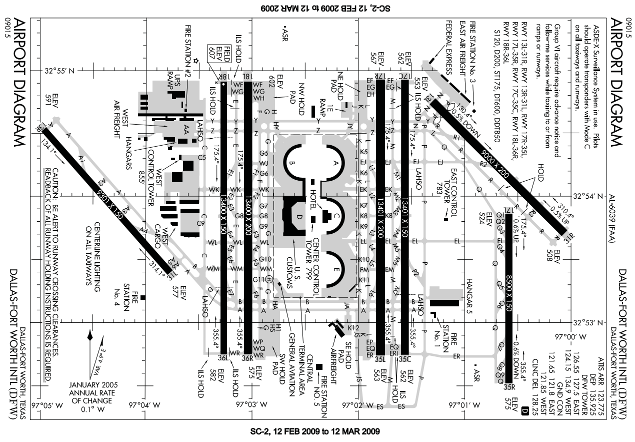

Figure 3: Diagram of Dallas Fort Worth Airport

Figure 4: The experimental setup with the participant (center) and the pseudopilots (left and right).

Figure 4 shows the entire experimental setup with the three work stations, including 2 psuedopilots and 1 participant. In addition it shows the video camera recording and the extra computers necessary to run the server and DFAS. Figure 5 shows a participant using the system. The participant's desk has reference charts on the left, flight strips on the bottom, and the dynamic display on the right. Interaction with the system was through standard oral communication messages.

In addition, we have begun analysis of our pilot study and present a some initial findings here. We evaluate the comprehensive system and looks at metrics over the AFAR system and the human participant as well. In the pilot study the agents were in two conditions, the first where transfer of control was determined a priori (Static) and second where it was calculated online using DFAS (Dynamic) as explained previously in Section 5.

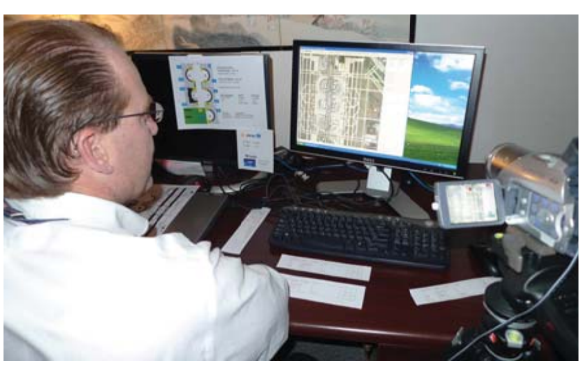

Figure 5: Ground Controller using AFAR.

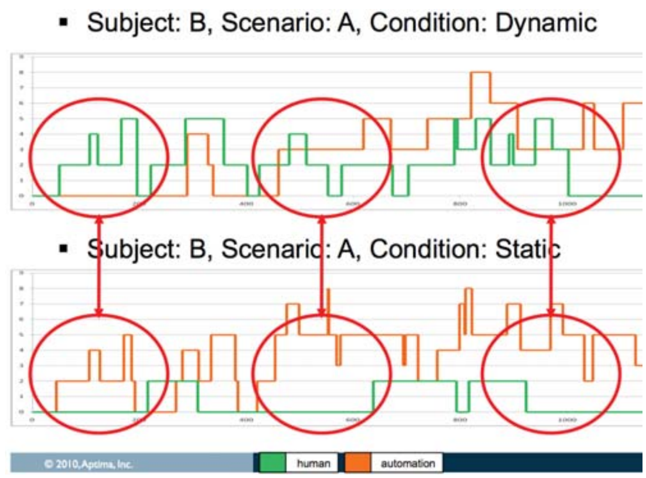

Figure 6: Step Diagram of Task Load Over Time.

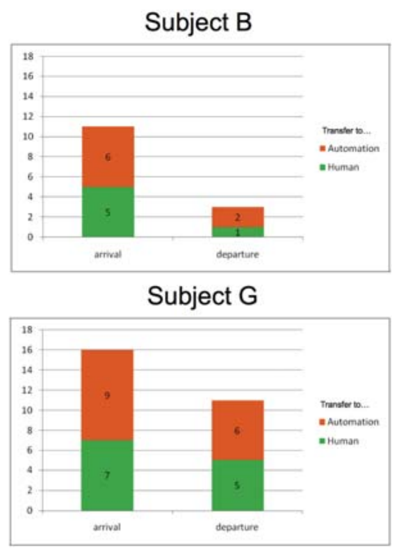

Figure 7: Total number of transfers of control.

Figure 6 shows the peformance of the DFAS algorithm in the AFAR System. Specifically, Figure 6 shows a step function of how many tasks occur over time for the human and the automation. The top graph shows Static configuration and the bottom graph shows the Dynamic (DFAS) condition. The circles are isolating points in time that highlight the differences between the static and the dynamic (DFAS) setting. DFAS achieves its desired outcome by balancing the workload between the human and automation. This results in shifting tasks to the human when the human is not busy (as seen in the circles), but also making sure that the human participant does not get overwhelmed.

In order to look further into how big of a role that the DFAS algorithm was playing, we also looked at how many transfers of control were executed by DFAS. Figure 7 shows the total number of allocations to automation and to the human participant. The departures required less transfers than the arrivals. However, these dynamic transfers are relatively balanced across the human and automation, suggesting that

Question	Mean[1]		Standard Deviation	
	Static	Dynamic	Static	Dynamic
1. I am confident that I can manage ground traffic without the automation. [2]	**18.25**	16.5	1.71	3.42
2. The performance of the automation enhanced my performance.	**16.25**	**16.25**	4.50	2.63
3. I am familiar with the operation of the automation.	**15**	13	4.16	6.48
4. I trust the automation.	**14.75**	**14.75**	3.77	4.35
5. The automation is reliable.	**17**	**17**	3.16	4.08
6. The automation is dependable.	**17**	**17**	3.16	4.08
7. The automation has integrity.	**17**	**17**	3.16	4.08
8. I am comfortable with the intent of the automation.	17	**18.25**	3.16	1.71
9. The automation is deceptive.	**18**	**18**	1.41	0.82
10. The automation behaves in an underhanded manner.	**18.5**	18.25	0.58	0.50
11. I am suspicious of the automation's intents.	18.25	**18.5**	0.96	0.58
12. I am suspicious of the automation's actions.	18.5	**18.75**	0.58	0.50
13. I am wary of the automation.	16.5	**18.5**	1.73	0.58
14. The automation's actions will have a harmful or injurious outcome.	14	**15.75**	4.62	3.95
15. I am confident in the automation.	10	**14.5**	6.53	5.80

Figure 8: Trust survey results.

DFAS did not solely take burden off the participant and give automation more functions to perform. Instead the aim of DFAS is to have automation play a role when it is better for overall human-agent performance.

In addition, we studied how these agent configurations affected the participant's view of the agent automation. Figure 8 shows the results from surveys that each of the participants completed looking at issues of trust in the automation that they were using. Bolded values are the maximum values. As seen in Figure 8, in most cases the DFAS condition was the same or better in terms of how much the participants trusted and understood the automated agent system. This is a promising result, given that a goal of our work is to build automated systems that humans will trust and use.

8. CONCLUSIONS

In this paper we have presented the AFAR project which aims to apply agents facilitate human-automation interaction in order to better define roles in the NextGen aviation environment. We believe that this technology will be beneficial to industry automation designers and help bring about more productive human-automation relationships. We have recently completed the building of the initial version AFAR system and conducted an initial pilot study. In the near future we plan to begin the full experiments and start showing the feasibility and usefulness of Dynamic Function Allocation Strategies in the NextGen environment, thereby benefitting NASA and the FAA as well.

9. ACKNOWLEDGMENT

The authors would like to acknowledge that this research was funded by NASA NRA contract number NNA08BC68C.

10. REFERENCES

[1] J. Boyan and M. Littman. Exact solutions to time-dependent MDPs. In *NIPS*, pages 1026–1032, 2000.

[2] J. M. Bradshaw, H. Jung, S. Kulkarni, S. Johnson, M. Feltovich, P. Allen, J. Bunch, N. Chambers, L. Galescu, R. Jeffers, N. Suri, W. Taysom, and A. Uszok. KAA: Policy-based explorations of a richer model for adjustable autonomy. In *Proceedings of the International Joint Conference on Autonomous Agents and Multi Agent Systems*, 2005.

[3] A. Degani and M. Heymann. Formal verification of human-automation interaction. In *Human Factors, 44*, 2002.

[4] M. A. Goodrich, T. W. McLain, J. D. Anderson, J. Sun, and J. W. Crandall. Managing autonomy in robot teams: observations from four experiments. In *SIGART Conference on Human-Robot Interaction, HRI*, 2007.

[5] Joint Planning and Development Office (JPDO). Concept of operations for the next generation air transportation system. 2007.

[6] L. Li and M. Littman. Lazy approximation for solving continuous finite-horizon MDPs. In *AAAI*, pages 1175–1180, 2005.

[7] J. Marecki, S. Koenig, and M. Tambe. A fast analytical algorithm for solving markov decision processes with real-valued resources. In *IJCAI*, January 2007.

[8] R. Parasuraman and V. Riley. Humans and automation: Use, misuse, disuse, abuse. In *Human Factors, 39*, pages 230–253, 1997.

[9] R. Parasuraman and C. D. Wickens. Humans: Still vital after all these years of automation. In *Human Factors, 50*, 2008.

[10] M. Puterman. *Markov decision processes*. John Wiley and Sons, New York, 1994.

[11] J. Reitsema, W. Chun, T. W. Fong, and R. Stiles. Team-centered virtual interactive presence for adjustable autonomy. In *American Institute of Aeronautics and Astronautics (AIAA) Space 2005*, 2005.

[12] P. Scerri, D. Pynadath, and M. Tambe. Towards adjustable autonomy for the real world. *Journal of Artificial Intelligence Research*, 17:171–228, 2002.

[13] N. Schurr, J. Marecki, and M. Tambe. Resolving inconsistencies in adjustable autonomy in continuous time (RIAACT): A robust approach to adjustable autonomy for human-multiagent teams. In *Proceedings of 7th International Conference on Autonomous Agents and Multiagent Systems*, Portugal, 2008. ACM.

[14] N. Schurr, J. Marecki, and M. Tambe. Improving adjustable autonomy for time-critical domains. In *Proceedings of 8th International Conference on Autonomous Agents and Multiagent Systems*, Hungary, 2009. ACM.

[15] N. Schurr, P. Patil, F. Pighin, and M. Tambe. Using multiagent teams to improve the training of incident commanders. In *AAMAS '06*, NY, USA, 2006. ACM.

[16] B. P. Sellner, F. Heger, L. Hiatt, R. Simmons, and S. Singh. Coordinated multi-agent teams and sliding autonomy for large-scale assembly. *Proceedings of the IEEE - Special Issue on Multi-Robot Systems*, July 2006.

[17] M. Sierhuis, J. Bradshaw, A. Acquisti, R. Hoof, R. Jeffers, and A. Uszok. Human-agent teamwork and adjustable autonomy in practice. In *In Proceedings of the seventh International Symposium on AI, Robotics and Automation in Space. Nara, Japan, 2003.*, 2003.

[18] P. Varakantham, R. Maheswaran, and M. Tambe. Exploiting belief bounds: Practical pomdps for personal assistant agents. In *AAMAS*, 2005.