DESIGN OF AN AUTOMATED FLIGHT PLANNING AND FLEET ALLOCATION TOOL THAT OPTIMISES THE DELIVERY OF WATER FROM WATER SOURCES

Master's Thesis

By Jesse Stewart

Completed November 2019

Under the supervision of Professor Karl Sammut

Submitted to the College of Science and Engineering in partial fulfilment of the requirements for the degree of Bachelor of Engineering (Robotics) (Honours), Master of Engineering (Electronics)

Flinders University, Adelaide

Declaration of Originality

I declare that this thesis does not incorporate without acknowledgement material previously submitted towards my degree, and to the best of my knowledge does not contain any material written by another person expect where appropriate reference is made.

Jesse Stewart

November 2019

Acknowledgements

I would like to sincerely acknowledge my academic supervisor Professor Karl Sammut, for the guidance and many opportunities provided since 2015, without whom I would not be completing my degree to become the engineer I am today.

Abstract

Wildfires are becoming an increasingly common occurrence across Australia with both fire intensity and surface area impact for each fire event continuing to grow. As the frequency of wildfires increase and rural firefighting resources are strained the benefit brought by aerial fire suppression is becoming more critical for the timely containment of wildfires. The highly dynamic environment encountered during aerial suppression operations requires a system that can adapt to the changing conditions with minimal impact to operational efficiency. Available aircraft must be assigned to currently active fire fronts, have their flight paths optimised to minimise the distance travelled, service the fire fronts while meeting the objectives of the suppression efforts, and utilise all available water sources to minimise their turnaround time. The challenges faced during management of aerial suppression resources amount to a fleet allocation and flight planning problem.

Building upon analysis of prior work in the fields of wildfire containment strategies, fleet management, vehicle routing, constraint-based planning, and flight planning, a fleet allocation and flighting planning system is proposed. The proposed system uses a layered approach allowing for a series of task specific algorithms to be integrated to form a unified solution. The fleet allocation layer implements a fire characterisation protocol to inform aircraft assignments to active fire fronts. The flight cost approximation layer generates an approximate flight cost by accounting for airspace to avoid built-up areas, represented by cost regions. Finally, the flight planning layer adopts an augmented greedy algorithm for the initial assignment of aircraft followed by an analysis of the expensive flight assignments, optimising where possible through assignment swapping.

The fleet allocation and flight planning system developed was implemented behind a prototype user interface to demonstrate the systems capabilities. The user interface implemented consists of a map as the primary interface window allowing for entities such as water sources, airbases, and fires to be introduced into the system through a mouse click. In addition to the testing and demonstration capability it affords, the user interface provides a unique and efficient tool for the generation of test data. The resulting system provides a

reliable means of allocating available aircraft following an assessment of the fire, generating an approximate flight cost likely to be incurred, and forming a flight plan for the distribution of aircraft to fires and water sources. Incorporated alongside, where possible, the system provides a means of strategically utilising short-term resources to maximise the efficiency of aerial suppression resources.

Ultimately, this project was able to identify through both an extensive research phase and vigorous development, the key components of an intelligent fleet allocation and flight planning system. Using the insights gained, a novel system was implemented for allocation and planning of aerial suppression aircraft. The solution was designed such that it could be utilised separate from the aerial firefighting example case and is domain independent. Deliberate decisions made throughout the development process has afforded the system the potential to form the basis for future intelligent fleet management applications.

Table of Contents

List of Figures

List of Abbreviations

AOB – Air Observer

AAS – Air Attack Supervisor

SEAT – Single Engine Air Tanker

SACFS – South Australian Country Fire Service

NOB – Nominated Operational Base

AGL – Above Ground Level

CAR – Civil Aviation Regulation

VRP – Vehicle Routing Problem

VRPPD - Vehicle Routing Problem with Pickup and Delivery

VRPTW - Vehicle Routing Problem with Time Windows

VRPMOB – Vehicle Routing Problem with Multiple Overlapped Batches

CVRP – Capacitated Vehicle Routing Problem

CSP – Constraint Satisfaction Problem

DCSP – Dynamic Constraint Satisfaction Problem

WPF – Windows Presentation Foundation

1 Introduction

1.1 Background

Aerial firefighting aircraft integrated into wildfire suppression efforts can often be the differentiating factor in the timely containment of wildfires. While aerial firefighting aircraft are frequently utilised for wildfire suppression to great effect, attacking a fire solely from the air has proven not to be an effective means of extinguishing wildfires. Aerial firefighting resources have been shown to provide the most significant impact when used to compliment ground-based operations. This includes tasks such as establishing containment lines, reducing fire front intensity, and rapid defence of vulnerable and critical ground targets, without which would considerably increase the strain placed on ground suppression operations [1].

Aircraft have long been associated with wildfire firefighting, both in reconnaissance and firebombing capacities. Having first emerged in the mid-1910s in the united states, the use of aircraft for monitoring and reconnaissance purposes was adopted in Australia in the 1930s. The ability to spot wildfires in their infancy and monitor their progress proved to be a valuable resource for the local firefighters. While the first recorded aerial suppression experiments in Australia date back to the 1930s, it was not until the conclusion of WW2 experiments continued using surplus military resources. Finally, following decades of experiments and regular feasibility studies states began to establish aerial firefighting units in the 1960s [2]. Initially, the idea was conceived to limit damage from fire fronts left inaccessible by grounds crews due to difficult terrain. The introduction of aerial firefighting provided a means of minimising wildfire damage where the only option prior would have been to let them burn. Following the initial development and the early years of operation, aerial firefighting aircraft soon became a mainstay in the firefighting arsenal for managing wildfires.

For the past twenty years research focus has shifted from the feasibility and integration of aerial firefighting resources to modernising the techniques and optimising resource usage, all contributing to improving the effectiveness of aerial firefighting operations. The effectiveness of aerial firefighting is a complex equation with countless factors impacting every fire event. Factors such as the distance between the fire and accessible water sources, time to replenish, retardant material, terrain, weather conditions, aircraft type, payload size, proximity to drop site the aircraft can achieve, pilot experience, fire size, fuel loading, fuel type, obstacles such as dense canopies, and timely follow-up by experience ground crews all effect the overall effectiveness of the aerial suppression effort [3].

1.1.1 Firefighting Aircraft

There are countless variations of firefighting aircraft in use around the world, both fixed wing and rotary wing, and many different sizes of each. The aircraft used and the firefighting specific equipment are driven by factors including short term repurposing of existing operational aircraft, military and commercial aircraft surplus to requirements, aircraft that capitalise on a local trained pilot community, and aircraft specifications suited to the region of operation.

Often where vast distances between operational bases, fire fronts and water sources are encountered fixed wing aircraft are favoured due to the range and transit speed they offer. Instances where greater manoeuvrability at the drop zone is required, the terrain is more challenging and conventional landing facilities are inconvenient to access during operations helicopters are the logical choice.

The Air Tractor 802A is an example of an aircraft used extensively around the world and regularly adapted for aerial firefighting. The Air Tractor line of aircraft was initially designed, marketed and operated as a crop-dusting aircraft for the agriculture market. With its payload capacity, handling during flight, minimal take-off and landing requirements, unsealed runway capability, fuel economy, and range, the aircraft possess the necessary qualities for firefighting aircraft. Today, the Air Tractor 802A has the largest capacity of any single engine tanker aircraft. It can be purchased from Air Tractor as the 802F, a ready to fly firefighting aircraft, and can switch between crop-dusting, seeding and firefighting duties to meet local demand [4].

A by-product of their widespread adoption is the large cohort of qualified, highly trained, and experienced agricultural pilots who fly the aircraft year-round. Other than many hours of experience flying the specific aircraft, agricultural flying provides the best simulation of the conditions and flying style encountered at active fire fronts. Flying at high speeds, at low altitudes, and in close proximity to ground based hazards, all while managing the deployment of a payload ensures pilots operating the aircraft at fire fronts have suitable experience in challenging conditions.

The benefit of an aircraft platform with the capability to transition between firefighting, crop-dusting, and seeding operations all with a common pilot base is the commercial prospects for aviation operators. Aviation operators can maximise aircraft utilisation across a year by contracting out aircraft to agriculture operations and government agencies during their peak times. Through the diversification of contract opportunities, greater numbers of Air Tractor 802 aircraft can be maintained for when they are needed most.

To assist during the conceptual phase of the project along with implementation decisions made during the project's development a select sample of firefighting aircraft were considered. The aircraft from the sample were used to inform the initial algorithm design, front-end prototype, and the operational constraints incorporated into the system. Aircraft were selected for the sample to provide a good insight into the aircraft used throughout Australia and many other countries around the world. The aircraft chosen were the two primary aerial firefighting platforms used by the South Australian Country Fire Service (SACFS) and an example of a smaller scale medium lift helicopter. The SACFS was chosen as it is a local organisation and provided a good snapshot of the firefighting procedures and resources used Australia wide. The SACFS was also one of few firefighting agencies operational procedures could be obtained from.   The primary aircraft used by the SACFS for immediate firebombing operations were the Air Tractor 802A and the Erikson Aircrane. The additional aircraft considered to round out the sample was the Sikorsky S-70 Firehawk, a common helicopter with capability to be retrofitted for aerial firefighting.

Air Tractor 802A/F - Single Engine Air Tanker (SEAT)

The Air Tractor 802A and its predecessors have been the primary Single Engine Air Tanker (SEAT) used by the SACFS for decades for aerial firefighting. Its payload capacity, long ranges and favourable inflight capabilities lends the platform to conditions found Australia wide. Due to the relative size of the country and the land area each states fire service must monitor and respond to fire events within, fixed wing platforms are ideal. The dual application nature of the platform described previously, and the extensive agricultural aviation industry reinforces the aircraft selection.

The aircraft is capable of deploying water, foam or chemical retardants to the fire front depending on the operation being undertaken. The payload drop is controlled through a computerised trap door mechanism. The computer-controlled trap door provides options as to how the payload is distributed. The trap door can be partially opened to control the density of the payload deposited, also known as the coverage level. It also allows for the payload to be partially deployed, allowing for split drops. The Air Tractor 802A in its firefighting configuration is shown below in Figure 1-1.

Figure 1-1: New South Wales Rural Fire Service Air Tractor 802A [5].

Table 1: Air Tractor 802A Performance Specifications [6]

Aircraft: Air tractor 802A – AT-802A	
Platform type	Fixed wing
Role	Air tanker
Water tank capacity	3100 L
Mid-mission refill capable	No
Fuel tank capacity	961 L
Fuel type	Jet-A1
Range	982 km
Working speeds	209 – 257 km/h
Cruise speed	356 km/h
Empty weight	2951 kg
Maximum take-off weight	7257 kg
Runway requirements (TO/L)	608 m
Landing speed	100 km/h

Erickson S-64E Aircrane - High Volume Helicopter

To supplement the SEAT aircraft during peak fire incident times and at large fire incidents the Erickson Aircrane is often utilised by the SACFS. In its base form the Erickson Aircrane is a multi-purpose twin jet engine heavy lift aircraft that operates with three crew. Originally adapted from the Sikorsky S-64 airframe, it provides a larger volume tanker option without being prohibited by short runways used by the aforementioned SEATs. The Erickson Aircrane in its firefighting configuration, with its accompanying 7500 litre tank fitted at the centre of the aircraft and its snorkel allowing for mid-air refilling of the tank is shown in Figure 1-2.

Figure 1-2: Erickson Aircrane with refilling snorkel suspended [7].

In addition to providing a larger tanker option rotary wing aircraft enable access to impromptu water sources such as lakes, damns, rivers, reservoirs, and even single use retardant refill tanks as shown in Figure 1-3. Another beneficial attribute includes its manoeuvrability around fire fronts resulting greater control during retardant drops.

Image removed due to copyright restriction. Original can be viewed online at: https://fireaviation.com/2018/01/10/erickson-receives-contract-to-build-two-aircranes-for-south-korea/

Figure 1-3: Erickson Aircrane refilling from onsite tank [8].

While the Erickson Aircrane brings many benefits, it also brings several negatives that require serious consideration. Many of these negatives are associated with increased operational costs and additional resource requirements. These costs include the maintenance of multiple complex mechanical assemblies per aircraft, consumables such as aviation fuel and oil, the personnel requirements before, during and post flight, and the cost of contracting the aircraft. It is these significant costs that may limit the dispatch of the helicopter to only critical fire fronts.

Table 2: Erickson Aircrane Performance Specifications [9].

Aircraft: Erickson S-64E Aircrane	
Platform type	Rotary wing
Role	Air tanker
Crew	3
Water tank capacity	7500 L
Mid-mission refill capable	Yes
Fuel tank capacity	4900 L
Fuel type	Jet-A1
Range	370 km
Working speeds	0-212 km/h
Empty weight	8724 kg
Maximum take-off weight	19050 kg
Runway requirements (TO/L)	Helipad
Landing speed	-

Sikorsky S-70 Firehawk - Medium Lift Helicopter

While fixed wing aircraft such as the SEATs used in Australia suit some conditions, they do not fit all conditions. Another common approach to firefighting aircraft is medium lift helicopters, using either Auxiliary water tanks such as the configuration shown in Figure 1-4, or suspended water vessels. The Sikorsky S-70 Firehawk is an ideal example of such an aircraft with a comparable water tank capacity to the seats while maintaining mid-flight refill capabilities and increased manoeuvrability and drop control.

The medium lift helicopter sized to provide comparable operational capability to the SEATs suffers the same negative qualities inherent to all helicopters. The costs associated with their consignment, consumables, operation and crew resources are all greater than the requirements for SEAT aircraft. While the medium lift helicopters provide similar payload capacities with similar airspeeds a major short coming is still the range of the aircraft.

Just as the SEATs have an additional use when not required for firefighting operations, medium lift helicopter fleets too have an important role when not engaged fighting fires. For example, in the case of the fleet of S-70 aircraft operated by the LA County Fire, firefighting is one aspect of their role. When not undertaking firefighting operations the fleet of aircraft is used conducting search, rescue, retrieval and surveillance operations. As with the SEATs, the broad range of potential applications enables a sizeable fleet of aircraft to be maintained such that it is an effective resource during wildfire incidents.

Image removed due to copyright restriction. Original can be viewed online at: https://fireaviation.com/2017/12/14/cal-fire-receives-final-approval-to-purchase-new-sikorsky-s-70i-helicopters/

Figure 1-4: A LA County Fire Department Sikorsky S-70 Firehawk water drop demonstration [10].

Table 3: Sikorsky S-70 Firehawk Performance Specifications [11].

Aircraft: Sikorsky S-70 Firehawk	
Platform type	Rotary wing
Role	Air tanker
Crew	2
Water tank capacity	3700 L
Mid-mission refill capable	Yes
Fuel tank capacity	1360 L
Fuel type	Jet-A1
Range	465 km
Working speeds	0-360 km/h
Empty weight	5350 kg
Maximum take-off weight	10000 kg
Runway requirements (TO/L)	Helipad
Landing speed	-

1.2 Problem Definition

Timely initial response by aerial suppression to wildfires either for the purpose of directly attacking the fire front or establishing a fire break can prove critical to wildfire containment. Equally, sustained aerial suppression operations stand to benefit from rapid aircraft turnaround times. Current ad-hoc flight planning and resource allocation methods coupled with the underutilisation of short-term resources provide suboptimal efficiency from aerial suppression resources. These inefficiencies result in greater fire damage, higher operational costs, and extended volunteer time.

1.3 Research Question

Building upon the problem introduced above, the thesis aims to explore the following topic:

Design of an automated flight planning and fleet allocation tool that optimises the delivery of water from water sources.

1.4 Research Objectives

- Investigate the current state of fleet and resource allocation algorithms, path planning algorithms, and flight planning algorithms as influenced by the Constraint Satisfaction Problem (CSP), the Vehicle Routing Problem (VRP), as well as other algorithms that suit the multi-depot, multi-vehicle, multi-goal planning algorithm structure.
- Propose a solution to the fleet allocation, flight cost approximation, and flight planning stages, in-particular how they combine to form a unified fleet allocation and flight planning system.
- Implement the solution proposed and configure it as the backend of a firefighting resource management user interface. The solution should be implemented using a layered approach with each subsequent layer linked appropriately. The intention behind the layered approach is to allow functionality to be inserted or deleted seamlessly.
- Evaluate the tool that is developed by applying the fleet allocation and flight planning tool to generated scenarios, assessing the systems output. The aim of the evaluation process it to verify the tool developed is providing an efficiency gain over the existing approach to aerial firefighting.

1.5 Methodology

The methodology followed over the course of the project:

Chapter 2: Previous Work

A review of the current practices in aircraft fleet management and flight planning is undertaken to identify the shortcomings of their application to aerial fire suppression. The knowledge gained from the review conducted is used to inform the proposal of a fleet allocation and fight planning tool specific to resource optimisation for aerial wildfire suppression.

Chapter 3: Fleet Allocation

A fleet allocation algorithm is designed to assess the active fires and assign an appropriate number of aircraft to each fire. A set of parameters are introduced to enable numeric classification of wildfires which can then be used to inform aircraft assignments.

Chapter 4: Approximating Flight Costs

A means of updating the approximate cost of potential flight assignments is designed to provide a better representation of the true costs associated with a potential flight assignment. Cost regions are introduced to enable identification of specific regions of airspace that should be avoided, indicating both the location and an accompanying importance.

Chapter 5: Flight Planning

A flight planning algorithm is designed to inform both the initial distribution of aircraft at the commencement of operations and continued assignment of aircraft during operations. A two-stage planning algorithm is used to establish an initial distribution of aircraft before searching for potential improvements to the assignments made.

Chapter 6: Prototyping Application

A prototyping application is developed to enable rapid testing and feature development during the integration of the fleet allocation, cost approximation, and flight planning algorithms. A simple two-dimensional plot is used to display the output of the algorithms to enable rapid development.

Chapter 7: Data Entry and User Interface

A user interface is designed to act as the data entry console for the end user, system monitoring tool, and test data generation tool. The user interface is constructed using the Windows Presentation Foundation framework, enabling rapid development using the provided display components and customisable modules such as a satellite map interface.

Chapter 8: Results and Discussion

Each algorithm developed is integrated to form a unified fleet allocation and flight planning tool. The algorithms are tested using simple data generated using the prototyping application tool and intricate data generated using the user interface. The performance of each is analysed and the strengths and limitations of the design identified.

Chapter 9: Conclusion

The results of fleet allocation and flight planning tool are summarised and used to reflect on the research objectives identified at the beginning of the thesis. Finally, the future work required to improve the implemented system is discussed and further application areas are introduced.

2 Previous Work

2.1 Effectiveness of aerial firefighting suppression of wildfires

The primary objectives that accompany the deployment of firefighting aircraft are limiting the perimeter growth of a wildfire and assisting ground suppression efforts. Operational effectiveness of aerial firefighting is a major consideration for both initial deployment and sustained engagement of aerial firefighting resources. An equally important consideration is what benefits simultaneous ground and aerial suppression efforts have over ground suppression efforts alone. Both of these are explored in [12], accompanied by an introduction to the considerations made when deploying aerial firefighting resources, how the decisions made during the deployment process impact the resulting effectiveness of the aerial suppression efforts, and their utilisation during the firefighting effort.

The application of aerial firefighting to aid in the suppression of wildfires with a key emphasis on the reduction of the wildfire containment time has been the subject of much research over the past twenty years. As the application of aerial firefighting is refined and the techniques for initial deployment, ongoing management during fire suppression, and determining the appropriate time to cease aerial fire suppression operations become better understood, the utilisation of aerial firefighting has become common place. The utilisation of aerial firefighting in Australia has increased considerably over the past fifteen years and has led to aerial firefighting becoming a critical component of the wildfire suppression arsenal across Australia.

Determining the effectiveness of aerial suppression on wildfire containment requires the identification of fire attributes that enable comparison between different scenarios. Throughout the literature reviewed two measures have been discussed, the size of the fire at the point of containment and the period of time required for containment to be achieved [12] [13] [14] [15] [16] . As a fire's containment time is defined as the period time from fire ignition to reaching its maximum size, the two measures mentioned are providing the same basic metric.

Of all the factors that contribute to both the efficiencies and inefficiencies of aerial firefighting, the factor with the greatest impact is the deployment time of the from the beginning of a new fire event. In situations where the decision to deploy aerial firefighting resources is made early in the suppression efforts, the aircraft will reach the fire front sooner and have a smaller fire front to cover. This leads to the maximum positive impact of the aerial firefighting resources possible. Combined with the targeted mop up process conducted by ground crews, fire containment time can be minimised.

While the timely deployment of aerial firefighting aircraft can provide the greatest increase in efficiency, any delays during the process can dramatically reduce the benefits of aerial firefighting. Sources of delay include the misjudgement of fires by incident managers at the fire front who opt not to engage aerial firefighting aircraft, incident managers personal views on benefits provided by aircraft, long transit distances for aircraft from their nominated operational bases, unprepared aircraft and resource mobilisation [12]. The outcome of any delay is the enlargement of the fire perimeter during the delay window prior to the initial attack. These places higher workloads on the available aircraft and the ground crews as a larger fire front must be attacked using the same resources.

Major delays incurred due to prohibitive weather conditions, aircraft operation time limits, or a choice not to deploy aircraft for the initial attack can be costly. Both from an aircraft operations standpoint but also the added burden placed on the ground crews. Deploying aircraft after a substantial delay and maintaining the initial operational objectives of attacking the fire front can result in a negligible benefit being provided to suppression efforts. All the while incurring the operational costs associated with the aircraft operations including fuel, maintenance, personnel cost, and resource movement. If the fire front is too large upon the arrival of firefighting aircraft, the objectives of the aircraft must be redefined from a direct attack role to an assistance role.

Another critical takeaway from the literature reviewed on the effectiveness of aerial firefighting is the usefulness operating independently of ground crews. Deployment of aerial firefighting resources as the only means of fire suppression is often ineffective at eliminating the fire completely, requiring timely follow up from ground crews performing localised suppression and mop up of spot fires as necessary [3].

While its effectiveness is questionable when deployed independently there are situations where standalone aerial fire suppression is the best course of action. In conditions that are too difficult or unsafe for ground crews, aerial suppression can be applied in an attempt to limit the spread of the fire and assist ground crews. These difficult and unsafe conditions can be simplified to three key areas, challenging terrain, unpredictable weather conditions and high fuel loads.

Fires inaccessible by ground crews due to challenging terrain are often situated such that when they spread, spread into other areas that cannot be accessed by ground crews. In this instance aerial firefighting can be used in an attempt to limit the spread and subsequent damage when conditions are at their worst. Unpredictable weather conditions have the potential to leave ground crews vulnerable to injury from an erratic fire front. High fuel loads are present in the form of dense layers of highly combustible fuel that result in extremely intense fires. Aerial firefighting can be used to reduce the intensity at strategic points allowing ground crews to attack the fire from the ground [12].

2.2 Current Aerial Firefighting Management Strategies

2.2.1 Current Aircraft Control, Flight Following and Altitude Allocation

Throughout wildfire suppression operations a substantial reliance is placed on having experienced personnel leading and managing the process. This is even more prevalent during aerial fire suppression where the reliance on experienced personnel is extended to pilots and aircraft management personnel. A pilot's intuition and experience flying their aircraft in busy, high stress environments where execution is crucial is a skill set that must be possessed prior to them undertaking aerial firefighting missions.

While pilot experience and aviation conventions provide the primary control of aircraft movement over the incident airspace, a basic framework is still required to ensure all personnel are operating with the same expectations. The management of aircraft has been separated by type treating operational flights, bombing flights, and non-operational flights differently. Operational flights consist of flights moving aircraft to, from, and around the fire incident. Firebombing flights define the movement of aircraft directly over the fire incident. Non-operational flights encompass flights conducted for the purposes of moving aircraft and other assets by air to support sustained aerial firefighting efforts [17].

Procedurally, operational and non-operational flights follow an almost identical set of rules. The expectation is the pilots plan and execute the flights using general aviation best practices and following the appropriate Civil Aviation Regulations (CAR). The difference between the two is the reporting structure. Each flight type requires the same information to be broadcast, however, the communication channels required differ. Operational flights require more broadcast channels and consist of more communication levels, all with the aim to appropriately coordinate and direct the available resources.

The primary airspace management technique employed by the SACFS is the nomination of a series of work cells, delineated by a series of altitudes with reference to the ground level over the fire incident. There are three nominated work cells, one each for firebombing aircraft, air attack supervisor aircraft, and air observation aircraft. Each work cell occupies a continuous 500ft block with the blocks assigned such that firebombing aircraft occupy 0 to 500ft Above Ground Level (AGL), air attack supervisor aircraft occupy 1000ft to 1500ft AGL, and air observation aircraft occupy 1500ft to 2000ft AGL [18]. During firefighting operations, the airspace directly above the fire incident becomes protected airspace and as such all aircraft not directly involved in the aerial firefighting efforts must not descend below 3000ft AGL. The work cells as used by the SACFS are shown in Figure 2-1.

Within these work cells a circuit direction is defined and CAR 163 is sighted to ensure appropriate separation between bombing aircraft is maintained. A circuit with a consistent direction is used to maintain a stable flow of aircraft within the incident airspace. The circuit direction is nominated by the AAS and considers the wind direction, aerial suppression objectives and how the seating in the AAS impacts it's visibility. CAR 163.1 is a regulation applicable to all aviation operations within Australia. It states, "The pilot in command of an aircraft must not fly the aircraft so close to another aircraft as to create a collision hazard" [19]. By defining a circuit direction and citing CAR 163.1, the onus is on the pilots operating inside the work cells to maintain safe separation to all other aircraft.

Image removed due to copyright restriction.

Figure 2-1: SACFS Incident airspace management.

2.2.2 Current Water Source Restrictions

Currently, only water sources that meet strict requirements for either the water source's static capacity or the minimum rate at which the water source can be replenished are considered during the planning phase. As outlined in the SACFS's operating procedures two forms of water sources are considered, a primary water source or a continuous water source. A primary water source must have a minimum static capacity of 45,000 litres. This number was derived from the requirements for two SEATs delivering four payloads per hour to operate for a two-hour period. A continuous water source must the have the supporting infrastructure to provide water at a rate of 30,000 litres every hour for an eight hour period [20].

The severe limitation on the water sources that can be used also influence where airbases can be established. According to SACFS's operating procedures, an airbase must have the capability to provide the prescribed amount of water to be classified as an airbase. The locations of continuous water sources currently in use are indicated on the map shown in Figure 2-2 by the pink aircraft symbols.

A consequence of limiting the types of water sources factored into the planning phase is the underutilisation of water sources such as dams, lakes, rivers, reservoirs, and portable tanks. Although these forms of impromptu and short-term water sources may not be accessible by all aircraft types in use they can still be used at the pilot's discretion. Not factoring them into the planning process severely limits the potential efficiency of the aerial firefighting resource.

Image removed due to copyright restriction.

Figure 2-2: SACFS recognised airbases.

2.3 Fleet Allocation

Fleet allocation describes how the available resources that compose a fleet are distributed upon their initial deployment. While this may not involve the direct assignment of individual vehicles to a target, it does specify the quantity of each vehicle to be assigned to each target. The same principles hold irrespective of the type of vehicles within a fleet, their starting locations, the goal, or overall application space of the fleet in question.

The work in [16] presents a scenario based standard response model for the stationing and deployment of wildfire suppression resources. Each system objective is represented by a function. The objectives include the number of resources stationed at a depot and how many resources are deployed to fires in the vicinity of the depot. The output of the system is a resource station allocation and deployment, found through minimising the sum of the objective functions. Weight values were used to introduce the preference of the personnel making the decisions, choosing what objective has a higher priority.

The analysis they carried out explored the deployment of firefighting resources from fifteen stations in the Amador-El Dorado unit. During the analysis one hundred fire scenarios were used, the information for which was derived from a stochastic simulation. The analysis focused on how the assignment of firefighting resources to fire stations effected the number of fires that did not receive the desired number of resources within the response time. In this instance the response time was defined as thirty minutes. A time chosen by the researchers as the maximum time for an initial attack likely to still result in timely containment of the fire.

The analysis was conducted by applying model developed to data recorded by the Californian Department of Forestry and Fire Protection from Californian wildfires. The standard response model developed formed a tractable model which was solved to find the optimal stationing of firefighting resources and their subsequent deployment to fires [21]. The standard response model implemented showed a trend to focus the available resources to the regions frequented by fire events, over stationing resources to achieve spatial diversity. However, it was noted this left some areas susceptible to sub-standard responses by firefighting resources. The number of fires contained within a sufficiently short period of time was sensitive to the number firefighting recourses allocated to surrounding stations. Therefore, a fine balance was required assigning the maximum firefighting resources to optimise timely containment frequency while not exceeding the budgetary constraints.

A similar approach presented in [22] further explores the two-step process of the station allocation at the beginning of the fire season and the deployment upon a fire event. It furthers the work presented in [16] by bother considering a more comprehensive range of potential fire locations, accounting for both ground based and aerial suppression resources and factoring in socioeconomic and policy constraints.

At the time the research was undertaken the distribution of aerial resources was approximately uniform across the country. When the method developed analysed the existing distribution it was found that 90.1% of fires received the predefined response. Taking into consideration the available budget and other constraints on the system, the optimal deployment of helicopters provided resulted in 92% of fires that receiving the predefined response, an increase of 1.1%. The optimal solution identified through their analysis prioritised locations where fires were most frequent, often in the vicinity of metro areas. In addition to metro areas, mountainous areas and challenging terrain was also a priority of the initial resource distribution, due to their limited accessibility by ground resources and the increased rate of perimeter growth.

To validate the number of scenarios used to solve the problem the upper and lower bounds for the objective values were estimated using the sample average approximation method. Three set of scenarios were used, with sizes of 30, 50 and 100. Across the varying sample sizes the optimal gap only decreased by 1%. From this they deduce that the 100 random scenarios are adequate and increasing the scenarios used would serve no purpose other than to increase the computational loads associated with solving the problem.

The work presented in [14] explored how the personnel making the decision during the firefighting effort effected how aerial firefighting resources are used, and the subsequent impact on fire containment efforts. To accomplish this the requests made by incident managers for various firefighting resources were compared to optimal response models for each fire. The models were developed from available data on the environmental conditions, weather conditions and behaviour of each fire.

Much of the data relate to resource requests and subsequent allocations at the fire front were obtained from interviews with experienced incident commanders. The authors highlighted the limited data sets available for fire events, a common obstacle to wildfire assessment and is explored further by [15]. This included the specific characteristics of the fire itself, the accompanying statistics and records of how suppression resources were applied at the fires and the short-term effect suppression resources had.

The researchers identified two primary limitations of the approach implemented. Firstly, the operational objectives at each fire front were assumed to be the same. In real world conditions the objectives at each fire front will vary as some fires will pose greater danger to people, receive different numbers of resources, require more resource to completely contain etc. Secondly, the method they used to approximate decisions by proxy measures made by incident managers. Observations and decisions made by humans at different locations, under different levels of stress, and in different conditions can only ever be approximate.

2.4 Flight Planning

Flight planning is the process of planning aircraft movement to determine the optimal means of aircraft traveling from their start location to their destination. Flight planning for an application such as aerial firefighting requires a unique approach to solve the problem, a problem that has not been explored in detail previously. As a result, no one field provides the background alone. Instead elements of constraint satisfaction problem, vehicle routing problem, and traveling salesman problem, along with other planning constructs that support multiple depot, multiple goal, multiple vehicle, and multiple trip planning each contribute filling gaps in knowledge.

One approach to solving a planning problem largely driven by constraints is the Constrain Satisfaction Problem (CSP), an approach becoming increasingly common for planning of multi-vehicle unmanned networks. The work presented in [23] used constraint satisfaction problem as one of their solutions for the assignment and management of a swarm of unmanned aerial vehicles. Their implementation used a set of variables which each possessed a domain of possible values and a set of constraints to limit the values variables take concurrently. Due to the mission objectives and the coordination required between every unmanned aircraft the solution used thirty-six variables and constraints to define the system. This illustrated the potential for rapid growth of a solution when multiple vehicles and multiple objectives are involved. The map of the CSP variables constraints used to illustrate the constraints is shown in Figure 2-3.

Image removed due to copyright restriction. Original can be viewed online at: https://www.sciencedirect.com/science/article/abs/pii/S2210650217306703?via%3Dihub (Pg. 485)

Figure 2-3: Graphical representation of the CSP model presented in [23].

The CSP implementation outlined in [24] introduces the concept of constraint scaling, the behaviour allowing conflicting constraints to be implemented without disrupting the system. Hard constraints are used to indicate constraints that should not be broken under any circumstances, medium constraints indicate constraints that should only be broken when necessary, and soft constraints indicate constraints which not optimal to do so can be broken to attain a more practical plan. The construct of constraint importance scaling introduced allowed for missions to be achieved safely that would not have been possible with a conventional CSP due to limited fuel resources.

A planning approach that provided a great deal more insight was the Vehicle Routing Problem (VRP). As the VRP is a generalisation of the Traveling Salesman Problem (TSP) there are large numbers of variations that have been developed for specific applications. Some variants of note for the aerial firefighting application being explored are VRP with Pickup and Delivery (VRPPD), VRP with Time Windows (VRPTW), Capacitated VRP (CVRP), and VRP with Multiple Trips (VRPMT). The vast array of VRP permutations demonstrate the adaptability of the basic VRP structure.

The work presented in [25] introduces the VRP with Multiple Overlapped Batches (VRPMOB) to solve a vehicle routing problem consisting of time windows, multiple deliveries, multiple backhauls, multiple trips per vehicle, and vehicle delay costs. The search algorithm implemented adopted a two-stage solution, first it finds an initial solution to establish a starting point before searching different neighbours to find potential performance gains.

The initial solution was implemented using a cheapest insertion algorithm, a more sophisticated implementation of the greedy algorithm. Each insertion was completed with a goal of minimising the increase in the solutions cumulative cost. Following the initial solution, they implemented a neighbourhood search algorithm to search the surrounding neighbours to identify any potential efficiency gains through assignment switching. They also identified the major challenge of their VRPMOB implementation, the use of a fixed fleet size. As the fleet size was fixed, any new vehicle assignments made were going to influence all subsequent vehicle assignments.

The work presented in [26] explores the VRP with time windows and multiple deliverymen. The assignment of delivery vehicles and deliverymen to the required routes were implemented using decision variables defining the system constraints and customer requirements. A static optimisation was then used to optimise the original vehicle and personnel allocations. Due to the systems comprehensive mathematical implementation, the computational requirements of the system limited the algorithms use. To reduce the computational requirements imposed by the problem Solom’s heuristic is introduced.

The work in [27] presents an implementation of the VRP algorithm specific to aircraft routing with refuelling. Two methods of solving the aircraft routing problem were covered, one using a mixed-integer linear program and a second using dynamic programming. The mixed-integer linear program decouples the refuelling decisions from the aircraft planning problem, separating them into a parallel computation process. The second planning algorithm tested used was based on Dijkstra's algorithm. The Dijkstra algorithm is augmented from its original format through additional label values used for fuel tracking.

The authors then proceeded to test the two approaches with large data sets to ascertain which algorithm better handled the planning problem. It was identified that the mixed-integer linear program approach did not scale well and routinely timed out before a solution was found. Conversely, the dynamic programming implementation scaled well and returned a feasible every simulation iteration irrespective of the number of aircraft, number of refuelling points, and number of destinations.

The work in [28] presents a VRP algorithm containing provisions for both multiple trips and time windows. As seen in previous works, a two-stage VRP algorithm was implemented. An initial solution was found using a sequential insertion algorithm due to the speed with which it enabled the initial list to be established. After the initial solution had been found an ant colony optimisation algorithm was applied to optimise the initial allocation of the vehicles. Also noteworthy was how the specific VRP considerations, backhaul, time windows, and multiple trips, were incorporated into the algorithm. The algorithm and all associated decision variables were configured allowing the different VRP considerations to be either included or omitted.

2.5 Summary

Throughout the review of the previous work several important factors have been identified worth considering during the design of the fleet allocation and flight planning tool.

Effectiveness of aerial firefighting suppression

Prior understanding of the effectiveness of aerial fire suppression on wildfire containment was vital before such a project was continued. Studies conducted both locally across Australia and abroad in the United States recognised the utilisation of aerial firefighting resources when deployed in a timely manner did expedite the containment time of wildfires. Based on the findings of these studies development of systems to increase the efficiency of these finite and extremely valuable resources is a worthwhile endeavour.

Current Aerial Firefighting Management Strategies

The management strategies in use today for aerial firefighting aircraft place serve restrictions on the resources sources that are factored into the planning process. The operating procedures in use by the SACFS highlight many of the areas where efficiency gains could be made, specifically from the utilisation of resource not currently considered. Currently all resources required for sustained aircraft operation including water, retardant and aircraft fuel can only be sourced from operational airbases. This is a consequence of the strict requirements placed on resources considered during the planning process.

Operating procedures sourced from the SACFS highlight the current methods of aircraft management over fire incidents. The primary separation is done by aircraft function and then supplemented by nominating a flight circuit and specifying pertinent CARs. As the project is focusing on the fleet allocation and flight planning of existing piloted aircraft, the same aircraft management methods will be adopted throughout the project.

The scarcity of procedural documentation from rural firefighting organisations may be attributed to the degree of operational information contained and the potential threat to operations should that information be used to interfere with operations. From the many requests made, the SACFS was the only organisation to make their documentation available upon request.

Fleet Allocation

Fleet allocation can take into consideration many constraints on the system such as spatial limitations, budgetary constraints, fleet constraints, and support resource limitations. While every paper reviewed during chapter 2.3 used some form of an economic value as a metric for a solutions efficiency, many papers reviewed considered how the containment time varied for each fleet allocation strategy considered. An approach of note was the two-stage fleet assignment, where stage one considered the allocation of vehicles to operating bases and stage two considered the deployment of vehicles to active fires. While the fleet allocation process will draw from the insights taken from the papers reviewed, less emphasis will be placed on the economic constraints and the budgetary implications of an aerial firefighting mission.

Flight Planning

The planning algorithms reviewed during the chapter highlighted the uniqueness of the flight planning problem as it pertains to aerial firefighting aircraft. However, the review of existing literature led to a series of algorithms being identified that once implemented could complement each other to provide the functionality required for the flight planning algorithm.

As most aircraft movements are based on constraints derived from the aircraft capabilities, fire characteristics and existing operational practices, the CSP was the first algorithm explored. The DCSP specifically showed promise as provisions for dynamic operation were incorporated. Use of CSP directly was ultimately avoided as its capability to handle either conflicting constraints or the need to prioritise constraints was not obvious. However, the relationships between constraints could be used to inform the algorithm design.

The VRP was another algorithm quickly identified due to its extensive history in vehicle planning. As the VRP looks to find the optimal set of routes for the available vehicles with the goal of minimising the cumulative distance travelled it in principle was suited to the task.

However, as aerial firefighting operations encompass multiple aircraft, multiple trips, multiple depots, multiple goals, capacity constrained vehicles, and time windows the existing implementations of the VRP were likely going to be difficult to apply. Variations of the VRP do exist to that consider features such as split loads, time windows, capacity constrained vehicles, and pickup and delivery, however, combining each of these features has potential to result in a convoluted solution.

While CSP and VRP show two potential avenues for the planning stage to take, each on their own will be challenging to wrangle to fulfil the planning requirements. These learnings suggest the best course of action is likely to integrate the favourable attributes from each algorithm to form a single algorithm tailored to the unique planning problem faced.

2.6 Research Proposal

After reviewing the current state of fleet allocation and flight planning utilised for aerial firefighting, a layered process is proposed to achieve an intelligent fleet allocation and flight planning tool. The tool aims to combine the fleet allocation, flight cost approximation and flight planning process in such a way to overcome the shortcomings identified in chapter 2.4. Implementing components as a series of subassemblies, as shown in Figure 2-4, will allow the tool to be reconfigured as required. Whether for experimentation during the development process, export for use in applications with minimal changes required, or expansion following completion of this project, it affords the tool flexibility.

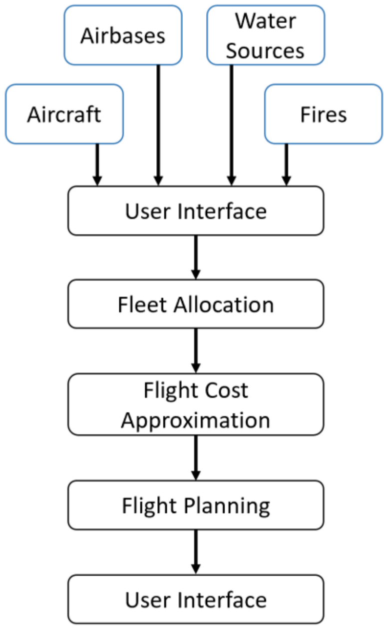

Figure 2-4: Proposed layered system diagram

The user interface layer is both the first and last layer of the system. Initially the user interface will be responsible for accepting user input, generating new data sets and loading previous data sets. The information entered either by the user or loading of a previous data set will form the basis of the subsequent layers. Following the fleet allocation, flight cost approximation and flight planning layers the user interface will become a display for continued monitoring during operations and reviewing the data generated throughout.

The fleet allocation layer will be responsible for characterising each fire and assigning the available aircraft to the fire fronts. The fire will be characterised using the information provided when a fire is first entered into the system along with the updated information provided for each fire. Fleet allocation will be achieved by allocating a number of aircraft from a pool of available aircraft proportional to the fire rating of each fire.

The flight cost approximation layer will be responsible for determining the shortest flight path between aircraft and either fires or water sources. Cost regions that are to be introduced to the system during the data entry phase will be used to augment the direct flight path between an aircraft and either a fire, water source or airbase. This will allow direct overflight of densely populated areas, significant or critical infrastructure, and hazardous weather cells to be avoided. On a case by case basis the urgency of the flight being undertaken, and the significance of the cost region will be used to determine if overflight of the cost region is permissible, and if so by how much.

The flight planning layer will be responsible for using the aircraft assignments and approximate flight costs calculated to plan the aircraft movements. For both the initial distribution of aircraft and during sustained operations the flight planning layer will aim to optimise resource usage by minimising cumulative distance travelled. During sustained operations checks will be incorporated to the assignment process to identify aircraft nearing completion of their assignment. By forecasting aircraft availability and incorporating additional aircraft into new plans the aircraft reassignment may be more efficient. The flight planning layer will also incorporate a redundancy check before a new flight assignment is made. This check will aim to ensure that an aircraft can safely return to their nominated operational base from any point over the course of the flight being assigned.

The proposed system will tackle the following short comings identified during the review in chapter 2.5:

- Utilisation of impromptu and short-term resources allowing for aircraft that can access them faster turnaround times.
- Prioritisation of active fire fronts through a structured classification process.
- A planning algorithm tailored to the dynamic environment encountered during aerial firefighting. Specifically accounting for the multiple aircraft, depots, goals, and trips along with the aircraft capacity constraints and time constraints on all operational aircraft.

2.7 Assumptions Made Prior To Development

Prior to embarking on the development phase of the project several assumptions were defined. Some assumptions defined aimed at improving the applicability of the tool developed to current procedures used during aerial firefighting. Other assumptions were defined to simplify the development process given the time constraints and complexity of the unconstrained problem.

- Direction of retardant drop made up to pilot therefore not factored in.

- Solution to be implemented should not use economics as the primary objective.
- Vehicles are only assigned for or incorporated into the plan if their current assignment is complete at the time of the new assignment.

In addition to the general assumptions listed above, a set of aircraft specific assumptions that have been made throughout the development process include:

- Each aircraft has a constant nominated operational base.
- Weather at all airbases will always be acceptable for both take-off and landing.
- Flying conditions over the fire front are always within safe operating limits.
- The SACFS mandated touchdown time will not impede the ability of a mission to be completed after it has been assigned.
- The retardant type used will have no impact on the systems operation.
- All aircraft considered have the same average airspeed.

3 Fleet Allocation

3.1 Objectives

Throughout the dry months and periods of serious fire danger, aircraft are stationed at nominated operational bases in a state of passive standby, a state from which stationed aircraft can be airborne within forty-five minutes from the of their engagement. On days of extreme fire danger, the standby state is upgraded to active standby, eliminating the response time and enabling immediate dispatch of all airborne resources.

While access to an aerial firefighting fleet that can be quickly mobilised is important, a means of directing aircraft to meet mission objectives is also critical. Therefore, the first step in the process of engaging firefighting resources was the allocation of aircraft to active fire fronts. The pool of available aircraft is a finite resource and, in the event multiple active fire fronts require attention simultaneously, a method of strategically allocating the available aircraft is required.

The goal of the fleet allocation process is to maximise the effective utilisation of aerial firefighting resources given the available aircraft, the fire front conditions and the fire locations. This chapter covers the implementation of the fire characterisation process and the use of the values obtained from the characterisation to inform the allocation of aircraft.

3.2 Fleet Allocation Considerations and Fire Characterisation

The aim of the fleet allocation procedure implemented for the thesis was assigning aircraft to the active fire front based on the priority of each fire. As such, a major component of the fleet allocation process was fire characterisation. Fire characterisation was achieved through the implementation of a list of parameters used to score the critical elements of an active fire. While there were countless parameters that could have been factored into the fire characterisation process, for the purposes of development and simplification for the thesis only the following parameters were considered:

- The fuel loading at the fire front.

- The potential for the fire to pose immediate danger to property and lives.
- The forecast contribution of aerial fire suppression.
- The monitored effectiveness of aerial suppression by ground crews.
- The accessibility of the fire ground by ground crews.

The parameters above were all existing parameters identified from the SACFS primary operating procedure, Aviation Operations Management [13]. The problem encountered using the parameters in their current form was the output of a parameter's assessment. The parameters considered used different output including; a binary output based on quantitative analysis, numeric values of varying scales, and some were void of any form of numeric description. This was addressed by using a consistent numeric scale for each parameter considered. The numeric scale used ranged from zero to nine, where zero was no influence and nine was considered a major influence. For implementation purposes the synthesis of the parameter values was done at the point of data entry, however, a standardised procedure would need to be adopted if standardisation of fire parameters was to be introduced.

3.3 Implementation

Using the fire characterisation parameters outlined in the previous section a process was developed determine the relative importance of each fire. This was accomplished by first updating the parameters with a set of weight values, normalising the values for all fires, scaling the normalised values to simplify aircraft allocation, and finally allocating aircraft to each fire proportional to its rating.

3.3.1 Weight Application and Initial Sum

The importance of the parameters defined were not necessarily equal, therefore, a method of defining the relative importance of each parameter was required. Furthermore, the importance of each parameter could not be static. Fire by nature is dynamic, therefore, the importance of each parameter needed to allow for change throughout the period of deployment as the fire evolved. Both concerns were addressed through the introduction of a series of weight values. The weights assigned to each parameter could be altered as required to better reflect the fire's state at that time.

The weighted parameters were calculated using the matrix operation shown in Equation 3. The matrix formed from the parameter weights was multiplied by the matrix formed from the parameter values, resulting in a matrix of weighted parameters.

$$\boldsymbol{w} = [w_{P1} \quad w_{P2} \quad w_{P3} \quad w_{P4} \quad w_{P5}] \tag{1}$$

$$\boldsymbol{P} = \begin{bmatrix} Parameter_1 \\ Parameter_2 \\ Parameter_3 \\ Parameter_4 \\ Parameter_5 \end{bmatrix} \tag{2}$$

$$\boldsymbol{P}_{weighted} = \boldsymbol{w} \times \boldsymbol{P} \tag{3}$$

These weight values were then used to establish a cumulative rating value for each fire by summing all weighted parameters together, shown in Equation 4.

$$fire\ rating = \sum \boldsymbol{P}_{weighted} \tag{4}$$

3.3.2 Normalising and Scaling

Following the establishment of the initial ratings for each fire, the ratings for all fires were normalised and scaled. First, ratings were normalised as per Equation 5.

$$rating_{norm} = \frac{rating - rating_{min}}{rating_{max} - rating_{min}} \tag{5}$$

Following the normalisation process the set of normalised fire ratings varied between zero and one. While the aircraft allocation function could have used the normalised values, the complexity of the aircraft allocation was greatly reduced by first scaling the list of normalised values. The scaling process as demonstrated in Equation 6, transformed the normalised values to a fraction. Following scaling of the normalised values the sum of the scaled ratings was one.

$$rating_{scaled} = \frac{rating_{norm}}{\sum ratings_{norm}} \tag{6}$$

3.3.3 Aircraft Allocation

The advantage of first establishing the scaled values was evident when determining the number of aircraft to allocate to each fire. The process became trivial and was achieved by multiplying the number of aircraft available, A_A, by the scaled rating of each fire, $rating_{scaled_i}$, as shown in Equation 7.

$$Aircraft_{fire_i} = A_A \times rating_{scaled_i} \quad (7)$$

A noteworthy potential limitation of the allocation method proposed was the rounding error that can occur. Anytime a double precision value used for computation was rounded and stored as an integer, potential existed for the value to round plus or minus one from the intended value. However, the failure mode of the rounding error was designed such that the risk of the dangerous outcome, exceeding the range of permissible values, was mitigated. Therefore, the result of rounding error was always an under assignment of aircraft by one or more, never an over assignment. This rounding error was more prevalent as the number of available aircraft approached the number of active fires, due to the fewer aircraft available to be assigned to each fire.

In the event the aircraft available are not all assigned, the initial list of fire ratings was consulted, and the remaining aircraft were added one at a time as the list was traversed from the highest to lowest fire rating. While this may not have been the most efficient method for utilising the unassigned aircraft, the list of fire ratings was the best source of information accessible to inform the allocation of un-allocated aircraft.

Following the allocation process the output of the fleet allocation stage, a list of aircraft to be assigned to each of the active fires was known. Importantly, the sum of all aircraft assigned was equal to the number of aircraft available. The number of aircraft assigned to each fire was also stored in each of the respective fire map objects for used in the following stages.

3.4 Adapting Allocation Parameters

The allocation method implemented above was designed to use observations and parameters currently employed by the SACFS. Providing both backwards compatibility with current SACFS procedures but also to constrain the complexity of the allocation stage for development purposes. However, the decision to implement these parameters using a fixed numerical scale was a deliberate one.

Currently the parameters considered during the allocation process are a combination of metrics specified in the SACFS operation manual and recorded observations made by firefighting personnel on the ground attending the fire front and at the SACFS control centre. The values are derived through visual inspection, environmental observations, personnel experience and the use of more sophisticated sensing equipment such as infrared cameras in places. As fire analysis continues to evolve and greater reliance is placed on computer analysis techniques, the numeric scale implemented will accept these values and continue to apply the fleet allocation process as informed by the individual fire ratings.

3.5 Summary

The fleet allocation process developed during this chapter forms the first layer of the fleet allocation and flight planning tool outlined in chapter 2.6. The chapter defines a set of parameters to describe each fire incident, which in turn enables all fires to be characterised and a rating assigned. The initial fire rating is then normalised and scaled to simplify the subsequent assignment process. The ratings obtained are then used to calculate how many of the available aircraft are assigned to each fire incident.

4 Flight Costs and No-fly Zones

4.1 Objectives

For the flight planning of both the initial aircraft distribution and the ongoing reassignment of aircraft a means of determining the minimum costs of a potential flight was required. As the primary heuristic used during the planning process was the cost of conducting a flight, a method of updating the flight cost to best represent how the aircraft was likely to be flown was critical. In this instance the cost of the flight was determined by the distance to be travelled, therefore altering the flight path could dramatically impact the planning stage.

4.2 Cost Regions

While no-fly zones accomplished the goal of avoiding certain airspace, they would not have easily allowed for external factors to influence the degree to which a no-fly zone was obeyed. Therefore, for the purposes of distinguishing between absolute and preferred no-fly zones, areas where direct overflight was a consideration were implemented using cost regions. Regions where overflight was unfavourable, the degree to which it should be avoided could then be quantified. Reasons for implementing cost regions included avoiding direct overflight of densely populated areas, avoiding direct overflight of significant landmarks and other critical structures, navigating around other fire incidents, and navigating around airspace or landmarks that could have posed a threat to aircraft such as dangerous weather cells or tall antennas.

Cost regions can be factored into the planning in one of two ways. Firstly, cost regions could be introduced manually through the user interface. Secondly, cost regions were enabled by default at locations that were operationally significant such as the area directly surrounding fires. Both methods are displayed in Figure 4-1. The transparent sand coloured circles show cost regions implemented through the user interface, while the transparent red region surrounding the fire marker indicates a cost region inherent to the fire object.

Figure 4-1: Introducing cost regions through the user interface.

4.3 Cost Region Weights

The implementation of a cost value for each cost region enabled the relationship between cost regions and the approaching aircraft to be dynamic. While there are many circumstances where encroaching on the defined cost regions may be detrimental to the aircraft or the mission, there are occasions where potential for varying degrees of overflight exist. Not adhering to cost regions established to protect the aircraft from hazards could prove catastrophic. However, situations where cost regions were established to minimise disruption on the ground, or the aircraft was performing a critical manoeuvre, partial or complete overflight had the potential to introduce considerable efficiency gains through the reduction of the flight distance. The use of cost region weights allowed the degree to which the cost region should have been observed and the required urgency of a flight before overflight to be considered.

For the purposes of simplifying the cost region implementation during the initial system development, how an aircraft approaches the cost region was reduced to two cases. If the planned flight was critical to the mission, the weight of the cost region was used to determine a percentage of the cost region's radius which could be overflown. In the event the aircraft was not undertaking a mission critical flight such as asset relocation or returning to its nominated operational base, the complete cost region was adhered to. The simplified process for navigating the weight of a cost region was implemented to best model how a piloted aircraft, with an experienced pilot at the controls, would approach the different type of cost regions given the operational setting.

The weight of all cost regions was either defined within the entities they are associated with, as found with the fire specific cost regions, or defined by the user at the time of defining the cost region using the user interface. The weight values used ranged between zero and one, defining the percentage of overflight that was permissible within a cost region for a mission critical flight. For hazardous instances such as fire, the cost region and exclusion zone around an active fire front was maintained by assigning the cost region for all fires a weight of one. Instances where cost regions were used to reduce the overflight frequency of populated areas a much lower weight was applied. If a zero weight was applied to a cost region, that cost region was ignored during the planning of mission critical flights.

To demonstrate the concept of the cost regions, the following two examples consider the same aircraft assigned a mission critical flight. The aircraft approaches a cost region assigned to the main streets of a small town designed to alleviate the noise and disturbances that come with firefighting aircraft. Due to the nature of the cost region its associated weight is zero, therefore the aircraft would not divert during its transit to or from a fire. The same aircraft then approaches a radio antenna with an accompanying cost region surrounding it. Due to the height of the antenna and the supporting guide wires the cost region has a weight of 0.75. Therefore, the aircraft undertaking a mission critical flight can encroach on the on the defined cost region by twenty five percent of the region’s radius.

4.4 Calculating Updated Flight Path

4.4.1 Check for intersection

The first stage of updating the flight costs for a planned assignment was to check the direct path for any intersections with active cost regions, from the starting location to the target location. If no intersection was found the direct path was already the best approximation of the flight cost. If a collision was detected, first the type flight being assigned and weight of the cost region the collision occurred in were checked. If the weight of the cost region was not zero, the aircraft was not undertaking a mission critical flight, or the cost region implemented posed a significant hazard to the aircraft, the approximate flight path of the aircraft required updating.

4.4.2 Calculate skirting points

As the updated flight path could only provide an approximation of the eventual flight path's cost, a method of avoiding the cost regions that itself used an approximation was satisfactory. In this instance the flight path was updated by augmenting the initial straight-line path through the additional of points to the route.

Working from the start location to the target location, first a point was found on the perimeter of the first cost region intersecting the line. The point on the perimeter was chosen such that the shortest line connecting the added point and the line intersecting the cost region was perpendicular to both the line intersecting the cost region and the tangent at the added point. The path between the new point and the goal was then checked for further intersections using the same process. If further intersections were found additional points were added with the procedure being repeated until no further intersections existed.

The new flight path was stored using a list of points. The first point in the list was the starting location of the aircraft. Additional points found to evade the cost regions were added to the list in the order they were found. When all intersections had been satisfied, the goal location was appended to the end of the list. The cost of the updated flight path could then be obtained by calculating the distance between each successive point within the list. While this method did not guarantee the optimal path with the lowest cost, it did result in a path that provides a better representation of the likely path of a piloted aircraft.

The process of augmenting the original direct path is shown in Figure 4-2. The red dashed line could not be traversed without an intersection with the cost region occurring. To avoid the cost region a point was added on the perimeter using the method described. The new path represented by the solid blue line was obtained by connecting the points on the list.

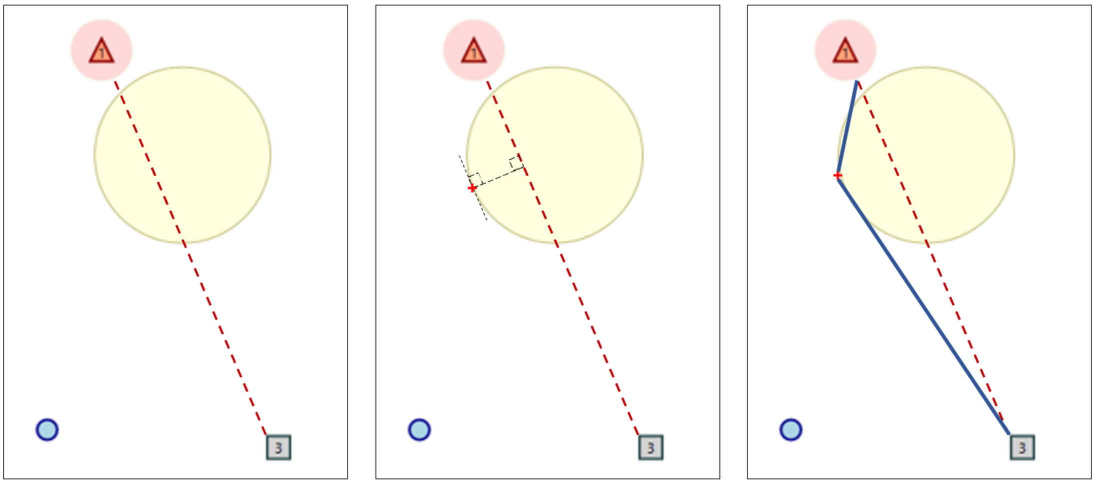

Figure 4-2: The process of adding points to augment the flight path.

4.4.3 Factoring in cost region weights

The weight values assigned to each cost region defined the degree of overflight permissible during a mission critical flight. Higher weight values were used to indicate the degree to which a cost region should be followed. As the weight of the cost region was reduced the overlap an aircraft could be increased. The decision of assigning weight values ranging from zero to one was deliberate, allowing the weight values to directly indicate a percentage of the cost regions radius which could be overlapped.

To achieve the overlap, instead of the point being placed on the perimeter of the cost region, the point was placed on the line connecting the intersecting line with the perimeter a distance proportional to the weight value. For example, for a weight of one the added point remained on the perimeter, for a weight of 0.5 the added point was placed halfway along the normal line, and finally, for a weight of 0.25 the added point was added three quarters of the way along the normal line. These cases are demonstrated in Figure 4-3. In the event weight of a cost region was zero the direct path would have been taken.

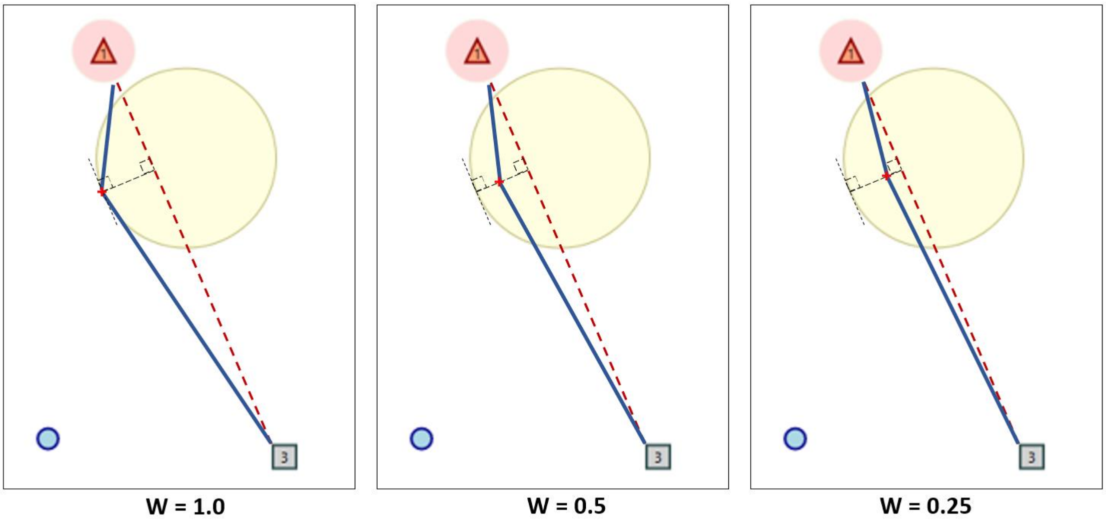

Figure 4-3: The degree of overflight determined by weight of cost region.

One caveat of the technique used was the order in which the checks for an intersection were required to take place in. The path required to be checked for intersections beginning at the current location of the aircraft and working towards the goal location. Failing to do so resulted in points being added in such a way as to clear the intersections while potentially creating further intersections.

While the flight cost was updated to better reflect the path the aircraft will have to take, all flight paths considered were only an approximation of the flight cost. Irrespective of whether the direct path was used, or the flight path was updated using the above process to avoid the cost regions. The process implemented was attempting to model how a pilot would fly the aircraft between its current location and the goal location. Due to the nature of having a pilot in the control loop, every aircraft will react to obstacles and manually defined cost regions differently.

4.5 Summary

The objective of this chapter is to implement a method of updating flight path costs to account for no-fly zones, avoidance of built up areas and aircraft hazards. The concept of cost regions is proposed to describe these instances, providing a dynamic version of a conventional no-fly zone that used a weighting value to determine how aircraft would interact with it. The approximate flight path for the aircraft is updated by adding points to the original direct route deliberately placed to route the aircraft around the cost regions. The weight value assigned to each of the cost regions encountered is used to determine the degree of overlap of the new flight path within the cost region. The updated path costs following the alterations made to the original direct routes formed the primary input to the flight planning system.

5 Flight Planning

5.1 Objectives

At this point, the number of aircraft assigned to each fire have been identified. However, the optimal airbases to assign aircraft from and how aircraft should be managed during assignments is unknown. The primary objective of this chapter is to design and implement a flight planning solution to optimise both the initial aircraft distribution and the continued flight planning during operations. The flight planning solution will be implemented as a stand-alone layer with minimal reliance on the proceeding or subsequent layers.

5.2 Planning Heuristic

Due to the limitation applied by the approximate flight costs used, a heuristic based algorithm was chosen for the flight planning. As the flight cost was approximate the use of a heuristic based planner was predicted to not detrimentally effect the performance of the system. For simplicity of implementation during the development phase the flight distance between the start and goal locations was chosen. This distance was described by either the direct path between the two locations or the updated path derived using the approximation procedure developed during chapter 4.

5.3 Initial Aircraft Distribution

The initial aircraft distribution accounts for cases where significant portions of the fleet were being assigned at the start of a new fire incident. This included mass deployment of available aircraft upon engagement of aerial firefighting aircraft, redistribution of aircraft following the initiation of a new fire event, or the distribution of aircraft upon the recommencement of a new day of aerial firefighting.

5.3.1 Initial Plan and Flight Assignments

Following the fleet allocation and path cost determination stages, the costs for every aircraft to travel to each fire was known. These flight costs were used to inform the planning process. Due to the potential complexity of the airbase, aircraft and fire structure, finding all permutations as a means of identifying the optimal set of flight assignments was unfavourable. For limited airbase, aircraft and fire numbers running all permutations would have required an insignificant computational overhead, however, as the number of factors acting on the system grows the number of potential permutations grows exponentially. A planning algorithm was therefore required that was going to provide a reasonable initial plan for the fleet while also scaling as the number of airbases, aircrafts and fires increased. The scalability constraint was introduced to maintain its multidisciplinary capabilities beyond aerial firefighting.

To achieve the initial distribution the greedy algorithm was implemented. As the name suggests, the greedy algorithm iterates through the list of flight costs, chooses the cheapest flight and makes the assignment. Following the initial assignment, the list of remaining flight costs was updated and the flight with the cheapest flight cost assigned. This process was repeated for the remaining aircraft until all available aircraft had been assigned to an active fire front.

The primary assumption made by the greedy algorithm is the locally optimal choice at each step will result in the globally optimal solution [29]. However, satisfying the locally optimal case in the greedy manner leaves potential for efficient assignments early on to result in expensive assignments to be made towards the end of the list. Considering these factors, the use of the greedy algorithm was not the most efficient choice and required further steps the improve the efficiency.

5.3.2 Flight Assignments Switching

As identified in chapter 5.3.1, the key downfall of the greedy algorithms use for the initial plan was the potential for high cost, inefficient allocations to occur towards the end of the assignment process. In an effort the reduce the high cost assignments that may have been made, the high cost assignments at all airbases were compared with low cost assignments at neighbouring airbases.

The high and low-cost allocations were compared in pairs and the cost of maintaining the existing aircraft assignments was compared with the cost of the assignments being switched. If the combined cost of the new potential assignments was reduced, the two assignments were switched. This method was aimed at remedying potentially expensive allocation made later in the greedy allocation process. While it was still not an optimal solution, if the greedy allocation was inefficient during later assignments, the overall cost was improved. In the event this process did not result in an improved cumulative collection of costs, the assignments found using the greedy algorithm were used.

5.4 Mid-Operation Aircraft Planning

During sustained firefighting operations the planning of aircraft continues to evolve. Upon the completion of an assignment, a new optimal water source, airbase or fire had to be chosen such that the aircraft could be reassigned. While this process was similar to the initial distribution, additional considerations had to be made for future aircraft movements and the suitability of new flight assignments had to be assessed.

5.4.1 Considering Future Aircraft Movements

When one or more aircraft became available for assignment, the future movement of the remaining aircraft in the fleet first had to be considered. The efficiency of new assignments was improved when multiple aircraft were replanned simultaneously. Therefore, no new assignments were made before first considering which aircraft were soon to complete their current assignment and become available and could be incorporated into the plan. Committing aircraft to new flight assignments without first looking forward to forecast aircrafts availability in the short term risked dramatically increasing the workload placed on available resources.

To enable the current status of an aircraft's assignment to be monitored a variable was introduced to all aircraft objects that indicated the completeness of its current assignment. The variable used a range of zero to one, showing the percentage of the assignment that had been completed. A threshold value was then introduced that defined how complete an aircraft's current assignment must be before it was factored into the next set of assignments. The limitation of this method was the completion value was based on the distance the aircraft had travelled of the planned trip. Therefore, the same completion value for two different aircraft could result in two significantly different wait times. However, the assumptions made previously that aircraft share the same speed therefore meant the distance travelled was suitable for all aircraft.

The key to considering future aircraft movements was to only forecast the future movements of all aircraft once for each set of aircraft assignments made. If the future movements were forecast repeatedly with hopes of finding further efficiency gains through simultaneously planning more aircraft, a loop was established until all aircraft became available again. As such it was detrimental to the utilisation of the aircraft and the systems overall efficiency.

5.4.2 Mid-operations Checks and Considerations

Following the consideration of future aircraft movements, the operational constraints placed on the aircraft were then checked. Operational factors such as the weather conditions, time of day, remaining flight time, and the status of each aircraft all had to be ok before the process could continue. If any of these conditions were violated, appropriate action had to be taken. The most common outcome of violation resulted in all new assignments being ceased and the aircraft directed to return to their NOB.

5.4.3 Return to Base Requirements

An aircraft's ability to safely return to its NOB from any point during or following its assignment was paramount. Therefore, prior to an aircraft receiving a new assignment the appropriate checks were made to ensure that the aircraft could safely return to its NOB following the completion of the assigned flight. This included considering the maximum duration a pilot can fly, the remaining flight time left in the day, the fuel required to return to base, and potential adverse changes in the weather.

While the safety margins factored into the flight assignments must not be so large as to adversely restrict the movement of the aircraft, allowing for appropriate safety margins was critical due to the higher risks associated with flying aerial firefighting missions. The safety factors considered when calculating the return to base plan also needed to be implemented such that they were dynamic. Factors such as the weather conditions and the aircrafts distance from its NOB following the completion of the planned assignment were all used to inform the safety margins factored in. In the event an aircraft could not safely return to its

NOB following the completion of the flight being assigned, the aircraft was directed to return to its NOB immediately.

5.5 Summary

The objective of this chapter is the design and implementation of a flight planning algorithm, building upon the fleet allocation and flight cost determination stages. Flight planning is achieved using the flight distance as the heuristic, and a heuristic based planner. To accomplish this a two-stage solution is implemented. First, an initial assignment of all available aircraft is found using the greedy algorithm. Secondly, the high cost assignments made from each airbase are tested against low cost assignments made from neighbouring airbases, checking if when the assignment is swapped the combined cost would be improved. The final flight plan developed for either initial aircraft distribution or mid-mission aircraft reassignment is then executed to either resupply the aircraft or drop the payload at the fire front.

6 Prototype Application

6.1 Objectives

All three components developed previously stood to benefit from a stripped back display and interface for rapid prototyping purposes. The objective of this chapter is to develop a simple application to combine the fleet allocation, flight cost approximation and flight planning components for testing purposes. The prototype application will consist of an entity generation engine to create a list of basic entities, a means of recording the test results, and a simple two-dimensional plot of the water sources, fires, airbases, and cost regions.

6.2 Simple Entity Generation

For the purposes of ongoing testing a means of generating simplified entities was required. Furthermore, two types of tests were required, a static data set that remained the same between each test and dynamic data set that was randomised upon each use. While the assumption could be made that random data sets alone were sufficient, there were some cases where maintaining single set of entities was beneficial to the testing procedure. To accommodate types of data sets a Boolean switch was implemented specifying the data set required at run time.

During the entity generation of a random data set, a series of rules were implemented to govern the generation of new entities ensuring the random set generated was a reasonable example of real life. The rule checks were applied to a prospective entity prior to its inclusion in the data set. The first check ensured there was no overlap of existing objects and the second check ensured adequate separation was maintained. If either rule was violated the prospective entity was discarded and a new random entity generated before the same checks were applied.

6.3 Two-Dimensional Plot

Following the generation of the data sets and the application of the three algorithms developed, the output of the test was plotted using a simple two-dimensional plot. The same display framework as the user interface adopts later, Windows Presentation Foundation (WPF), was used to create the plot window.

The example plot window shown in Figure 6-1, shows the output following the generation of a random set of entities before applying the planning algorithms. The airbases are represented by the grey squares, the water sources are represented by the blue circles, the fires are represented by the red triangles, and the cost regions are represented by the opaque circles. The circles surrounding the fire objects represent the cost regions inherent to fire entities.

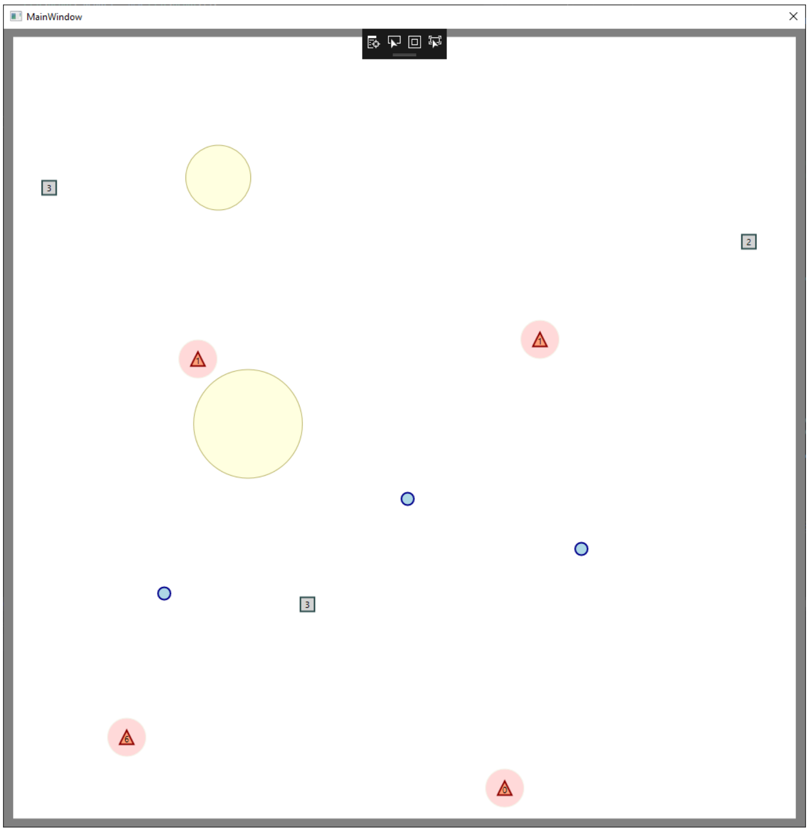

Figure 6-1: An example of the prototype application's plotting window.

Within the plotted objects some contain a number, shown in Figure 6-2. The number displayed inside the grey airbase objects indicate the number of aircraft stationed there at the start of operations. The numbers displayed inside the red fire objects indicate the number of aircraft that were assigned to that fire.

Figure 6-2: Entity markers used. Airbases (grey squares), fires (red triangles), and water sources (blue circles).

6.4 Summary

The prototype application developed during this chapter forms the basis of the intermediate test bed used to continue development of the algorithms. Simplified data sets are generated containing either static or randomised objects. The fleet allocation, flight cost approximation, and flight planning algorithms are then applied before plotting the output to a simple two-dimensional plot.

7 Data Entry and User Interface

7.1 Objectives

Given a primary consideration throughout has been the development of a tool that can be integrated into the current firefighting process, providing a low friction method of inputting information into the system is required. The objective of this chapter is to develop a user interface that will allow for both data entry and system monitoring. The data entry component must allow for introduction of fire, water, cost region, airbase, and aircraft entities into the system while also providing provisions for updating the details for each entity. The monitoring component should allow for the current status of water sources, fires, and airbases to be monitored throughout the duration of the mission, including a plot of current aircraft locations and their movements.

7.2 The User Interface Design

The user interface was developed using Windows Presentation Forms as the development framework. The use of an existing graphics framework that utilised the same backend langue of C# used for the algorithm implementation expediated the time from design and ideation to an operational platform. In addition to the C# backend, the mixed C# and XAML front end allowed for quick UI graphics to be implemented in a bespoke manner.

The application layout used a three-column design consisting of two information panels on each side of the window and a single map panel at the centre of the application. The central map panel contained a satellite map allowing for users to use known landmark cues to identify points of interest. While a human operator interacting with a map contained some error, the detail achievable through manual interaction with the map far outperformed the limitations encountered, largely due to the utilisation of piloted aircraft. The basic user interface layout and its three-panel structure is shown in Figure 7-1.

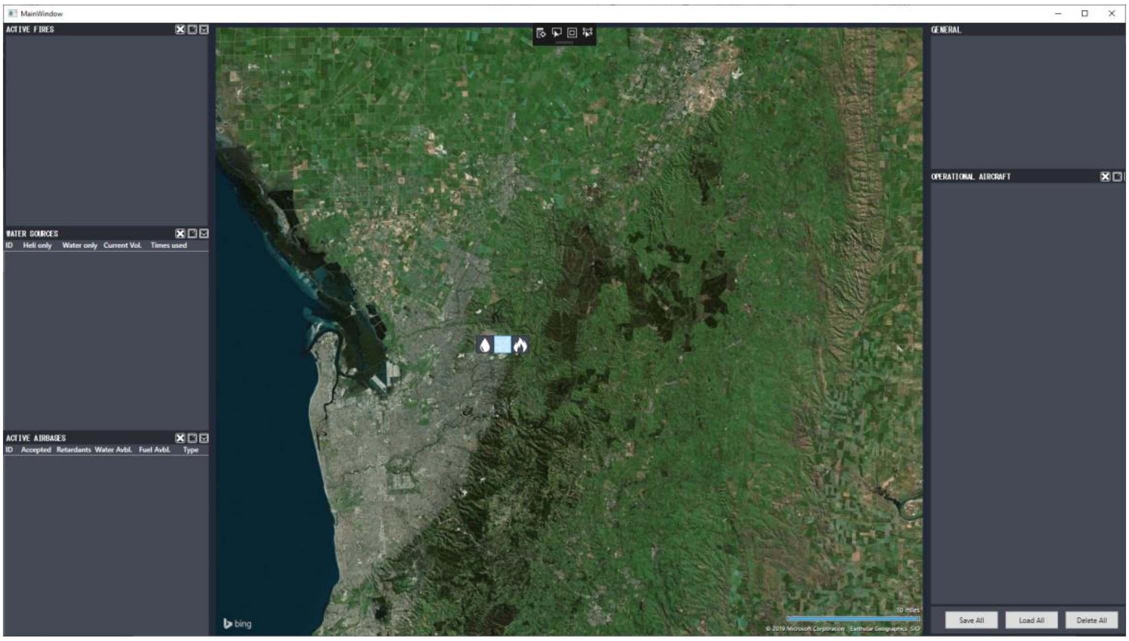

Figure 7-1: Blank user interface design showing the three panel design, two information panels separated by a signal map panel.

7.3 Entity Generation

Entity generation was the name given to the creation of airbases, water sources, fires and cost regions from the user interface. Entities were added through a series of pop-up windows displayed on a left mouse click, the location of which informed the placement of the entity on the map. Within the first pop-up the selection of the available entities was displayed, shown in Figure 7-2.

Figure 7-2: Entity options provided at the first pop-up.

Upon selection of the type of entity to be added, additional pop-up windows were triggered specific to the type of entity. The example of the new airbase and new water source pop-up fields are shown below in Figure 7-3. A populated map with all possible entities is shown in Figure 7-4.

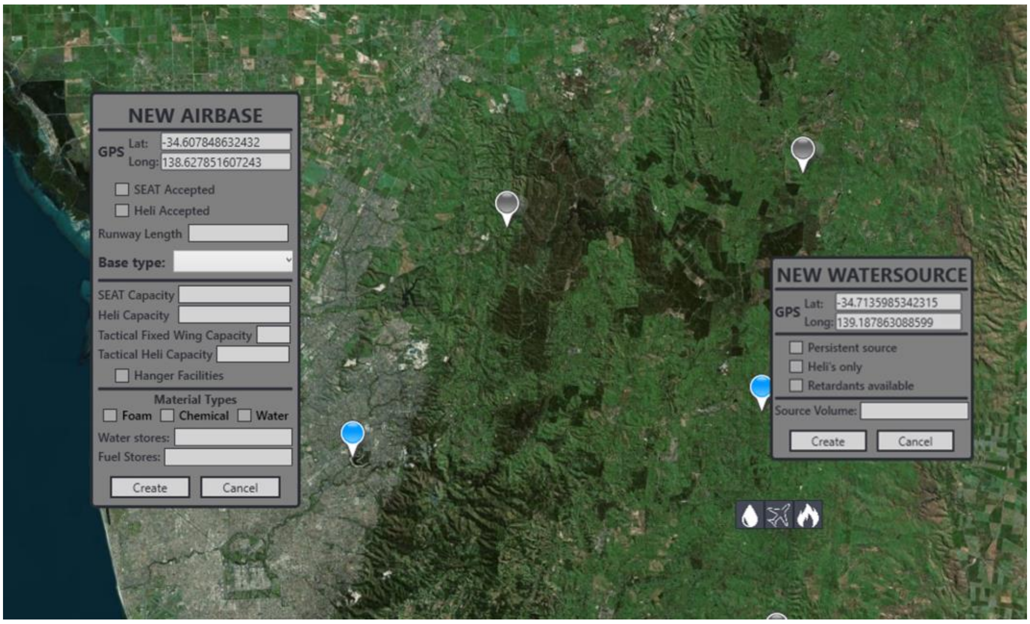

Figure 7-3: New airbase and water source pop-up fields.

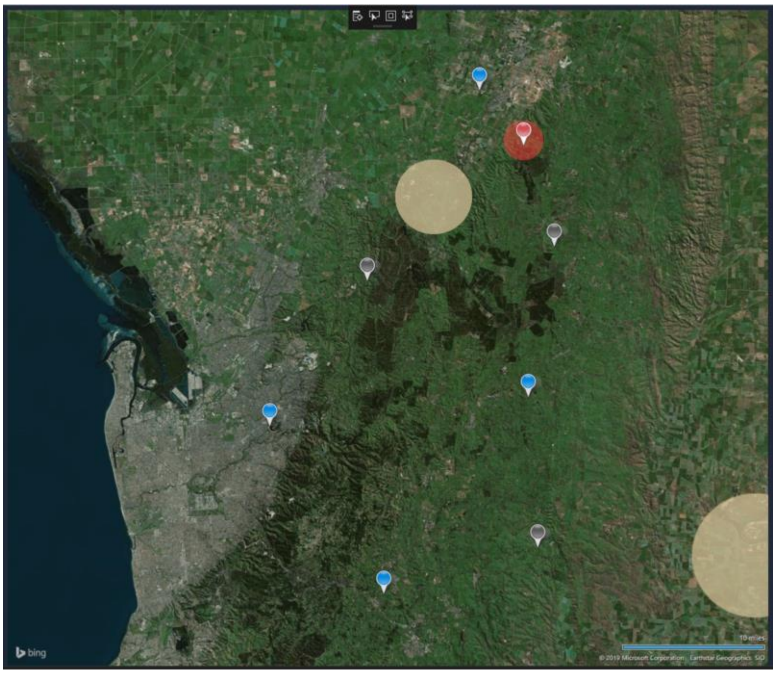

Figure 7-4: The user interface's map panel populated with airbases (grey), water sources (blue), fires (red), and cost regions (opaque circles).

7.4 Editing Parameters

An important function of the user interface was real time updates of entity data as well as the update of any mission critical information. While data fields such as water source volumes, fuel stores, stationed aircraft, and the number of times a resource had been serviced were updated automatically, a method for manually updating entity and mission data was required.

To accomplish this the pop-up fields used for entity creation were reused to provide data editing capabilities to the fields of interest. Access to the pop-up field for data editing was accessed by double clicking the specific entity of interest in the side panels of the user interface. Once the desired fields had been edited and confirmed, they were used to update the status of the system.

7.5 Monitoring Aircraft Location

As the primary component of the user interface was the satellite map, a suitable platform for monitoring aircraft status during operations was available. The monitoring functionality was designed to provide a visual representation of the aircrafts movement as the mission progressed. In addition to aircraft movement, it enabled the aircrafts current operational status and current assignment to be monitored as they unfolded. While the functionality was integrated to the user interface and the aircraft entities themselves, the visualisation of aircraft movement was not implemented during simulation.

7.6 User Interface Utilisation

While the initial objective driving the design of the user interface was developing a means of demonstrating the system and how operators could interact with the system, it also proved to be a beneficial tool during the development phase of the project. The functionality incorporated into the user interface provided a convenient means of creating datasets suitable for testing the fleet allocation, flight cost approximation and flight planning algorithms. The manual nature of this data generation method enabled datasets to be created that were designed to test the extreme cases with potential to break the algorithms.

7.7 Summary

The objective of this chapter is to create a user interface that enables data to be entered by the end user of the tool, while also enabling the generation of data sets for development purposes during the project. The user interface is implemented using the WPF framework, making use of the predefined user interface functions to expediate its implementation. The user interface is then used to as a data generation tool to compile a series of test data sets.

8 Results and Discussion

8.1 Overview

The purpose of this chapter is to present the implemented system as proposed in chapter 2.6 and examine the systems performance. The first section explores the testing of the three individual sub-components that form the system. This algorithm testing process was conducted using data sets generated using the user interface developed in chapters 6 and 7. The findings from the testing of the sub-components were used to inform the final variant of the system applied at the back end of the user interface.

8.2 Sub-component Testing

Each of the algorithms were developed in isolation to enable the problem to be broken down into manageable sections. The primary tool used throughout testing was the prototype application implemented in chapter 6 which allowed for simpler entities to be used. For each permutation used the number of airbases, fires, waters, and aircraft were specified at run time. A Boolean value was used to determine whether the entities would load from a pre-defined list suitable for debugging purposes or be generated at random suitable for general testing.

8.2.1 Fleet Allocation

Upon each run the parameters associated with a random fire entity generated were randomised. Following one execution of the test application all data inputted to the prototype application or output from the prototype application was recorded for demonstration purposes. Table 4 shows the data from the fire characterisation process, including each of the fire assessment parameter values along with the parameter weights implemented. The bottom three rows of the table show the initial fire rating from the sum of the weighted parameter values, the normalised rating, and the scaled rating used later for the aircraft allocation.

Table 4: Fire characterisation parameters.

Fire Parameter (weight)	Fire 1	Fire 2	Fire 3	Fire 4
Fuel Load (0.3)	9	9	6	3
Endangerment rating (0.4)	3	8	5	5
Forecast effectiveness (0.2)	9	5	7	7
Observed effectiveness (0.0)	8	8	8	8
Ground access (0.1)	9	8	6	6
Initial rating	6.6	7.7	5.8	4.9
Normalised rating	0.607	1	0.321	0
Scaled rating	0.3148	0.5185	0.1667	0

Following the fire characterisation and the calculation of the scaled fire ratings the fleet assignments were made. The scaled rating allowed for fleet assignments to be made through the direct multiplication of the available aircraft and the scaled ratings. Following the initial aircraft assignment, a check was made to ensure all available aircraft were assigned and rounding error didn't result in one or more aircraft not being assigned. In the event any aircraft were left unassigned, the list of fires was cycled through in order of fire rating severity and remaining aircraft allocated at each step. The resulting aircraft allocation from the fire parameters shown in Table 4 and ten available aircraft are shown below in Table 5.

Table 5: Fire assignments based following fire assessment.

	Fire 1	Fire 2	Fire 3	Fire 4
Assigned Aircraft (10)	3	5	2	0

The map generated randomly and used for demonstration above is shown below in Figure 8-1. The map shows the rules implemented in chapter 6, aimed at eliminating randomly generated map objects overlapping or residing in too close a proximity to other map objects achieve the desired goal. By preventing overlapping objects and maintaining spatial separation the map generated was more indicative of real-world scenarios.

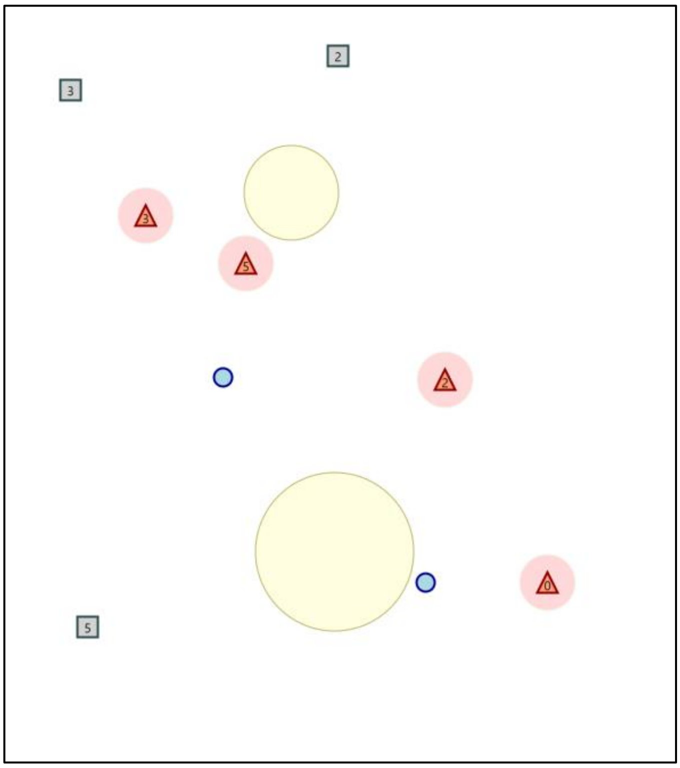

Figure 8-1: A randomly generated set of entities.

The checks incorporated to eliminate the risk associated with over or under assignment of aircraft performed well throughout. Following hundreds of test cases passed through the fleet allocation algorithm, no faults occurred through over assignment and no new inefficiencies were introduced from under assignment.

One limitation of the fleet allocation algorithm identified during the testing phase was the absence of forward planning. The allocation procedure implemented considered the available aircraft to be all aircraft stationed at each airbase. No considerations were made for strategically retaining aircraft at airbases to enable rapid response to new fire incidents. Long term the tool would benefit from a more complete assessment of the conditions and risk before aircraft assignments are made. An assessment should factor in the number of aircraft permitted to be assigned to a single fire, the number of aircraft that should be retained at an airbase, and the number of aircraft permitted to undertake a long-range transit out of their operational zone.

Another limitation identified following the implementation of the fleet allocation process is the reliance currently placed on the observations of experienced personnel at the fire front, observations from the AAS or AOB aircraft, or experience personnel stationed at the operations command centre. As the observations and qualitative assessments are being performed by different personnel with varying experience levels a degree of variance is unavoidable. Until assessment methods provide quantitative results the variance must be accepted and minimised using duplicate assessments by multiple personnel.

8.2.2 Flight Cost Approximation

The flight cost approximation process calculated the distance from every airbase to every fire location from the data plotted in Figure 8-1. The distance values calculate shown below in Table 6 were returned using a dimension of pixel widths. This was an artefact of the graphics implementation in the prototype application. While appropriate code could have been incorporated to rectify this, the purpose of the prototype application was rapid testing and pixel width was therefore sufficient.

Table 6: Distance from every airbase to every fire.

	Fire 1	**Fire 2**	**Fire 3**	**Fire 4**
Airbase 1	128	216	415	217
Airbase 2	364	347	216	409
Airbase 3	219	208	299	496

Throughout the testing phase a common observation was the newly calculated flight paths encroaching on the cost regions defined, as shown in Figure 8-2. As only one point was being used to route the path past the cost region, the flight paths were overlapping the cost region. This was more prevalent for larger cost regions, the existing flightpath was centred on the cost region or fire, water source or airbase locations near the cost region in question. While aircraft remain piloted the overlaps occurring are not likely to impact the system as the system only requires an approximation. However, if semi-autonomous or completely autonomous aircraft were to adopt the planner for logistics purposes, inaccuracies present would greatly affect the capability of the tool developed.

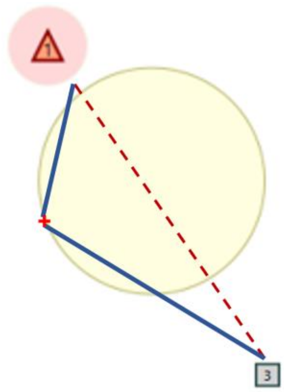

Figure 8-2: Calculated flight path overlap.

Ultimately, updating aircraft flight paths to avoid direct overflight of sensitive areas or avoid threats to the aircraft was not going to provide an exact flight cost for use in the subsequent planning stage. As the aircraft considered are pilot by humans there was no means of obtaining an exact cost for a potential flight. With the current technologies in place, the system could only do its best to emulate the path the aircraft would likely take and how pilots would navigate no-fly zones.

Beyond the tool's example application of water delivery optimisation during firefighting efforts, integrating no-fly zones into the planning process was critical. For example, if the intended use of the tool being developed was extended to defence applications incorporating no-fly zones would become essential during the planning process. Avoiding airspace behind enemy lines, avoiding unprotected and high-risk airspace, and avoiding locations that could give enemies an insight into the movement of a fleet and its objectives must all be factored in to establishing the cost of aircraft movement.

8.2.3 Flight Planning

The flight planning when tested on single or small collection runs operated as expected. Shown in Table 7 is the result of the flight planning as applied to the example data from the previous analysis. The neighbour search as implemented proved less impactful than first predicted. Often the final flight assignments were the result of the initial solution derived using the greedy algorithm. One potential cause could be the small sample sets used for the initial integration and test phases. The use of a smaller data set in turn meant each object had less neighbours, therefore, less potential for efficiencies to be found.

Table 7: Aircraft planning fire assignments.

Aircraft # (airbase)	Fire #
1 (1)	1
2 (1)	1
3 (1)	1
4 (2)	2
5 (2)	2
6 (2)	2
7 (2)	3
8 (2)	3
9 (3)	2
10 (3)	2

Efforts were being made to run tests that would facilitate data collection allowing for a more detailed analysis process, however, the resulting complexity of the layered system made the process difficult. The most basic comparison planned was developing the functionality to calculate every possible permutation of the plan, allowing the worst, best, and range of costs to be found. The hurdle was writing a program that would automate every assignment permutation to be run for every layer for a given data set. The challenges faced during the extended testing and attempt at better analysing the planner were also experience during the implementation of a simulation, discussed in chapter 8.5.

The first limitation encountered during the implementation of the flight planning process was encountered previously during the development of the fleet allocation and flight cost approximation stages, the requirement for a pilot in the loop. The movement of the available aircraft can be optimised as best they can using the flight costs approximated in chapter 4, however, the flight paths considered during the planning phases are unlikely to be followed exactly due to the pilots in the loop. While this needed to be considered during the evaluation of the methods used, an approximation of the flight path was still the best source of information that could have been used to inform the flight planning.

The second limitation encountered that became evident during the testing process was the predictability of the flight planning. For operations such as aerial firefighting, freight management, humanitarian endeavours, and other applications where malicious interference is unlikely, the planning priority of improving efficiency to maximise resource utilisation and reduce costs was the preferred choice. However, for planning applications where adversaries maybe working to disrupt the mission, the predictability that accompanies optimal or close to optimal solutions would be detrimental to the overall goal.

The flight planning approach implemented used the cost of the flight path as the planning heuristic, and the goal of the planning process was to establish a plan that was as close to the optimal solution as possible. In the events when fewer numbers of depots, vehicles, targets and resources are used, the closer the flight plans of the aircraft would likely be to the optimal solution. Adversaries that stand to gain from thwarting the mission objectives can implement similar planning processes and attempt to pre-empt plans being developed. With only a basic understanding of current aircraft positions and a general mission objective, stress points in an optimal solution plan could be identified and exploited to the adversary's benefit.

8.3 User Interface Analysis

The user interface design implemented in chapter 7 achieved the requirement of a primary console to enable user interaction with the system. It had the functionality to generate fully specified fire, water source, airbase, cost region, and aircraft entities along with the provisions to maintain current system information relevant to the operation such as fire suppression objectives and weather.

Although the user interface implemented meets the immediate requirements outlined for the project's development, there are several limitations present effecting both the interfaces long term usability and the interfaces technical adaptability. These were all associated with the user interfaces application using the WPF framework.

Limitations specific to the end user include the applications dependence on the Windows operating system and the ability to only run one instance of the application for the entire operation. Limitations specific to the technical adaptability of the tool was the limitation placed on the programming languages used for the front end and back end development. The implementation of the WPF used a combination of C# and XAML for the display specific code. This in turn heavily informed the selection of the programming language used to implement the backend algorithms due to the challenges porting data between C# and C++ for example.

A potential alternative to the use of the WPF framework would have been the use of a webpage-based user interface. It would have addressed every limitation identified from the use of WPF. The user interface could be used independent of operating system with any current web browser. It would also permit an unlimited number of user interfaces to be open, either for the purpose of interacting with the system or for users interested in monitoring the system in operation.

A webpage-based implementation would have enabled the selection of programming language used for the implementation of the algorithms to be made purely based on language requirement dictated by the algorithms. With a webpage-based design, porting data between the webpage and backend algorithms would have been a trivial process and existing frameworks exist for this purpose.

8.4 System Review

The greatest strength of the project was the layered architecture which allowed for nimble development while also considering longer term implementation and revision. By treating each layer as a self-contained component with unified inputs and outputs the layers could be integrated quicker, swapped out with a new revision of itself allowing continued operation of the main system, and layers could be bypassed completely for test and debugging purposes. This was closely followed by the interaction design and implementation. While a featured user interface was not a requirement for an implementation project such as this, it proved to be an invaluable tool for generating data sets, testing and rapid prototyping.

The notable limitation of the system developed was the lack of testing enabling the performance of the system to be quantified. While the combined system was tested extensively with randomised data throughout the development, an effective method for running large numbers of test and translating the results into usable data was missing, as expanded upon in chapter 8.5.

8.5 Simulation

As identified numerous times a major limitation was the difficulty establishing a repeatable simulation method to characterise the performance of the system. The initial plan was to compare the greedy and neighbour search combination to a series of other planning algorithms. The aim was to also develop a means of determining every possible planning permutation from a data set to enable an extensive analysis to be conducted.

The biggest obstacle encountered during the process of building the simulation was the method used to store the information at the airbases, fires and water sources. Each object had at minimum of two levels containing aircraft assignments, therefore, to access every permutation a convoluted means of iterating through all permutation was required. A major change to introduce if the system was to be redesigned would be the simplification of the assignment storage. The use of a more common data structure for which simple methods of traversing exist is paramount.

8.6 Summary

The purpose of the chapter is to demonstrate the implementation of the tool developed and present its performance, strengths, and limitations. The results and limitation are then used to identify the weakness of the system developed. Overall the system met the initial objectives of implementing fleet allocation and flight planning tool that optimised the utilisation of resources, however, the testing process uncovered several limitations. The primary limitation identified is the sub-optimal method used to implement the data storage, making the process of automation and simulation extremely challenging.

9 Conclusion

Aerial suppressions application in containing wildfires is becoming increasingly common as both the frequency and severity of wildfires increase. With the cumulative costs associated with the resource's operation being so high, along with the personal and financial cost that accompany extensive fire damage, efficiency use of aerial suppression is critical. While the importance of timely attendance by aerial suppression is well documented, the systems and procedures in places are largely dependent on experience firefighting personnel making to make the mission critical decisions.

The purpose of this research was to determine the requirements that would enable the management and planning of firefighting aircraft to be conducted in a more rigorous and repeatable manner. These requirements were identified during the review of current practices adopted by aerial fire suppression, vehicle route planning algorithms, and resource optimisation algorithms. From these requirements a layered fleet allocation and flight planning system that would integrate with current practices was proposed. The system was successfully implemented behind the user interface that was developed.

Fleet Allocation

The fleet allocation algorithm must have access to up to data information on all the fires considered during the allocation process. Without current fire information the aircraft assignments made will not be efficient rendering all future steps superfluous.

Flight Cost Approximation

The cost regions identified must be incorporated into the planning and the resulting flight path adjusted to compensate. If the cost regions are not accounted for and a direct flight path is used all subsequent flights will be negatively influenced.

Flight Planning

The flight planning algorithm must have a complete and up to date list of constraints and characteristics of every component in the system, for example a detailed aircraft model.

Without a detailed model of every system component, the planning process may be underestimating or overestimating the capabilities of aircraft, airbases and water sources.

User Interface

Regular use of the user interface must be adopted by the system users. Failure of the operators to update the information in the user interface will result in the system poorly reflecting the current operational state.

Conclusive Remarks

An investigation of the current procedures adopted by aerial suppression resources and previous work specific to planning of these resources informed the direction for the project. However, no previous work explores the allocation and flight planning of firefighting aircraft in such detail. Therefore, the solution developed throughout the project resulted in a unique intelligent fleet allocation and flight planning tool being developed. With future work, the solution implemented could form the basis of a system revolutionising the aerial fire suppression field through the synchronised use of both manned and unmanned vehicles.

9.1 Future Work

Throughout this project, the scope of the project limited some details that could be incorporated into the solution in a timely manner. Furthermore, several gaps in the implementation were identified during the design and integration phases. This section provides an overview of key improvements to explore further:

Fleet Allocation

The fleet allocation algorithm implemented for this project was void of any form of budgetary constraints. The deployment of aerial suppression resources is resource intensive from the time aircraft are requested, and every additional aircraft adds to the operational cost. Therefore, implementing budgetary constraints in the fleet allocation phase would improve the tools adaptability to real world scenarios due to the importance placed on the economic constraints of operations. In addition to the exploration of incorporating budgetary constraints, investigating how a fleet can be managed in its entirety could prove beneficial.

Actions such as limiting the aircraft deployed to the bare minimum required to achieve the operational objective, or intentionally retaining several aircraft at NOBs following a risk assessment of likely future fire events that will require a rapid response should be considered.

Fire Characterisation

The fire characterisation implemented for this project required decisions and judgements made by experienced personnel to be synthesised to a form that could be readily processed. The system would benefit from exploring more sophisticated technologies to assess the active fires, removing the human influence from the system.

Fire Management strategies

The current means of fire management implemented uses the rating derived for each fire to allocate resources, prioritising fires where the requirement for suppression resources is greatest. It doesn't consider how containing each fire will impact the future effectiveness of aerial suppression. A means of incorporating a more comprehensive analysis to inform the fleet allocation could greatly improve the efficiency of aerial fire suppression.

User Interface

The current user interface relies on an operator to maintain the information with the system. Integrating information currently utilised by other systems firefighting organisations use would alleviate the additional workload placed on operators while also ensuring data provided to the system is up to date.

9.2 Alternate Applications

The primary focus of the project was the development of an intelligent fleet allocation and path planning tool. However, for the purposes of demonstrating the system throughout its development as well as providing a method of tracking the progress of the project a defined use case was beneficial. The demonstration case chosen was aerial firefighting, specifically exploring how the fleet allocation and flight planning tool developed could be applied to improve the efficiency of resource usage. In the aerial firefighting example, the primary resource of concern was water. The choice of an aerial firefighting demonstration application was a deliberate one. Enabling the tool developed to be implemented to a dynamic and suitably complex application, while minimising the challenges introduced when working with IP sensitive or classified information.

There are many potential use cases for the type of tool that was developed. Two examples are briefly described below to demonstrate how the tool was designed to remain application agnostic.

9.2.1 Aerial Monitoring and Data Collection

One potential application of a fleet allocation and flight planning tool is the management and planning of aircraft undertaking reconnaissance, surveying, and data collection flights. These flight roles are the crucial during the undertaking of military intelligence gathering, aerial surveillance, search and rescue, and boarder security missions. Due to careful design choices considered throughout the project, the fleet allocation and flight planning tool developed can be adapted to alternate use cases. The conversion process would become especially trivial for scenarios utilising multiple aircraft or vehicles, multiple depots, multiple targets, and a resource which usage benefits from optimisation.

9.2.2 Protection of High Security Assets

Another potential use case for the fleet allocation and flight planning tool developed is the management and planning of a fleet of autonomous vehicles, entrusted with the task of maintaining surveillance over high value or high security assets. These assets may be desirable due to their worth, equipment stored within, their use if captured, or their alibility to provide a military advantage, either due to additional artillery assets or intelligence that can be gained. Such assets could employ a fleet of autonomous vehicles to monitor their surrounds where permeant measures cannot be established for their protection.

One such example could be that of a naval ship docked in a civilian harbour. Prior to the ship's arrival in the harbour thorough sweeps of the surrounding infrastructure, the dock above and below the surface, the sea floor surrounding the dock can be employed to minimise the risk of the asset being compromised. However, without ongoing monitoring over the course of the vessel being docked the asset is still at risk from new threats.

A fleet of autonomous platforms that could consist of flying drones, autonomous surface vessels and autonomous underwater vessels could be deployed to monitor the asset in the uncontrolled environment over the course of the docking period. The fleet allocation and flight planning tool proposed could be adapted to maximise the surrounding area covered by the surveillance vehicles, prioritising areas at greater risk or not recently inspected.

10 Bibliography

[1] Bushfire Cooperative Research Centre, “Assessing Aerial Suppression Drop Effectiveness,” Bushfire Cooperative Research Centre, Melbourne, 2009.

[2] J. Kightly, “50 Years of Firebombing Operations,” National Aerial Firefighting Centre, 19 January 2017. [Online]. Available: http://www.nafc.org.au/?page_id=382. [Accessed 14 September 2019].

[3] Bushfire Cooperative Research Centre, “Effectiveness and efficiency of aerial firefighting in Australia,” Bushfire Cooperative Research Centre, Melbourne, 2009.

[4] Field Air, *Guide to Air Tractor Aircraft,* Ballarat: Field Air, 2017.

[5] Bidgee, “Kennedy Aviation (VH-XAY) Air Tractor AT-802,” 22 1 2014. [Online]. Available: https://commons.wikimedia.org/wiki/File:Kennedy_Aviation_(VH-XAY)_Air_Tractor_AT-802_taxiing_after_refilling_with_fire_retardant_at_Wagga_Wagga_Airport.jpg. [Accessed 19 August 2019].

[6] Air Tractor, “Air Tractor AT-802A,” Air Tractor, 2019. [Online]. Available: https://airtractor.com/aircraft/at-802a/. [Accessed 12 March 2019].

[7] E. Valle, “Elvis Erickson S-64 Air-Crane Helictoper,” 18 September 2012. [Online]. Available: https://commons.wikimedia.org/wiki/File:N179AC-Elvis-739.jpg#globalusage. [Accessed 19 August 2019].

[8] B. Gabbert, “Erickson Aircrane reloads,” 12 August 2014. [Online]. Available: https://fireaviation.com/2018/01/10/erickson-receives-contract-to-build-two-aircranes-for-south-korea/. [Accessed 02 August 2019].

[9] Erickson Incorperated, “Technical and performance information,” March 2016. [Online]. Available: https://www.fire-brigade.asn.au/operations/Erickson-Inc-S64.pdf. [Accessed 11 April 2019].

[10] T. Bell, *A Los Angeles County Fire Department Sikorsky S-70 Firehawk helicopter demonstrates a water drop,* California: Fire Avitation, 2013.

[11] Lockheed Martin, “Sikorsky S-70 Firehawk,” 4 2019. [Online]. Available: https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/firehawk/8370_Firehawk_Brochure_LR.PDF. [Accessed 15 9 2019].

[12] M. Plucinski, G. McCarthy, J. Hollis and J. Gould, “The effect of aerial suppression on the containment time of Australian wildfires estimated by fire management personnel,” *International Journal of Wildland Fire,* no. 21, pp. 219-229, 2012.

[13] South Australian Country Fire Service, “SOP 8.3 – Aviation Operations Management,” SACFS, Adelaide, 2018.

[14] M. Hand, H. Katuwal, D. E. Calkin and M. P. Thompson, “The influence of incident management teams on the deployment of wildfire suppression resources,” *International Journal of Wildland Fire ,* vol. 26, pp. 615-629, 2017.

[15] M. P. Thompson, P. Freeborn, J. D. Rieck, D. E. Calkin, J. W. Gilbertson-Day, M. A. Cochrane and M. S. Hand, “Quantifying the influence of previously burned areas on

suppression effectiveness and avoided exposure: a case study of the Las Conchas Fire," *International Journal of Wildland Fire,* vol. I, no. 25, pp. 167-181, 2016.

[16] R. G. Haight and J. S. Fried, "Deploying Wildland Fire Suppression Resources with a Scenario-Based Standard Response Model," *INFOR: Information Systems and Operational Research,* vol. 45, no. 1, pp. 31-39, 2007.

[17] South Australian Country Fire Service, "Aviation Operations Standard Operating Procedure 8.5 - Flight Following," SACFS, Adelaide, 2018.

[18] South Australian Country Fire Service, "SACFS INCIDENT AIRSPACE MANAGEMENT," SACFS, Adelaide, 2018.

[19] Commonwealth of Australia Consolidated Regulations, "CIVIL AVIATION REGULATIONS 1988 - REG 163," - - -. [Online]. Available: http://www5.austlii.edu.au/au/legis/cth/consol_reg/car1988263/s163.html. [Accessed 3 June 2017].

[20] South Australian Country Fire Service, *Aviation Operations Standard Operating Procedure 8.12 - Airbases and Helibases,* Adelaide: SACFS, 2018.

[21] Y. Lee, J. S. Fried, H. J. Albers and R. G. Haight, "Deploying initial attack resources for wildfire suppression: spatial coordination, budget constraints, and capacity," *Canadian Journal of Forest Research,* vol. 43, no. 1, pp. 56-65, 2012.

[22] Y. Lee, B. Lee and K. H. Kim, "Optimal Spatial Allocation of Initial Attack Resources for Firefighting in the Republic of Korea Using a Scenario Optimization Model," *Journal of Mountain Science,* vol. 11, no. 2, pp. 323-335, 2014.

[23] C. R. Atencia, J. Del Ser and D. Camacho, "Weighted stratergies to guide a multi-objective evolutionary algorithm for multi-UAV mission planning," *Swarm and Evolutionary Computation,* vol. 44, no. 1, pp. 480-405, 2017.

[24] J. Munoz-Morera, F. Alarcon, I. Maza and A. Ollero, "Combining a hierarchical task network planner with a constraint satisfaction solver for assembly operations involving routing problems in a multi-robot context," *International Journal of Advanced Robotic Systems,* pp. 1-13, 2018.

[25] M. Yu and X. Qi, "A vehicle routering problem with multiple overlapped batches," *Transportation Research,* vol. E, no. 61, pp. 40-55, 2013.

[26] J. De Le Vega, P. Munari and R. Morabito, "Robust optimization for the vehicle routing problem with multiple deliverymen," *Central European Journal of Operations Research,* vol. 27, no. 4, pp. 905-936, 2017.

[27] T. E. Kannon, S. G. Nurre, B. J. Lunday and R. R. Hill, "The aircraft routing problem with refueling," *Optimization Letters,* vol. 9, no. 8, pp. 1609-1624, 2015.

[28] J. O. Onh and SupraYogi, "Vehicle Routing Problem with Backhaul, Multiple Trips and Time Window," *Jurnal Teknik Industri,* vol. 13, no. 1, pp. 1-10, 2011.

[29] F. Liu, "A greedy algorithm for solving ordinary transportation problem with capacity constraints," *Operations Research Letters,* vol. 45, no. 4, pp. 388-391, 2017.

[30] G. Xu and X. Zhong, "Real-time wildfire detection and tracking in Australia using geostationary satellite: Himawari-8," *Remote Sensing Letters,* vol. 8, no. 11, pp. 1052-1061, 2017.

[31] Bushfire Cooperative Research Centre, “Assessing Aerial Suppression Drop Effectivenss,” Bushfire Cooperative Research Centre, Victoria, 2009.