Determining Optimal Asset Location for Rapid and Efficient Wildfire Suppression: A Simulation-Based Approach

Daniel Hulse*, Seydou Mbaye†, Misty D. Davies‡
Intelligent Systems Division, NASA Ames Research Center, Moffett Field, CA, 94035, USA

The impact of wildfire incidents has been growing in recent years, posing a serious threat to communities at the urban-wildland interface. To address this problem, there have been growing calls to use UAVs to increase the capacity of responsible agencies to quickly and effectively suppress fires and to reduce risks associated with firefighting. One of the opportunities associated with UAVs is the ability to rapidly and autonomously operate from limited-access air bases where fires are expected to burn. This study provides an approach to determine where these air bases should be placed in order to most rapidly extinguish fires, given provided fuel distributions. This approach uses an integrated simulation of fire propagation and UAV-based suppression actions to determine how much of a given environment was burned over a range of scenarios. It then uses an optimization method to explore the space and determine the location with the least burned area. Results show the approach to efficiently and effectively provide optimal bases for single-base placements over a range of scenarios, though future work is required to adequately calibrate the model and study how it can be used in multiple-base placement problems.

I. Introduction

Wildfires are becoming an increasing threat in the United States of America [1]. In addition to threatening the natural ecosystem [2], wildland fires also endanger human lives [3], and their management comes with high economic cost [4, 5]. A report from the National Interagency Fire Center indicated an average 61,410 annual wildfires between 2013 and 2022 (burning an average of 7.2 million acres per year). Additionally, per the same report 68,988 wildfire incidents were recorded in 2022, marking a net increase of about 10,000 from 2021 [6]. The resulting destruction from wildfires can lead to significant economic losses. A 2017 report from the National Institute for Standards and Technology indicated that the estimated economic cost of managing wildfires ranges from \$71.1 billion to \$347.8 billion annually [5].

Given the danger they present [7] and how rapidly they can spread, a rapid and efficient response is essential to contain wildfires and thus mitigate risks to people and property. However, wildland firefighting is also highly dangerous job. The National Institute for Occupational Safety and Health (NIOSH) reported over 400 fatalities among on-duty wildfire fighters between 2000 and 2019 [8]. The same NIOSH report also notes that wildfire firefighters can be at risk of facing sudden cardiac deaths as a result of long exposure and physical exertion [8]. The inherent dangers of firefighting further limit the ability of firefighters to perform unsafe firefighting actions, such as flying aircraft through smoke or creating fire breaks in areas that are too close to advancing flames.

While much of wildland firefighting occurs on the ground, aerial support often plays an important role (such as aerial spotting, moving personnel, and/or wildfire suppression) in modern firefighting operations, especially for larger, more remote, and higher-risk fires. Aerial wildfire suppression operations are specifically conducted through specialized aircraft that can drop water or fire retardant to slow or contain advancing flames. Because wildland fires can change over time due to changes in environmental conditions (e.g., wind, humidity, etc.) [9], it is critical to have aerial resources strategically coordinated with ground operations. To enable the analysis and planning of wildfire management operations, simulation can be used to determine where the fire is likely to spread to (and, thus, where to conduct operations) [10].

The use of drones in wildfire management has been extensively documented in the literature [11–13]. Firefighting experts have additionally stated that unmanned aircrafts have the potential to significantly help efficiently mitigate wildland fires while decreasing risks to humans [14]. However, there are still questions about how and where they should be deployed, both in the near and long-term. Some commonly-envisioned (and, in many cases, deployed) near-term

*Software Systems AST, NASA Ames Research Center, daniel.e.hulse@nasa.gov.
†Researcher, NASA Ames Research Center, seydou.mbaye@nasa.gov
‡Associate Chief for Aeronautics Systems - Intelligent Systems (TI) Division, Airspace Operations and Safety Program (AOSP) Technical Advisor for Assurance and Safety, AIAA Fellow, misty.d.davies@nasa.gov

roles for UAVs in wildland firefighting include surveillance and small item delivery. In the long term, however, there is a broader opportunity to use UAVs to automate dangerous tasks currently performed by human pilots, such as performing water and retardant drops [15–18]. Automating wildland fire mitigation in this way presents a number of opportunities and challenges for fire protection agencies. One of these opportunities is that UAVs need not deploy from fully-staffed airports, but may rapidly deploy from limited-access air bases scattered throughout the wildland area, closer to where fires are expected to start. One challenge related to this opportunity is selecting where to best locate these air bases given limited resources.

To solve this problem, this work develops a simulation-based approach to optimize base placements provided given environmental information, including fuel distributions and potential ignition scenarios. This approach uses a high-level integrated simulation of fire propagation and air-based fire suppression to quantify firefighting metrics (e.g., burn area or containment) from a given fire scenario. An optimization algorithm is then used to evaluate and ultimately select potential air base locations that minimize the firefighting metrics. The rest of this paper is organized as follows: Section II provides background on optimal wildfire management as well as more general resilience optimization methodologies. Section III then details the optimization and simulation methodology, with Section IV presenting and discussing the results of an initial demonstration, including its assumptions and limitations. Finally, Section V summarizes findings and highlights avenues for future work.

II. Background

This section provides some relevant background on the simulation and optimization framework used in this paper, as well as previous efforts in Wildfire Management Optimization.

A. Resilience Optimization

Designing resilient systems means ensuring that they will efficiently function under hazardous scenarios (e.g., disasters) in dynamic environments [19]. The optimization of resilience is the use of dedicated algorithms to improve this type of design problem by formulating it as a mathematical optimization problem. Some of the major difficulties in resilience optimization include ensuring that resilience is appropriately balanced against other considerations (such as design and operational cost), accurately and efficiently modelling the effects of variables, and managing the computational cost of algorithms as they optimize objectives over large sets of scenarios. Many frameworks and algorithms have thus been formulated for resilience optimization (such as [20–23]), often with domain-specific considerations in mind.

In prior work, the authors of this study provided an overall framework for understanding resilience optimization problems [24] and studied how various optimization approaches in the literature fit within this framework [25]. To summarize these works, resilience optimization problems can be formulated as an optimization of the design to a wide range of scenarios (known as Resilience-based Design Optimization, or RDO), an optimization of the contingency management policy in a particular scenario (known as Resilience Policy Optimization, or RPO), or as an integrated problem of optimizing both the design and contingency management policy (known as Integrated Resilience Optimization, or IRO). Typically, while solving RPO problems may be helpful for understanding what can be done in-time to various catastrophes, their applicability at design-time is limited because it is not certain at design time what specific catastrophe(s) the system undergo (and thus need to mitigate). Additionally, while solving RDO problems can be used to optimize different design features over sets of hazardous scenarios, it is often difficult to fully understand how these features will be leveraged in particular scenarios without an idea of what the corresponding (scenario-specific) resilience policies will be. The IRO formulation solves this problem by solving both formulations simultaneously. However, this comes at the expense of computational resources. Thus, various architectures have been explored to improve the computational performance of algorithms.

The approach outlined in Section III.B is a formulation of wildfire base placement as a Resilience-based Design Optimization (RDO) problem. The main difference between this problem (given in Equation 3) and the general RDO formulation is the lack of a design cost (assumed constant here).

B. Wildfire Mitigation Optimization

From a strategy perspective, wildfire management can be difficult because wildfire behavior is subject to a range of uncertainties [26] that make it difficult to predict. At the same time, wildfire management authorities have limited resources which they may deploy to fight fires. Thus, strategic decision-making is needed to ensure that resources are effectively used to safely mitigate wildfires.

Several studies have thus considered the optimization of wildfire response as a resource allocation problem. Rideout et al. [27] proposed a methodology leveraging Integer Linear Programming with a goal to optimize allocation of resources, such as manpower or engines, for multiple wildfire events. Their framework, which is based on a "least cost plus loss," defines what fire event to prioritize for resource deployment, contingent upon available budget. Griffith et al. [28] combined Monte Carlo tree search and mixed-integer optimization for dynamic resource allocation in a tactical wildfire management scenario. Heyns et al. [29] further presented an optimization framework able to define optimal locations for camera towers. Their model outperforms traditional methods, and feedback from experts showed that the results are implementable and replicable [29]. Chan et al. [30] additionally developed a novel algorithm, Firefly, which relies on drones that monitor firelines and inform resource allocation. Their framework considers fire spread and zone prioritization as unknowns that can be uncovered over a time period.

The common thread of these aforementioned studies is that they all specifically deal entirely or in part with the use of UAVs to help with allocating and deploying resources to a given wildfire incident. This is a resilience policy optimization problem, because it deals with the in-time management of a given wildfire event, rather than the optimization of infrastructure across a range of events. This paper, on the other hand, deals with the placement and use of air bases for UAV-based wildfire suppression over a range of potential fires which could occur, which is a resilient design optimization formulation.

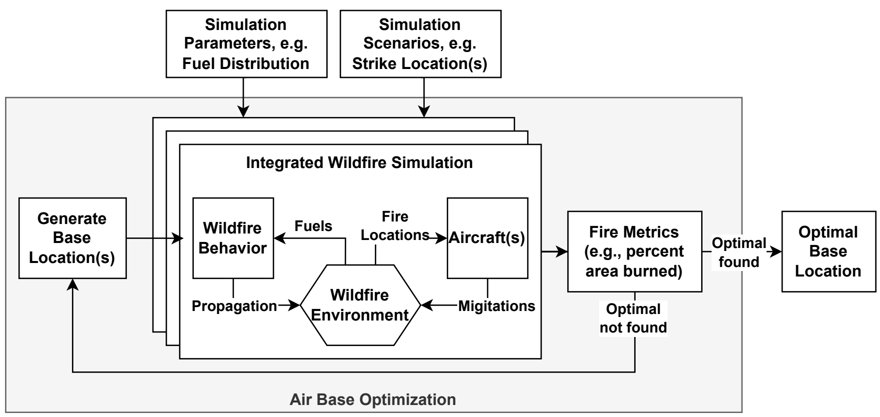

Fig. 1 Proposed simulation and optimization methodology.

III. Methodology

To systematically allocate UAV air bases for wildfire suppression, this paper proposes the approach depicted in Fig. 1, which optimizes an integrated simulation of wildfire propagation and suppression over a range of wildfire scenarios. The goal of this approach is to automatically generate optimal air base placements, given an input map of fuel distributions and environmental conditions, over a range of likely wildfire scenarios. To use this approach, one would (1) build the model behaviors and ensure that they match with the expected performance of the UAVs, (2) identify likely scenarios and assumptions for simulation parameters, including fuel distributions, (3) define wildfire risk management objectives and constraints, and then (4) implement and run an optimization algorithm to minimize the given objectives while satisfying the given constraints. The next subsections describe the details and assumptions built into the simulation used here, as well as the formulation of air base placement as a resilience-based design optimization problem.

A. Integrated Wildfire Propagation and Suppression Model

The integrated wildfire propagation and suppression model is used to determine the effects of wildfire suppression actions (resulting from a particular base placement) on fire propagation in different wildfire scenarios. In previous work, the authors developed a similar combined fire propagation and suppression model [31] in the fmdtools simulation package (see: [32]) to understand the effects of different firefighting concepts of operations on wildfire suppression [**?**]. While this model was able to provide some insight into how communications could affect wildfire suppression, it had a number of inherent limitations, namely: (1) it was slow to execute and difficult/unwieldy to parameterize and thus optimize, (2) there was little-to-no ability to tailor it to a specific real-world use-case (e.g., a wildland map or set of operators), and (3) it lacked important behaviors having to do with human/aircraft interactions such as distributed situational awareness and airspace de-confliction.

As a part of this work, a much more flexible model is being developed, called the Aerial Disaster Response Model (`aerialdrm`). The goal of `aerialdrm` is to enable the adaptable analysis of aircraft over a range of aerial disaster response missions. To enable this, the `aerialdrm` project leverages the improvements in `fmdtools` 2.0 (see: [33]), many of which were specifically developed for the lightweight and parameterizable modelling of environments (e.g., maps of terrain, fuels, or urban infrastructure), and operators/communications [34]. The next subsections describe aspects of this model which are being developed to model wildfire response scenarios, including wildfire propagation and aircraft-based suppression.

1. Wildfire Modelling

Fuel	Ignition Time	Ext. Time
Grass	25 min	45 min
Scrub	37 min	100 min
Forest	50 min	200 min

(a) Fuel Properties

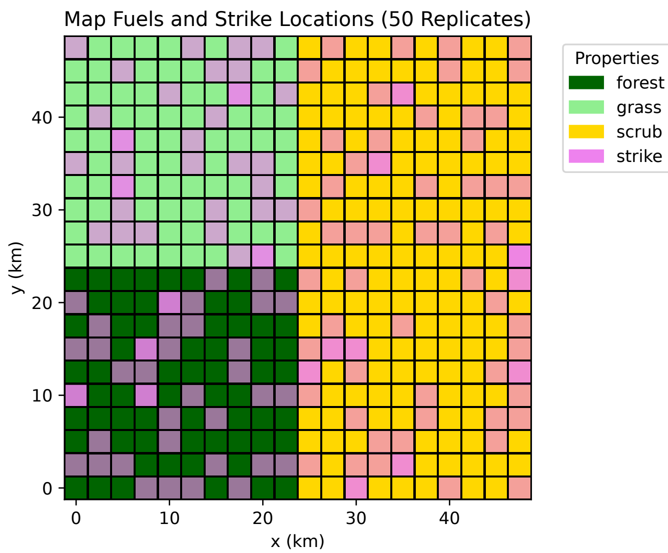

(b) Fire Propagation Map and Strike locations over given scenarios

Fig. 2 Environmental fuel properties.

The `aerialdrm` simulation represents fire propagation using a grid with a variety of different fuels (forest, grass, and scrubland) with properties that determine the ignition and extinguishment time of the fire in a given grid point. These fuel distributions may be treated as an input, or set based on preset maps. To start the fire, a lighting strike may be placed at any given grid point (or set of points), which sets the `burning` status of that point to `True`. For demonstration purposes, this map was set to the map shown in Figure 2b, with strikes randomly placed at the points shown (over 20 scenarios).

To propagate the fire, the ignition time of points adjacent to existing burning points are incremented down at each timestep until the increment reaches 0, at which point their `burning` status is also set to `True`. When these points are `burning` they also increment their burn time state until it reaches 0, at which point the grid point is `extinguished`.

2. Aircraft Modelling

Aircraft in the model are deployed out of base locations to perform retardant drops. As a part of their overall missions, aircraft proceed through a sequence of modes shown in Figure 3a. In short, the aircraft starts at a base, then flies to a fire, mitigates a fire, flies back to base, and then resupplies at the base. To further explain these modes:

1) When flying to the fire (the `fly_to_fire` mode), the closest fire that may spread to a nearby grid point is selected, which the aircraft flies to at a rate determined by its speed ($6.6\frac{km}{min}$).
2) To mitigate the fire (the `mitigate_fire` mode), the aircraft transitions the grid point from `burning` to `extinguished` states. This is a somewhat over-optimistic assumption given the grid point size and is expected to change in future work, but was used here for modelling simplicity.

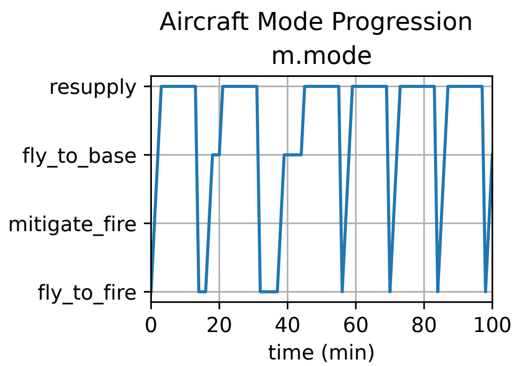

(a) **Example aircraft procession through modes and states over time.**

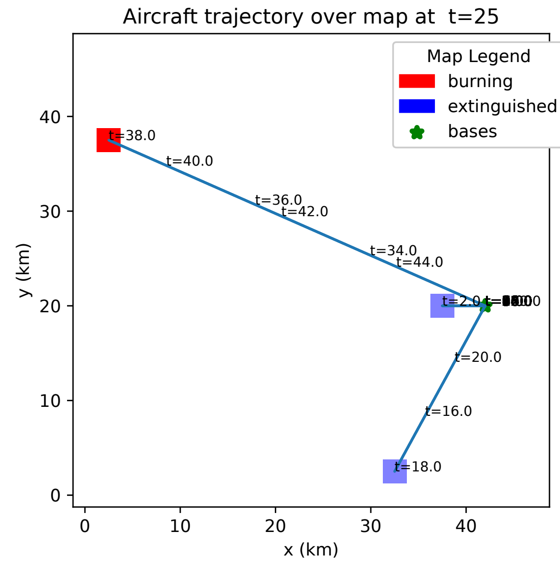

(b) **Example aircraft trajectory superimposed on a given wildfire state.**

Fig. 3 Aircraft behavior over the course of a simulation.

3) The aircraft then flies back to the base that it operates out of in the `fly_to_base` mode. Aircraft may only operate out of a single base.
4) The aircraft then enters the `resupply` mode, during which it is at the base for a set amount of time, and after which it can then fly to the fire again.

The flight behavior associated with these modes in a controlled scenario (in which there is no fire propagation) is shown in Figure 3b, where the trajectory of the aircraft is superimposed on a fire map showing a time-step in the middle of operations. As shown, the aircraft first mitigates nearby fires (which are mitigated by $t = 25$ minutes), after which it flies to base and then mitigates the fire that is further away (not yet mitigated at $t = 25$ minutes). In this sense, base placement not only determines the distance the aircraft must travel to mitigate fires, but also determines which fires are most likely to be mitigated first.

3. Integrated Simulation

In the integrated simulation, the fire propagates as the aircraft performs its mitigation actions. The structure of this simulation is shown in the central portion of Figure 1. As shown, the overall simulation is composed of the wildfire behavior, aircraft(s), and a shared wildfire environment. This simulation is parameterized in terms of fuel distributions (giving one the ability to change the map properties) as well as the number and location of air bases, and simulation/strike scenarios. It further can be run to various end-conditions, such as a specified timestep, or the containment of the fire.

The overall behavior of the integrated model is shown in Figure 4. In this version of the simulation, there are three lightning strikes and one base. As shown, the aircraft is able to mitigate two of the three before they are able to spread further. However, there is one left which spreads to adjacent points. Given this, the aircraft continues to mitigate this fire until it is put out at $t = 147$ minutes. While the integrated simulation can be given a range of fuel distributions (as well as different grid scales and sizes), this version was used for demonstration purposes to enable the intuitive verification of expected results (which should be easy to predict and interpret when the fuels are uniform and distinct).

4. Metrics

To quantify the overall impact of a fire in this integrated simulation, a method was implemented to calculate the percentage of the area that was burned by the end of the simulation. This can be calculated using the equation:

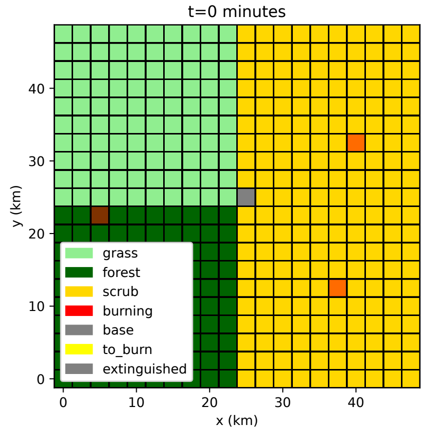

(a) Start of simulation: fires started by 3 lightning strikes.

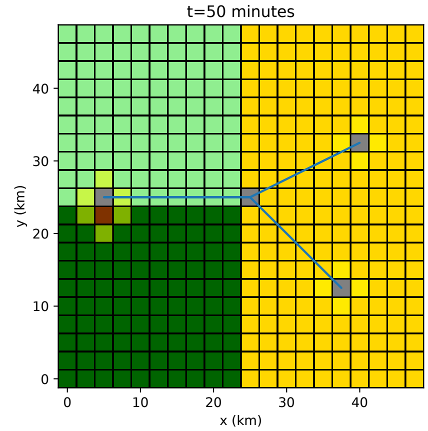

(b) UAV is able to mitigate two fires before they spread.

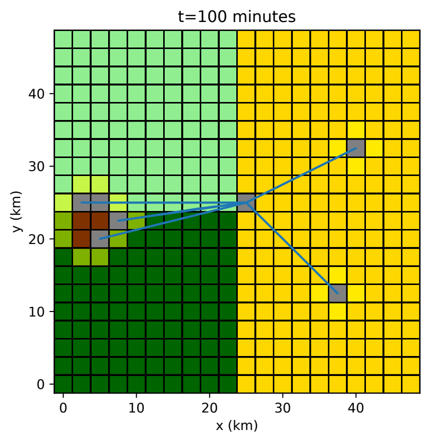

(c) UAV must now contain remaining fire.

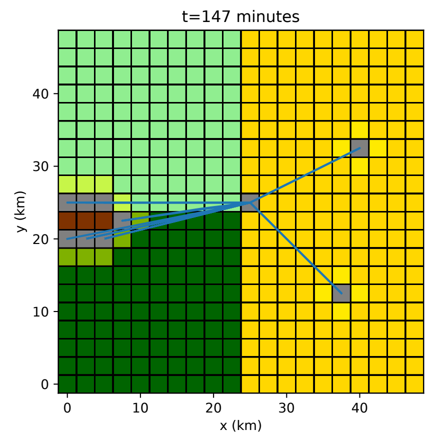

(d) Fire contained.

Fig. 4 Combined propagation and mitigation of fire over time in a three-fire scenario.

$$p_b = \frac{100}{n} \sum_{i=0}^{n} \mathbb{1}_{B \vee E}(i) \tag{1}$$

where:

- p_b is the percent burned,
- n is the number of grid points,
- i is the grid point index,
- B is the set of burning points, and
- E is the set of extinguished points.

This metric is calculated at the end of the simulation when the fire is contained, or a maximum time ($t = 400$ minutes for this simulation) has been reached. To quantify the percent burned over a range of scenarios (in this case, strike locations), the expectation of this metric is taken using the Monte Carlo method (i.e., averaging the metric over all randomly-generated scenarios) using:

$$\mathop{\mathbb{E}}_{s}\{p_b\} = \frac{1}{n_s} \sum_{s=0}^{n_s} p_b(s) \tag{2}$$

where:

- s is a scenario (e.g., individual set of initial strikes),
- S is the set of scenarios, and

- n_s is the number of scenarios.

This metric may be used to quantify how large fires are expected to grow over a range of potential fire scenarios that could occur.

B. Optimization

The integrated fire suppression model can then be used to optimize the placements of base(s) to minimize expected fuel burn given the provided distributions of fuels. The next subsections discuss the formulation of this problem as well as an algorithm which may be used to solve it.

1. Formulation

The optimization of air base locations is a Resilience-based Design Optimization problem with coupled control parameters (see: [35]), a formulation which is used very often in the optimization of infrastructure to natural disasters (e.g., [36, 37]). In this type of problem, a design (in this case, the placement of the base) is to be used and thus optimized over a wide range of circumstances. This design also implicitly determines the policy for how each individual scenario is mitigated–in this case, what fires are selected over time to be mitigated based on air base proximity. Assuming that there is not a design or operational cost associated with placing a base in a particular location, this may be formulated as as the optimization problem:

$$\begin{aligned} &\min_{\vec{x}} \mathop{\mathbb{E}}_{s} p_b(\vec{x}) \\ &s.t.\ \vec{x}_{min} \leq \vec{x} \leq \vec{x}_{max} \end{aligned} \tag{3}$$

where $\vec{x}$ is the design vector (in this case, air base x and y-coordinates) and $\vec{x}_{min}$ and $\vec{x}_{max}$ are the bounds of the grid.

2. Algorithm

In this work, the DIRECT algorithm was used to solve the optimization problem. DIRECT (DIviding RECTangles) is a global algorithm that functions by iteratively dividing and sampling rectangles in the design space. Based on the merit of the given sample points, further rectangles are divided around the best sample point(s) until the solution reaches a desired accuracy. An overview of the DIRECT algorithm is provided in Ref [38].

The DIRECT algorithm was chosen to solve this problem because of (1) the non-differentiable (and somewhat noisy) nature of the objective and variables, (2) The low number of variables ($2n$, where n is the number of bases), and (3) The authors' observation of its empirical robustness on this problem compared to other algorithms such as Nelder-Mead and Powell's method. The implementation chosen was the locally-biased variant of DIRECT [39] that is built in to `scipy.optimize` [40].

IV. Results

To demonstrate the ability of this simulation-based approach to optimize air base placements, this work considered the simple problem of placing a single base on the fuel distribution, and subject to 50 different 3-strike wildfire scenarios. The provided fuel distribution and initial strike locations across the given scenarios are shown in Figure 2b. To optimize the placement of the base, the DIRECT algorithm was used on this problem over 100 iterations. The progression of the DIRECT algorithm on this problem is shown in Figure 5. As shown in (Figure 5a), the algorithm successively explored smaller regions of the design space, eventually narrowing into a region of the scrub-land northeast of the center of the map. This exploration yields an optimal value over time shown in 5b in which there is first a major jump between values, and then there is more time taken refining the solution within the grassland. This shows the effectiveness of the algorithm and the efficiency of the overall approach on this simple problem–not only is optimal base placement able to substantially improve firefighting performance, but an entire optimization (over 50 scenarios) is able to be performed in under 15 minutes.

Running the optimization algorithm on this problem yielded substantially improved results from the simulation over a naive (centered) base placement. Figure 6 shows the average percentage burned areas returned by the simulations over the scenarios which were optimized over. In the naive placement (where the base is placed in the center), as shown in Figure 6a, there is a high frequency of burned area throughout the grassland area (often in the range of 16-18%). This means that fires are often spreading uncontrolled in the grassland area, before spreading into the scrubland, where the fire may continue to propagate (more slowly) but is more often extinguished.

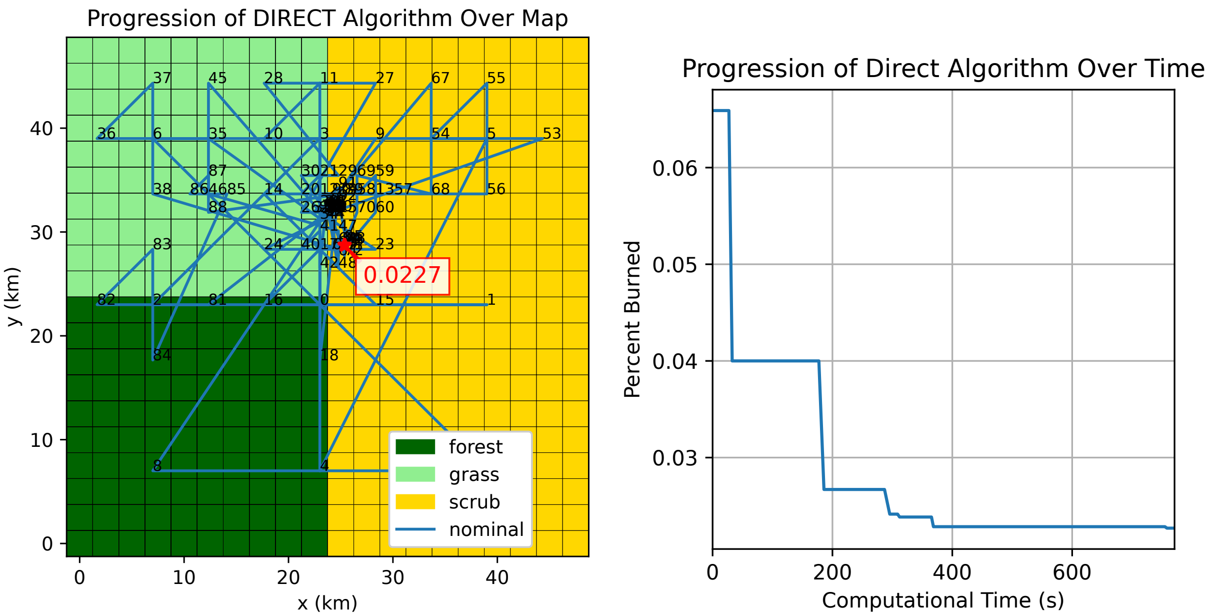

(a) Exploration of Direct Algorithm over map.

(b) Progress of optimization over time.

Fig. 5 Progression of DIRECT Algorithm.

In the optimized base allocation, on the other hand, there is much less overall burned area, as shown in Figure 6b. Note also the change in overall scale of burn frequency, with no single pixel burning over 16% of the time. Additionally, while some fires do escape the grassland, the frequency is greatly reduced. This is expected, because the the base is placed further north in the map, where it can more easily (and is more likely to prioritize) mitigate fires in the grassland and prevent their escape into/through the scrub-land. The resulting area with notable frequent spread is the south-eastern scrub-land areas that are farthest away from the base. This is expected because this is the (non-forest) area that is farthest from the base and thus is likely to be mitigated last.

Table 1 Results of Optimization Algorithm for Base Placement Problem

	Base Location (km)	Avg. Percent Burned Area
Unoptimized	23.75, 23.75	6.98
Optimized	25.37, 28.73	2.27

The results of the optimization are summarized in Table 1. As shown, the optimized base placement improves the average percentage of the area burned substantially, from 7% to 2.3%. This optimal placement thus results in a substantial decrease from the un-optimized center-point solution, because it causes the aircraft to prioritize (and more easily mitigate) wildfires in the top half of the map, where the average ignition time is lower (being both grassland and scrub). This is expected because faster-spreading fires more readily influence the objective (percent burned area). Interestingly, even though the grassland has a lower ignition time than scrub time, the optimal point is closer to the scrub land, which seems to be a reflection of the higher forest ignition time, as well as the fact that grass fires cannot propagate into the forest (since the forest ignition time is higher than grass extinguish time).

This shows the potential value of simulation-based air base placement: by enabling the quantification of the severity of a range of different fire scenarios, it can enable one to best target particularly wildfire-prone areas while ensuring all areas are adequately covered.

A. Discussion and Limitations

While the results demonstrate the potential of the overall simulation-based approach, there are a range of limitations which hinder its practical real-world use, including modelling and simulation assumptions and the formulation of the

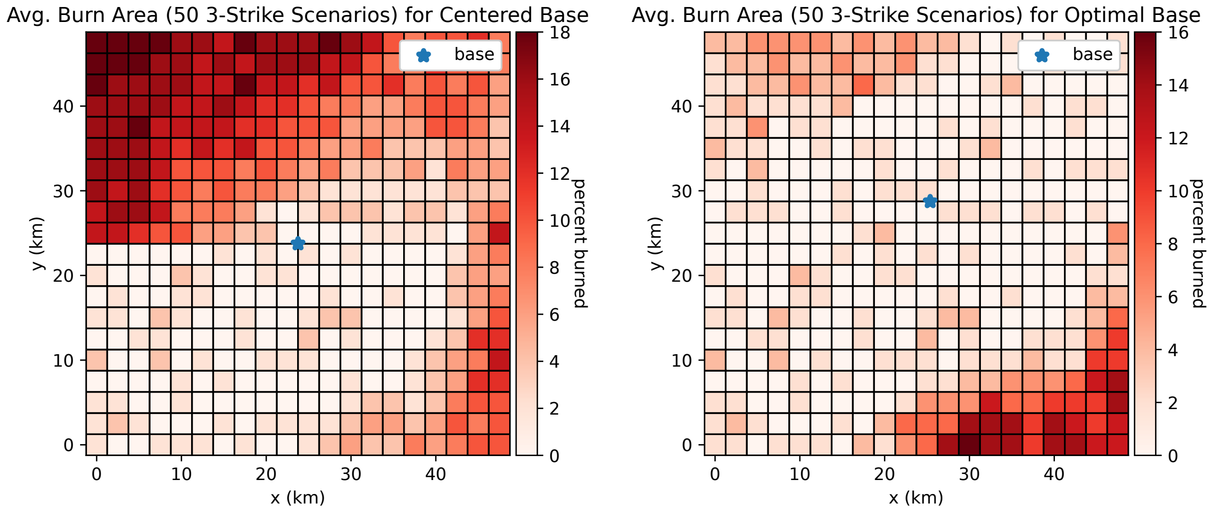

(a) Unoptimized base configuration and results.

(b) Optimal base configuration and results.

Fig. 6 Average burned areas given unoptimized (left) and optimized (right) base placements.

optimization problem.

As of yet, modelling assumptions have not been calibrated against real-world fire propagation and scenarios, and thus should not be taken as a realistic representation of wildfire propagation or suppression behavior. Some of the major issues currently present in this fire model are that (1) wildfire mitigation actions undertaken by the Aircraft are 100% effective and require very little time (1 minute/timestep), (2) because the grid size is limited, fires cannot propagate outside the grid, which artificially limits fire spread, and (3) the environment is only subject to three-strike fires with no directional wind. While, in some ways, the scenario difficulty (three-strike fire scenarios) counteracts the mitigation effectiveness assumption for the purpose of optimization (while increasing the number of random strike locations to sample from), in order to be used more predicatively, the model itself should be better calibrated to replicate real-world fire behavior. Thus, future work should calibrate the parameters and assumptions of this model to more justifiably match real-world scenarios.

Additionally, the way the problem was formulated imposes some limitations on the real-world usefulness of this approach. Formulating the optimization problem as minimizing percent burned area neglects the risk-related considerations (e.g., will the fire destroy homes or infrastructure) that often drive firefighting decisions. Additionally, the problem formulation (Resilient Design Optimization with Coupled Control) uses fairly simplistic assumptions for firefighting policy, since the aircraft always mitigate the nearest fire location first, rather than prioritizing spread rate. This may lead to base placement results which overstate the need to move the base in order force in-time firefighting decisions about which fires to fight first. To resolve this problem, it may be necessary to either improve the wildfire prioritization method or formulate the optimization problem as a nested base placement design problem and wildfire management problem. This would ensure that the evaluation of base locations is based on optimal (or, at least, more realistic) utilization of the UAVs. Future work should explore alternative formulations to this problem to study how they may affect optimal base placement decisions.

Finally, the problem considered here was relatively simple, looking at only the placement of a single base. This made it a good candidate for solution with the DIRECT algorithm, which performs well on low-dimensional (< 6-variable) problems. However, in reality stakeholders may be more interested in placing multiple bases around a larger area. Future work should thus extend this method to multi-base allocation problems, as well as investigate suitable algorithms for these problems.

V. Conclusion

This work marks a first step towards understanding how to place bases for UAV-based wildfire suppression. Towards this goal, this work demonstrated the use of parameterized simulation for modelling fire propagation which may be provided different fuel distributions. It further demonstrated the use of the DIRECT algorithm in optimizing the base location to minimize fuel burn over a range of scenarios. As shown, the overall approach is able to substantially improve the simulated wildfire response, and the optimization algorithm performs well on the multi-modal nature of the problem. Nevertheless, it should be stressed that this is in-progress work with a wide range of limitations, including (as discussed in Section IV.A), modelling assumptions and determining the impact of base placement on wildfire prioritization decisions.

Future work should resolve these limitations while providing an expanded study and set of capabilities. Specifically, extending the study into the placement of more than one base could generate more useful insight into how to allocate bases throughout a firefighting area. Additionally, future work should apply this method on a range of different maps, including maps generated from real-world fuel distributions. Finally, future work should also expand the number of metrics/objectives under consideration such that the model can be used for more holistic decision-making (e.g., based on expected harms to people and property) rather than solely prioritizing minimizing fuel burn.

Acknowledgments

This research was funded by the System-Wide Safety project in the NASA Aeronautics Research Mission Directorate. The findings herein represent the research of the authors and do not necessarily represent the view of the U.S. Government or NASA. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the U.S. Government.

References

[1] Sanderfoot, O. V., Bassing, S. B., Brusa, J. L., Emmet, R. L., Gillman, S. J., Swift, K., and Gardner, B., "A review of the effects of wildfire smoke on the health and behavior of wildlife," *Environmental Research Letters*, Vol. 16, No. 12, 2022, p. 123003. https://doi.org/10.1088/1748-9326/ac30f6.

[2] Robinne, F.-N., Hallema, D. W., Bladon, K. D., Flannigan, M. D., Boisramé, G., Bréthaut, C. M., Doerr, S. H., Di Baldassarre, G., Gallagher, L. A., Hohner, A. K., et al., "Scientists' warning on extreme wildfire risks to water supply," *Hydrological Processes*, Vol. 35, No. 5, 2021, p. e14086. https://doi.org/10.1002/hyp.14086.

[3] Butler, C., Marsh, S., Domitrovich, J. W., and Helmkamp, J., "Wildland firefighter deaths in the United States: A comparison of existing surveillance systems," *Journal of occupational and environmental hygiene*, Vol. 14, No. 4, 2017, pp. 258–270. https://doi.org/10.1080/15459624.2016.1250004.

[4] Brown, E. K., Wang, J., and Feng, Y., "US wildfire potential: a historical view and future projection using high-resolution climate data," *Environmental Research Letters*, Vol. 16, No. 3, 2021, p. 034060. https://doi.org/10.1088/1748-9326/aba868, URL https://dx.doi.org/10.1088/1748-9326/aba868.

[5] Thomas, D., Butry, D., Gilbert, S., Webb, D., and Fung, J., "The Costs and Losses of Wildfires: A Literature Review," 2017. https://doi.org/10.6028/NIST.SP.1215.

[6] National Interagency Coordination Center, "Wildland Fire Summary and Statistics Annual Report 2022," , 2022. URL https://www.nifc.gov/sites/default/files/NICC/2-Predictive%20Services/Intelligence/Annual%20Reports/2022/annual_report.2.pdf.

[7] Stawovy, L., and Balakrishnan, B., "The threat of wildfires and pulmonary complications: A narrative review," *Current Pulmonology Reports*, Vol. 11, No. 4, 2022, pp. 99–105. https://doi.org/10.1007/s13665-022-00293-7.

[8] "Fighting Wildfires," , May 2023. URL https://www.cdc.gov/niosh/topics/firefighting/default.html.

[9] Lagouvardos, K., Kotroni, V., Giannaros, T. M., and Dafis, S., "Meteorological Conditions Conducive to the Rapid Spread of the Deadly Wildfire in Eastern Attica, Greece," *Bulletin of the American Meteorological Society*, Vol. 100, No. 11, 2019, pp. 2137 – 2145. https://doi.org/10.1175/BAMS-D-18-0231.1, URL https://journals.ametsoc.org/view/journals/bams/100/11/bams-d-18-0231.1.xml.

[10] Sullivan, A. L., "Wildland surface fire spread modelling, 1990–2007. 3: Simulation and mathematical analogue models," *International Journal of Wildland Fire*, Vol. 18, No. 4, 2009, pp. 387–403. https://doi.org/10.1071/WF06144.

[11] Hossain, F. A., Zhang, Y. M., and Tonima, M. A., "Forest fire flame and smoke detection from UAV-captured images using fire-specific color features and multi-color space local binary pattern," *Journal of Unmanned Vehicle Systems*, Vol. 8, No. 4, 2020, pp. 285–309. https://doi.org/10.1139/juvs-2020-0009, URL https://doi.org/10.1139/juvs-2020-0009.

[12] Thul, L., and Powell, W., "An Information-Collecting Drone Management Problem for Wildfire Mitigation," 2023. https://doi.org/10.48550/arXiv.2301.07013.

[13] Akhloufi, M. A., Couturier, A., and Castro, N. A., "Unmanned Aerial Vehicles for Wildland Fires: Sensing, Perception, Cooperation and Assistance," *Drones*, Vol. 5, No. 1, 2021. https://doi.org/10.3390/drones5010015, URL https://www.mdpi.com/2504-446X/5/1/15.

[14] Mbaye, S., Jones, G., and Davies, M. D., "Analysis of Input from Wildfire Incident Experts to Identify Key Risks and Hazards in Wildfire Emergency Response," *AIAA SCITECH 2023 Forum*, 2023. https://doi.org/10.2514/6.2023-2710, URL https://arc.aiaa.org/doi/abs/10.2514/6.2023-2710.

[15] Bradshaw, H., "Drone swarms could stop wildfires, researchers say," *BBC*, 2024. URL https://www.bbc.com/news/articles/cg3e4pw294po.

[16] Cambell, M., "Firefighting drones prove they can help contain wildfires faster and safer," *Spectrum News*, 2022. URL https://spectrumlocalnews.com/tx/south-texas-el-paso/weather/2022/09/29/fire-fighting-drones-prove-they-can-help-contain-wildfires-faster-and-safer.

[17] Ausonio, E., Bagnerini, P., and Ghio, M., "Drone Swarms in Fire Suppression Activities: A Conceptual Framework," *Drones*, Vol. 5, No. 1, 2021. https://doi.org/10.3390/drones5010017, URL https://www.mdpi.com/2504-446X/5/1/17.

[18] Aydin, B., Selvi, E., Tao, J., and Starek, M. J., "Use of Fire-Extinguishing Balls for a Conceptual System of Drone-Assisted Wildfire Fighting," *Drones*, Vol. 3, No. 1, 2019. https://doi.org/10.3390/drones3010017, URL https://www.mdpi.com/2504-446X/3/1/17.

[19] Ehlen, M., Sanyal, P., and Pepple, M., "A Framework for Guiding Homeland Security Programs from Risk to Resilience," 2015. URL https://www.osti.gov/biblio/1648691.

[20] Wang, Z., Wen, Y., Wang, Z., and Zhi, P., "Resilience-based design optimization of engineering systems under degradation and different maintenance strategy," *Structural and Multidisciplinary Optimization*, Vol. 66, 2023. https://doi.org/10.1007/s00158-023-03671-0.

[21] FENG, Q., HAI, X., SUN, B., REN, Y., WANG, Z., YANG, D., HU, Y., and FENG, R., "Resilience optimization for multi-UAV formation reconfiguration via enhanced pigeon-inspired optimization," *Chinese Journal of Aeronautics*, Vol. 35, No. 1, 2022, pp. 110–123. https://doi.org/https://doi.org/10.1016/j.cja.2020.10.029, URL https://www.sciencedirect.com/science/article/pii/S1000936120305616.

[22] Dui, H., Wang, X., and Zhou, H., "Redundancy-Based Resilience Optimization of Multi-Component Systems," *Mathematics*, Vol. 11, No. 14, 2023. https://doi.org/10.3390/math11143151, URL https://www.mdpi.com/2227-7390/11/14/3151.

[23] Kong, J., Zhang, C., and Simonovic, S. P., "Optimizing the resilience of interdependent infrastructures to regional natural hazards with combined improvement measures," *Reliability Engineering & System Safety*, Vol. 210, 2021, p. 107538. https://doi.org/https://doi.org/10.1016/j.ress.2021.107538, URL https://www.sciencedirect.com/science/article/pii/S095183202100096X.

[24] Hulse, D., and Hoyle, C., "Understanding Resilience Optimization Architectures: Alignment and Coupling in Multilevel Decomposition Strategies," *Journal of Mechanical Design*, Vol. 144, No. 11, 2022, p. 111704. https://doi.org/10.1115/1.4054993, URL https://doi.org/10.1115/1.4054993.

[25] Hulse, D., Biswas, A., Hoyle, C., Tumer, I. Y., Kulkarni, C., and Goebel, K., "Exploring Architectures for Integrated Resilience Optimization," *Journal of Aerospace Information Systems*, Vol. 18, No. 10, 2021, pp. 665–678. https://doi.org/10.2514/1.I010942, URL https://doi.org/10.2514/1.I010942.

[26] Thompson, M. P., and Calkin, D. E., "Uncertainty and risk in wildland fire management: A review," *Journal of Environmental Management*, Vol. 92, No. 8, 2011, pp. 1895–1909. https://doi.org/https://doi.org/10.1016/j.jenvman.2011.03.015, URL https://www.sciencedirect.com/science/article/pii/S0301479711000818.

[27] Rideout, D., Wei, Y., and Kirsch, A., "Optimal allocation of initial attack resources to multiple wildfire events," *International Journal of Safety and Security Engineering*, Vol. 1, 2011, pp. 312–325. https://doi.org/10.2495/SAFE-V1-N3-312-325.

[28] Griffith, J. D., Kochenderfer, M. J., Moss, R. J., Misic, V. V., Gupta, V., and Bertsimas, D., "Automated Dynamic Resource Allocation for Wildfire Suppression," 2017. URL https://api.semanticscholar.org/CorpusID:43398511.

[29] Heyns, A., du Plessis, W., Kosch, M., and Hough, G., "Optimisation of tower site locations for camera-based wildfire detection systems," *International Journal of Wildland Fire*, Vol. 28, 2019, pp. 651–665. https://doi.org/10.1071/WF18196.

[30] Chan, H., Tran-Thanh, L., and Viswanathan, V., “Fighting Wildfires under Uncertainty - A Sequential Resource Allocation Approach,” *International Joint Conference on Artificial Intelligence*, 2020. URL https://api.semanticscholar.org/CorpusID:220480938.

[31] Andrade, S., Hulse, D., Irshad, L., and Walsh, H. S., “Supporting Hazard Analysis for Wildfire Response Using fmdtools and MIKA,” 2022. URL https://ntrs.nasa.gov/citations/20220014099.

[32] Hulse, D., Walsh, H., Dong, A., Hoyle, C., Tumer, I., Kulkarni, C., and Goebel, K., “fmdtools: A fault propagation toolkit for resilience assessment in early design,” *International Journal of Prognostics and Health Management*, Vol. 12, No. 3, 2021. https://doi.org/10.36001/ijphm.2021.v12i3.2954.

[33] NASA, “fmdtools Version 2.0.5,” GitHub repository, November 2024. URL https://github.com/nasa/fmdtools.

[34] Irshad, L., and Hulse, D., “Modeling Distributed Situation Awareness in Resilience-Based Design of Complex Engineered Systems,” *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Vol. 87295, American Society of Mechanical Engineers, 2023, p. V002T02A050. https://doi.org/10.1115/DETC2023-116689, URL https://ntrs.nasa.gov/citations/20220008569.

[35] Hulse, D., and Hoyle, C., “Understanding Resilience Optimization Architectures: Alignment and Coupling in Multilevel Decomposition Strategies,” *Journal of Mechanical Design*, Vol. 144, No. 11, 2022, p. 111704. https://doi.org/10.1115/1.4054993.

[36] Miller-Hooks, E., Zhang, X., and Faturechi, R., “Measuring and maximizing resilience of freight transportation networks,” *Computers & Operations Research*, Vol. 39, No. 7, 2012, pp. 1633–1643. https://doi.org/10.1016/j.cor.2011.09.0.

[37] Wang, Z., Wang, Q., Zukerman, M., Guo, J., Wang, Y., Wang, G., Yang, J., and Moran, B., “Multiobjective path optimization for critical infrastructure links with consideration to seismic resilience,” *Computer-Aided Civil and Infrastructure Engineering*, Vol. 32, No. 10, 2017, pp. 836–855.

[38] Jones, D. R., and Martins, J. R., “The DIRECT algorithm: 25 years Later,” *Journal of global optimization*, Vol. 79, No. 3, 2021, pp. 521–566. https://doi.org/10.1007/s10898-020-00952-6.

[39] Gablonsky, J. M., and Kelley, C. T., “A locally-biased form of the DIRECT algorithm,” *Journal of Global optimization*, Vol. 21, 2001, pp. 27–37. https://doi.org/10.1023/A:1017930332101.

[40] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, Vol. 17, 2020, pp. 261–272. https://doi.org/10.1038/s41592-019-0686-2.