INVESTIGATION OF INTELLIGENT RESOURCE MANAGEMENT FOR AVIATION COMMUNICATIONS

Eric Knoblock[1*], Michael Gasper[1], Rafael Apaza[2], Hongxiang Li[3], Zhe Wang[3], Ruixuan Han[3], Nora Schimpf[3], Nathaniel Rose[4] and Neil Adams[4]

[1]Communications and Intelligent Systems Division, NASA Glenn Research Center, Cleveland, OH, USA
[2]Aeronautics Mission Office, NASA Glenn Research Center, Cleveland, OH USA
[3]Department of Electrical and Computer Engineering, University of Louisville, Louisville, KY, USA
[4]HX5, LLC,, Cleveland, OH, USA
**eric.j.knoblock@nasa.gov*

Keywords: Aeronautics, spectrum, machine learning, AAM, resource management

Abstract

The emergence of new aerial vehicles into the airspace as part of new initiatives, such as Advanced Air Mobility (AAM), will place growing demand for spectrum resources to support airspace operations. The traditional approach of using fixed channel allocations within standard service volumes will not allow for dynamic and efficient distribution of resources based on airspace demand; consequently, a new approach to aviation spectrum management will be required to meet the anticipated needs of airspace users. The National Aeronautics and Space Administration (NASA) is investigating the application of advanced concepts to implement a novel spectrum management approach that allows for the intelligent utilization of aviation spectrum throughout the airspace while maintaining the quality of service prescribed by aeronautical standards. This technical investigation evaluates the dynamic assignment of resources for both air-ground and air-air communication links applicable to both the emerging AAM initiative as well as the existing air traffic management system. The performance of the proposed spectrum management concepts will be evaluated using a custom modeling and simulation capability that is currently under development. The implementation of these approaches is anticipated to facilitate increased spectrum utilization efficiency and enhanced airspace capacity, which will better serve the needs of future applications.

1 Introduction

New airspace concepts, such as AAM, are currently under development and will require novel air traffic management approaches that employ modern technological advancements to enable a more collaborative and inter-networked environment. Critical information exchanges in this future environment will require diverse air-ground and air-air communications among airspace users [1]. While specific communications requirements remain in formulation, it is expected that the static use of existing spectrum allocations will not meet the anticipated demand.

To address this challenge, researchers at the NASA Glenn Research Center and the University of Louisville are investigating approaches to modernize aviation spectrum management using advanced techniques in wireless communications and artificial intelligence (AI), with the goal of improved efficiency of aviation spectrum utilization and improved safety of future operations. Accordingly, this research effort is conducted in three inter-related development areas. The first focuses on the development of algorithms to predict the demand for communication resources throughout the airspace. The second area focuses on the development of learning-based resource allocation methods for both air-ground and air-air communications use cases. Lastly, the third area focuses on the development of a capability to evaluate the performance of the integrated prediction and learning-based resource allocation functionality via modeling and simulation in realistic, operational scenarios.

The remainder of this paper is organized as follows. Section 2 describes the prediction methods. Section 3 discusses learning-based resource allocation in air-ground and air-air communications, and Section 4 provides an overview of the modeling and simulation capability. Finally, Section 5 provides concluding remarks.

2 Demand Prediction

As a key component of the intelligent spectrum management concept, the prediction function considers information such as convective weather patterns, aircraft traffic flows, and airspace restrictions to determine the demand for communications resources (e.g., spectrum) throughout the airspace. The resulting “prediction products” are then used by the allocation module to dynamically assign resources for air and ground elements based on the demand. The scope of this effort includes investigation of both learning-based and non-learning-based methods of prediction. However, initial efforts have focused on learning-based predictions, and more specifically, neural network models that process time-dependent data, e.g., recurrent neural networks (RNNs).

To determine the resource demands at various levels of airspace, this research is currently evaluating three types of predictions: aircraft trajectory, airspace density, and communications demand. Training data for the neural network models are derived from historical data captures, including weather patterns and aircraft trajectories, which are available from various data repositories, including the Sherlock Data Warehouse. In addition to historical data sources, efforts are underway to develop custom aircraft tracks and airspace structures, such as those envisioned for future aeronautical concepts.

An important aspect of the prediction effort is the determination of performance metrics (such as prediction accuracy, tolerable latency, and prediction look-ahead windows, for example) to meet the service demands of the scenarios under consideration. Determining these metrics serves to support the analysis and will be captured as the prediction and allocation algorithms reach greater maturity. The following sections discuss the various prediction methods, relevant datasets, model formulations, and research completed up to this time.

2.1. Flight Trajectory Prediction

Flight trajectory prediction can provide essential knowledge supporting resource allocation. Accurate trajectory predictions localize an aircraft to the regions of airspace it may traverse, providing supplemental knowledge to determine communication resource requirements. Aircraft position can furthermore determine communication channel metrics such as link path loss, which may inform power budgeting and resource allocation.

Trajectory prediction formulations consider an aircraft's 4D coordinates (latitude, longitude, altitude, and time) and its deviations from a last-filed flight plan. Flight plan information is frequently supplemented with atmospheric data products due to the strong correlation between en-route flight reroutes and convective weather. Atmospheric data products are composed of both direct (e.g., wind vectors, humidity, air temperature) and airspace-tailored (e.g., Echo Top, Vertically Integrated Liquid) weather data, available through NOAA and MIT CIWS/CoSPA databases. Supervised deep learning offers an immediate approach to leveraging this data; hybrid-recurrent neural networks are typically developed, which are able to interpret and reduce atmospheric data into lower-dimensional features and synthesized with flight plan information to generate predictions along a sequence of flight points [2].

This effort has investigated trajectory prediction as a sequence-to-sequence problem, mapping complete flight plans and supplemental atmospheric data onto complete trajectory predictions. Initial research compared atmospheric data products and deep learning models using the hybrid-recurrent structure in Fig. 1 [3]. Echo Top was found to be a holistically accurate weather feature; attention mechanisms offered some potential benefit, and were cited for further investigation. More recent efforts, however, identify limitations as flight data is generalized across the national airspace [4]. Improving general accuracy across known and unknown flight routes remains a priority for future trajectory work, and may be relieved with additional supplemental data. Current models may significantly benefit from knowledge of airspace activity and restrictions, which heavily impact takeoff and descent. This supplemental data may take the form of NOTAMs, nearby aircraft position histories, and overall center traffic levels, and may require alternate prediction paradigms or separate models to effectively predict along each flight phase (takeoff, en-route, descent).

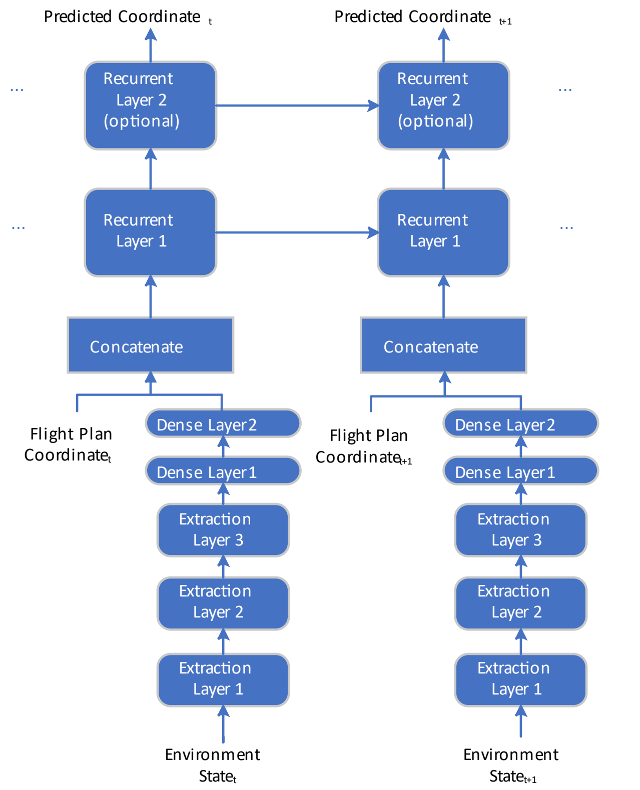

Fig 1. General Structure of Hybrid-Recurrent Architecture for Trajectory Prediction

2.2. Airspace Density Prediction

Airspace density (i.e., aircraft count) prediction considers airspace demand as aggregate forecasts, which may directly support estimates of communication resource requirements. Existing research models regions of airspace as a state transition matrix between each sector, formulating changes in sector density as a conservation-of-flow problem [5]. Airspace statistical data and density histories, such as instantaneous sector counts and daily sector transition totals available via NASA's Sherlock Data Warehouse, can be combined with NOTAMs and convective weather information to train a predictive model. While foundational research forecasts sector densities using Linear Dynamic System Models, Recurrent Neural Networks and other machine learning approaches will be necessary to accommodate airspace dynamism and expanded supplemental data.

2.3. Communications Demand Prediction

Communications demand prediction considers multiple resources (channel access duration, bandwidth, power budget, etc.) and forecasts their requirements using trends within the airspace. As communication technologies are subject to change over the span of this research, forecasting should remain agnostic of underlying communication systems. This may be achieved by separating prediction work into the forecasting of communication events within the airspace and subsequent inference of resource requirements.

The predominant challenge in this area of demand prediction is sufficiently representing communication events in terms of available data. Though airspace density is an immediate representation of demand, it does not account for the airspace complexities (such as jet route structure) impacting communication needs. Current research represents air-ground communications in terms of air traffic controller workload, as defined in prior research on airspace sectorization [6] [7]. This workload is defined in four categories:

- **Monitoring workload**, approximating a sector baseline based on its current airspace density and average transit time;
- **Coordination workload**, representing handoffs between airspace sectors and classes;
- **Conflict resolution workload**, representing aircraft adjustments to maintain minimum separation; and
- **Maneuvering workload**, representing aircraft adjustments to mitigate external factors such as convective weather.

Even with Sherlock data available to represent workloads, airspace events and workload still are not fully captured. Workload distribution varies significantly between sectors due to their nonuniform structure, and sector demands are often dependent on those of its neighbors. Federated learning approaches may provide solutions for these complexities. Data balancing techniques within federated learning, such as the Astraea framework, may help to train a global model to recognize general workload trends [8]. Capturing more sector-specific relationships and complex data dependencies may be possible with federated personalization techniques such as knowledge distillation and mixed models [9].

3 Learning-based Resource Allocation

Aviation spectrum, along with other communication resources such as time and power, must be optimally allocated to air-ground and air-air communication links to achieve the maximum spectrum utilization efficiency. As a result, spectrum optimization is usually boiled down to joint resource allocation problems.

In this section, we discuss various learning-based resource allocation approaches. First, a learning-based graph coloring approach is introduced for channel allocation in air-ground communications. Then, we discuss a deep reinforcement learning (DRL) based framework for use cases in air-ground and air-air communications.

3.1. Graph Coloring based Resource Allocation

The use of graph coloring in the context of resource allocation has been a widely-studied problem for spectrum management applications [10] [11] [12] [13] [14], and this investigation focuses on the use of a graph coloring approach to address the channel assignment problem in AAM scenarios. Graph coloring is a method that assigns colors to vertices of a graph such that no two vertices linked by an edge share the same color, and Fig. 2 provides an example of how aircraft proximity can be translated to an intersection graph whose vertices and edges are defined by regions of airspace where an aircraft may experience detrimental interference. Aircraft that are sufficiently separated in space would not present a risk of interference and would thus not be joined by an edge in the graph representation.

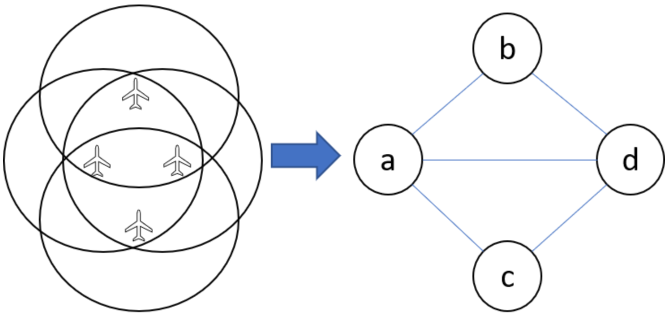

Fig 2. Example areas of potential interference around aircraft and translating it into a graph

The team has chosen to implement a graph convolutional neural network inspired by [15] [16]. As an input to the network, an adjacency matrix is formed based on the relative positioning of aircraft within a region of airspace. Given the adjacency matrix as an input, the neural network model returns a coloring solution that attempts to satisfy the graph coloring problem, or in the absence of a solution, minimizes the number of conflicts through an energy-based loss function.

For this study, a graph G is represented by N x N adjacency matrix $\mathbf{A}$ describing the relation of nodes and edges in the graph. The adjacency matrix serves as the input to the neural network. The output of the network is an N x C matrix $\mathbf{C}$ that contains the set of probabilities that a node N would be assigned any color C such that $\sum_C \mathbf{C}_{n,C} = 1$. Ideally, the probability of a single color should approach 1 in a well-trained network. In this context, reducing the loss to zero indicates a coloring solution with no conflicts. However, in cases where there may be no valid solution, the network will minimize conflicts in accordance with the loss function.

The graph network implemented in this work utilizes two topology adaptive graph convolutional layers from [17], with dimensional embedding and hidden layer sizes both set to 60. The input layer uses a ReLU activation function, and the output uses a Softmax activation to generate the output probabilities for each color. Network training is performed using a slightly modified version of the physics-inspired loss function from [15]. The loss is calculated as

$$Loss = \sum_{i}^{C} p_i^T \mathbf{A} \, p_i$$

where p_i is a length N column of $\mathbf{C}$, and the probability of assigning the i^{th} color to each node in the graph is represented by the adjacency matrix $\mathbf{A}$.

3.2. *Deep Reinforcement Learning based Resource Allocation*

Another method under investigation for intelligent resource allocation is the use of Deep Reinforcement Learning (DRL) as described in following sub-sections.

3.2.1. *Deep Reinforcement Learning:* The dynamic resource allocation problem in aeronautical communications can be modeled mathematically as a Markov Decision Process (MDP), which can be efficiently solved by DRL. By combining deep learning and reinforcement learning, DRL can automatically train the neural network to make optimal decisions. A DRL model includes several essential elements: agent, action, state, environment, reward, and policy.

- Agent: It interacts with the environment by gathering information, taking actions, and receiving reward.
- Action: It is the execution of a decision made by the agent. An action set includes a list of discrete/continuous possible actions.
- State: It is the immediate observation of the environment at a specific moment.
- Environment: It is the airspace system including air and ground elements, and any other relevant information such as available channels and their channel state information. The environment transits to a new state after any agent takes an action.
- Reward: It is the feedback of the environment to the agent. The reward function is designed to evaluate the state-action pair executed in the environment. Depending on the specific optimization problem, the reward can be either positive or negative.
- Policy: It is the state-action mapping function (i.e., brain of the agent). In DRL, a customized neural network is selected as the policy that proactively learns from the trail-and-error experiences to maximize the state-action value (Q-value). Ideally, the policy achieves the maximized accumulated reward after training.

Fig. 3 illustrates the DRL-based optimization diagram to support general air operations. The agent is a resource manager (RM) that utilizes the designated policy to make decisions. In AAGN/AACN, the environment information includes aircraft's geographical locations, channel state information (CSI), desired communication QoS, etc. At each step, the agent observes the current state of the environment. Based on that, the agent employs the policy to make decisions and take actions. Then, the environment transits to a new state. After that, the agent receives its reward from the updated environment.

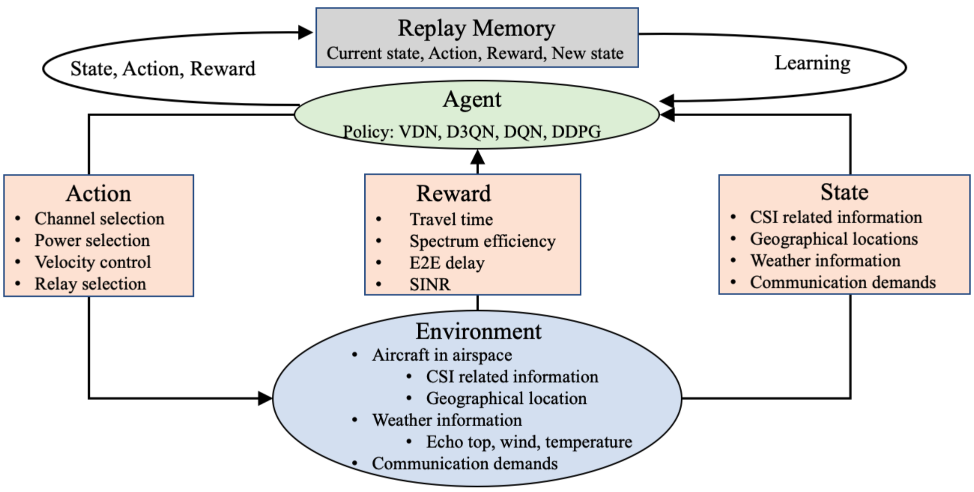

Fig 3. DRL Optimization diagram

3.2.2. *Case Study: DRL-based Resource Allocation:* Generally, aeronautical communication networks can be categorized as air-ground communication networks (AGCN) and air-air communication networks (AACN). In this Section, we first describe our recent works in AGCN and AACN, then we will discuss a new problem formulation that integrate both networks.

- AGCN: As in every aviation system, AVs need to maintain reliable bidirectional communications with ground control stations (GCS) to support aircraft safety-related operations. Specifically, communication refers to control non-payload communication (CNPC), which includes command and control data, air traffic control relay data, and sense and avoid data. Losing CNPC can lead to fatal consequences.

For AGCN, one typical application is urban air mobility (UAM) which is envisioned by NASA and the Federal Aviation Administration (FAA). UAM provides a safe and fast air transportation paradigm to support cargo and passenger mobility in populated areas which can significantly reduce traffic congestion and satisfy increasing mobility demands.

Unlike the typical data-driven application where spectrum utilization efficiency aims to maximize achievable data rate, SUE in UAM application should focus on spectrum availability and reliability so that more AVs can provide safe transportation services. In [18], the International Telecommunications Union (ITU) gives the general SUE definition, where $SUE = \frac{M}{B \cdot S \cdot T}$. The notation M represents the communication-aided useful effect, and B, S, and T represent the frequency bandwidth, the geometric space, and the time, respectively. In UAM applications, the completion of delivery is the communication-aided useful effect, and SUE should be defined as the completed deliveries per unit time divided by the product of the frequency bandwidth and the geometric space.

Given the above SUE definition, our previous work [19] introduced the mission completion time as the SUE metric. By minimizing the mission completion time, the spectrum resources can be used to support other AVs to execute delivery tasks which can inherently improve the SUE in UAM applications. Specifically, in [19], an air transportation system was considered where multiple AVs transport passengers/cargo from different sources to destinations along pre-defined paths. During the flight, the minimum communication QoS must be guaranteed at all times to ensure flight safety. The objective is to minimize the mission completion time by jointly optimizing AVs' velocity and spectrum allocation. To solve the optimization problem, we first formulate it as a Markov Decision Process (MDP) and then propose a multi-agent DRL algorithm that incorporates Value Decomposition Networks (VDN) with Dueling Double Deep Q Networks (D3QN). Simulation results reveal that the DRL-based algorithm outperforms non-learning-based solutions.

- AACN: AACN is built upon a group of aircraft that are able to conduct wireless communication. It allows multiple air-air (A2A) communication pairs to exchange information at the same time. AACN has multiple applications, one of them is to support the multi-access edge-computing (MEC). Each aircraft is equipped with computing capabilities that is able to collect various sensory data and send the data to the end users with limited local computing capabilities. To conduct different sensing applications, there are various communication requirements (e.g., data rate, SINR, packet loss rate) and frameworks for demand (e.g., single-hop, multi-hop).

Any two aircraft nodes that are within the communication range can establish a direct A2A link, named as single-hop communication. If two aircraft are too far away to send information to each other directly, other aircraft can proactively serve as relays to support the communication. This kind of structure is named as multi-hop communications. We have been investigated on DRL-based resource allocation in single-hop [20], [21]and multi-hop A2A communication scenarios [22]. We consider limited available resources (including spectrum, power, and relays) in model development and training. In the spectrum scarcity AACN scenario, a frequency channel may be used by multiple A2A links, resulting in co-channel interference.

In single-hop communications, we consider joint channel selection and power control to maximize the weighted sum spectral efficiency (WSSE) [20], [21]. We design a distributed and dynamic DRL-based algorithm to solve this joint optimization problem. Specifically, two policies including Deep Q-Network (DQN) and DQN + Deep Deterministic Policy Gradient (DDPG) are employed to exploit the local information from its neighbours and learn to make optimized decisions. DQN can work with discrete action space, while DDPG can work with continuous space that is suitable for power allocation. The distributive structure makes it scalable to large networks, which can be widely applied in MEC application.

In multi-hop communications, packets could be sent from different sources and received by different destinations in the same network, bringing us a Multiple Sources and Multiple Destinations (MSMD) routing problem. Thus, the spectrum access and routing decisions are jointly considered to guarantee a reliable and efficient A2A communication. In paper [22], E2E delay is one of the important indicators to evaluate the A2A communication performance. E2E delay is the accumulated delay at each hop, which contains queuing delay and transmission delay. Queuing delay is the wait duration from the time when one packet arrives at a node and when it is served. Transmission delay is the duration to deliver one packet from one node to another. DQN policy is utilized to find an optimal routing and channel selection strategy that minimizes the E2E communication delay.

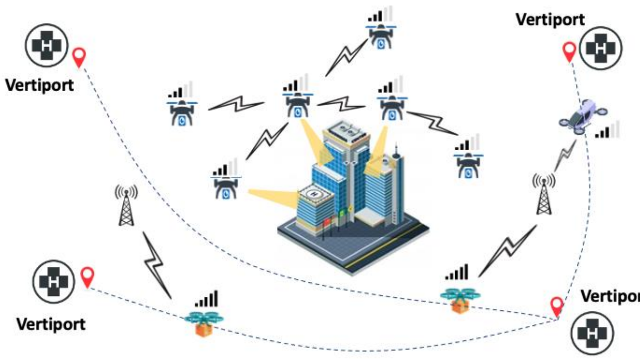

Fig 4. RL-based Resource Allocation Use Case

In this case study, we consider a scenario where two types of AVs coexist in the airspace, as shown in Fig. 4. On one hand, there are multiple moving AVs that perform UAM applications transporting cargo and passengers from sources to destinations. During the flight, these AVs need to connect with ground control stations to ensure safety, composing AGCN. Their objective is to minimize the mission completion time. On the other hand, a swarm of AVs is operated in the network for other applications e.g., surveillance, sensing, or relay tasks. Within a swarm, effective communication among these AVs is required to coordinate and achieve the common objective, which imposes AACN between AVs. Therefore, finding out the best routing strategy to minimize the packet transmission delay is their objective. By considering both types of AVs, such a use case imposes a hybrid communication network that includes both AGCN and AACN.

Let **Q**, **A**, **P**, and **L** denote the set of variables for AVs trajectories, channel allocation, transmitting power, and relay selection. The channel allocation is represented by a binary indicator where $a_{n,k}(t) = 1$ if channel k is allocated to AV n at time t. Otherwise, $a_{n,k}(t) = 0$. For AACN, we assume there is a total of M sources-destination pairs, and **L** contains all the next relay selections for M pairs. For AGCN, we assume a set of UAM AVs, denoted by $\mathcal{N}$ in the system. In this hybrid communication network, each AV communication QoS should be above a threshold to ensure the reliability of communication links. Specifically, we consider the signal-to-interference-plus-noise ratio (SINR) as the communication QoS metric which is a function of the channel allocation AVs' trajectories **Q**, channel allocation **A**, transmitting power **P**, and relay selection **L**.

For AGCN, Let T_n be the mission completion time of AV n. For AACN, let M and D_m denote the set of source-destination pairs and the E2E delay of AV m. Particularly, D_m is defined as the summation over each hop delay that includes queuing delay and transmission delay. The overall SUE maximization problem can be formulated as a joint mission completion time and E2E delay minimization problem:

$$P1: \min_{\mathbf{P},\mathbf{A},\mathbf{Q},\mathbf{L}} \alpha \sum_{n\in\mathcal{N}} T_n + \beta \sum_{m\in M} D_m$$

s.t. $$C1: \gamma_n(t) \geq \gamma_{qos}, \forall n \in \mathcal{N}$$

$$C2: \gamma_m(t) \geq \gamma'_{qos}, \forall m \in M$$

$$C3: a_{n,k}(t), a_{m,k}(t) = \{0,1\}$$

$$C4: \sum_{n\in\mathcal{N}} a_{n,k}(t) \geq 1, \forall n \in \mathcal{N}$$

$$C5: \sum_{m\in M} a_{m,k}(t) \geq 1, \forall m \in M$$

$$C6: p_n(t) \leq p_{max}, \forall\, p_n \in \mathbf{P}$$

$$C7: p_m(t) \geq p'_{max}, \forall p_m \in \mathbf{P}$$

$$C8: \| \mathbf{q}_n(t) - \mathbf{q}_n(t-1)\|_2 \leq V, \forall n \in \mathcal{N}$$

$$C9: \| \mathbf{q}_n(t) - \mathbf{q}_{n'}(t-1)\|_2 \geq d_{min}, n \neq n'$$

where α and β represent the weight of two metrics. Notation $\gamma_n(t)$ represents AV n' SINR at time t and constraints $C1$ guarantee their minimum communication QoS requirement γ_{qos} can be satisfied at any time for all AVs. Similarly, constraint $C2$ guarantees the communication quality of AVs in the swarm. Constraint $C3$ specifies the binary channel allocation constraint. Constraints $C4$ and $C5$ indicate that each AV will have spectrum channels to support its communication. Constraints $C6$ and $C7$ are the maximum power constraint for two types of AV. Constraint $C8$ is the velocity constraint for UAM AV where V is AV's maximum moving velocity. Constraint $C9$ enforces collision avoidance for AVs. Note that problem ($P1$) is a non-convex, multi-stage combinatorial optimization problem, and finding the optimal solution via conventional optimization techniques suffers from the curse of dimensionality. Recently, DRL has drawn significant attention to solve such problems. Based on the MDP, DRL utilizes deep neural networks to handle the large-scale optimization space, which makes it a viable solution for finding the optimal solution.

4 Modeling and Simulation

To evaluate the performance of the integrated prediction and learning-based allocation functionality, the team is developing a modeling and simulation capability that consists of computational tools and airspace environment visualization. The following paragraphs provide an overview of the features of the airspace modeling toolset.

4.1. Toolset Architecture

The modeling and simulation toolset implements a modular software architecture that allows for integration with external elements (e.g., flight simulators, live aircraft data feeds) that may be used to assist in the modeling and evaluation of airspace scenarios. The high-level software architecture is shown in Fig. 5.

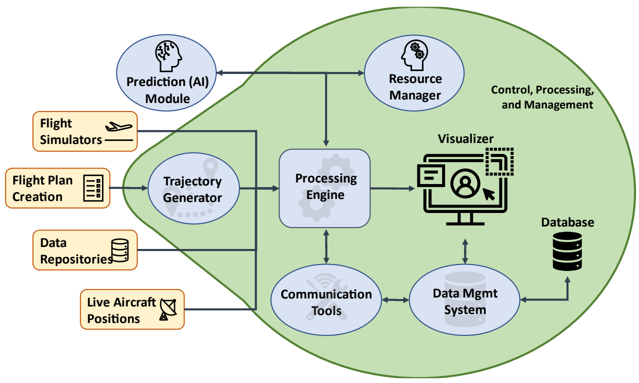

Fig 5. High level overview of the modelling and simulation modularization, including connections to external interfaces

The processing engine evaluates simulation objects and updates their status based on pre-defined timesteps or real-time data streams. The properties and locations of the simulation objects are updated at each step and then redrawn within the visualizer. After, the processing engine waits for the next time step and repeats the process. Data processing at each time interval utilizes various software packages and tools. Geospatial calculations determine aircraft locations and also the length of the communications path between transceivers. Additional functionality allows for the determination of link budgets, signal-to-interference noise ratios, and link service quality.

Aircraft flight track data imported into the simulation is pre-processed before being rendered by the visualizer. The trajectory calculations include interpolation of aircraft position throughout its flight path, determination of the aircraft's current location in terms of airspace sector and heading, and determination of link budgets and interference information. The pre-processed aircraft trajectory data can then be stored in a database and later retrieved for simplified processing and improved performance in subsequent simulations.

Live data feeds are processed similarly; although, the processing is done live, and the simulation updates after each time interval. Interfacing with the X-Plane flight simulator software is an example of live data processing where the simulation tool receives live updates based on the simulated aircraft trajectory, which is then rendered by the visualizer.

4.2. Toolset Features

The custom software toolset allows a user to create custom airspace scenarios to evaluate the performance of advanced learning-based methods by integrating with external machine learning models. Simulation results are then presented graphically, allowing the user to evaluate numerical results or observe the scenario playback in the visualizer. Details on many of the key features of this software are described below. These features allow for rapid development of novel airspace scenarios, providing user-friendly interfaces to create the necessary components for the evaluation of resource management algorithms in an arbitrary airspace configuration.

4.2.1. *Link Analysis:* As the processing engine updates active objects, the data is further processed by a set of communication tools that perform link budget calculations. Given a ground-to-aircraft link at a specific frequency, the controller determines all possible interfering links and calculates the signal-to-interference ratio as an indicator of service quality. The intended link (ground-to-air) is compared to the aggregation of all interfering co-channel and adjacent channel aircraft.

4.2.2. Visualization: Simulation visualization allows for the dynamic display of all objects within a scenario, and graphical rendering of the visualizer is built upon the NASA WorldWind Software Development Kit [23]. The interface allows the user to easily manipulate and observe the simulation while in progress, as well as develop custom scenarios, including definition of custom airspace volumes, ground stations, flight paths, communication system parameters, and channel allocation method. Once the simulation is complete, the results are captured as time-series data values that can be plotted within the tool for quick evaluation or can be exported to a set of formatted data files for more detailed analysis in external software packages.

4.2.3. *Custom Airspaces:* As a part of simulation scenario development, custom airspace configurations can be constructed via the graphical placement of a number of reconfigurable polygon objects that (in combination) define the desired airspace volume geometries or corridor structures. Ground station locations are similarly placed, thereby creating an association (or link) between an airspace and a service volume that define the areas of coverage for the customized scenario. These user-defined constructs can be stored in a database for future use, and in subsequent simulations, the user has the option to either use an archived configuration, edit an existing configuration, or create a new airspace.

4.2.4. *Flight Track Development:* Aircraft objects within the simulation are defined via flight tracks, which consist of a number of timestamped points in three-dimensional space with associated velocity information. The scenario development function allows a user to 1) define aircraft parameters, such as ascent rate, cruising altitude, and velocity, 2) select any number of intermediary waypoints that define the trajectory, and 3) complete and generate a full flight track from departure point to destination point for use in simulation. Additional options include the ability to import existing flight tracks, which can be either user-defined, generated via an external flight simulator, or from actual commercial flight tracks. Simulated flight data is accomplished via an interface to the X-Plane flight simulation software. Actual flight data may be 1) historical data acquired from repositories such as the Sherlock Data Warehouse or 2) live data streams of aircraft tracks from the FAA's System-Wide Information Management Flight Data Publication Service (SFDPS) [24].

4.2.5. Line-of-Sight Evaluation: The line-of-sight (LOS) evaluation capability allows a user to select any two points in 3D space on the visualizer map and determine if there are any obstructions between the objects, either due to terrain or interfering structures such as buildings. Terrain data is included as a part of WorldWind SDK, and the LOS evaluation tool uses the highest-available resolution of terrain within a particular region. Building data for urban areas is obtained via querying the OpenStreetMaps API. This is accomplished via the user selecting a set of rectangular regions in space which defines the area of interest that is used as a query for any available information on obstructions from buildings and other structures. The data returned from the API are rendered as simple polygons within the visualizer. The LOS capability can be used for a variety of purposes, including site planning and radio coverage analysis for critical areas. The LOS evaluation can also be executed as part of the link analysis function within a full simulation.

4.2.6. *Convective Weather Modeling:* Convective weather modeling allows for the rendering of 3D cloud echo top with dynamic resolution depending on the level of zoom. Historical echo top data can be quickly loaded into the tool and displayed at any timestamp throughout simulation scenario. Weather data may be included as an airspace environment data item that is provided to learning-based models to aid in predictions of aircraft trajectories. Future development efforts include rendering of weather data from the FAA SWIM feed.

4.2.7. *Machine Learning Connectivity:* This capability allows a user to establish an interface between the toolset and external learning-based models created using Python machine learning libraries. As part of that exchange, the toolset generates a string of data that represents the state of the simulation at each time step and transmits that data over a TCP socket connection to the receiving Python software. The data may then be processed within a series of Python scripts, including learning-based resource management models, where the resulting resource allocation information is returned to the toolset to further evaluation within the simulation.

5 Conclusions

Spectrum depletion is a concern for aviation and will serve as an impediment to the modernization of the evolving air transportation system. This research investigates the use of advanced concepts to implement an intelligent resource allocation capability for enhanced spectrum utilization efficiency. This proposed new approach may better meet the anticipated spectrum needs of future applications. Follow-on efforts will focus on implementation and evaluation of the concepts, as well as potential technology infusion opportunities.

6 References

[1] R. D. Apaza, E. J. Knoblock and H. Li, “A New Spectrum Management Concept for Future NAS Communications,” in 2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC), 2020.

[2] Y. Pang, H. Yao, J. Hu and L. Yongming, “A Recurrent Neural Network Approach for Aircraft Trajectory Prediction with Weather Features From SherlocK,” in AIAA Aviation 2019 Forum.

[3] N. Schimpf, E. J. Knoblock, Z. Wang, R. D. Apaza and H. Li, “Flight Trajectory Prediction Based on Hybrid- Recurrent Networks,” in 2021 IEEE Cognitive Communications for Aerospace Applications Workshop (CCAAW), 2021.

[4] N. Schimpf, E. J. Knoblock, Z. Wang, R. D. Apaza and H. Li, “A Generalized Approach to Aircraft Trajectory Prediction via Supervised Deep Learning,” IEEE Transactions on Aerospace and Electronic Systems, Under Review.

[5] B. Sridhar, T. Soni, K. Sheth and G. Chatterji, “An Aggregate Flow Model for Air Traffic Management,” in AIAA Guidance, Navigation, and Control Conference and Exhibit.

[6] A. Yousefi, Optimum Airspace Design with Air Traffic Controller Workload-based Partitioning, George Mason University, 2005.

[7] O. Hakan and Y. Kadriye, “A new approach to air traffic controller workload measurement and modelling,” Aircraft Engineering and Aerospace Technology: An International Journal, vol. 83, pp. 35-42, 2011.

[8] D. Moming, L. Duo, C. Xianzhang, T. Yujuan, R. Jinting, Q. Lei and L. Liang, “Astraea: Self-Balancing Federated Learning for Improving Classification Accuracy of Mobile Deep Learning Applications,” in 2019 IEEE 37th International Conference on Computer Design (ICCD), 2019.

[9] K. Viraj, K. Milind and P. Aniruddha, “Survey of Personalization Techniques for Federated Learning,” in 2020 Fourth World Conference on Smart Trends in Systems, Security and Sustainability (WorldS4), 2020.

[10] C. Lee, S.-M. Oh and A.-S. Park, “Interference avoidance resource allocation for D2D communication based on graph-coloring,” in 2014 International Conference on Information and Communication Technology Convergence (ICTC), 2014.

[11] D. Tsolkas, E. Liotou, N. Passas and L. Merakos, “A graph-coloring secondary resource allocation for D2D communications in LTE networks,” in 2012 IEEE 17th international workshop on computer aided modeling and design of communication links and networks (CAMAD), 2012.

[12] X. Cai, J. Zheng and Y. Zhang, “A graph-coloring based resource allocation algorithm for D2D communication in cellular networks,” in 2015 IEEE International Conference on Communications (ICC), 2015.

[13] S. Basloom, A. Nazar, G. Aldabbagh, M. Abdullah and N. Dimitriou, “Resource allocation using graph coloring for dense cellular networks,” in 2016 International Conference on Computing, Networking and Communications (ICNC), 2016.

[14] E. Pateromichelakis and K. Samdanis, “A graph coloring based inter-slice resource management for 5G dynamic TDD RANs,” in 2018 IEEE International Conference on Communications (ICC), 2018.

[15] M. J. A. Schuetz, J. K. Brubaker and H. G. Katzgraber, “Combinatorial Optimization with Physics-Inspired Graph Neural Networks,” Nature Machine Intelligence, vol. 4, no. 4, pp. 367-377, 2022.

[16] M. J. A. Schuetz, J. K. Brubaker, Z. Zhu and H. G. Katzgraber, “Graph Coloring with Physics-Inspired Graph Neural Networks,” in arXiv preprint arXiv:2202.01606, 2022.

[17] J. Du, S. Zhang, G. Wu, J. M. Moura and S. Kar, “Topology adaptive graph convolutional networks,” in arXiv preprint arXiv:1710.10370, 2017.

[18] International Telecommunication Union, “SM.1046 : Definition of spectrum use and efficiency of a radio system,” 2017. [Online]. Available: https://www.itu.int/rec/R-REC-SM.1046-3-201709-I/en.
[19] R. Han, H. Li, E. J. Knoblock and R. D. Apaza, “Dynamic Spectrum Allocation in Urban Air Transportation System via Deep Reinforcement Learning,” in 2021 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC), 2021.
[20] Z. Wang, H. Li, E. J. Knoblock and R. D. Apaza, “Joint Spectrum Access and Power Control in Air-Air Communications - A Deep Reinforcement Learning Based Approach,” in 2021 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC), 2021.
[21] Z. Wang, H. Li, E. J. Knoblock and R. D. Apaza, “Learning based edge computing in air-to-air communication network,” in 2021 ACM/IEEE 6th Symposium on Edge Computing (SEC), 2021.
[22] Z. Wang, R. Han, H. Li, E. J. Knoblock, R. D. Apaza and M. R. Gasper, “Deep Reinforcement Learning based Routing in an Air-to-Air Ad-hoc Network,” in 2022 40th Digital Avionics Systems Conference (DASC), 2022.
[23] National Aeronautics and Space Administration, [Online]. Available: https://worldwind.arc.nasa.gov/. [Accessed 19 January 2022].
[24] Federal Aviation Administration, “SWIM Flight Data Publication Service (SFDPS),” 2021. [Online]. Available: https://www.faa.gov/air_traffic/technology/swim/sfdps/. [Accessed 25 January 2022]