

Multi-UAV Planning for Cooperative Wildfire Coverage and Tracking with Quality-of-Service Guarantees

Esmaeil Seraj · Andrew Silva · Matthew Gombolay

Received: 05.15.2022 / Accepted: 05.24.2022

Abstract In recent years, teams of robot and Unmanned Aerial Vehicles (UAVs) have been commissioned by researchers to enable accurate, online wildfire coverage and tracking. While the majority of prior work focuses on the coordination and control of such multi-robot systems, to date, these UAV teams have not been given the ability to reason about a fire's track (i.e., location and propagation dynamics) to provide performance guarantee over a time horizon. Motivated by the problem of aerial wildfire monitoring, we propose a predictive framework which enables cooperation in multi-UAV teams towards collaborative field coverage and fire tracking with probabilistic performance guarantee. Our approach enables UAVs to infer the latent fire propagation dynamics for time-extended coordination in safety-critical conditions. We derive a set of novel, analytical temporal, and tracking-error bounds to enable the UAV-team to distribute their limited resources and cover the entire fire area according to the case-specific estimated states and provide a probabilistic performance guarantee. Our results are not limited to the aerial wildfire monitoring case-study and are generally applicable to problems, such as search-and-rescue, target tracking and border patrol. We evaluate our approach in simulation and provide demonstrations of the proposed framework on a physical multi-robot testbed to account for real robot dynamics and restrictions. Our quantitative evaluations validate the performance of our method accumulating $7.5\times$ and $9.0\times$ smaller tracking-error than state-of-the-art model-based and reinforcement learning benchmarks, respectively.

To appear in the Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS). This work was sponsored by ONR under grant N00014-18-S-B001, MIT Lincoln Laboratory grant 7000437192, Lockheed Martin Corporation under grant GR00000509, and GaTech institute funding.

E. Seraj
Georgia Institute of Technology, Atlanta, GA, 30332, USA
E-mail: eseraj3@gatech.edu

A. Silva
Georgia Institute of Technology, Atlanta, GA, 30332, USA
E-mail: andrew.silva@gatech.edu

M. Gombolay
Georgia Institute of Technology, Atlanta, GA, 30332, USA
E-mail: matthew.gombolay@cc.gatech.edu

Keywords Multi-Robot Coordination · Multi-Robot Systems · Planning and Scheduling · Aerial Systems · Cooperating UAVs

1 Introduction

Unmanned Aerial Vehicles (UAVs) have gained significant interest across research communities due to their agility, cost-efficiency, and ability to access and monitor remote areas. Multi-robot systems, such as teams of UAVs, have been commissioned as Mobile Sensor Networks (MSNs) in applications which require observing (i.e., monitoring and surveilling) the physical world and gathering information [1, 2]. MSNs have numerous applications in surveillance and disaster response [3, 4, 5, 6], environment monitoring and agriculture [7, 8], wildfire fighting [9, 10, 11], search-and-rescue [4], manufacturing [12, 13], homeland security and border patrol [14, 15].

While multi-robot systems are capable of executing time-sensitive, complex missions that are intractable for a single robot, it is challenging to efficiently coordinate such systems and to optimize the collaborative behavior among robots [16, 17]. This coordinated planning problem becomes even more challenging when robots have to collaborate in a dynamic environment, such as in monitoring an active wildfire. Dynamic environments are restless; meaning regardless of robots' collective actions, the states of the environment continually change. As such, classical multi-agent planning and scheduling approaches designed for static environments will fail in such domains, since coordination plans made for a timestep may be inapplicable to the next step due to the environment dynamicity [18, 19].

In this paper, we study the problem of coordinated planning of multi-UAV teams deployed as MSNs for cooperative surveillance and tracking of a dynamic environment. In our work, we enable UAVs with the ability to estimate the changing states of their environment and leverage these estimated information to reason about their cooperation plan for collaborative monitoring. Particularly, we consider safety-critical and time-sensitive scenarios where only a limited number of UAVs are available. Therefore, in our problem, not only it is important to reason robustly under environmental dynamicity via active planning with limited resources, it is exigent to have probabilistic guarantees that the UAV team can execute the coordinated plan. In our approach, the UAV agents actively evaluate a probabilistic error-bound that provides such guarantee and attempt to revise their coordinated plan when a performance guarantee cannot be provided due to changes in environment states. While the problem of multi-agent planning and scheduling has been studied extensively in prior work [18, 19, 20, 21, 22, 23], many of these works either lack the ability to actively plan in a dynamic environment, or do not provide a probabilistic performance guarantee that assures the teams' success in executing a coordinated plan.

1.1 Motivating Application: Aerial Wildfire Monitoring

We adopt the application of aerial wildfire monitoring as a running case-study and motivate our paper in this important safety-critical problem. In the application of wildfire fighting, human firefighters on the ground need online and dynamic observation of the firefront (i.e., the moving edge of fire) to anticipate a wildfire's unknown characteristics, such as size, scale, and propagation velocity, in order to plan their strategies accordingly. To support human firefighters, teams of UAVs can be deployed as MSNs to estimate the states of fire across thousands of acres and provide human firefighters with such information [24, 25, 5].

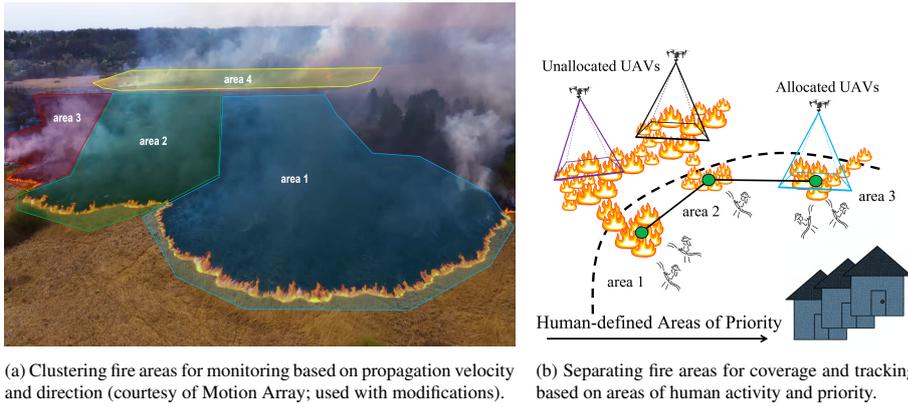


Fig. 1: Figure 1a demonstrates an example of separating wildfire areas for coverage and tracking based on fire propagation velocity and direction. The firefront in each area is moving in a different direction and with a varying velocity and, therefore, at least one separate UAV is required to monitor and track the firefront in each area. In a similar perspective, Figure 1b shows an example of disjoint coverage areas that are separated according to areas of human (i.e., firefighters) activity and priority. In this case, the autonomous UAVs are teaming with humans to safely manage their human collaborators by providing for them *high-quality* information regarding their time-varying proximity to fire (i.e., areas of human operation). When we have access to additional UAV resources that may not be needed in the prioritized areas, we can deploy such UAVs (i.e., unallocated UAVs) to monitor the rest of the wildfire areas (not specified or prioritized).

Nevertheless, coordinated control of UAV teams in this volatile setting provides particular challenges [26, 27, 28, 29, 30].

Fighting wildfires safely and effectively requires accurate online information on firefront location, size, shape, and propagation velocity [31, 32, 33, 24]. To provide firefighters with this realtime information, researchers have sought to utilize satellite feeds to estimate fire location information [34, 35, 36]. Unfortunately, the resolution of these images is too low for more than simple detection of a wildfire's existence [36]. Firefighters need frequent, high-quality images of the wildfire to make strategic plans [24, 10, 37]. Here, we define *high-quality information* as local, high-resolution images (or other sensory information) that are captured from a close by distance with respect to the areas prioritized by humans.

In large scale dynamic field coverage and tracking applications, such as aerial wildfire monitoring, UAVs need to be distributed effectively to cover the entire area of interest [9, 11]. In many cases, the region that need to be covered and tracked by the UAV team can be sub-divided into smaller areas according to different characteristics, priorities or needs. Figure 1a depicts an example of separating wildfire areas for coverage and tracking based on fire propagation velocity and direction. The firefront in each area is moving in a different direction and with a varying velocity and therefore, at least one separate UAV is required to monitor and track the firefront in each area. In a similar perspective, Figure 1b shows an example of disjoint coverage areas that are separated according to regions of human (i.e., firefighters) activity and priority. In this case, the autonomous UAVs are teaming with humans to safely manage their human collaborators by providing for them *high-quality* information regarding their time-varying proximity to fire (i.e., areas of human operation). Accordingly, when the UAV resources are limited, a multi-agent planning algorithm is required such that the real-time cooperative field coverage and tracking performance can be guaranteed with as few UAV agents as possible.

1.2 Contributions

To overcome the aforementioned limitations, herein we develop a probabilistically-guaranteed framework for real-time, large-scale coordinated planning and cooperative coverage of dynamic environments. We ground our approach in application to aerial wildfire monitoring and develop an algorithmic framework in which we explicitly infer the latent fire propagation dynamics and leverage the estimated information in real-time to plan for a guaranteed quality of service. In our predictive mechanism, we define the term “service” as actively estimating the states of a propagating wildfire via limited number of UAVs and providing human firefighters on the ground with probabilistic guarantees regarding their proximity to a fire. We design our planning and coverage (i.e., monitoring) strategies upon a multi-step Adaptive Extended Kalman Filter (AEKF) predictor and the FARSITE wildfire propagation mathematical model [38].

A primary motivation of our approach is the need for computationally lightweight, yet probabilistically-guaranteed, algorithms to scale to large numbers of UAVs and coverage areas. A central contribution of our work is a set of novel, analytical temporal and tracking-error residual bounds that allow UAVs to enable probabilistically-guaranteed coordination in monitoring and tracking dynamic targets (i.e., firespots). We derive these bounds in three different scenarios: (1) stationary fire, (2) moving fire without considerable grow (i.e., spawn), and (3) quickly moving, multiplying fire. For each case, we derive analytic and probabilistic temporal upper-bounds to facilitate performance guarantee for collaborative field coverage. The primary contributions of our work are:

1. An algorithmic planning framework based on Kalman estimation and uncertainty propagation to learn and leverage in real-time a predictive mechanism that enables probabilistic guarantees for tracking dynamic targets (i.e., firefronts).
2. Introducing the Uncertainty Residual Ratio (URR) tracking-error bound as well as a set of novel, analytical temporal upper-bounds that allow UAVs to enable probabilistically-guaranteed coordination in monitoring and tracking dynamic targets (i.e., firespots)
3. An efficient and scalable collaborative field coverage algorithm by leveraging our analytical temporal bounds for centralized planning and distributed execution.
4. Presenting quantitative and experimental results in simulation and on physical robots that show the effectiveness of our URR bound and cooperative coordinated multi-UAV planning framework by demonstrating efficient scalability and significantly outperforming two prior SOTA baselines [10, 37] in a wildfire coverage and tracking task.

We note that, despite the utility of UAVs as the robot agents and the aerial wildfire monitoring application in our paper, the proposed coordinated, collaborative planning and area coverage algorithms as well as the derived set of analytical temporal and the URR tracking-error bounds are broadly applicable to other safety-critical applications that require mobile sensors to accurately track moving targets in their environment. See Section 1.3 for a detailed discussion of the applicability of our proposed framework. We investigate the cooperative field monitoring problem by considering the wildfires in three separate scenarios of *stationary*, *moving* and *moving-spreading* target points (e.g., firespots).

In the application of wildfire fighting, human teams (i.e., firefighters) fight a wildfire until the fire is extinguished or working condition becomes too harsh to continue relative to the cost of abandoning. However, determining when a situation is too dangerous relative to the costs of abandoning the effort is a non-trivial ethical dilemma (e.g., how to measure the lives of immovable patients at a hospital in the fire’s path relative to those of the firefighters) complicated by a lack of information regarding the fire’s propagation characteristics. We

take the perspective that this decision should remain in the hands of human professionals. As such, in our case-study, we focus 1) on developing a tight, probabilistic, upper-bound to reason about the minimum number of robots required to ensure high-quality information for the human decision-makers (Section 5) and 2) measuring the tightness of our test by the number of robots required to satisfy this bound relative to other state-of-the-art approaches (Section 8).

We empirically evaluate the performance, feasibility, and scalability of our framework in various experiments, alongside a state-of-the-art model-based [10] and a reinforcement learning (RL) benchmark [37] for UAV-based aerial field coverage and coordinated planning. Our experiments demonstrate a promising utility of our approach, accumulating $7.5\times$ and $9.0\times$ smaller tracking-errors than the two benchmark methods. Moreover, we assess the feasibility of our framework through implementation on physical robots in a mock wildfire monitoring scenario. The results of our experimental evaluations are presented in a supplementary video, available on <https://youtu.be/zTR07cK1wRw>.

1.3 General Applicability of the Proposed Framework

While grounded in an application to wildfire monitoring, the proposed coordinated planning and collaborative field coverage frameworks for teams of autonomous robots are broadly applicable to other domains and problems in which one is attempting to perform intelligent tracking and monitoring of many targets with few robotic sensors in large-scale operations [39]. Such applications include search-and-rescue [40, 41, 42], tracking of wildlife poachers [43, 44, 45], oil spill surveillance in oceans [46, 47, 48], border patrol and protection [49, 39] and even air traffic control for urban air mobility [50, 51, 52]. Here, we discuss these potential applications.

Search and Rescue Missions– Considering a scenario in which a natural disaster has struck a city, a team of UAVs can be used to keep monitoring various distant search sites for detecting survivors. Here, robots need to make sure to visit each search site frequently enough to not miss any signs of motion; here, our analytical temporal and measurement-uncertainty bounds for the case of static targets can provide a probabilistic guarantee for surveillance quality. This example can be extended to other search-and-rescue scenarios where humans (groups or individuals) are lost in forests or mountains and various sites need to be searched persistently or even for animal and wildlife monitoring and control, such as tracking specific animal species.

Border Patrol and Protection– Similar to the search-and-rescue example, border patrol and protection can be a relevant application for our framework in which aerial surveillance and tracking of multiple moving target (e.g., cars) can be performed by hovering UAVs. Various distant areas each including multiple targets can be monitored with a small team of UAVs which estimate the states (i.e., position and velocity) of specific targets. For this application, our analytical bounds for dynamic targets can provide the minimum number of UAVs, required to ensure a high tracking quality of targets. A related example would be the tracking of wildlife poachers.

Oil Spill Detection and Surveillance– Detecting and surveilling oil spills in the oceans are also a closely relevant application for our frameworks. Similar to the aerial wildfire detection and monitoring, teams of autonomous UAVs can be deployed for large-scale oil spill surveillance, including multiple distant and dynamically growing spills, similar to the third case of our wildfire case-study, the moving-spreading fire areas.

2 Related Work

Recent advances in UAV technology have opened up the possibility of providing real-time, high-quality fire information to firefighting teams across thousands of acres [5, 24, 25]. However, coordinated control of UAV teams in this volatile setting provides particular challenges, such as decentralized controller design, task scheduling problems, large-scale communication, smoke detection and image stabilization, etc. [26, 27, 28, 29, 30].

Among the vision-based fire coverage methods [25, 53, 54, 55], Merino *et al.* [56] investigated the utility of visual and infrared cameras on UAVs to monitor the evolution of the firefront shape and then attempted to extend the proposed approach to a collaborative scheme performed by multiple UAVs. However, Yuan *et al.* [57] note that vision-based approaches for fire monitoring and tracking struggle with image stabilization and camera obstruction, e.g., due to smoke, which produce significant errors in sensor measurements. Without addressing these errors properly, relying on such systems can be fatal.

Our work is closely relevant to the cooperative multi-robot target tracking literature. A comprehensive review of the taxonomy and recent approaches on multi-robot target detection and tracking is provided by Robin and Lacroix in [58]. An intensity function-based algorithm is presented in [59] to generate a control law for tracking dynamic targets with a group of UAVs. The problem of dynamic distributed task allocation in ground-robot teams for coordinated target tracking is studied in [60] through a bio-inspired approach. Mottaghi *et al.* [61] developed a particle filter-based approach to create a potential field to track a moving fire with multiple robots. Hausman *et al.* [62] developed a probabilistic localization and control method for a UAV team with fixed number of robots which seeks to minimize the expected future uncertainty of the target position. Both mutual information and EKF's covariance are investigated as measures of uncertainty.

A review of the recent methods on the automation methods for wildfire remote sensing application via UAVs is presented in [63], introducing and discussing three key metrics, i.e., situational awareness, decisional ability, and collaboration ability, in the recent relevant literature. Recent work by Sujit *et al.* [33] proposed a cooperative approach to detect and monitor multiple spots of fire (i.e. hotspots) using two groups of detector and service UAV agents. Afghah *et al.* [5], proposed a distributed leader-follower coalition formation model to cluster a set of UAVs into multiple coalitions that collectively cover a designated area. Kumar *et al.* [64] proposed a cooperative control algorithm to first cooperatively track the firefront shape for accurate situational awareness and then, autonomously fight the fire using fire suppressant fluid. Ghamry *et al.* [65] proposed a distributed algorithm with multiple stages such as search, confirmation and propagation monitoring for a team of UAVs to evenly distribute in a leader-follower manner to track an elliptical fire perimeter. Pham *et al.* [10] designed a fire heat-intensity based distributed control framework for a team of UAVs to be capable of closely monitoring a wildfire in open space in order to track its development. Harikumar *et al.* [66] proposes a search and dynamic formation control framework for a multi-UAV system to efficiently search for a dynamic target in an unknown environment.

The majority of prior work on distributed control of UAVs, e.g. for MSN applications, focuses on maximizing the coverage of a particular area of interest [67, 68], such as an area of wildfire [11, 5]. Previous studies typically sought to maximize coverage (of a fire area) either through density-function based approaches [69, 70, 71, 9] or by maximizing the average pixel density across a terrain [72, 11]. In the latter case, UAV-based wildfire coverage studies tend to adopt a fire-intensity function for their coverage problem formulation [10, 67, 72, 9, 69, 70, 71] to model areas of fire according to an intensity model.

Many of the aforementioned approaches require an accurate fire-shape function [64] to work, or assume an enlarging elliptical perimeter for burning area to be monitored and tracked by moving UAVs [36, 33, 65, 66]. However, assuming a shape model for large-scale wildfires (when robot help is required) is not accurate [73]. Other approaches such as [10] use an “interestingness” function in order to enforce UAVs to look for areas of fire with predefined specifications. However, this function requires accurate, online measurements of heat intensity over the entire wildfire, which are often unavailable.

Previous image pixel density-based or intensity function-based approaches to distributed control of UAVs for field coverage have typically sought to maximize the fire coverage, as their sole objective [10, 5, 67, 68]. However, these approaches do not explicitly reason (i.e., through tracking and filtering) about a fire’s state (i.e., position, velocity, scale and associated uncertainty), nor do they develop a predictive model for fire propagation by key parameters (e.g., fire-spread rate due to available fuel/vegetation and wind). Other approaches such as [24] and [25] utilized Kalman filter to track fire position on input image data, using frames centroids as model inputs. Nevertheless, unlike these estimation-based approaches, we develop a predictive system with upper-bounds on the measurement-uncertainties that provides probabilistically-guaranteed performance.

Additionally, the fire-intensity based coverage approaches, such as [10] and [11], tend to generate a density-function reflecting the heat-intensity in different parts of a terrain. Density-functions and Voronoi Tessellation based coverage approaches are amongst the most popular methods for static or dynamic field coverage in which robots are distributed and assigned to different parts of a map to maximize coverage [67, 72, 9, 69, 70, 71]. Accordingly, fire-intensity based coverage approaches [10, 11] typically tend to model the location of a firefront as the “coolest” or “hottest” part of a fire visible via infrared sensors [10] according to the wildfire scenario, but this assumptions are not always accurate [74] for the applications of continuous monitoring and tracking. These methods utilize a fire heat intensity model to generate a time-varying density function and use artificial potential field to create driving force to control UAVs towards areas of minima (coolest part of fire) or maxima (hottest part of fire) [10]. By using artificial potential fields, such methods generate control inputs for each UAV at the low actuator level and thus, are different from our high-level coordinated planning for tracking moving targets. Intuitively, our high-level planner deals with the question of “what to do next?” while the mentioned low-level works deal with the problem of “how to get there?”. Moreover, most of these methods are designed to cover as much of an area as possible with limited resources without systematically reasoning about the minimum required number robots for the task.

More closely along the line of our work, Bailon-Ruiz *et al.* [75] proposed a model-based planning algorithms to monitor a propagating wildfire using a fleet of UAVs. The approach tailors a variable neighborhood search to plan surveillance trajectories for a fleet of fixed-wing aircrafts according to a given fire propagation model and a given wind model. The fire state predictions are performed for hours into the future by using the integrated models and are updated through aircraft observations. Despite the similarities to our work, however, we did not consider the approach in [75] as a baseline for comparison, since the method assumes both a fire propagation model and a wind model to be given for making predictions and plans on the scale of hours into future and is designed for fixed-wing aircrafts with constrained motion dynamics. In our approach, we enable UAVs to actively infer fire propagation model parameters through environment observations and plan accordingly, whereas, in [75], the state estimation is performed for extended temporal windows into the future. An active state estimation framework may be able to better cope with changes in fire dynamics over time. In addition, the method in [75] directly plans trajectories for fixed-

wing UAVs with constrained differential-drive dynamics which are not applicable to our presumed omni-directional UAVs. The focus of our method, on the other hand, is to provide a performance-guaranteed plan and an upper-bound on the number of UAV agents needed to monitor the fire areas without losing the track quality.

Furthermore, data-driven and learning-based approaches such as Reinforcement Learning (RL), have also been broadly applied to the problem of dynamic field coverage to enable collaborative monitoring of wildfires [37, 76, 77, 78, 79, 80]. The problems of high-dimensional state-space and imperfect sensory information, which are common in the aerial wildfire monitoring application, are tackled in [81] and [82] for teams of fixed-wing aircrafts using two deep RL methods. A similar problem was investigated and tackled in [83] where a Q-learning [84] was leveraged to learn one surveillance policy for a group of independent Q-learners. Ure *et al.* [85] proposed a decentralized approach for the problem of multiple learning and collaborating agents in fire monitoring cases where agents estimate different models from their local observations, but they can share information by communicating model parameters. In [86, 87], a graph-based actor-critic [84] method is introduced to learn efficient communication protocols for a group of cooperating heterogeneous agents (i.e., sensing and manipulating agents) performing wildfire fighting.

While RL-based approaches occasionally show promising results in specific contexts, such as for small-scale fires [37, 81], there are key limitations in terms of lack of formal guarantees on the boundedness of error, non-standard reward-function specification, scalability and inadaptability to domain shift, which are all limiting factors in application to safety-critical domains such as wildfire fighting.

3 Preliminaries: FARSITE Fire Propagation Mathematical Model

Our coordinated multi-UAV planning framework relies on a dynamic target's approximate motion model to reason about the quality of service in robots' target-tracking mission. This approximate model is leveraged in Kalman filter to infer its latent parameters through environment observations. For our aerial wildfire coverage and firefront tracking case-study, we adopt the Fire Area Simulator (FARSITE) wildfire propagation mathematical model [38].

In this work, we leverage the simplified FARSITE model [10, 9, 11] to predict firefront propagation and use dynamic observations to infer latent parameters of the model. In our framework, UAVs quantify their estimation uncertainty by explicitly inserting the inferred parameters into the model and following a nonlinear uncertainty propagation law. The FARSITE wildfire propagation model used here can be replaced with any other parameterized model, such as the correctable fire simulation model introduced in [88].

Considering q_t^i as the location of i -th firespot on firefront at time t and \dot{q}_t^i as a firefront's growth rate at location q_t^i (i.e., fire propagation velocity), the wildfire propagation dynamics can be written as in Equation 1, where δt is the time-step and $\dot{q}_t^i = \frac{d}{dt}(q_t^i)$ is a function of fire spread rate (R_t , i.e., fuel and vegetation coefficient), wind speed (U_t), and wind azimuth (θ_t), which are available to our system through weather forecasting equipment.

$$q_t^i = q_{t-1}^i + \dot{q}_{t-1}^i \delta t \quad (1)$$

In Equation 1, by ignoring the superscript i and without losing generality, the firefronts growth rate, \dot{q}_t , can be estimated for each propagating spot on firefront by Equations 2-3,

where \dot{q}_t^x and \dot{q}_t^y are first-order firefront dynamics for X and Y axes [38, 9].

$$\dot{q}_t^x = C(R_t, U_t) \sin(\theta_t) \quad (2)$$

$$\dot{q}_t^y = C(R_t, U_t) \cos(\theta_t) \quad (3)$$

In above equations, $C(R_t, U_t)$ is a firespot's velocity which is a function of the fuel/vegetation coefficient, R_t , and mid-flame wind speed, U_t . We use equations from Finney [38] to calculate C as in Equation 4, in which $LB(U_t) = 0.936e^{0.256U_t} + 0.461e^{-0.154U_t} - 0.397$ and $GB(U_t) = LB(U_t)^2 - 1$.

$$C(R_t, U_t) = R_t \left(1 - \frac{LB(U_t)}{LB(U_t) + \sqrt{GB(U_t)}} \right) \quad (4)$$

4 Problem Formulation and Algorithmic Overview

4.1 Problem Statement

The proposed coordinated planning and collaborative coverage framework with quality-of-service guarantees is described and formulated in this section and is presented in Algorithm 1. We assume that firefront loci are detected through vision or thermal cameras [56, 57]. Moreover, as discussed in Figures 1a and 1b, the UAV team is required to collectively monitor and track the fire (or generally moving targets) within several disjoint areas, such as areas of human-activity or human-defined areas of priority where high-quality information is indeed required. As such, we also assume that locations of these areas are known a priori.

Upon receiving the request to monitor and track the fire in a total of N_h areas of wildfire, an available UAV, d , from a set of homogeneous UAVs, $d \in \{\text{UAV}\}^{N_d}$ (N_d is total number of UAVs), is required to provide, online (or at within reasonable time frames) and high-quality information regarding the firefronts in the specified areas. The information here is referring to the estimated state information of a firespot, such as its location and velocity. To this end, the UAV d conducts a tour on firespots, $(q_t^1, \dots, q_t^{N_q}) \in \{Q_t\}^{N_q}$, where N_q is the total number of detected firespots in all of N_h areas. This tour is performed to generate a graph, G_t , on firespots and accordingly a path, \mathcal{P}_t^g , on this graph for the UAV to follow as its tour-paths. The graph, G_t , is created by considering the firespots, $q_t^{N_q}$, as its nodes and the straight path between them as its edges. The path, \mathcal{P}_t^g , must be created and updated such that the UAV, d , is capable of visiting all N_h areas frequently enough such that the estimated state information, i.e., firespot locations and velocities, would remain up-to-date between the UAV visits. A list of inputs and outputs are summarized in Algorithm 1.

Since the velocity of the assigned UAV is limited to v_{max}^d and fire's propagation parameters (i.e., R_t , U_t and θ_t in FARSITE model in Equation 1) alter stochastically over time, it will be challenging for a single UAV to monitor all firespots within all N_h areas without losing the track quality, particularly if the monitoring areas are distant. As such, more UAVs need to be recruited and the monitoring tasks need to be divided among them. Note that, UAV resources are limited and thus, in large-scale operations such as wildfire fighting it is not feasible nor is efficient to simply use as many UAVs as possible to increase the quality of the estimated information. Accordingly, the number of required UAV agents need to be systematically bounded to efficiently monitor fire propagation within all N_h human-defined areas of priority, while not losing the fire tracking quality. For this purpose, we propose

Table 1: Summary of key nomenclature used in our paper.

Notation	Domain	Definition and Properties
N_q	\mathbb{N}^+	Total number of detected firespots (or “targets” in general)
N_d	\mathbb{N}^+	Total number of available UAVs (or “robots” in general)
N_h	\mathbb{N}^+	Total number of areas to be monitored and tracked
$\hat{\cdot}$	—	Accent used for estimated variables
q_t^i	$\mathbb{R}^{1 \times 2}$	Location of i -th firespot at time t (i.e., $q_t = [q_t^x, q_t^y]$)
p_t^d	$\mathbb{R}^{1 \times 3}$	Position of d -th UAV at time t (i.e., $p_t = [p_t^x, p_t^y, p_t^z]$)
v_{max}^d	\mathbb{R}	Maximum linear velocity of UAV d
\mathcal{M}_t	\mathbb{R}	Fire propagation model (or in general, a target’s motion model)
\mathcal{O}_t	\mathbb{R}	UAV’s observation model (value at time t)
T_{UB}	\mathbb{R}	The upper-bound time required for a UAV to complete a tour in a designated area
URR_t^q	\mathbb{R}	The uncertainty residual ratio
$S_{t t}$	$\mathbb{R}^{m \times m}$	The measurement residual covariance matrix; m is the number of state variables
\mathcal{A}_t^d	$\mathbb{N}^{1 \times 2}$	Assigned UAV-path pairs

leveraging online inference of wildfire dynamics and quantifying UAVs’ accumulated estimation uncertainties about the fire model’s parameters in order to efficiently coordinate the UAV team. For the readers’ convenience, we provide Table 1 listing the key variables and notations used throughout the article.

4.2 Algorithmic Overview

Our solution to the aforementioned problem in Section 4.1 is overviewed here. Algorithm 1 presents our multi-UAV coordinated planning framework for active monitoring of dynamic environments, such as wildfire areas. Figure 2 represents a flowchart diagram of the core analytical process for guaranteeing performance in our proposed algorithm.

Initially areas that need to be monitored are prioritized and divided from the rest of the map (Line 3). A selected UAV travels to these areas and gathers sensing information by flying over firespots and creating a graph, G_t , with detected firespots as its vertices and a tour-path, \mathcal{P}_t^g , by connecting the vertices on the generated graph as its edges (Lines 4-5). The graph G_t is an undirected graph generated by connecting each point to its closest neighboring point (i.e., minimum spanning tree) and thus, the initial, quickly-generated graph, G_t , and tour-path, \mathcal{P}_t^g , are not efficient and therefore, need to be modified. In our approach, generating the revised graph G_t' and the path $\mathcal{P}_t^{g'}$ on detected firespots within specified areas includes leveraging a Close-Enough Traveling Salesman Problem (CE-TSP) step to account for multiple fire-spots observable within UAVs’ Field-of-View (FOV) (Line 6). In CE-TSP, a UAV only needs to get “close enough” to fire-spots for state estimation rather than going to the exact location of each fire-spot. Next, the UAV must determine through a *feasibility test* that whether monitoring all dynamic fire-spots within the N_h areas can be guaranteed by taking the generated path, such that the track of none of the spots is lost.

The performed feasibility test (Lines 9-13) is based on an analytical, probabilistic bound on the tracking error, which evaluates the measurement uncertainty residual about the state estimate, \hat{q}_t^i , of a firespot, q_t^i , at time t , where $i = 1, \dots, N_q$. The tracking-error residual bound which we refer to as URR_t^q , is evaluated through determining an upper-bound on the time, T_{UB} , required to cover each firespot in the N_h human-defined vicinities of priority

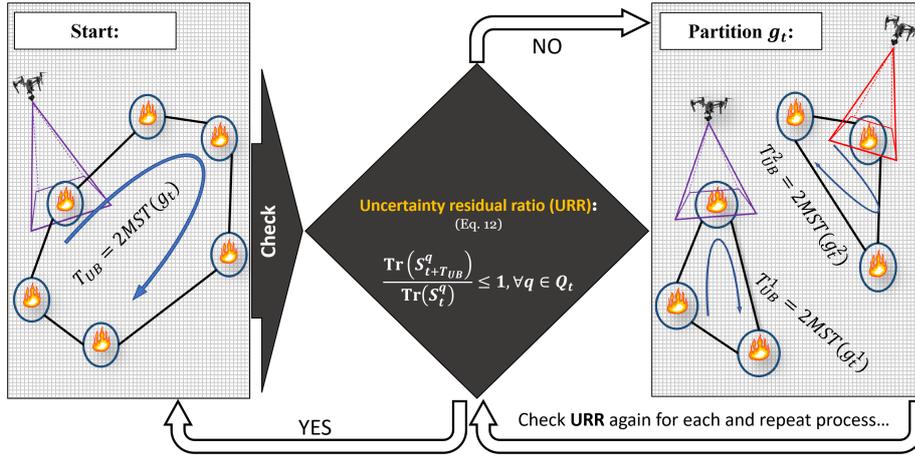


Fig. 2: A flowchart diagram representing the URR checking process as the analytical safety condition. If the URR bound is not satisfied for a node in g_t , the graph is partitioned into two sub-graphs g_t^1 and g_t^2 (i.e., through k -means clustering), and accordingly into two paths, and another UAV is called-in from the available UAVs in the team. The process is repeated until URR is satisfied for all sub-graphs.

(Line 7). In our feasibility test, T_{UB} is compared to the maximum time allowable for each fire track to propagate before measurement, so that the post-measurement residual is less than the residual after the last time a given firespot was observed (Line 35). As shown in Figure 2, if the test is satisfied, the aforementioned process continues and a near-optimal tour is computed via a k -opt procedure. If the test fails, the UAV divvies up the responsibilities for covering the fire locations by partitioning graph G_t^q into $\langle g_t^1, g_t^2 \rangle$, recruiting another UAV to assist and repeating this process until URR_t^q bound is satisfied for all sub-graphs of G_t^q . The partitioning process can be performed through clustering approaches such as k -means. UAVs' status are updated at the end for allocated/unallocated (i.e., 1/0 respectively) UAVs.

In the meantime, the unallocated team of UAVs can be utilized for coordinated, distributed field coverage (see Section 6 and Algorithm 2), for monitoring the rest of the wildfire area (i.e., areas that are not specified or prioritized). To this end, unallocated UAVs generate a search graph and cluster the hotspots to generate a time-optimal path by leveraging the estimated temporal upper-bound T_{UB} .

5 Coordinated Planning for Monitoring with Guaranteed Quality-of-Service

Our objective is to propose a performance-guaranteed coordinated planning framework for collaborative wildfire tracking. Particularly, our approach consists of a probabilistically-guaranteed, realtime coordination algorithm that provides human firefighters with real-time information about fire states within the specified areas. We then develop a predictive distributed coverage method for large-scale dynamic field coverage based on our probabilistic bounds to improve the performance in both wildfire area coverage and firefront tracking.

5.1 Online Inference of Wildfire Dynamics

To perform a real-time inference of wildfire dynamics, we utilize the Adaptive Extended Kalman Filter (AEKF) to predict a distribution over the states of observed firespots and

Algorithm 1: Stages of the proposed multi-UAV coordinated dynamic field monitoring framework (pre-specified areas of priority) with guaranteed performance.

input : Obtain the fire-map $\{Q_t\}^{N_q}$ and total number of areas to cover N_h , list of all UAVs $\{\text{UAV}\}^{N_d}$, set of UAV poses, velocities and observation model $\langle \{p_t\}^d, v_{max}^d, \mathcal{O}_t \rangle$, fire propagation model \mathcal{M}_t

output: Assigned UAV-path pairs $\{\mathcal{A}_t^d\}$

```

1 // main loop //
2 while MissionDuration do
3   Update fire-map and priority areas:  $\langle \{Q_t\}^{N_q^1}, \{Q_t\}^{N_q^0} \rangle \leftarrow \text{UpdateMap}(\{Q_t\}^{N_q})$ 
4   Estimate fire states and dynamics:  $\{\hat{q}_t\} \leftarrow \text{Sense}(\{Q_t\}^{N_q}, \{p_t\}^d, \mathcal{M}_t, \mathcal{O}_t)$ 
5   Form a graph and compute path:  $\langle G_t, \mathcal{P}_t^g \rangle \leftarrow \text{ComputePath}(\{\hat{q}_t\})$ 
6   Modify graph of estimated spots:  $\langle G_t', \mathcal{P}_t^{g'} \rangle \leftarrow \text{CETSP}(G_t, \mathcal{P}_t^g)$ 
7   Compute  $T_{UB}$  for  $d \in \{\text{UAV}\}^d$  to cover  $G_t'$ :  $T_{UB}^{g'} \leftarrow \text{TUB}(\mathcal{P}_t^{g'}, \{\hat{q}_t\}, v_{max}^d)$ 
8   Create a set from graphs and their corresponding  $T_{UB}$ s:  $\{\mathcal{G}\} \leftarrow \langle G_t', T_{UB}^{g'} \rangle$ 
9   while  $\exists g \in \{\mathcal{G}\}$  s.t. False  $\leftarrow \text{FeasibilityTest}(\{\mathcal{G}\}, \{\hat{q}_t\})$  do
10     $\langle \{g'_1, g'_2\}, \{\mathcal{P}'_{g_1}, \mathcal{P}'_{g_2}\} \rangle \leftarrow \text{CETSP}(\text{Cluster}(\{\mathcal{G}\}^g, 2))$ 
11     $\langle T_{UB}^{g'_1}, T_{UB}^{g'_2} \rangle \leftarrow \text{TUB}(\{\mathcal{P}'_{g_1}, \mathcal{P}'_{g_2}\}, \{\hat{q}_t\} \in g, v_{max}^d)$ 
12     $\{\mathcal{G}\} \leftarrow \langle \{g'_1, g'_2\}, \{T_{UB}^{g'_1}, T_{UB}^{g'_2}\} \rangle$ 
13  end
14  Assign UAVs to paths:  $\{\mathcal{A}_t^d\} = \langle \mathcal{P}_t^{g'}, d \rangle, \forall d \in \{\text{UAV}\}^{N_d}$ 
15 end
16 // inner functions //
17 def Sense( $\{Q_t\}^{N_q}, \{p_t\}^d, \mathcal{M}_t, \mathcal{O}_t$ ): // dynamic fire state estimations
18    $\hat{q}_{t+1} = \arg \max_{q_{t+1}} \rho(q_{t+1}, p_{t+1}, \mathcal{M}_t, \mathcal{O}_t)$ 
19 def ComputePath( $\{\hat{q}_t\}$ ): // generate and optimize path
20    $G_t, \mathcal{P}_t^g \leftarrow \text{MST}(\{\hat{q}_t\})$  // MST(.): minimum spanning tree graph
21   while TimeAvailable do
22      $\mathcal{P}_t^g \leftarrow \text{k-opt}(G_t, \mathcal{P}_t^g)$  // k-opt(.): to optimize path
23   end
24 def TUB( $\mathcal{P}_t^g, \{\hat{q}_t\}, v_{max}^d$ ): // temporal service upper-bound
25   Determine fire scenario according to estimated states  $\hat{q}_t^i$ 
26   switch
27     Case (1):  $T_{UB}^{C1} \leftarrow \frac{2\text{len}(\mathcal{P}_t^g)}{v_{max}^d}$  // See Equation 14 for details
28     Case (2):  $T_{UB}^{C2} \leftarrow \frac{8\zeta^\alpha(|G_t|-1)\text{len}(\mathcal{P}_t^g)}{v_{max}^d(1-4\zeta^\alpha(|G_t|-1))}$  // See Equation 15 for details
29     Case (3):  $T_{UB}^{C3} \leftarrow \frac{-\beta + \sqrt{\beta^2 - 4\gamma\delta}}{2\gamma}$  // See Equation 21 for details
30 def FeasibilityTest( $\{\mathcal{G}\}, \{\hat{q}_t\}$ ):
31   if  $\frac{\text{Tr}(S_{t+T_{UB}|t})}{\text{Tr}(S_{t|t-1})} \leq 1, \forall g \in \{\mathcal{G}\}$ , return True else, return False

```

accordingly, a measurement covariance for each firespot (i.e., each discretized point of fire-front) through non-linear error propagation [89, 90]. In a linear system with Gaussian noise, the Kalman filter is optimal [91]. In a system that is nonlinear, the Kalman filter can be used for state estimation, but other methods such as the particle filter may give better results, normally at the price of significant additional computational effort [91]. In our work, we develop computationally lightweight algorithms to scale to large numbers of UAVs and coverage areas. As such, we leverage a Kalman filter as it provides such low-complexity algorithm while also performing satisfactorily.

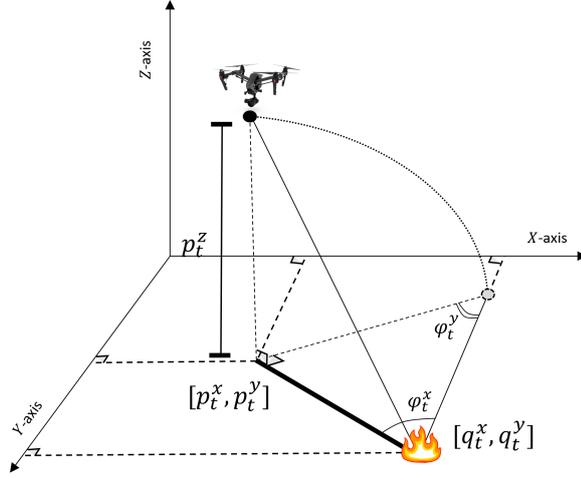


Fig. 3: The observation model of a hovering UAV with position $p_t = [p_t^x, p_t^y, p_t^z]$, with respect to a target (e.g., firespot) on the ground with location $q_t = [q_t^x, q_t^y]$. The angle parameter, φ_t , represents the angle by which the UAV observes the target on the ground. We note that, although the angle-parameters, φ_t^y and φ_t^x , are complementary angles, we utilize them both since we ultimately leverage these angles for tracking 2D pose uncertainty in $\langle x, y \rangle$ (and not just angle uncertainties).

Considering $q_t = [q_t^x, q_t^y]$ as the location of firefronts on the ground at time, t , and $p_t = [p_t^x, p_t^y, p_t^z]$ as the UAV pose, we seek to estimate a distribution over a firespot's location one step forward in time, \hat{q}_{t+1} , given the current firefront distribution and the previous measurement (shown as $q_{t|t-1}$ in our notations throughout the paper), fire propagation model with parameters at time, t , \mathcal{M}_t , and UAV observation model of the field, \mathcal{O}_t . Considering the parameters in fire propagation model (i.e., R_t, U_t and θ_t in FARSITE model in Equation 1), the firespot state estimation problem can be formulated as maximizing the joint probability density function (PDF), ρ , in Equation 5. Equation 5 describes a joint probability density function for estimating the next-step position of a fire-spot, given the current firefront distribution and the previous measurement, fire propagation model, and UAV observation model. Therefore, maximizing this PDF means finding a value of the random variable (the value estimate) that can occur with the highest probability.

$$\hat{q}_{t+1} = \arg \max_{q_{t+1}} \rho(q_{t|t-1}, p_{t|t-1}, \mathcal{M}_t(R_{t|t-1}, U_{t|t-1}, \theta_{t|t-1}), \mathcal{O}_t(q_{t|t-1}, p_{t|t-1})) \quad (5)$$

The PDF in Equation 5 is calculated through the AEKF estimator. The process model, \mathcal{M}_t , predicts states given its parameters (i.e., R_t, U_t, θ_t) by leveraging the fire propagation model (Section 3). The observation model, \mathcal{O}_t , maps the predicted state estimates to observation space. As demonstrated in Figure 3, considering a dynamic target on the ground, $q_t = [q_t^x, q_t^y]$, moving according to a nonlinear model, \mathcal{M}_t (i.e., Eq. 1), at each time, t , the observation mapping along X and Y axes, \mathcal{O}_t^x and \mathcal{O}_t^y , of a flying perception agent with pose, $p_t = [p_t^x, p_t^y]$ and altitude p_t^z , with respect to the target's location can be shown as in Equations 6-7.

$$\mathcal{O}_t^x : \{q_t, p_t\} \rightarrow \varphi_t^x : \|q_t^z - p_t^z\|_2 = \|q_t - p_t\|_2 \tan \varphi_t^x \quad (6)$$

$$\mathcal{O}_t^y : \{q_t, p_t\} \rightarrow \varphi_t^y : \|q_t - p_t\|_2 = \|q_t^z - p_t^z\|_2 \tan \varphi_t^y \quad (7)$$

Moreover, we incorporate process and observation noises to account for error induced by methodological uncertainties (i.e., fire propagation model and UAV observation model inaccuracies) and other sensor errors (i.e., camera and weather forecasting equipment), which are modeled as white noise with covariances Λ_t and Γ_t , respectively. We leverage an adaptive update framework [90] for process and observation noise covariances in standard EKF equations, which introduces *innovation*- and *residual*-based updates for Λ_t and Γ_t , respectively. The adaptive updates, as introduced in [90], and shown in Equations 8-9, relax the assumption of constant process and observation covariances, Λ_t and Γ_t , and reduce sensitivity to their initialization by adaptively estimating these parameters based on innovation and residual to improve the dynamic state estimation accuracy of the standard EKF.

$$\Lambda_t = \gamma \Lambda_{t-1} + (1 - \gamma) \left(K_t \tilde{d}_t \tilde{d}_t^T K_t^T \right) \quad (8)$$

$$\Gamma_t = \gamma \Gamma_{t-1} + (1 - \gamma) \left(\tilde{y}_t \tilde{y}_t^T + H_t P_{t|t-1} H_t^T \right) \quad (9)$$

In above equations, γ is the update step-size, \tilde{y}_t is the measurement residual, and \tilde{d}_t is the measurement innovation. The measurement residual is defined as the difference between actual measurement and the predicted measurement using the posterior, while the measurement innovation is defined as the difference between the actual measurement and its predicted value. Moreover, $K_t = P_{t|t-1} H_t^T S_t^{-1}$ is the near-optimal Kalman gain, in which, $P_{t|t-1}$ is the predicted covariance estimate, H_t is the observation Jacobian matrix and S_t is the covariance residual. Through Equations 8-9, Λ_t and Γ_t are adaptively estimated as the Kalman filter leverages its observations to improve the predicted covariance matrix, $P_{t|t-1}$. Accordingly, to derive the AEKF uncertainty (i.e., measurement-residual) propagation equations, we define the process state vector as $\vec{\Theta}_t = [q_t^x, q_t^y, p_t^x, p_t^y, p_t^z, R_t, U_t, \theta_t]^T$ and $\vec{\Phi}_t = [\varphi_t^x, \varphi_t^y, \hat{R}_t, \hat{U}_t, \hat{\theta}_t]^T$ as the mapping vector¹ through which the state estimates are translated into a unified *angle*-parameters, as firespots on the ground are perceived by hovering UAVs (Figure 3) [9]. Eventually, the model and observation measurement uncertainties propagated by AEKF estimation can be shown as in Equations 10-11.

$$P_{t|t-1} = F_t P_{t-1|t-1} F_t^T + \Lambda_t \quad (10)$$

$$S_{t|t} = H_t P_{t|t-1} H_t^T + \Gamma_t \quad (11)$$

Here, $P_{t|t-1}$ is the predicted covariance estimate and $S_{t|t}$ is the innovation (or residual) covariance matrix. Moreover, F_t and H_t are the process and observation Jacobian matrices which include the gradients of FARSITE model equations in Equations 2-3 with respect to all state variables in $\vec{\Theta}_t$ and the gradients of UAV's observation model in Equations 6-7 with respect to all observation space variables in $\vec{\Phi}_t$, respectively (see [9] for derivations).

5.2 Guaranteeing the Quality of Service

Here, we first present the CE-TSP solution for generating more efficient tour-paths for UAV agents. Next, we introduce and elaborate on our analytical condition for performance guar-

¹ Note that the state variables in $\vec{\Theta}_t$ and $\vec{\Phi}_t$ include firespot and UAV positions, q_t and p_t , as well as R_t, U_t and θ_t which are the FARSITE wildfire propagation model parameters. While these variables are generally application-dependent, we emphasize that R_t, U_t and θ_t are specific to FARSITE (see Section 3) model and can be replaced with other measurable model parameters, in case FARSITE is replaced with other fire propagation models such as the correctable fire simulation model in [88].

antee, i.e., the Uncertainty Residual Ratio (URR) bound, and then, we derive the analytical temporal upper-bounds for a probabilistically-guaranteed coordination based on URR.

5.2.1 Close-enough Traveling Salesman Problem (CE-TSP)

When a UAV agent is tasked to monitor (i.e., estimate the firefront states and closely track each firespot), the first step in our coordinated planning module is to generate a search graph, G_t , with firefront points within the specified areas as the vertices and distances in between as the edges. To this end, we leverage the Close-Enough Traveling Salesman Problem (CE-TSP) with Steiner zone [92] variable neighborhood search where the agent only gets “close enough” to each fire-point instead of visiting their exact locations².

In the application of aerial wildfire propagation monitoring, at each time-step, a UAV can observe multiple firespots (i.e., nodes) within its FOV simultaneously. As such, it is not required for the UAV to travel to the exact location of each firefront point. Instead, the UAV first identifies the overlapping Δ -disks between k firespots as their corresponding Steiner zone [92] and then chooses the centroid node from the identified Steiner areas as a new vertex in its search graph instead of the original k points. Through this process the original graph G_t is modified into a revised graph G'_t (Line 5 in Algorithm 1). The variable k can be tuned empirically and based on the background application.

5.2.2 Analytical Performance Guarantee Condition: Uncertainty Residual Ratio (URR)

When the modified search graph, G'_t , is generated, the UAV performs tours on this graph by visiting each graph-node. Upon reaching a new node, the UAV updates the firespot’s state estimate and calculates two quantities: (1) the analytical upper-bound time, T_{UB} , required to complete a tour on G'_t and arrive back at the current node considering the new state estimates (described in Section 5.2.3) and (2) the current measurement residual covariance through Equation 11. Leveraging these two parameters, the UAV propagates the uncertainty residual for T_{UB} steps into the future by repeatedly applying AEKF’s prediction step (i.e., Equation 10) T_{UB} times, and then performing an update step at the end (i.e., Equation 11). In this way, we emulate gathering a measurement again and compute the innovation covariance at the end of the prediction steps. We introduce the Uncertainty Residual Ratio ($URR_t^{\hat{q}}$) in Equation 12, in which $\text{Tr}(\cdot)$ represents the trace operation.

$$URR_t^{\hat{q}} = \frac{\text{Tr}(S_{t+T_{UB}|t})}{\text{Tr}(S_{t|t-1})} \quad \text{and} \quad URR_t^{\hat{q}} \leq 1, \forall q \in \{Q_t\}^{N_q} \quad (12)$$

The $URR_t^{\hat{q}}$ bound is an indicator of the scale to which the UAV agent is capable of tracking the firespot, q_t , without losing tracking information, while performing tours on the search graph, G'_t . An unsatisfied URR bound (i.e., a URR value greater than one) indicates growing uncertainty and demonstrates a quickly growing (or propagating) wildfire, about which the UAV will not be capable of providing online state-estimates and without losing any track information, while completing a tour on G'_t . Similarly, a URR smaller than one for all firespots in the current graph indicates that online information can be provided by the UAV while keeping track of all propagating firespots. We reiterate that, the goal of the

² Note that, this step might not be required in applications other than aerial wildfire monitoring, in which instead of N_h fire areas, N_h specific moving points/targets need to be monitored. Accordingly, the CE-TSP step in our framework can be replaced with a regular TSP [93].

UAV team here is to provide firefighters with the online, high-quality information about the propagating firefront within the specified areas such that these information can be used in real-time for strategizing firefighting plans.

To maintain control over the measurement uncertainty, we posit that the UAV observers would want the measurement uncertainty residual with respect to a target on the ground not to increase from $t = t_0$ to $t = t_0 + kT_{UB}$ for any positive integer constant k if the UAV observes the target from the same relative position. The reason is that the measurement uncertainty residual as computed by the EKF in our formulation is only dependent on the relative distance between the observer and the target and is independent of time (please refer to the provided Appendix for a mathematical proof and a detailed discussion). As such, the bound in URR ratio determines if the current uncertainty residual (the denominator of Eq. 12) is greater or smaller than the uncertainty residual at time $t + t_{UB}$ (the numerator of Eq. 12) when the UAV completes the tour and revisits the current target. Accordingly, if the ratio is greater than one, it means that the uncertainty residual is increasing, indicating the UAV is falling behind on monitoring the target and the tracking quality is getting worse.

As demonstrated in flowchart diagram in Figure 2, if the URR bound is not satisfied for a graph node ($URR_t^q > 1$ for $q \in \{Q\}^{N_q}$), the generated search graph, G'_t , is partitioned into two smaller sub-graphs $\langle g'_1, g'_2 \rangle$ (i.e., through k -means clustering), and accordingly two paths $\langle \mathcal{P}'_{g_1}, \mathcal{P}'_{g_2} \rangle$. Therefore, the UAV divides up the monitoring task into smaller portions and then recruits another UAV from the available UAVs in the team to collaborate by taking over one of the sub-graphs for monitoring. This process is repeated for all sub-graphs and all of their nodes until URR_t^q is satisfied for all $q \in \{Q_t\}^{N_q}$ and t . We note that, in our formulation, the UAVs only begin the surveillance and tracking the URR bound when they arrive to the determined sub-graph. As such, to minimize the time it takes a UAV to arrive to its assigned sub-graph we solve an optimization problem. Each UAV is assigned to one partition by solving a minimization problem with UAVs as variables, partitions as domains, and distance to the partition centroid as constraints. This way, the closest available UAV is selected for a sub-graph. Additionally, as described in Section 5.2.1, in cases where several firespots can be observed in a UAV's FOV at once, the URR bound is computed for the centroid node from the identified Steiner areas.

The introduced URR bound in Equation 12 depends on the upper-bound traverse time, T_{UB} , which itself is dependent on two major factors: (1) the maximum linear velocity of the UAV (v_{max}^d) and (2) the propagation rate of the fire (Or in general, the motion velocity of any dynamic target which is subject to monitoring). As such, the relations between T_{UB} and the aforementioned two factors need to be derived analytically for all possible scenarios. The following section is dedicated to analyzing such scenarios in the face of the aerial fire monitoring domain and the respective analytical temporal bounds are presented.

5.2.3 Probabilistic Temporal Upper-Bound for Service (T_{UB})

In this section, we derive a probabilistic upper-bound, T_{UB} , on the time required by a UAV to service (i.e., visit once and estimate states) each fire location. The T_{UB} is used in Equation 12 to determine whether a UAV can be probabilistically guaranteed to service each fire location fast enough to ensure a bounded track residual for each firespot. As described in Section 5.2.2, the T_{UB} is dependent on two major factors: (1) the maximum linear velocity of the UAV (v_{max}^d) and (2) the propagation velocity and growth rate of firespots. As such, we derive three bounds, one for each of the following possible scenarios: (1) stationary target

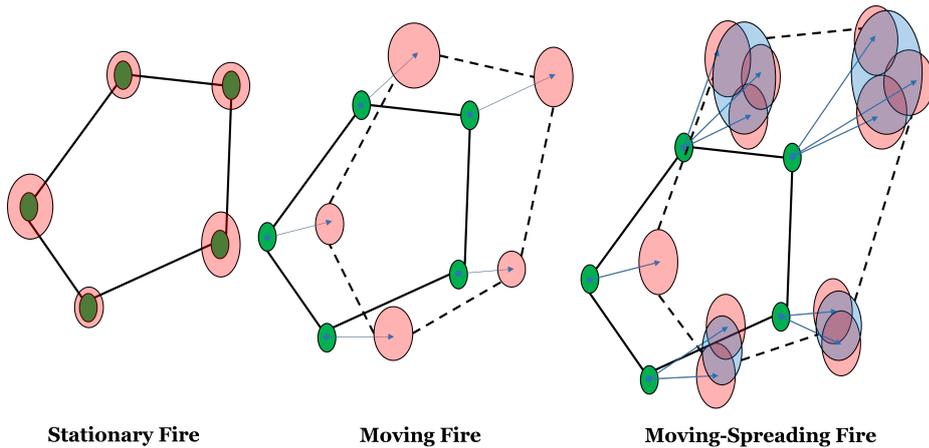


Fig. 4: The three considered scenarios in which, a fire (i.e., target) could be stationary (left), moving (middle) or moving-and-spreading (right). Green dots show a firespot’s current location and red circles represent the distribution over the firespot’s next-step location. Blue circles show the Steiner zones. Note that the black, solid and dashed lines represent a UAV’s current and next-step tour paths, respectively.

points³, (2) moving target points, and (3) moving-spreading target points. Figure 4 depicts the three mentioned scenarios subject to our study.

For the derivations, we reason about the velocity of firespots (e.g., target points), ζ , at the α confidence level, such that, $Pr[\zeta < \zeta^\alpha] = 1 - \alpha$ where smaller α means more conservative. In other words, the probability of the fire being quicker than ζ^α is α . As such, we need to derive an α -dependent probability for our upper-bound service time T_{UB} , such that it is lower-bounded by the *actual maximum time*, T^* , a UAV can go without visiting a $q \in \{Q\}^{N_q}$ and not losing its track quality. Accordingly, and assuming conditional independency between firespots’ velocities given latent fire dynamics⁴, the probability that our bounds are “correct” (meaning that $T_{UB} \geq T^*$) can be shown as in Equation 13. For our experiments, we set $\alpha = 0.05$. While, specific assumptions made for each Case are discussed in the respective Section, we note that the accuracy of the presented bounds depend on the level of accuracy of the utilized model for fire propagation⁵, or generally for applications other than wildfire monitoring, the accuracy of the approximate motion model used for a moving target subject to monitoring.

$$Pr[T_{UB} \geq T^*] \leq 1 - (1 - \alpha)^{N_q} \quad (13)$$

Case 1: Stationary Target Points – In the first scenario all firespots are almost stationary (i.e., stationary customers in the travelling salesman problem). In this case, the search graph, G'_t , is a static graph, and therefore, we utilize a minimum spanning tree (MST) to obtain the initial path, $\mathcal{P}_t^{g'}$. The MST path (i.e., the Hamiltonian cycle computed from the MST) is not the optimal path, but it is fast to generate, and thus, it ensures a quick, sound way to obtain a path that leads to guaranteed service. Once an initial consensus is reached upon providing the updated firefront state estimates according to the MST path, the assigned UAVs can then

³ The scenario designs are motivated such that they expand the applicability of our framework to domains other than wildfire monitoring, and thus, here we use the term *target points* instead of *firespots*.

⁴ We explicitly estimate the latent fire dynamics in our AEKF model.

⁵ Here, we used the FARSITE model which can be replaced with any other parameterized fire propagation model, such as the correctable fire simulation model introduced in [88].

spend time improving the efficiency of their tours through the k -opt algorithm. Accordingly, we derive the upper-bound time for a UAV with maximum linear velocity v_{max}^d to complete a tour on the static search graph in Case 1, as the Hamiltonian cycle of the generated MST path, shown by its spatial cost $\Delta\mathcal{T}^{MST}$ as introduced in Proposition 1. The $\Delta\mathcal{T}^{MST}$ is the length of the generated MST path.

Proposition 1 *The upper-bound time for a UAV to complete a monitoring tour, $T_{UB}^{(C1)}$, in Case 1, the stationary target points, can be calculated as shown in Equation 14.*

$$T_{UB}^{(C1)} = \frac{2}{v_{max}^d} \Delta\mathcal{T}^{MST} \quad (14)$$

Case 1 is presented to develop a tight bound employed under nominal fire conditions and a foundation to derive the traverse temporal upper-bounds for the other cases. Moreover, while there are many approaches for approximating the TSP with guarantees, such as Christofides ($O(n^3)$), we choose other approaches with lower complexity, such as the double tree ($O(n^2)$) algorithm. This is because we consider large-scale scenarios with thousands of discrete firespots (discretization of hundreds of square-miles), which signifies the importance of the lower upper-bound on complexity.

Case 2: Moving Target Points – In the second scenario firespots, move but do not grow (i.e., spawn) considerably. As shown in the middle plot in Figure 4, the upper-bound service time, T_{UB} , for a UAV now depends not only on the maximum linear velocity of the UAV, v_{max}^d , and the number of nodes N_q (i.e., length of discretized firespot list $\{Q\}^{N_q}$), but it also depends on the velocity of the propagating firespots. We derive the upper-bound on the time it takes a UAV to revisit each firespot location (i.e., each node of the search graph) for Case 2 as introduced in Proposition 2. In Equation 15, $\widehat{q}_t^\xi|_\alpha$ is the estimated linear velocity of the spot q_t in the $\xi \in \{x, y\}$ direction evaluated at confidence interval defined by α .

Proposition 2 *The upper-bound time for a UAV to complete a monitoring tour, $T_{UB}^{(C2)}$, in Case 2, the moving target points, can be calculated as shown in Equation 15.*

$$T_{UB}^{(C2)} = \frac{8\zeta^\alpha (N_q - 1) \Delta\mathcal{T}^{MST}}{v_{max}^d (1 - 4\zeta^\alpha (N_q - 1))} \quad (15)$$

In Equation 15, ζ^α is defined as in Equation 16 and we assume a worst-case velocity for all firespots propagating at the fastest fire's rate of speed in the x- and y-directions, as represented by ζ^α . Tuning α enables control of the degree of confidence in our system at the cost of making the UAV coordination problem more difficult. We emphasize that our independence and uniformity of fastest and universal velocity assumptions for firespots pose a worst-case scenario assumption. In other words, the actual service time, T^* , will not become greater than our upper-bound time, T_{UB} , if these assumptions are relaxed.

$$\zeta^\alpha = \arg \max_{q, q'} \sqrt{\left(\widehat{q}_t^x|_\alpha\right)^2 + \left(\widehat{q}_t^y|_\alpha\right)^2} \quad (16)$$

Proof To arrive at the bound in Equation 15, we start by considering the temporal cost of traveling to each firespot under the stationary case, $T_{UB}^{(C1)}$. Moving firespots (even with approximately similar linear velocities), will lead to graph nodes moving possibly in different directions and therefore, the edges of the initial MST graph will shrink or expand. As demonstrated in the middle plot in Figure 4, although the blue arrows, representing velocity and direction of each moving node, have the same lengths (i.e., equal velocities), the next-step tour of the UAV (the dashed black line) is longer than the current tour (solid black line). In the worst case of firespots moving in opposite directions, each edge of the search graph expands according to two times the velocity of the fastest firespot, $2\zeta^\alpha$. Also, if the graph has N_q nodes for the UAV to consider, this yields $(N_q - 1)$ MST edges and $2(N_q - 1)$ Hamiltonian paths in-between the nodes. We note that the total expansion (shrinkage) of the graph is a function of the time the fire is able to expand (shrink), which is a self-referencing relation, as shown in Equation 17. Next, by factoring in the universal velocity of the nodes, ζ^α , and replacing $T_{UB}^{(C1)}$ in Equation 17 by its value in Equation 14 and then solving for $T_{UB}^{(C2)}$, we arrive at Equation 15, as demonstrated below.

$$T_{UB}^{(C2)} = 4\zeta^\alpha (N_q - 1) \left(T_{UB}^{(C1)} + T_{UB}^{(C2)} \right) \quad (17)$$

$$T_{UB}^{(C2)} (1 - 4\zeta^\alpha (N_q - 1)) = 4\zeta^\alpha (N_q - 1) T_{UB}^{(C1)} \quad (18)$$

$$T_{UB}^{(C2)} (1 - 4\zeta^\alpha (N_q - 1)) = \frac{8\zeta^\alpha (N_q - 1)}{v_{max}^d} \Delta\mathcal{T}^{MST} \quad (19)$$

$$T_{UB}^{(C2)} = \frac{8\zeta^\alpha (N_q - 1) \Delta\mathcal{T}^{MST}}{v_{max}^d (1 - 4\zeta^\alpha (N_q - 1))} \quad (20)$$

We note that in Case 2, the UAV's FOV becomes irrelevant as the tour is based on UAVs' inference of the future distribution of the fire, and the location is not growing so that it would escape the FOV after T_{UB} steps into the future. Moreover, in the aerial firefront tracking application, when a fire moves to an area that is already burnt or otherwise is observed to dissipate, the fire is pruned from consideration, and the UAV path is updated as there is no fuel for the fire to actually move there.

Case 3: Moving-Spreading Target Points – In the third scenario that we consider, the firespots move and grow quickly, as in a propagating and spreading wildfire. In this case, single nodes of fire expand over time and escape the UAV's FOV. As such, we must consider the time it takes for a spawning point, q_t , to grow large enough to escape a UAV's FOV, given the maximum leaner velocity, v_{max}^d and current FOV width, w_t , of the UAV. The FOV width for a UAV is directly related to the UAV's current altitude, p_t^z , and the camera half-angle, ϕ , and can be calculated as $w_t = 2p_t^z \tan \phi$. We derive the upper bound for UAV's traversal time allowed to maintain the track quality of each firespot in Case 3 as introduced in Proposition 3. Equation 21.

Proposition 3 *The upper-bound time for a UAV to complete a monitoring tour, $T_{UB}^{(C3)}$, in Case 3, the moving-spreading target points, can be calculated as shown in Equation 21.*

$$T_{UB}^{(C3)} = \frac{\mathcal{V} + \text{sqr}t \left((1 - \mathcal{V})^2 - \frac{64\mathcal{V}(N_q - 1)\Delta\mathcal{T}^{MST}(\zeta^\alpha)^2}{v_{max}^d w_t (1 - 4\zeta^\alpha (N_q - 1))} \right) - 1}{4\mathcal{V}\zeta^\alpha (w_t)^{-1}} \quad (21)$$

We refer to the term, \mathcal{V} , in Equation 21 as the *velocity ratio constant*, which captures the ratio between the velocities of the firespots and the UAV and is computed $\mathcal{V} = \frac{2\zeta^\alpha N_q}{v_{max}^d}$.

Proof To arrive at the bound for Case 3 as in Equation 21, we start by quantifying the maximum planar width and length of the enlarging area. For a specific firespot, q_t , the time-varying planar width and length of the enlarging area along X and Y axes can be computed as $\mathcal{W}(t) \leq 2 \widehat{q}_t^x \Big|_\alpha T_{UB}^{(C3)}$ and $\mathcal{L}(t) \leq 2 \widehat{q}_t^y \Big|_\alpha T_{UB}^{(C3)}$ at the α confidence level, as firespots are now allowed to spread for a maximum of $T_{UB}^{(C3)}$ units of time before the UAV revisits them. Assuming a vertical scanning pattern, we round up the maximum possible $\mathcal{W}(t)$ and thus, the total number of passes the UAV would need to take from left to right can be calculated as $n(t) \leq \left\lceil \frac{2 \widehat{q}_t^x \Big|_\alpha T_{UB}^{(C3)}}{w_t} \right\rceil$, and the total distance traversed for each pass is obtained by $d(t) = 2 \widehat{q}_t^y \Big|_\alpha T_{UB}^{(C3)}$. Accordingly, the total pass traversed is given by $d^{tot}(t) = n(t)d(t)$. Therefore, the time it takes one firespot to escape the FOV of the UAV can be calculated as in Equation 22.

$$\tau_q(t) = \frac{d^{tot}(t)}{v_{max}^d} = \frac{n(t)d(t)}{v_{max}^d} = \frac{\left\lceil \frac{2 \widehat{q}_t^x \Big|_\alpha T_{UB}^{(C3)}}{w_t} \right\rceil 2 \widehat{q}_t^y \Big|_\alpha T_{UB}^{(C3)}}{v_{max}^d} \quad (22)$$

In order to account for all of the firespots, we compute the summation over all τ_q in Equation 22. However, the center of each spreading fire point also moves; thus, $T_{UB}^{(C2)}$ from Equation 15 must also be added to this summation, as shown in Equation 23.

$$T_{UB}^{(C3)} = T_{UB}^{(C2)} + \sum_{q \in \{Q\}} \frac{2}{v_{max}^d} \widehat{q}_t^y \Big|_\alpha T_{UB} \left\lceil \frac{2 \widehat{q}_t^x \Big|_\alpha T_{UB}^{(C3)}}{w_t} \right\rceil \quad (23)$$

To solve Equation 23 for $T_{UB}^{(C3)}$ and find the final upper-bound, we make two simplifying assumptions. First, we remove the ceiling operator and add one to the term inside the operator to achieve continuity, as shown in Equation 24.

$$\left\lceil \frac{2 \widehat{q}_t^x \Big|_\alpha T_{UB}^{(C3)}}{w_t} \right\rceil \leq \frac{2 \widehat{q}_t^x \Big|_\alpha T_{UB}^{(C3)}}{w_t} + 1 \quad (24)$$

Second, we adopt a similar approach to Case 2 by assuming the area required to cover to account for the growth of each fire location is upper-bounded by N_q times the area of growth for a hypothetical fire growing quickest (Equation 25).

$$\sum_q \widehat{q}_t^y \Big|_\alpha T_{UB} \left\lceil \frac{2 \widehat{q}_t^x \Big|_\alpha T_{UB}^{(C3)}}{w_t} \right\rceil \leq N_q \zeta^\alpha \left\lceil \frac{2 \zeta^\alpha T_{UB}^{(C3)}}{w_t} \right\rceil \quad (25)$$

With these two conservative assumptions and replacing $T_{UB}^{(C2)}$ from Equation 15, the upper-bound service time in Equation 23 can be revised as shown below, in Equation 26-27.

$$T_{UB}^{(C3)} = T_{UB}^{(C2)} + \frac{2N_q \zeta^\alpha T_{UB}^{(C3)}}{v_{max}^d} \left[\frac{2\zeta^\alpha T_{UB}^{(C3)}}{w_t} + 1 \right] \quad (26)$$

$$T_{UB}^{(C3)} - \frac{2N_q \zeta^\alpha T_{UB}^{(C3)}}{v_{max}^d} \left[\frac{2\zeta^\alpha T_{UB}^{(C3)}}{w_t} + 1 \right] = \frac{8\zeta^\alpha (N_q - 1) \Delta \mathcal{T}^{MST}}{v_{max}^d (1 - 4\zeta^\alpha (N_q - 1))} \quad (27)$$

We note that Equation 27 is in the form of a general quadratic equations (i.e., $z - az(bz + 1) = \delta$) in which $z = T_{UB}^{(C3)}$. This form can be reorganized into $\gamma z^2 - \beta z + \delta = 0$, where $\gamma = ab = \frac{4N_q(\zeta^\alpha)^2}{w_t v_{max}^d}$ and, $\beta = 1 - a = 1 - \frac{2N_q \zeta^\alpha}{v_{max}^d}$ and, $\delta = \frac{8\zeta^\alpha (N_q - 1) \Delta \mathcal{T}^{MST}}{v_{max}^d (1 - 4\zeta^\alpha (N_q - 1))}$. The upper-bound traversal time, $T_{UB}^{(C3)}$, for Case 3 can be obtained from the general form of solutions to quadratic equations, $T_{UB}^{(C3)} = \frac{-\beta + \sqrt{\beta^2 - 4\gamma\delta}}{2\gamma}$, in which replacing γ , β and δ results in Equation 28. Finally, plugging in the velocity ratio constant, \mathcal{V} , into the Equation 28 and rearranging the terms obtains $T_{UB}^{(C3)}$ as presented in Equation 21.

$$T_{UB}^{(C3)} = \frac{\left(\frac{2N_q \zeta^\alpha}{v_{max}^d} - 1 + \sqrt{\left(1 - \frac{2N_q \zeta^\alpha}{v_{max}^d}\right)^2 - \frac{128N_q(N_q-1)\Delta \mathcal{T}^{MST}(\zeta^\alpha)^3}{w_t(v_{max}^d)^2(1-4\zeta^\alpha(N_q-1))}} \right)}{8(w_t v_{max}^d)^{-1} N_q (\zeta^\alpha)^2} \quad (28)$$

6 Coordinated Distributed Field Coverage Module (No Specified Areas)

UAV agents in the team which are not assigned to monitor the firefronts within specified (human-defined) areas of priority directly (i.e., unallocated UAVs), can be used to explore the rest of the wildfire. See Figure 1b as an example illustration. When we have access to additional UAV resources that may not be needed for the performance-guaranteed coverage of the prioritized areas, we can deploy such UAVs (i.e., unallocated UAVs) to monitor the rest of the wildfire areas. To this end, in this section we propose a coordinated, distributed field coverage framework for multiple UAVs to collectively cover and monitor the rest of wildfire areas, other than areas specified by human-teams. We show that how as a corollary of deriving our analytical upper-bound service times, T_{UB} , these temporal values can be leveraged to design a simple yet effective coverage strategy.

To design our proposed coordinated, collaborative field coverage algorithm we focus on settings with *Centralized Planning and Distributed Execution* (CPDE) paradigm. In other words, communication between UAV agents is not restricted during the planning phase which is done by a centralized computer; however, during execution of the assigned plans, each UAV only performs locally and can communicate with its neighboring UAVs. We note that the CPDE paradigm is a standard and widely-used setting for multi-agent planning [94, 95, 96]. The steps to our distributed coverage module for non-prioritized areas are summarized in Algorithm 2 and are elaborated here.

After detecting the fire map and hotspots, a set of firespots based on the aforementioned Steiner zone variable neighborhood search method are generated (see Section 5.2.1). The set of nodes are partitioned (i.e., K-means clustering) according to the number of available (i.e., unallocated) UAVs. Each UAV is assigned to one partition by solving a simple constraint

Algorithm 2: Stages of the proposed distributed field coverage module when no area is prioritized for surveillance.

input : Obtain the fire-map $\{Q_t\}^{N_q}$, list of all UAVs $\{UAV\}^{N_d}$, set of UAV poses, velocities, etc.
output: Assigned UAV-path pairs $\{\mathcal{A}_t^d\}$

- 1 **objective:** Unallocated UAVs, $\{UAV\}^{N_d^0}$, to cover unspecified areas, $\{Q_t\}^{N_q^0}$
- 2 // main loop //
- 3 Initialize $t \leftarrow 0$
- 4 **while** *MissionDuration* **do**
- 5 Update fire-map and priority areas: $\langle \{Q_t\}^{N_q^1}, \{Q_t\}^{N_q^0} \rangle \leftarrow \text{UpdateMap}(\{Q_t\}^{N_q})$
- 6 Update UAVs status: $\langle \{UAV\}^{N_d^1}, \{UAV\}^{N_d^0} \rangle \leftarrow \text{UpdateDroneStats}(\{UAV\}^{N_d})$
- 7 CoordinatedCoverage $(\{Q_t\}^{N_q^0}, \{UAV\}^{N_d^0}, t)$ // pass in the time t
- 8 Update time: $t \leftarrow t + 1$
- 9 **end**
- 10 // inner functions //
- 11 **def** CoordinatedCoverage $(\{Q_t\}^{N_q^0}, \{UAV\}^{N_d^0}, t)$:
- 12 $\langle \{T_{UB_{i1}}^{g'}, \dots, T_{UB_i}^{g'}\}, \{\mathcal{P}'_{i1}, \dots, \mathcal{P}'_i\} \rangle \leftarrow \text{CSP}(\text{TUB}(\text{Cluster}(\{Q_t\}, \text{len}(\{UAV\}^{N_d^0}))))$
- 13 **if** $(t \geq \min(T_{UB}^i)) \parallel (i \in \{UAV\}^{N_d^0} \text{ is deployed for human support})$
- 14 $\langle \{\mathcal{P}'_{i1}, \dots, \mathcal{P}'_{i_{new}}\} \rangle \leftarrow \text{Cluster}(\{Q_t\}, \text{len}(\{UAV\}_{new}^{N_d^0}))$ // re-clustering
- 15 Repeat line 12 for the new assignments

satisfaction problem (CSP) with UAVs as variables, partitions as domains, and distance to the partition centroid as constraints (Line 12 in Algorithm 2). Note that the partitioning step and solving the CSP problem is done centrally, during the centralized planning phase.

After assigning UAVs to a partition, each UAV agent starts to execute the following tasks distributedly (Line 12 in Algorithm 2). First, an optimal path for coverage and tracking is found by applying the k-opt algorithm. The upper-bound time, T_{UB} , is then calculated for a firespot in the center of a UAV's FOV. T_{UB} provides the UAV with an estimate of how long it will take a fire to escape its FOV as determined by fire propagation velocity. When a route is identified in this way, UAVs can apply this reasoning for the next T_{UB} time-steps before recalculating a new path (Line 13 in Algorithm 2). After T_{UB} has passed, the partitions are revised, and an optimal path is recalculated, since fire locations have likely changed significantly during this time (Lines 14-15 in Algorithm 2). Once a UAV is recruited for help to monitor firefronts within human-specified areas, one of the unallocated UAVs is dispatched and the process is repeated from the central planning phase. In scenarios where the centralization of the planning step does not impose an issue, leveraging this method provides a simple, low-complexity approach that plans for multiple UAVs at the high-level (without considering the low-level control inputs) to efficiently cover a dynamic field.

We emphasize that the relation between our coordinated, distributed field coverage algorithm described here and the coordinated planning framework (described in Section 5) is that the latter coordinates UAV agents to efficiently track the firefronts within specified disjoint areas of fire (i.e., areas prioritized by human firefighters) by using as few UAVs as possible, while guaranteeing the performance. The algorithm described in this section utilizes the remaining UAVs in the team (if any) to collectively surveil the rest of the wildfire area, if needed. In general, this module may not be needed, or it can be replaced with any other surveillance methods depending on the problem setting (e.g., if centralized planning is possible or if it needs to be fully distributed). For instance, if the centralized planning is

not feasible in an application, the coverage method in this section can be replaced with the fully decentralized controller in [9], or any other similar methods from the literature.

7 Empirical Evaluation

In this section, we empirically evaluate both our multi-UAV coordination and planning framework (Section 5) as well as the coordinated, distributed field coverage module introduced in Section 6 in an aerial firefront tracking and wildfire area monitoring case-study. In our experiments, robots in a team of homogeneous, autonomous UAVs (e.g., omnidirectional multi-rotor aircrafts such as quadcopters) are tasked to coordinate together to: (1) track and monitor the firefront within multiple specified human-defined vicinities of priority and provide real-time fire states and tracking information and (2) cooperatively cover and surveil the entire wildfire area. For the field coverage we test the performance of our approach in comparison with two state-of-the-art model-based and reinforcement learning benchmarks (see Section 7.1). Moreover, we demonstrate the feasibility of our framework through implementation on physical robots in a multi-robot testbed (Section 8).

7.1 Baselines

The first benchmark to which we compare is a recent distributed control framework for dynamic wildfire tracking as proposed by [10]. The distributed control framework includes two controller modules in which one is responsible for field coverage and the other for path planning (i.e., in-flight collision avoidance, maintaining a safe altitude, and moving towards new desired poses). The two controllers are defined as the negative gradients (gradient descent) of objective functions to maximize the area-pixel density of the UAV's fire observations and to maintain safe flight parameters with potential field-based criteria, respectively.

Second, we compare our approach to a reinforcement learning (RL) benchmark method proposed in [37]. Each UAV is controlled by an independent agent, and all agent policies are identical (scalable RL algorithm). The policy network architecture consists of three fully connected layers with ReLU activations, following prior work [37]. The agent receives as input an uncertainty map over the wildfire, as described in Section 5.2.2, and outputs a direction to move for the next time-step. The reward function is adapted from prior work, though elements that involve collision are removed, as we assume UAVs can occupy the same space. We follow hyperparameter settings given in prior work, and perform a sweep over hidden-layer dimensions, as they are not explicitly defined in prior work.

7.2 Simulation Environment and Results

7.2.1 Evaluating the Algorithms

To evaluate the efficacy of our proposed multi-UAV planning framework for firefront tracking, we performed a comprehensive simulation to determine the number of UAVs needed to satisfy the uncertainty residual ratio (given in Equation 12) for a range of numbers of disjoint areas specified for firefront tracking. We performed the evaluation for all three mentioned wildfire scenarios for ten trials and calculated the mean and standard error (SE) for each.

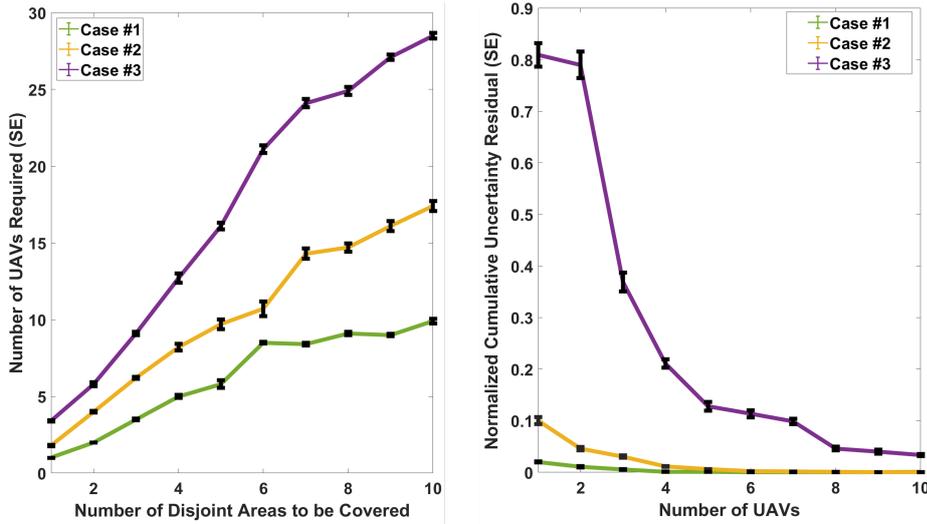


Fig. 5: The left-side figure depicts a quantitative evaluation of our analytical URR bound and shows that increasing the number of disjoint areas to be monitored resulted in a rise in the number of UAVs required to guarantee the performance. This also held as the wildfire propagation scenario changed and a fire propagated more aggressively. The right-side figure demonstrates the efficacy of the coordinated, distributed coverage algorithm, in which less cumulative uncertainty residuals (Y-axis) is better. This figure shows that UAVs easily covered the fire map in the case of stationary wildfire, while in the case of a fast-growing fire, more UAVs were required to cover the fire map efficiently.

The objective was to determine the number of required UAVs to guarantee the performance in each case. The results for this simulations are presented in Figure 5 (left-side figure).

For the aforementioned evaluations, we simulated one to ten (i.e., $N_h \in \{1, 2, \dots, 10\}$) distant areas of fire, assuming each corresponds to one prioritized area. Each area included 20-30 (i.e., random integer in the range $[20, 30]$) randomly placed fire-spots. Areas were randomly initialized within a 500-by-500 terrain. A UAV team was randomly positioned within a distant location at a corner of the terrain. A total number of 30 UAVs were assumed as the goal was to determine the number of required UAVs to guarantee the quality of service in each case. We chose the fire propagation velocity to be 0, 0.5, and 1 for wildfire Cases 1, 2, and 3, respectively. Moreover, we chose the spawning rate of the fire for Case 3 to be at most three (i.e., each fire can produce up to three more fires). Additionally, the team of UAVs were homogeneous in their dynamics and motion characteristics where the maximum linear velocity, v_{max}^d , of the UAVs was set to 500 for all UAVs.

To further investigate the UAV-team behavior during our multi-UAV planning framework in the above experiment, we demonstrated the computed T_{UB} times by the UAVs in Figure 6. Figure 6 shows the computed upper-bound times, T_{UB} , that it takes UAVs to complete a determined tour with respect to the three wildfire scenarios. The plots in Figure 6 are averaged across UAVs and across ten separate experiment trials. As shown, the computed upper-bound service times initially start with a high value, relative to the wildfire scenario since, at first, only one UAV is assigned to monitor and track the entire firefront points. As the time proceeds and the algorithm determines that more UAVs are required for a guaranteed tracking of the firespots, the computed T_{UB} times become smaller until the values converge to a reasonable time respective to the fire scenario.

Additionally, to demonstrate that our derived upper-bound times, T_{UB} , are tight with respect to the actual maximum time it takes UAVs to complete a tour, T^* , we computed the

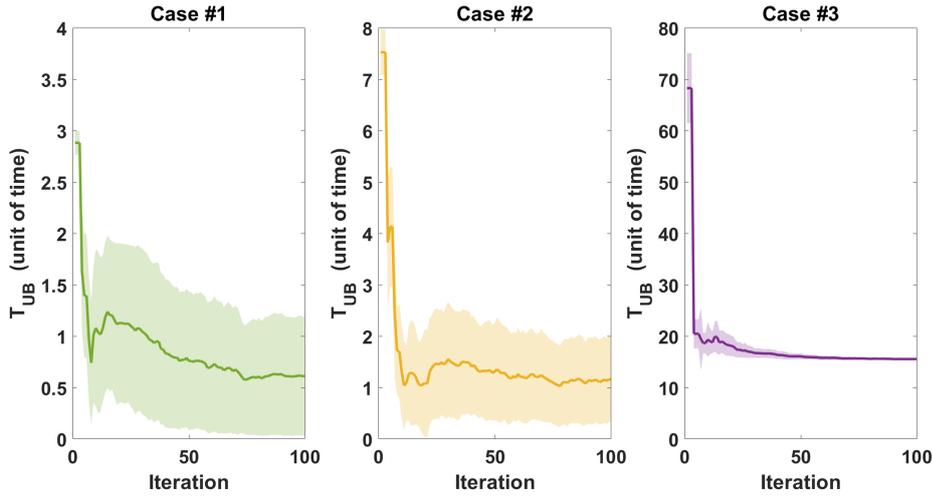


Fig. 6: The computed upper-bound times, T_{UB} , it takes a UAV to complete a determined tour over the course of a simulation with respect to the three wildfire scenarios. As shown, the computed upper-bound service times initially start with a high value, relative to the wildfire scenario, and gradually decrease until converging to a reasonable value, as the algorithm determines that more UAVs are required for a guaranteed tracking of the firespots. Plots are averaged across UAVs and across ten separate experiment trials.

$\frac{T_{UB}}{T^*}$ ratio for all fire scenarios in an experiment. This empirically-evaluated ratio evaluates the tightness of the derived upper-bound times where for a tight bound, the value of this ratio must remain close to one. Figure 7a shows the mean value of this ratio (\pm standard error), computed over 5000 trials for each fire scenario. As shown, the value of the ratio slightly increases as the fire scenario gets more intense. The upper-bounds for Case 1 (i.e., Equation 14), Case 2 (i.e., Equation 15), and Case 3, (i.e., Equation 21) are $1.1\times$, $1.18\times$, and $1.48\times$ greater than the actual maximum times, T^* , for the respective cases. Results show that our upper-bounds become more conservative for more aggressive fire scenarios. We note that for this experiment, the environment parameters are as described above, in Section 7.2.1, and the actual maximum times, T^* , were computed directly in simulation and without considering the MST paths designed in our algorithm.

We also evaluated our coordinated, distributed field coverage algorithm within a similar framework by calculating the cumulative uncertainty residual while tasking an increasing number of UAVs to surveil a large propagating wildfire map (i.e., no specified areas). The uncertainty residual at each timestep was calculated by inspecting the ground-truth fire map for firespots that were not covered by any UAVs, and the respective cumulative error residual obtained from AEKF was calculated. The uncertainty residual measure was an indicator of how successful the team of UAVs were in cooperatively covering all firespots and increased by the number of nodes not covered by any UAV. Figure 5 (right-side) shows the results for the evaluation of our coordinated, distributed coverage algorithm. All wildfire environment and UAV characteristics and parameters for experiments in Figure 5 (right-side) were similar to the description presented above and once again, the experiments were conducted for all three aforementioned wildfire scenarios.

We further investigated the efficacy of our algorithm by: (1) assessing the boundedness of the UAVs' measurement residuals through cumulative uncertainties and (2) evaluating the performance in an evolving wildfire scenario. For these experiments, we initialized two distinct fire areas in a large 500×500 terrain, each with around 30 initial firespots. Figure 7b

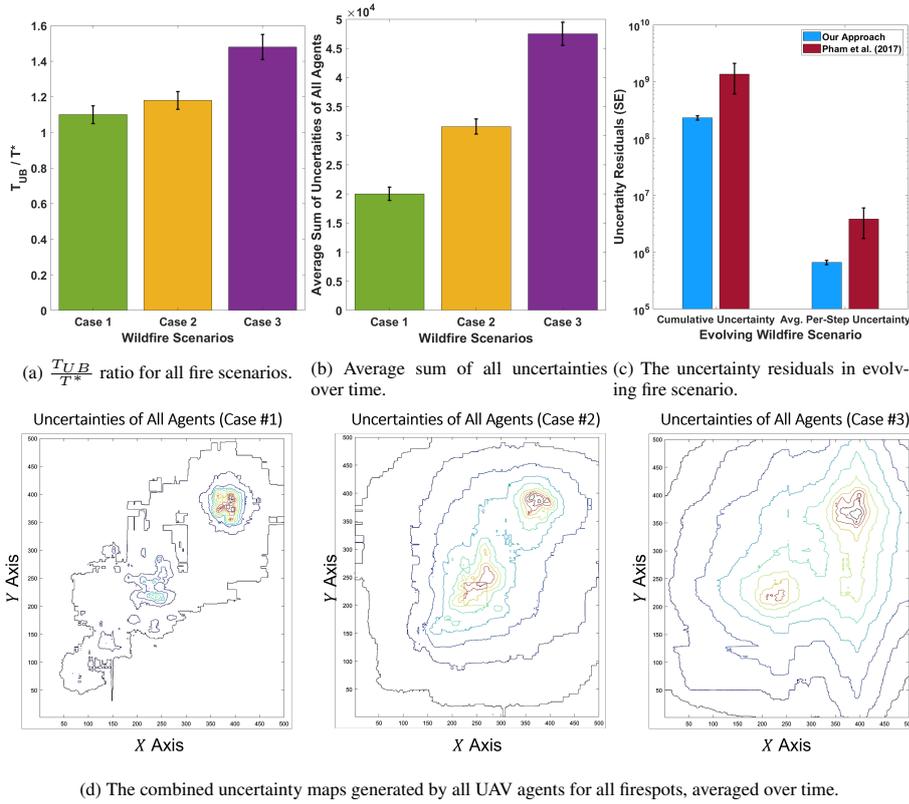


Fig. 7: Figure 7a shows the $\frac{T_{UB}}{T^*}$ ratio for all fire scenarios computed over 5000 trials for each case. This ratio evaluates the sanity of the derived upper-bound times, T_{UB} , and the actual maximum time it takes UAVs to complete a tour, T^* . As shown, for a tight upper-bound, the value of this ratio must remain close to one. Figure 7b shows the sum of all UAVs' uncertainties for all of the measurements, averaged over time. As shown, the measurement uncertain residuals are bounded. Figure 7c compares the measurement uncertainty residuals for our approach and [10] in the evolving fire scenario (i.e., fire starts in Case 1 and evolves into Case 2 and 3 at certain time-steps). As shown, our method outperforms the baseline in this more challenging scenario. Figure 7d shows the combined uncertainty maps generated by summing all measurement residuals of all UAVs for all firespots and averaging over time.

shows the sum of all agents' uncertainties for all measurements, averaged over time and across ten trials. As shown, the measurement uncertain residuals are bounded while the values increase as the wildfire scenario becomes more aggressive. Figure 7d illustrates the combined uncertainty maps in the 500×500 terrain for each of the three wildfire scenarios by summing all the measurement residuals of all agents for all firespots and averaging over time. Figure 7c compares the measurement uncertainty residuals for our coordinated field coverage approach and [10] in an evolving fire scenario. For this experiment, firespots start in Case 1 (i.e., stationary) and evolve into Case 2 and 3 at $t = 50$ and $t = 150$, respectively for a total of $t = 350$ time-steps. A total of five UAVs initiated at a distant location are tasked to cover and track the wildfire area at all times. As shown in Figure 7c, our method outperforms the baseline in this challenging scenario by accumulating $5.8 \times$ less uncertainty residual than the baseline over the simulation time and on average per time-step.

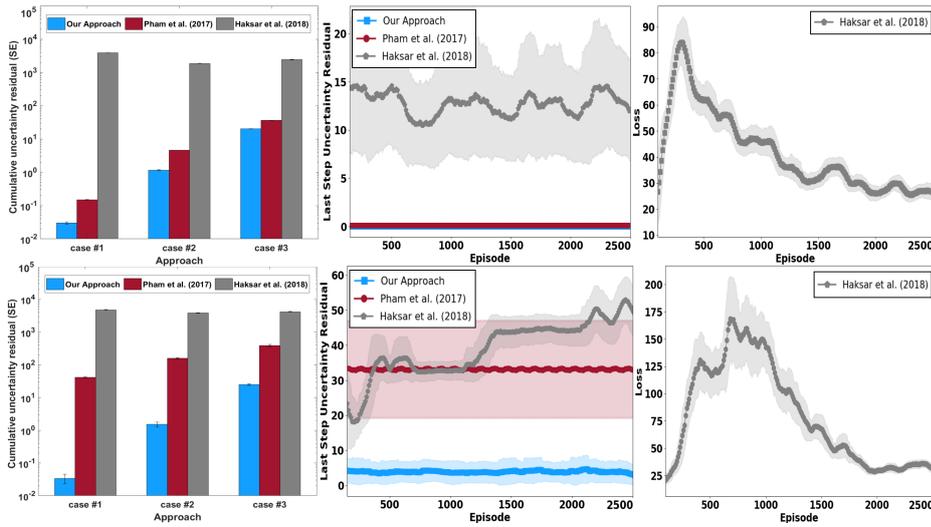


Fig. 8: This figure presents a comparison of our predictive coordinated controller to prior works for both our fire environment (bottom row) and that of [37] (top row). First column figures show the cumulative uncertainty residual during 100 simulation time steps (10 trials). Middle column figures show the uncertainty residual at the end of last simulation step over the course of 2500 episodes of 100 time-step simulation. The third column figures show the RL agent’s loss over the course of training.

7.2.2 Baseline Comparison

Furthermore, we also evaluated our method by assessing its performance in comparison with two state-of-the-art benchmark studies ([10] and [37]) for wildfire coverage. In our simulations, we tested all three algorithms in two different fire environments: (1) our fire environment where one area of fire including 40 randomly-placed firespots was initialized and a total of five randomly positioned UAVs within a distant location were assigned to cover the fire area within a 500-by-500 terrain. (2) the fire environment utilized in [37] where a total of 16 firespots all within a 4-by-4 square in the center of a 50-by-50 terrain were initialized and a total of 10 UAVs positioned around this initial square were assigned to cover the fire area. For both environments, UAVs were required to cover the entire fire map as much as possible, while maintaining an altitude that was both safe and conducive to high-quality imaging. Moreover, the fire model parameters, R_t , U_t and θ_t , were initialized as in [10] for comparison. The results for these evaluations are presented in Figure 8 in which top and bottom rows correspond to (1) our fire environment and (2) that of [37], respectively.

7.3 Discussion

We investigated the sensitivity of our algorithm to the fire scenario, number of separate areas to be monitored, and the number of UAVs available, as shown in Figure 5. The left-hand figure shows that increasing the number of disjoint areas to be monitored resulted in a rise in the number of UAVs required to guarantee the performance. This also held as the wildfire propagation scenario changed and a fire propagated more aggressively. As represented in Figure 5 (right-side figure), the UAV team was able to easily cover the fire map in the case of stationary wildfire, while in the case of a fast-growing fire, more UAVs were required to cover the fire map efficiently (i.e., with low cumulative measurement uncertainty residual).

As shown in Figure 6, the computed upper-bound service times, T_{UB} initially start with a high value and gradually decrease until converging to a reasonable value, as the algorithm determines that more UAVs are required for a guaranteed tracking of the firespots. Note that the initial and the converged values of the T_{UB} times have higher values for wildfire Cases 2 and 3 since the moving and moving-spreading firefronts result in expanded search graphs.

The comparison between our approach and control-theoretic [10] and reinforcement learning-based [37] benchmarks shows that we are able to achieve a $7.5\times$ and $9.0\times$ reduction in error residual for the most challenging cases, as depicted in Figure 8. For all cases, our approach achieved significantly lower uncertainty residual and cumulative uncertainty in both fire environments. These results demonstrate that our framework not only provides probabilistic bounds on guaranteeing performance, but also achieves an empirically more-optimal solution for maintaining a tight track on wildfire propagation. Herein, we declare that the main objective for these benchmark comparisons is to evaluate the soundness and feasibility of our coordinated field coverage approach in Section 6 as compared to standard state-of-the-art approaches for this purpose and to evaluate whether our *more-informed* approach can fulfill the expectations by producing comparable and/or better results.

Moreover, we note that the reward-function specification problem and consequently scalability issue in RL-based methods are present, as the RL agent fails to achieve good performance even after convergence (third column from left in Figure 8). Possible reasons for this include an over- or under-specified reward function from prior work, which emphasized fixed-wing aircraft flight patterns and did not explicitly encourage maximal coverage.

We note that, as discussed in Section 5.2.2, an unsatisfied URR bound means that the uncertainty residual may grow with each tour of the UAVs, which is an indicator that the UAVs may not be capable of monitoring the targets without losing the track quality. Therefore, a URR greater than one only means that the tracking quality cannot be guaranteed in the respective application. However, the strictness of this bound and whether a UAV can continue the tracking task with small growth in the uncertainty (e.g., 10% growth) depend on the rate of this growth over time and if such growth is acceptable in the respective application.

Additionally, the URR upper-bound is designed to recruit UAVs into the monitoring team. However, the URR bound also can be readily used for dismissing UAVs in a simple procedure such as: if URR bound is satisfied with n UAVs then check URR with $n - 1$ UAVs (graphs and tours must be recalculated with $n - 1$) and set $n = n - 1$ if URR is still satisfied. Otherwise, n UAVs will remain. Note that, the URR bound is checked by each UAV and for all the firespots in their assigned region. To select a UAV to be dismissed when the URR bound is satisfied with $n - 1$ UAVs, we can choose from regions that are close to each other. This can be achieved, for instance, by calculating the relative distances between the region centroids. When a UAV is dismissed in this way, the reclustering process of the regions is performed through a centralized clustering step such as applying a k -means with $n - 1$ desired clusters. We then assign UAVs to their new regions.

8 Demonstration: Multi-robot Testbed

In order to account for vehicle dynamics and motion constraints, we implemented and tested our coordinated planning and distributed field coverage modules in the Robotarium multi-agent robotic platform [97], at the Georgia Institute of Technology. The simulated growing wildfires using the introduced FARSITE [38] model are projected on the arena as the regions to be covered by the robot team. The fire simulation parameter setup for our experiments are similar to the empirical evaluation case. In Figures 9 and 10, the sub-figures (1)-(4) illustrate

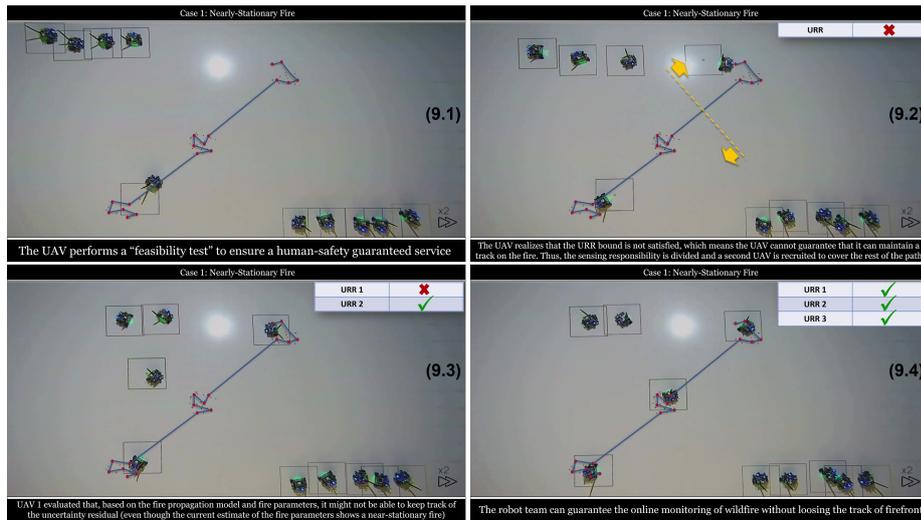


Fig. 9: This figure presents example demonstrations of the proposed coordinated planning algorithm for human safety, implemented on physical robots. In Figure (9.4) all URRs are satisfied and the algorithm determines that a team of three robots can collectively provide the estimated information for the three areas without losing the track of any spot. Link to video: <https://youtu.be/zTR07cK1wRw>.

initial to final robots standings in our experiments for (1) coordinated planning algorithm for tracking firefronts within specified separate areas and (2) the coordinated, distributed field coverage method, respectively. The boxes around each robot shows their respective altitude-dependent FOV. All experiments are repeated for the three wildfire cases and video recordings of these experiments can be found on <https://youtu.be/zTR07cK1wRw>.

For the coordinated planning algorithm with URR condition for performance guarantee, we specified ten robots and tested the feasibility of the algorithm for all three wildfire cases. In the first case, only three robots were needed to guarantee the quality of service at all times, while this number was six and nine for the second and third scenarios, respectively. Figure 9 presents example demonstrations of these experiments. Figure 9.1 shows the beginning of the experiment where a single robot is required to track and monitor three distant areas of fire (i.e., simulated firespots projected on the arena). The robot path is also shown. In Figure 9.2 the robot determines it cannot guarantee to provide real-time information about all three areas since the URR (shown on top left) is not satisfied and thus, the map is partitioned and another robot is summoned. The same process is repeated in Figures 9.3 – 9.4 until all URRs are satisfied and the three robots can collectively provide the estimated information for the three areas without losing the track of any spot. Video recordings of these experiments for all three wildfire scenarios can be found on the *first* part of the provided supplementary video.

Next, we tested the coverage performance using five and three robots, again for all three wildfire scenarios. Figure 10 presents four sample demonstrations of our experiments for the proposed coordinated field coverage algorithm, implemented on physical robots. Each sub-figure 10.1 – 10.4 includes two images: (1) moderate moving fire (left-side) and (2) fast moving-spreading fire (right-side). Figure 10.1 represents the initial step in the algorithm in which *unallocated* robots are selected. Figures 10.2 and 10.3 show the team of robots swarming towards the wildfire area and beginning the surveillance. Figure 10.4 shows that the robot team can successfully cover both wildfire cases and monitor the dynamic spots,

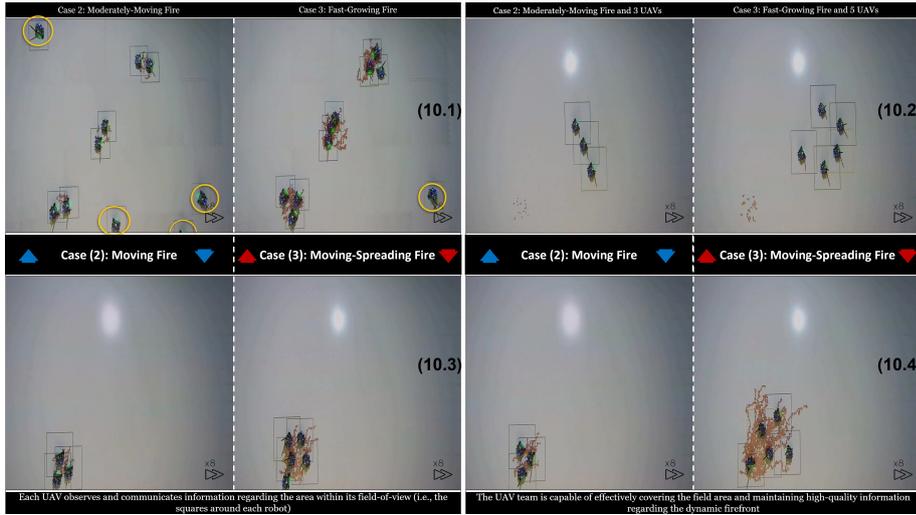


Fig. 10: This figure presents four sample demonstrations of our experiments for the proposed coordinated field coverage algorithm, implemented on physical robots. Each sub-figure 10.1 – 10.4 includes two images: (1) moderate moving fire (left-side) and (2) fast moving-spreading fire (right-side). Figure 10.1 represents the initial step in the algorithm in which *unallocated* robots are selected. Figure 10.4 shows that the robot team can successfully cover both wildfire cases and monitor the dynamic spots, after ten minutes of running the algorithm. Link to video: <https://youtu.be/zTR07cK1wRw>

after ten minutes of running the algorithm. Video recordings of these experiments can be found on the *second* part of the provided supplementary video.

9 Limitations and Future Work

As stated in Section 5.1, our coordinated planning algorithm is model-based. In our framework, UAVs quantify their estimation uncertainty by explicitly inserting the inferred model parameters from the environment into the model and following a nonlinear uncertainty propagation law. According to our experiments, a mismatch between the actual model used by the algorithm and the ground truth model can lead to significantly increasing the actual measurement uncertainty residuals and consequently an unreliable coverage plan. While this poses a limitation on the framework, in such cases, the best practice is to use adaptive estimation methods Adaptive EKF (AEKF) where the covariances can change adaptively. Nevertheless, the results of such methods can still be unreliable.

We performed an experiment to evaluate the performance of the AEKF in our framework for cases when there exists a mismatch between the process model and the actual fire propagation model. To impose a model mismatch, we altered the $LB(U_t)$ equation in the actual FARSITE model (introduced in Section 3) to be $LB_{new} = \mathcal{X}_1 \cdot LB(U_t) + \mathcal{X}_2$, where $\mathcal{X}_1 \sim \mathcal{N}(10, 3)$ and $\mathcal{X}_2 \sim \mathcal{N}(0, 3)$. This alteration would also alter the values of $GB(U_t)$ and $C(R_t, U_t)$. However, we did not revise the derivatives in the process Jacobian matrix, F_t , to account for this alteration. The rest of the environment parameters were exactly the same as in Section 7.2.1. We repeated the simulation for ten separate trials and recorded the cumulative uncertainties over each run. As a result, our experiments demonstrate that the described mismatch between the actual FARSITE model and the process model in the adopted AEKF, on average, caused 1.48 ± 0.4 times higher cumulative uncertainties across

wildfire scenarios, as compared to the case without a model mismatch. This increase in uncertainty residual demonstrates that even an adaptive estimation method such as AEKF cannot completely compensate for a model mismatch.

We discuss that in our approach both UAVs or a ground *control room* can be in charge of running the algorithm. By assuming local communication between neighboring UAVs, without losing generality, we can presume either case to be practical, assuming required hardware and computational resources are available. By assuming an on-board computer, each UAV can perform the URR bound check in Equation 12 and communicate its local belief of its assigned path to a ground control room, which then recruits additional UAVs to the task. We note that distributed communications could result in bottlenecks in our coordinated field coverage method; however, we rely on the rich literature on algorithms available to manage wireless communication grids to address this problem (see [98, 99]).

In Section 5.2.3, we made a worst-case scenario assumption regarding the independence of firespots and uniformity of fastest and universal velocity for all firespots. While these assumptions are made to generalize the applicability of the approach, in future work, the independence assumption can be relaxed for the wildfire monitoring application, given that an accurate model of the correlations between firespots are in hand. Accordingly, the calculated upper-bound service time, T_{UB} , will become smaller and closer to the actual service time, T^* , making the generated monitoring plans more optimal.

10 Conclusion

We have introduced a novel analytical measurement-residual bound on fire propagation uncertainty, allowing high-quality planning for real-time wildfire monitoring and tracking, while also providing a probabilistic guarantee on the quality of service. Our approach outperformed prior work for distributed control of UAVs for wildfire tracking, as well as a reinforcement learning baseline. Our quantitative evaluations validate the performance of our method accumulating $7.5\times$ and $9.0\times$ less error residual than the model based benchmark [10] and the learning-based benchmark [37] respectively when covering a large aggressive wildfire, over a course of 2500 episodes of 100-steps long simulations. Our method also outperformed the best-performing baseline [10] in a challenging evolving fire scenario by accumulating $5.8\times$ less uncertainty residual than the baseline over the simulation time and on average per time-step method. See Section 7 and 7.3 for the evaluations and a discussion on the presented results. Physical implementation of our framework on real robots in a multi-robot testbed demonstrate and validate the feasibility of our approaches.

We note that in this work we generalize by considering large-scale human-defined areas of priority for monitoring and tracking moving targets (such as firespots) and investigating three different scenarios of stationary targets (e.g., the traveling salesman problem and search-and-rescue), moving targets (e.g., border patrol and the tracking of wildlife poachers) and moving-spreading targets (e.g., wildfire monitoring and oil spill surveillance). We also note that due to the modular design of our framework here, any distributed control algorithm for dynamic field coverage, such as our introduced previous work [9], can replace the proposed coordinated coverage module in section 6 to enable *unallocated* UAVs to monitor the unspecified areas of wildfire.

Appendix

A.1 Time Independency of the EKF's Measurement Residual

Our analytical URR bound in Equation 12 depends on the state-estimation measurement residual computed at different time-steps. To maintain control over the measurement uncertainty, we posit that the UAV observers would want the measurement uncertainty residual with respect to a target on the ground not to increase from $t = t_0$ to $t = t_0 + kT_{UB}$ for any positive integer constant k if the UAV observes the target from the same relative position. Therefore, we examine the time-dependency of the propagated error through our EKF formulation. To this end, we follow the mathematical proof and discussions provided in [16] and [9, 100]. We state that the measurement uncertainty about the states of a dynamic point q_t , observed by a flying UAV is independent of time and is only a function of distance between the observer and the point. In the following, we mathematically proof this point.

First, we present how the uncertainty residual is quantified by an EKF. The total uncertainty residual propagated by EKF is composed of a model and an observation measurement uncertainties, both of which follow the general nonlinear uncertainty propagation law, shown in Equations 29-30, where $\Sigma_{t|t-1}$ is the predicted covariance estimate, $\Lambda_{t|t}$ is the innovation (or residual) covariance, F_t and H_t are the process and observation Jacobian matrices, and Q_t and Γ_t are the process and observation noise covariances, respectively.

$$\Sigma_{t|t-1} = F_t \Sigma_{t-1|t-1} F_t^T + Q_t \quad (29)$$

$$\Lambda_{t|t} = H_t \Sigma_{t|t-1} H_t^T + \Gamma_t \quad (30)$$

Considering Equations 29-30, the gradients in the process, F_t , and observation, H_t , Jacobian matrices are responsible for alterations in the uncertainty values. To compute these gradients, we calculate the derivatives of fire's propagation model, \mathcal{M}_t , and UAV's observation model, \mathcal{O}_t , with respect to the state variables. As discussed in Section 5.1 and considering the introduced state vectors, we first derive the process and observation Jacobian matrices (F_t and H_t) as follows in Equations 31-32, respectively. In Equations 31-32, $t' = t - 1$.

$$\frac{\partial \mathcal{M}_t}{\partial \mathcal{S}_i} \Big|_{\hat{\Theta}_{t|t'}} = \begin{matrix} & q_{t'}^x & q_{t'}^y & p_{t'}^x & p_{t'}^y & p_{t'}^z & R_{t'} & U_{t'} & \theta_{t'} \\ \begin{matrix} q_t^x \\ q_t^y \\ p_t^x \\ p_t^y \\ p_t^z \\ R_t \\ U_t \\ \theta_t \end{matrix} & \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & \frac{\partial q_{t'}^x}{\partial R_{t'}} & \frac{\partial q_{t'}^x}{\partial U_{t'}} & \frac{\partial q_{t'}^x}{\partial \theta_{t'}} \\ 0 & 1 & 0 & 0 & 0 & \frac{\partial q_{t'}^y}{\partial R_{t'}} & \frac{\partial q_{t'}^y}{\partial U_{t'}} & \frac{\partial q_{t'}^y}{\partial \theta_{t'}} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{matrix} \quad (31)$$

$$\frac{\partial \mathcal{O}_t}{\partial \Phi_i} \Big|_{\hat{\Phi}_t} = \begin{matrix} & q_t^x & q_t^y & p_t^x & p_t^y & p_t^z & R_t & U_t & \theta_t \\ \begin{matrix} \varphi_t^x \\ \varphi_t^y \\ \hat{R}_t \\ \hat{U}_t \\ \hat{\theta}_t \end{matrix} & \begin{pmatrix} \frac{\partial \varphi_t^x}{\partial q_t^x} & \frac{\partial \varphi_t^x}{\partial q_t^y} & \frac{\partial \varphi_t^x}{\partial p_t^x} & \frac{\partial \varphi_t^x}{\partial p_t^y} & \frac{\partial \varphi_t^x}{\partial p_t^z} & 0 & 0 & 0 \\ \frac{\partial \varphi_t^y}{\partial q_t^x} & \frac{\partial \varphi_t^y}{\partial q_t^y} & \frac{\partial \varphi_t^y}{\partial p_t^x} & \frac{\partial \varphi_t^y}{\partial p_t^y} & \frac{\partial \varphi_t^y}{\partial p_t^z} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{matrix} \quad (32)$$

In Eq. 31-32, we define the process state vector as $\vec{\Theta}_t = [q_t^x, q_t^y, p_t^x, p_t^y, p_t^z, R_t, U_t, \theta_t]^T$ and $\vec{\Phi}_t = [\varphi_t^x, \varphi_t^y, \hat{R}_t, \hat{U}_t, \hat{\theta}_t]^T$ as the mapping vector. As such, we calculate the partial derivatives in Eq. 31 by using Eq. 1-2 and applying the chain-rule to compute the derivatives of q_t^x and q_t^y with respect to parameters R_{t-1} , U_{t-1} , and θ_{t-1} . The partial derivatives are then derived as in Eq. 33-35, where $\mathcal{D}(\theta)$ is $\sin \theta$ and $\cos \theta$ for X and Y axis, respectively.

$$\frac{\partial q_t}{\partial \theta_{t-1}} = C(R_t, U_t) \frac{\partial \mathcal{D}(\theta)}{\partial \theta} \delta t \quad (33)$$

$$\frac{\partial q_t}{\partial R_{t-1}} = \left(1 - \frac{LB(U_t)}{LB(U_t) + \sqrt{GB(U_t)}} \right) \mathcal{D}(\theta) \delta t \quad (34)$$

$$\frac{\partial q_t}{\partial U_{t-1}} = \frac{R_{t'} \left(LB(U_{t'}) \frac{\partial GB(U_{t'})}{\partial U_{t'}} - GB(U_{t'}) \frac{\partial LB(U_{t'})}{\partial U_{t'}} \right)}{\left(LB(U_{t'}) + \sqrt{GB(U_{t'})} \right)^2} \mathcal{D}(\theta) \delta t \quad (35)$$

To compute the partial derivatives in the observation Jacobian matrix in Equation 32, we first need to derive the relation between the angle parameters, φ_t^x and φ_t^y , and the UAV pose. The angle parameters contain information regarding both firefront location $[q_t^x, q_t^y]$ and UAV coordinates $[p_t^x, p_t^y, p_t^z]$. According to Fig. 3, by projecting the looking vector of UAV to planar coordinates, the angle parameters are calculated as shown in Equations 36-37 for X and Y axes respectively, where $q_t = [q_t^x, q_t^y]$ and $p_t = [p_t^x, p_t^y]$.

$$\varphi_t^x = \tan^{-1} \left(\frac{p_t^z}{\|q_t - p_t\|} \right) \quad (36)$$

$$\varphi_t^y = \tan^{-1} \left(\frac{\|q_t - p_t\|}{p_t^z} \right) \quad (37)$$

The partial derivatives in the observation Jacobian matrix H_t for X-axis, presented in Eq. 32, are derived as in Eq. 38-40 and for Y-axis derivatives, we can derive as in Eq. 41-43.

$$\nabla_{q_t} \varphi_t^x = \frac{1}{1 + \left(\frac{p_t^z}{\|q_t - p_t\|} \right)^2} \left(\frac{-p_t^z (q_t - p_t)}{\|q_t - p_t\|^3} \right) = \left[\frac{\partial \varphi_t^x}{\partial q_t^x}, \frac{\partial \varphi_t^x}{\partial q_t^y} \right] \quad (38)$$

$$\nabla_{p_t} \varphi_t^x = \frac{1}{1 + \left(\frac{p_t^z}{\|q_t - p_t\|} \right)^2} \left(\frac{p_t^z (q_t - p_t)}{\|q_t - p_t\|^3} \right) = \left[\frac{\partial \varphi_t^x}{\partial p_t^x}, \frac{\partial \varphi_t^x}{\partial p_t^y} \right] \quad (39)$$

$$\frac{\partial \varphi_t^x}{\partial p_t^z} = \frac{1}{1 + \left(\frac{p_t^z}{\|q_t - p_t\|} \right)^2} \left(\frac{1}{\|q_t - p_t\|} \right) \quad (40)$$

$$\nabla_{q_t} \varphi_t^y = \frac{1}{1 + \left(\frac{\|q_t - p_t\|}{p_t^z} \right)^2} \left(\frac{(q_t - p_t)}{p_t^z \|q_t - p_t\|} \right) = \left[\frac{\partial \varphi_t^y}{\partial q_t^x}, \frac{\partial \varphi_t^y}{\partial q_t^y} \right] \quad (41)$$

$$\nabla_{p_t} \varphi_t^y = \frac{1}{1 + \left(\frac{\|q_t - p_t\|}{p_t^z} \right)^2} \left(\frac{-(q_t - p_t)}{p_t^z \|q_t - p_t\|} \right) = \left[\frac{\partial \varphi_t^y}{\partial p_t^x}, \frac{\partial \varphi_t^y}{\partial p_t^y} \right] \quad (42)$$

$$\frac{\partial \varphi_t^y}{\partial p_t^z} = \frac{1}{1 + \left(\frac{\|q_t - p_t\|}{p_t^z} \right)^2} \left(\frac{-\|q_t - p_t\|}{(p_t^z)^2} \right) \quad (43)$$

Now, considering EKF's covariance propagation Equations in Equations 29-30 as well as the gradients in process Jacobian matrix F_t as calculated in Eq. 33-35, we can see that the gradients in process Jacobian matrix are only functions of fire propagation model parameters (e.g., the FARSITE model in this case) such as fuel coefficient, R_t and wind velocity and direction, U_t and θ_t . Consequently, while these parameters do not vary significantly with time, the uncertainty drop due to process model is time-invariant. We note that FARSITE [38] assumes locality in time (i.e., within seconds or few minutes), making the assumption of time-invariant fire parameters fairly acceptable [101]. Moreover, the gradients in the observation Jacobian matrix, Eq. 38-43, are only functions of the Euclidean distance between the UAV pose and firespot coordinates. We also know that, since at the time of visiting a firespot the planar displacement between UAV and fire locations are approximately zero and the only distance between the two equals to the UAV altitude. Accordingly, both F_t and H_t are locally time-invariant and the total measurement uncertainty residual variations between two different time-steps (e.g., $t = t_0$ and $t = t_0 + kT_{UB}$) is not a function of time and is only a function of the UAV observer's altitude.

References

1. C. Caillouet, F. Giroire, and T. Razafindralambo, "Optimization of mobile sensor coverage with uavs," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2018, pp. 622–627.
2. Y. Zhao, X. Wang, C. Wang, Y. Cong, and L. Shen, "Systemic design of distributed multi-uav cooperative decision-making for multi-target tracking," *Autonomous Agents and Multi-Agent Systems*, vol. 33, no. 1, pp. 132–158, 2019.
3. I. F. Akyildiz and I. H. Kasimoglu, "Wireless sensor and actor networks: research challenges," *Ad hoc networks*, vol. 2, no. 4, pp. 351–367, 2004.
4. M. Liu, H. Gong, Y. Wen, G. Chen, and J. Cao, "The last minute: Efficient data evacuation strategy for sensor networks in post-disaster applications," in *2011 Proceedings IEEE INFOCOM*. IEEE, 2011, pp. 291–295.
5. F. Afghah, A. Razi, J. Chakareski, and J. Ashdown, "Wildfire monitoring in remote areas using autonomous unmanned aerial vehicles," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2019, pp. 835–840.
6. P. Doherty, J. Kvarnström, and F. Heintz, "A temporal logic-based planning and execution monitoring framework for unmanned aircraft systems," *Autonomous Agents and Multi-Agent Systems*, vol. 19, no. 3, pp. 332–377, 2009.
7. M. Li, Y. Liu, and L. Chen, "Nonthreshold-based event detection for 3d environment monitoring in sensor networks," *IEEE Transactions on Knowledge and Data Engineering*, vol. 20, no. 12, pp. 1699–1711, 2008.
8. K. Ota, M. Dong, Z. Cheng, J. Wang, X. Li, and X. S. Shen, "Oracle: Mobility control in wireless sensor and actor networks," *Computer Communications*, vol. 35, no. 9, pp. 1029–1037, 2012.
9. E. Seraj and M. Gombolay, "Coordinated control of uavs for human-centered active sensing of wildfires," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 1845–1852.
10. H. X. Pham, H. M. La, D. Feil-Seifer, and M. Deans, "A distributed control framework for a team of unmanned aerial vehicles for dynamic wildfire tracking," in *Intelligent*

- Robots and Systems (IROS)*, 2017 IEEE/RSJ International Conference on. IEEE, 2017, pp. 6648–6653.
11. H. X. Pham, H. M. La, D. Feil-Seifer, and M. C. Deans, “A distributed control framework of multiple unmanned aerial vehicles for dynamic wildfire tracking,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2018.
 12. M. J. Bays and T. A. Wettergren, “A solution to the service agent transport problem,” in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 6443–6450.
 13. M. Mozaffari, W. Saad, M. Bennis, and M. Debbah, “Efficient deployment of multiple unmanned aerial vehicles for optimal wireless coverage,” *IEEE Communications Letters*, vol. 20, no. 8, pp. 1647–1650, 2016.
 14. A. Ahmadzadeh, G. Buchman, P. Cheng, A. Jadbabaie, J. Keller, V. Kumar, and G. Pappas, “Cooperative control of uavs for search and coverage,” in *Proceedings of the AUVSI conference on unmanned systems*, vol. 2, 2006.
 15. F. Xia, Y.-C. Tian, Y. Li, and Y. Sung, “Wireless sensor/actuator network design for mobile control applications,” *Sensors*, vol. 7, no. 10, pp. 2157–2173, 2007.
 16. E. Seraj, L. Chen, and M. C. Gombolay, “A hierarchical coordination framework for joint perception-action tasks in composite robot teams,” *IEEE Transactions on Robotics*, 2021.
 17. S. Konan, E. Seraj, and M. Gombolay, “Iterated reasoning with mutual information in cooperative and byzantine decentralized teaming,” *arXiv preprint arXiv:2201.08484*, 2022.
 18. M. Brenner and B. Nebel, “Continual planning and acting in dynamic multiagent environments,” *Autonomous Agents and Multi-Agent Systems*, vol. 19, no. 3, pp. 297–331, 2009.
 19. C. Undeger and F. Polat, “Multi-agent real-time pursuit,” *Autonomous Agents and Multi-Agent Systems*, vol. 21, no. 1, pp. 69–107, 2010.
 20. H. Ravichandar, K. Shaw, and S. Chernova, “Strata: unified framework for task assignments in large teams of heterogeneous agents,” *Autonomous Agents and Multi-Agent Systems*, vol. 34, pp. 1–25, 2020.
 21. D. Sarne and B. J. Grosz, “Determining the value of information for collaborative multi-agent planning,” *Autonomous agents and multi-agent systems*, vol. 26, no. 3, pp. 456–496, 2013.
 22. L. Panait and S. Luke, “Cooperative multi-agent learning: The state of the art,” *Autonomous agents and multi-agent systems*, vol. 11, no. 3, pp. 387–434, 2005.
 23. M. N. Prasad and V. R. Lesser, “Learning situation-specific coordination in cooperative multi-agent systems,” *Autonomous Agents and Multi-Agent Systems*, vol. 2, no. 2, pp. 173–207, 1999.
 24. J. V. R. de Sousa and P. V. Gamboa, “Aerial forest fire detection and monitoring using a small uav,” *KnE Engineering*, pp. 242–256, 2020.
 25. Z. Xing, Y. Zhang, C.-Y. Su, Y. Qu, and Z. Yu, “Kalman filter-based wind estimation for forest fire monitoring with a quadrotor uav,” in *2019 IEEE Conference on Control Technology and Applications (CCTA)*. IEEE, 2019, pp. 783–788.
 26. R. Beard, D. Kingston, T. W. McLain, and D. Nelson, “Decentralized cooperative aerial surveillance using fixed-wing miniature uavs,” 2006.
 27. M. McIntire, E. Nunes, and M. Gini, “Iterated multi-robot auctions for precedence-constrained task scheduling,” in *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2016, pp. 1078–1086.

28. E. Nunes and M. Gini, "Multi-robot auctions for allocation of tasks with temporal constraints," in *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
29. H.-L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE transactions on robotics*, vol. 25, no. 4, pp. 912–926, 2009.
30. R. S. Allison, J. M. Johnston, G. Craig, and S. Jennings, "Airborne optical and thermal remote sensing for wildfire detection and monitoring," *Sensors*, vol. 16, no. 8, p. 1310, 2016.
31. J. R. Martinez-de Dios, B. C. Arrue, A. Ollero, L. Merino, and F. Gómez-Rodríguez, "Computer vision techniques for forest fire perception," *Image and vision computing*, vol. 26, no. 4, pp. 550–562, 2008.
32. D. Stipaniev, M. tula, L. Krstini, Damir and, T. Jakov, and Bugari, "Advanced automatic wildfire surveillance and monitoring network," in *6th International Conference on Forest Fire Research, Coimbra, Portugal.(Ed. D. Viegas)*, 2010.
33. P. Sujit, D. Kingston, and R. Beard, "Cooperative forest fire monitoring using multiple uavs," in *Decision and Control, 2007 46th IEEE Conference on.* IEEE, 2007, pp. 4875–4880.
34. K. Fujiwara and J.-i. Kudoh, "Forest fire detection in 2001 using three-dimensional histogram," in *Geoscience and Remote Sensing Symposium, 2002. IGARSS'02. 2002 IEEE International*, vol. 4. IEEE, 2002, pp. 2057–2059.
35. J.-I. Kudoh and K. Hosoi, "Two dimensional forest fire detection method by using noaa avhrr images," in *Geoscience and Remote Sensing Symposium, 2003. IGARSS'03. Proceedings. 2003 IEEE International*, vol. 4. IEEE, 2003, pp. 2494–2495.
36. D. W. Casbeer, D. B. Kingston, R. W. Beard, and T. W. McLain, "Cooperative forest fire surveillance using a team of small unmanned air vehicles," *International Journal of Systems Science*, vol. 37, no. 6, pp. 351–360, 2006.
37. R. N. Haksar and M. Schwager, "Distributed deep reinforcement learning for fighting forest fires with a network of aerial robots," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1067–1074.
38. M. A. Finney, "Farsite: Fire area simulator-model development and evaluation," *Res. Pap. RMRS-RP-4, Revised 2004. Ogden, UT: US Department of Agriculture, Forest Service, Rocky Mountain Research Station. 47 p.*, vol. 4, 1998.
39. M. Freed, W. Fitzgerald, and R. Harris, "Intelligent autonomous surveillance of many targets with few uavs," in *Proceedings of the Research and Development Partnering Conference, Department of Homeland Security, Boston, MA*, 2005.
40. S. Waharte and N. Trigoni, "Supporting search and rescue operations with uavs," in *2010 International Conference on Emerging Security Technologies*. IEEE, 2010, pp. 142–147.
41. P. Rudol and P. Doherty, "Human body detection and geolocalization for uav search and rescue missions using color and thermal imagery," in *2008 IEEE aerospace conference*. Ieee, 2008, pp. 1–8.
42. T. Tomic, K. Schmid, P. Lutz, A. Domel, M. Kassecker, E. Mair, I. L. Grixia, F. Ruess, M. Suppa, and D. Burschka, "Toward a fully autonomous uav: Research platform for indoor and outdoor urban search and rescue," *IEEE robotics and automation magazine*, vol. 19, no. 3, pp. 46–56, 2012.
43. M. A. Olivares-Mendez, C. Fu, P. Ludivig, T. F. Bissyandé, S. Kannan, M. Zurad, A. Annaiyan, H. Voos, and P. Campoy, "Towards an autonomous vision-based unmanned aerial system against wildlife poachers," *Sensors*, vol. 15, no. 12, pp. 31 362–31 391, 2015.

44. L. F. Gonzalez, G. A. Montes, E. Puig, S. Johnson, K. Mengersen, and K. J. Gaston, "Unmanned aerial vehicles (uavs) and artificial intelligence revolutionizing wildlife monitoring and conservation," *Sensors*, vol. 16, no. 1, p. 97, 2016.
45. E. Bondi, A. Kapoor, D. Dey, J. Piavis, S. Shah, R. Hannaford, A. Iyer, L. Joppa, and M. Tambe, "Near real-time detection of poachers from drones in airsims." in *IJCAI*, 2018, pp. 5814–5816.
46. V. V. Klemas, "Coastal and environmental remote sensing from unmanned aerial vehicles: An overview," *Journal of coastal research*, vol. 31, no. 5, pp. 1260–1267, 2015.
47. M. M. Marques, V. Lobo, R. Batista, J. Almeida, M. de Fátima Nunes, R. Ribeiro, and A. Bernardino, "Oil spills detection: Challenges addressed in the scope of the seagull project," in *OCEANS 2016 MTS/IEEE Monterey*. IEEE, 2016, pp. 1–6.
48. J. Allen and B. Walsh, "Enhanced oil spill surveillance, detection and monitoring through the applied technology of unmanned air systems," in *International oil spill conference*, vol. 2008, no. 1. American Petroleum Institute, 2008, pp. 113–120.
49. C. C. Haddad and J. Gertler, "Homeland security: Unmanned aerial vehicles and border surveillance." Library of Congress Washington DC Congressional Research Service, 2010.
50. P. D. Vascik, H. Balakrishnan, and R. J. Hansman, "Assessment of air traffic control for urban air mobility and unmanned systems," 2018.
51. D. P. Thippavong, R. Apaza, B. Barmore, V. Battiste, B. Burian, Q. Dao, M. Feary, S. Go, K. H. Goodrich, J. Homola *et al.*, "Urban air mobility airspace integration concepts and considerations," in *2018 Aviation Technology, Integration, and Operations Conference*, 2018, p. 3676.
52. P. D. Vascik and R. J. Hansman, "Scaling constraints for urban air mobility operations: Air traffic control, ground infrastructure, and noise," in *2018 Aviation Technology, Integration, and Operations Conference*, 2018, p. 3849.
53. F. De Vivo, M. Battipede, P. Gili, A. J. Yezzi, E. Feron, and E. Johnson, "Real-time fire segmentation via active contours for uav integrated wildfire propagation prediction," in *2018 AIAA Information Systems-AIAA Infotech@ Aerospace*, 2018, p. 1488.
54. P. Chamoso, A. González-Briones, F. De La Prieta, and J. M. Corchado, "Computer vision system for fire detection and report using uavs." in *RSFF*, 2018, pp. 40–49.
55. L. Merino, F. Caballero, J. R. Martínez-De-Dios, I. Maza, and A. Ollero, "An unmanned aircraft system for automatic forest fire monitoring and measurement," *Journal of Intelligent & Robotic Systems*, vol. 65, no. 1-4, pp. 533–548, 2012.
56. L. Merino, F. Caballero, J. R. M. de Dios, I. Maza, and A. Ollero, "Automatic forest fire monitoring and measurement using unmanned aerial vehicles," in *Proceedings of the 6th International Congress on Forest Fire Research. Edited by DX Viegas. Coimbra, Portugal*. Citeseer, 2010.
57. C. Yuan, Y. Zhang, and Z. Liu, "A survey on technologies for automatic forest fire monitoring, detection, and fighting using unmanned aerial vehicles and remote sensing techniques," *Canadian journal of forest research*, vol. 45, no. 7, pp. 783–792, 2015.
58. C. Robin and S. Lacroix, "Multi-robot target detection and tracking: taxonomy and survey," *Autonomous Robots*, vol. 40, no. 4, pp. 729–760, 2016.
59. B. Jung and G. S. Sukhatme, "Cooperative multi-robot target tracking," in *Distributed Autonomous Robotic Systems 7*. Springer, 2006, pp. 81–90.
60. L. Jin, S. Li, H. M. La, X. Zhang, and B. Hu, "Dynamic task allocation in multi-robot coordination for moving target tracking: A distributed approach," *Automatica*, vol. 100, pp. 75–81, 2019.

61. R. Mottaghi and R. Vaughan, "An integrated particle filter and potential field method applied to cooperative multi-robot target tracking," *Autonomous Robots*, vol. 23, no. 1, pp. 19–35, 2007.
62. K. Hausman, J. Müller, A. Hariharan, N. Ayanian, and G. S. Sukhatme, "Cooperative multi-robot control for target tracking with onboard sensing," *The International Journal of Robotics Research*, vol. 34, no. 13, pp. 1660–1677, 2015.
63. R. Bailon-Ruiz and S. Lacroix, "Wildfire remote sensing with uavs: A review from the autonomy point of view," in *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2020, pp. 412–420.
64. M. Kumar, K. Cohen, and B. Homchaudhuri, "Cooperative control of multiple uninhabited aerial vehicles for monitoring and fighting wildfires," *Journal of Aerospace Computing, Information, and Communication*, vol. 8, no. 1, pp. 1–16, 2011.
65. K. A. Ghamry and Y. Zhang, "Cooperative control of multiple uavs for forest fire monitoring and detection," in *Mechatronic and Embedded Systems and Applications (MESA), 2016 12th IEEE/ASME International Conference on*. IEEE, 2016, pp. 1–6.
66. K. Harikumar, J. Senthilnath, and S. Sundaram, "Multi-uav oxyrrhis marina-inspired search and dynamic formation control for forest firefighting," *IEEE Transactions on Automation Science and Engineering*, 2018.
67. S. G. Lee, Y. Diaz-Mercado, and M. Egerstedt, "Multirobot control using time-varying density functions," *IEEE Transactions on Robotics*, vol. 31, no. 2, pp. 489–493, 2015.
68. X. Lin and C. G. Cassandras, "An optimal control approach to the multi-agent persistent monitoring problem in two-dimensional spaces," *IEEE Transactions on Automatic Control*, vol. 60, no. 6, pp. 1659–1664, 2014.
69. W. Li and C. G. Cassandras, "Distributed cooperative coverage control of sensor networks," in *Proceedings of the 44th IEEE Conference on Decision and Control*. IEEE, 2005, pp. 2542–2547.
70. L. Zuo, M. Yan, Y. Guo, and W. Ma, "An improved kf-rbf based estimation algorithm for coverage control with unknown density function," *Complexity*, vol. 2019, 2019.
71. M. Santos, S. Mayya, G. Notomista, and M. Egerstedt, "Decentralized minimum-energy coverage control for time-varying density functions," in *2019 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*. IEEE, 2019, pp. 155–161.
72. M. Schwager, B. J. Julian, M. Angermann, and D. Rus, "Eyes in the sky: Decentralized control for the deployment of robotic camera networks," 2011.
73. D. Morvan, "Physical phenomena and length scales governing the behaviour of wildfires: a case for physical modelling," *Fire technology*, vol. 47, no. 2, pp. 437–460, 2011.
74. E. Pastor, L. Zárate, E. Planas, and J. Arnaldos, "Mathematical models and calculation systems for the study of wildland fire behaviour," *Progress in Energy and Combustion Science*, vol. 29, no. 2, pp. 139–153, 2003.
75. R. Bailon-Ruiz, S. Lacroix, and A. Bit-Monnot, "Planning to monitor wildfires with a fleet of uavs," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4729–4734.
76. J. Xiao, G. Wang, Y. Zhang, and L. Cheng, "A distributed multi-agent dynamic area coverage algorithm based on reinforcement learning," *IEEE Access*, vol. 8, pp. 33 511–33 521, 2020.
77. A. A. Adepegba, S. Miah, and D. Spinello, "Multi-agent area coverage control using reinforcement learning," in *The Twenty-Ninth International Flairs Conference*, 2016.

78. E. Seraj, X. Wu, and M. C. Gombolay, "Firecommander: An interactive, probabilistic multi-agent environment for joint perception-action tasks," *arXiv preprint arXiv:2011.00165*, 2020.
79. R. Zanol, F. Chiariotti, and A. Zanella, "Drone mapping through multi-agent reinforcement learning," in *2019 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2019, pp. 1–7.
80. E. Seraj, "Embodied team intelligence in multi-robot systems," in *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 2022, pp. 1869–1871.
81. K. D. Julian and M. J. Kochenderfer, "Distributed wildfire surveillance with autonomous aircraft using deep reinforcement learning," *Journal of Guidance, Control, and Dynamics*, vol. 42, no. 8, pp. 1768–1778, 2019.
82. ———, "Autonomous distributed wildfire surveillance using deep reinforcement learning," in *2018 AIAA Guidance, Navigation, and Control Conference*, 2018, p. 1589.
83. A. Viseras, M. Meissner, and J. Marchal, "Wildfire front monitoring with multiple uavs using deep q-learning," *IEEE Access*, 2021.
84. R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
85. N. K. Ure, S. Omidshafiei, B. T. Lopez, A.-a. Agha-Mohammadi, J. P. How, and J. Vian, "Online heterogeneous multiagent learning under limited communication with applications to forest fire management," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 5181–5188.
86. E. Seraj, Z. Wang, R. Paleja, D. Martin, M. Sklar, A. Patel, and M. Gombolay, "Learning efficient diverse communication for cooperative heterogeneous teaming," in *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 2022, pp. 1173–1182.
87. E. Seraj, Z. Wang, R. Paleja, M. Sklar, A. Patel, and M. Gombolay, "Heterogeneous graph attention networks for learning diverse communication," *arXiv preprint arXiv:2108.09568*, 2021.
88. E. Beachly, C. Detweiler, S. Elbaum, B. Duncan, C. Hildebrandt, D. Twidwell, and C. Allen, "Fire-aware planning of aerial trajectories and ignitions," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 685–692.
89. R. E. Kalman and R. S. Bucy, "New results in linear filtering and prediction theory," *Journal of basic engineering*, vol. 83, no. 1, pp. 95–108, 1961.
90. S. Akhlaghi, N. Zhou, and Z. Huang, "Adaptive adjustment of noise covariance in kalman filter for dynamic state estimation," in *2017 IEEE Power & Energy Society General Meeting*. IEEE, 2017, pp. 1–5.
91. D. Simon, *Optimal state estimation: Kalman, H infinity, and nonlinear approaches*. John Wiley & Sons, 2006.
92. X. Wang, B. Golden, and E. Wasil, "A steiner zone variable neighborhood search heuristic for the close-enough traveling salesman problem," *Computers & Operations Research*, vol. 101, pp. 200–219, 2019.
93. D. L. Applegate, R. E. Bixby, V. Chvatal, and W. J. Cook, *The traveling salesman problem: a computational study*. Princeton university press, 2006.
94. L. Kraemer and B. Banerjee, "Multi-agent reinforcement learning as a rehearsal for decentralized planning," *Neurocomputing*, vol. 190, pp. 82–94, 2016.
95. J. Foerster, I. A. Assael, N. De Freitas, and S. Whiteson, "Learning to communicate with deep multi-agent reinforcement learning," in *Advances in neural information pro-*

- cessing systems, 2016, pp. 2137–2145.
96. E. Seraj, v. Azimi, C. Abdallah, S. Hutchinson, and M. Gombolay, “Adaptive leader-follower control for multi-robot teams with uncertain network structure,” in *2021 American Control Conference (ACC)*. IEEE, 2021.
 97. D. Pickem, P. Glotfelter, L. Wang, M. Mote, A. Ames, E. Feron, and M. Egerstedt, “The robotarium: A remotely accessible swarm robotics research testbed,” in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 1699–1706.
 98. G. Danoy, M. R. Brust, and P. Bouvry, “Connectivity stability in autonomous multi-level uav swarms for wide area monitoring,” in *Proceedings of the 5th ACM Symposium on Development and Analysis of Intelligent Vehicular Networks and Applications*. ACM, 2015, pp. 1–8.
 99. M. Li, K. Lu, H. Zhu, M. Chen, S. Mao, and B. Prabhakaran, “Robot swarm communication networks: architectures, protocols, and applications,” in *2008 Third International Conference on Communications and Networking in China*. IEEE, 2008, pp. 162–166.
 100. E. Seraj, A. Silva, and M. Gombolay, “Safe coordination of human-robot firefighting teams,” *arXiv preprint arXiv:1903.06847*, 2019.
 101. P. Delamatar, A. Finley, and C. Babcock, “Downloading and processing noaa hourly weather station data,” *dim (st)*, vol. 1, no. 30538, p. 12, 2013.