The Journal of Supercomputing (2023) 79:13864–13888
https://doi.org/10.1007/s11227-023-05225-z

Surveillance mission scheduling with unmanned aerial vehicles in dynamic heterogeneous environments

Dylan Machovec[1] · Howard Jay Siegel[1,2] · James A. Crowder[3] · Sudeep Pasricha[1,2] · Anthony A. Maciejewski[1] · Ryan D. Friese[4]

Accepted: 21 March 2023 / Published online: 30 March 2023
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

Abstract

In this study, we design, evaluate, and compare multiple heuristic techniques for mission scheduling of distributed systems comprising unmanned aerial vehicles (UAVs) in energy-constrained dynamic environments. These techniques find effective mission schedules in real-time to determine which UAVs and sensors are used to surveil which targets. We develop a surveillance value metric to quantify the effectiveness of mission schedules, incorporating the amount and usefulness of information obtained from surveilling targets. We use the surveillance value metric in simulation studies to evaluate the heuristic techniques with a reality-based randomized model. We consider two comparison heuristics, three value-based heuristics, and a metaheuristic that intelligently switches between the best value-based heuristics. Additionally, preemption and filtering techniques are applied to further improve the metaheuristic. We show that, for all scenarios that we consider, the novel modified metaheuristics find solutions that are the best on average compared to all other techniques that we evaluate.

Keywords Heuristic methods · Scheduling · Unmanned aerial vehicles · Dynamic mission planning

1 Introduction

Unmanned aerial vehicles (UAVs) are used in many environments to gather information, such as in active battlefields scenarios. An example of such a scenario is shown in Fig. 1. In this example, seven UAVs form a distributed system that is being used to gather information about nine targets. We assume that a UAV can only surveil a single target at a time, so this is an *oversubscribed* scenario

✉ Dylan Machovec
dylan.machovec@gmail.com

Extended author information available on the last page of the article

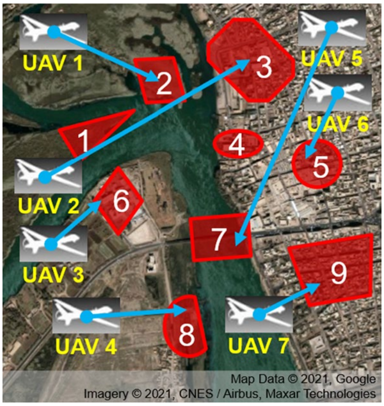

Fig. 1 An example scenario with seven UAVs and nine targets. Arrows drawn from a UAV to a target signify that the target is currently under surveillance by the UAV

because there are more targets than UAVs and it will not be possible to surveil all targets simultaneously with the fleet of UAVs. To gather as much useful information about the targets as possible, it is necessary to conduct mission planning and scheduling to determine how the UAV fleet should effectively surveil the targets.

As both the number of UAVs that are active simultaneously and the number of targets available in an environment increase, it becomes necessary to reduce the amount of human control and human scheduling required to operate them effectively [1]. This can be accomplished by designing and deploying heuristic techniques that can find effective mission scheduling solutions.

In this study, our focus is on the design of mission scheduling techniques capable of working in dynamic environments that are suitable for determining effective mission schedules in real-time. Because scheduling problems of this type are, in general, known to be NP-Hard, finding optimal solutions is not feasible [2]. Due to this, we consider fast heuristic techniques to find mission schedules and evaluate their performance compared to other heuristics. These techniques find mission schedules for our scenarios in less than a second on average.

To effectively compare and evaluate these techniques, we measure system-wide performance using a new metric called *surveillance value*. Surveillance value measures the overall performance of all information gathered by the UAVs, based on parameters such as the number of surveils that occur, the quality of information gathered by each surveil (e.g., image resolution), the importance of each target, and the relevance of the information obtained for a specific target. The novel contributions in this work include:

- The design of new mission scheduling heuristics that are used to dynamically determine which UAVs and sensors should be used to surveil each target;
- The modification of the heuristics using preemption and filtering techniques that enable more efficient utilization of the UAVs in a fleet;

- The construction of a model for a distributed system of UAVs surveilling targets in energy-constrained, dynamic scenarios where the value of each surveil can be quantified;
- The development of a novel system-wide performance measure of the information gathered by all UAVs considering multiple factors;
- The creation of a model for randomly generating scenarios defined by a set of UAVs and targets for the purpose of evaluating mission scheduling techniques, such as the heuristics considered in this work;
- The analysis and comparison of heuristics across many varied simulated scenarios.

This paper is organized as follows. In Sect. 2 the system model and environment are described. The methods used by both the novel mission scheduling heuristics and the comparison heuristics evaluated in this study are presented in Sect. 3. Section 4 contains the specific process used to generate the scenarios we use in our simulations. In Sect. 5, we show the results of the simulations and use the results to analyze and compare the behavior and performance of the heuristics. Related work is discussed in Sect. 6 and finally, in Sect. 7, we conclude and discuss possible future work.

2 System model

2.1 Overview

The distributed system considered in this study consists of a heterogeneous set of UAVs (with varying sensor and energy characteristics) and a heterogeneous set of targets (with varying surveillance requirements). These sets of UAVs and targets are dynamic, meaning that UAVs and targets can be added or removed from the sets at any time. Additionally, specific characteristics of the UAVs and targets can change dynamically at any time. We make a simplifying assumption that every UAV is always close enough to every target and has an unobstructed view of every target so that any sensors available to a UAV can surveil any target at any time.

Because UAVs cannot stay airborne indefinitely, this work considers mission scheduling strategies for a single day. At the end of the day, all UAVs would be able to return to their base of operations to refuel or recharge. In this study, a UAV can only surveil a single target at any given time. The problem space we explore primarily consists of oversubscribed systems, which means that there are fewer UAVs than targets. This will prevent all available targets from being surveilled simultaneously. While the majority of the environments in this study are oversubscribed, the techniques we design and evaluate are still applicable to undersubscribed systems.

Table 1 UAV characteristics

Characteristics	UAV 1	UAV 2	UAV 3	UAV 4
Total energy	1.0	0.5	0.8	0.8
Sensor type	VIS	SAR ǀ IR	VIS ǀ IR	SAR ǀ LIDAR
Energy Consumption/hour	0.05	0.15 ǀ 0.08	0.1 ǀ 0.05	0.15 ǀ 0.05
Sensor quality	7	7 ǀ 5	9 ǀ 3	7 ǀ 4

Table 2 Target Characteristics

Characteristics	Target 1	Target 2	Target 3	Target 4	Target 5	Target 6
Priority	3	4	5	6	8	10
Surveillance time	0.97 h	3.02 h	1.77 h	2.73 h	1.27 h	2.52 h
Allowed sensors	SAR	VIS ǀ IR	SAR ǀ IR ǀ LIDAR	VIS ǀ IR ǀ LIDAR	VIS ǀ SAR ǀ IR ǀ LIDAR	VIS ǀ IR
Sensor affinity	6	7 ǀ 4	2 ǀ 8 ǀ 5	3 ǀ 1 ǀ 8	9 ǀ 2 ǀ 5 ǀ 4	8 ǀ 6

2.2 Target and UAV characteristics

In our environment, each UAV has a single energy source with a fixed amount of *total energy* available to it. In our subsequent discussions, we normalize this value so that the maximum amount of energy available to any UAV is less than or equal to 1.0. Every UAV is equipped with one or more sensors that can be used to surveil targets. The *sensor types* considered in this work are visible light sensors (VIS), infrared light sensors (IR), synthetic-aperture radars (SAR), and light detection and ranging systems (LIDAR). Each UAV cannot have more than one sensor of a given type, which is a simplifying assumption in this work. The heuristics presented in this study will function in environments with multiple sensors with the same type. Each sensor available to a UAV also has an associated *sensor quality* value ranging from 1 (worst) to 10 (best) and a *rate of energy consumption*, which is normalized to the total energy available to the UAV and ranges from 0.0 to 1.0 normalized units of energy per hour. All sensors available to a UAV use the same energy source. The energy needed for the UAV's fixed flight plan is not included in the energy available to the sensors in our model. An example of UAV characteristics is shown in Table 1 for a fleet of four UAVs.

Targets represent locations of interest to be potentially surveilled by UAVs. A *priority* value is assigned to each target, which represents the overall importance of surveilling the target. Priority values are positive integers between 1 and 10, where higher numbers represent more important targets. Each target has a *surveillance time*, which specifies the number of hours that a UAV should spend surveilling the target in a single surveil. Because the kind of information that is useful for each target may vary, targets have a set of *allowed sensor types*, which constrains which sensors can surveil the target. Each of these allowed sensor types has an associated *sensor affinity* value, which range from 1 (worst) to 10 (best) and measures how useful or relevant the information gained from that sensor type is for the target. Table 2 contains target characteristics for an example set of six targets. Additionally, each target has a set of *surveillance intervals*, representing the time intervals in which a single surveil of the target should occur.

2.3 Dynamic events

Characteristics of the UAVs and targets in a scenario can change dynamically during the day. For example, changes in the weather may affect the quality of information collected by certain sensor types, which can be modeled in this study by a change in the sensor affinity value for targets affected by this change in weather. We model this as a change in sensor affinity because the weather would be local to one or more targets and would affect the sensors of all UAVs surveilling that target.

The dynamic changes that we model for UAVs are: (a) adding and removing UAVs from the scenario, (b) removing sensor types from UAVs, and (c) modifying the sensor quality for sensors of the UAVs. The set of targets can also dynamically change: (a) new targets can be added or removed from the scenario, (b) priority of targets can be adjusted, (c) time that a target should be surveilled in a single surveil can be modified, (d) allowed sensor types can be added or removed from the target, and (e) sensor affinities for each allowed sensor type can be altered.

In this study, we assume that any dynamic changes are unexpected and that the techniques that assign UAVs to surveil targets have no information about (a) when the changes will happen, (b) which UAVs and targets will have their characteristics changed, and (c) which specific characteristics will be changed. Additionally, the specific intervals of time during the day when each target can be surveilled are not known in advance.

2.4 Surveillance value

To evaluate the performance of different techniques for assigning UAVs to targets, it is necessary to measure the worth of individual surveils by a UAV on a target. For a UAV (u), target (t), and used sensor type (s) the value of a surveil ($\sigma(u, s, t)$) is given by the product of the priority (ρ), sensor affinity (α), and sensor quality (γ):

$$\text{value}(\sigma(u, s, t)) = \rho(t) * \alpha(t, s) * \gamma(u, s). \tag{1}$$

The total *surveillance value* earned over an interval of time is then defined by the sum of values earned by all surveils performed by UAVs in that interval of time:

$$\text{surveillance value} = \sum_{\substack{\sigma \in \textit{surveils} \\ \textit{performed}}} \text{value}(\sigma). \tag{2}$$

If a surveil is not fully completed, then a partial value will be earned for that surveil, which is directly proportional to the fraction of the surveil that was completed. A partial surveil can occur when the target's surveillance interval ends, the UAV runs out of energy, there are dynamic changes in the environment that stop the surveil, or a heuristic preempts the surveil.

In the case where a characteristic of the target or UAV that affects the value of the surveil changes during the surveil, then the value of the surveil before the change is calculated as a partial surveil that ends when the change occurs. The value of the surveil after the change is similarly calculated as a partial surveil using the remaining time until the end of the surveil or until another characteristic that would affect the value changes. For example, if during a five-hour surveil the priority of a target is doubled after three hours, then the value earned for the first three hours would be calculated as 60% of a full surveil using the initial priority and the value earned for the last two hours would be calculated as 40% of a full surveil using the doubled priority.

2.5 Problem statement

The goal of our proposed scheduling heuristics is to maximize surveillance value obtained over a day. This problem is constrained by the total energy available to each UAV. In this study, this constraint is only applied to the energy consumed by a UAV's sensors. Additionally, each UAV can only surveil one target and only operate one of its sensors at any time; similarly, at any point in time, each target can only be surveilled by one UAV. These are simplifying assumptions used in this study.

3 Mission scheduling techniques

3.1 Mapping events

Mapping UAVs to targets refers to the process of determining which UAVs will surveil which targets, which sensors will be used for surveils, and when the surveils will occur. When preemption is not considered, a UAV is an *available UAV* to be mapped if it is not currently surveilling targets and has energy remaining, and a target is an *available target* to be mapped if it is not currently being surveilled and it is eligible for being surveilled based on its surveillance intervals. A sensor of a UAV is said to be a *valid sensor type* for a given target if that sensor type is also in the target's list of allowed sensor types. If a UAV has a valid sensor type for a target, it is called a *valid UAV* for that target. Only valid UAVs are considered for mapping to a given available target.

The instant when a mapping of available UAVs to available targets occurs is called a *mapping event*. At a mapping event, a mission scheduling technique is used to assign available UAVs to surveil available targets based on the current state of the system. In this study, all techniques presented are real-time heuristics to allow mapping events to be completed in less than a second on average for the problem sizes we consider. There are different techniques for deciding when a mapping event should be initiated, e.g., at fixed time intervals or due to changes in the environment. In this study, we consider the case where mapping events occur with a fixed time interval. Most of our simulations use a fixed time interval of five minutes. We examined the impact of this interval as a part of our simulations and found that other time intervals do not significantly improve performance. In a real-world implementation, this interval of time can be derived based on empirical evaluations of the characteristics of the actual system.

3.2 Comparison techniques

3.2.1 Random

At a mapping event, the *Random* technique considers available targets in a random order. For each target, a random available valid UAV and a random valid sensor type of that UAV are selected. The selected UAV and sensor type are assigned to surveil the target. This results in both the target and the UAV becoming unavailable for new assignments until this new surveil completes. If there is no available UAV that has a valid sensor type for the target, then no UAV is assigned to the target. This repeats with the next target in the random ordering until there are no more assignments of UAVs to targets possible in the current mapping event.

3.2.2 Random best sensor

The *Random Best Sensor* heuristic is similar to the Random technique, except that it uses knowledge about the sensor quality of UAVs and the sensor affinity of targets to make decisions that are likely to result in higher surveillance value. Like the Random heuristic, available targets are considered in a random order and a UAV with a valid sensor type for this target is selected at random. Instead of selecting a random valid sensor type from the UAV, this heuristic chooses the sensor type with the maximum product of the UAV's sensor quality and the target's sensor affinity. Because both values are directly used along with the target's priority in the calculation for the value of a surveil, this strategy will often select higher value surveils compared to the Random heuristic. Next, the same process used by the Random heuristic occurs: the UAV is assigned to surveil the target with this sensor type and the heuristic continues with the next randomly ordered target until no more assignments are possible.

3.3 Value-based heuristics

3.3.1 Overview

The value-based heuristics in this study are designed to search through valid combinations of UAVs, targets, and sensor types to greedily assign UAVs to surveil targets based on the surveillance value performance measure. A valid combination is represented by an available target, a valid available UAV for that target, and a valid sensor type of the UAV for the target.

3.3.2 Max value

At a mapping event, the *Max Value* heuristic starts by finding a valid combination of a UAV, target, and sensor type that results in the maximum possible value for a single surveil. If there are multiple valid combinations with the same maximum possible value, then one of these combinations is selected arbitrarily. The heuristic then assigns the UAV from the selected combination to surveil the selected target with the selected sensor type. This process of finding the maximum value combination and starting a surveil based on the combination repeats until no more assignments of available UAVs to available targets are possible in the current mapping event.

3.3.3 Max value per time

The *Max Value Per Time* heuristic is identical to Max Value except for one difference. Instead of selecting the valid combination that results in the maximum possible value for a surveil, Max Value Per Time instead selects the valid combination that results in the maximum possible value divided by surveillance time of the target (based on a complete surveil of the target).

3.3.4 Max value per energy

The *Max Value Per Energy* heuristic is identical to Max Value except that instead of selecting the valid combination that results in the maximum possible value for a surveil, Max Value Per Energy instead selects the valid combination that results in the maximum possible *value per energy (VPE)*, equal to value divided by the projected energy consumed by the UAV for that surveil. The projected energy consumption can be easily calculated from the energy consumption rate of the selected sensor type and the surveillance time of the selected target (based on a complete surveil of the target). This general concept of performance per unit time and performance per unit of energy has been applied to similar problems in high-performance computing environments, e.g., [3, 4].

3.4 Metaheuristic

The value-based heuristics described in Sect. 3.3 are designed to perform well in specific situations and using the wrong heuristic for a scenario could result in poor performance. Because there may be insufficient information to predict which heuristic should be used, we design a metaheuristic to intelligently combine the best performing value-based heuristics. This does not include the Max Value Per Time heuristic because in the scenarios we consider, Max Value Per Time never performs better than either Max Value or Max Value Per Energy on average. The

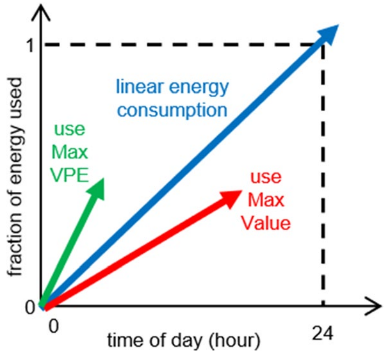

Fig. 2 A visualization of the decision-making process employed by the Metaheuristic. If the current energy consumption during the day of a UAV is below the linear energy consumption line, then Max Value is used for the UAV at the current mapping event. Otherwise, Max Value Per Energy (Max VPE) is used

Metaheuristic uses a two-phase process to find good surveillance options. The general strategy employed by this metaheuristic is illustrated in Fig. 2. The metaheuristic keeps track of the historical rate of energy consumption of each UAV and will use either Max Value or Max Value Per Energy depending on whether the historical rate of energy consumption is above or below a linear rate of energy consumption that would result in running out of energy at the end of the 24-h period we consider.

In the first phase, the Metaheuristic selects a candidate target and valid sensor type for each UAV. The fraction of the day that has passed (δ) and the fraction of the UAV's energy that has been consumed (ε) are used to determine if the strategy used by the Max Value or Max Value per Energy heuristic would be most effective. If $\delta > \varepsilon$, energy is being consumed slowly and Max Value is used. Otherwise, the UAV has been consuming energy at a relatively high rate and the strategy from Max Value Per Energy can be used to make energy-efficient decisions. Based on this choice, either the valid combination using the UAV that results in the maximum possible value or the maximum possible value divided by energy consumed is selected as the best candidate combination for the current UAV. The first phase ends when every UAV has a candidate combination selected. Note that multiple UAVs can select the same target as their candidate.

The second phase is used to determine which UAV from the first phase should be assigned to its candidate target and sensor type. Unlike the first phase, it is unnecessary to use strategies from multiple value-based heuristics in the second phase. This is because energy is a constraint for individual UAVs and not for the overall system. At the system level, all that is relevant to maximizing surveillance value is the value of each surveil. Thus, we choose the UAV with a candidate combination that results in the maximum possible value earned by its corresponding surveil. This chosen UAV is assigned to surveil its target. This process of selecting candidates in the first phase and making an assignment of the best candidate in the second phase is repeated until no more assignments are possible in the current mapping event. An example of a mapping produced by the Metaheuristic is shown in Fig. 3.

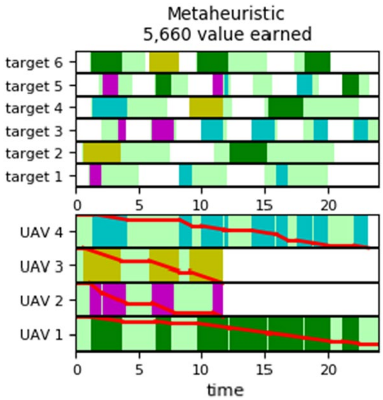

Fig. 3 The mapping produced by the Metaheuristic for our example scenario in Tables 1 and 2. This mapping earns a total value of 5,660 during the 24 h we consider. The red lines represent the percentage of remaining energy for each UAV. Light green regions represent the surveillance intervals for each target and when a UAV has energy remaining. White regions represent when a UAV or target is not available. Dark green (UAV 1), purple (UAV 2), yellow (UAV 3), and blue (UAV 4) regions represent when a surveil is active using each specific UAV

3.5 Heuristic modifications

3.5.1 Overview

We consider two modifications that can be applied to any of the previously described heuristics. One of these modifications allows the heuristics to preempt surveils and the other limits the options available to a heuristic to improve the chance that the heuristic makes optimal decisions.

3.5.2 Preemption

Preemption is a modification to a heuristic that increases the number of choices available. Specifically, in addition to assigning available UAVs to available targets, the heuristic can stop any surveil that is currently in progress, which will cause the affected UAV and target to become immediately available. Because we assume that all UAVs are always able to begin surveilling any target immediately (see Sect 2.1), we do not consider any overhead time due to preempting the surveils.

This modification adjusts the greedy heuristics described in this section so that the set of available targets and UAVs includes all targets and UAVs as long as the new combination that will preempt an existing surveil is better in terms of the metric used by the heuristic (e.g., a surveil is better for Max Value if it earns more value).

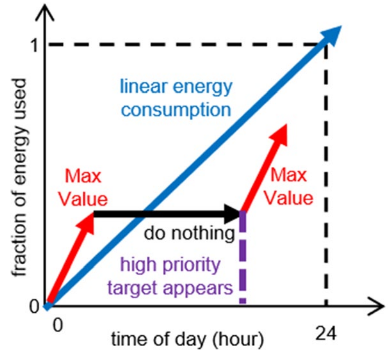

Fig. 4 A visualization of the decision-making process employed by the filtering technique. This modification is designed to prevent UAVs from using all of their energy quickly, which is useful when high priority targets would be available late in the day

3.5.3 Filtering

Filtering is a modification to a heuristic that reduces the number of choices available. Specifically, this modification is designed to control the rate of energy consumption of each UAV so that it will run out of energy close to the end of the day. An example of the benefits of filtering is shown in Fig. 4. In this example, using the Max Value heuristic on its own would use up all the energy of the UAV before the first 12 h have passed. After the first surveil using Max Value has ended, the filtering technique detects that the UAV has been consuming energy quickly and will run out of energy before the end of the day and begins aggressively removing options from the heuristic, resulting in the UAV doing nothing. Near the end of the day, a high priority target arrives and there is still energy remaining to surveil the new target.

The first step of filtering is to calculate the fraction of the day that has passed (δ) and the fraction of the UAV's energy that has been consumed (ε). We define the threshold factor (F) to be the unitless value ε / δ. The base VPE threshold (τ) is equal to the average VPE of all surveils so far for the UAV. Finally, the VPE threshold (T) is equal to $\tau \times F$. At a mapping event, only surveils with VPE > T are considered for each UAV.

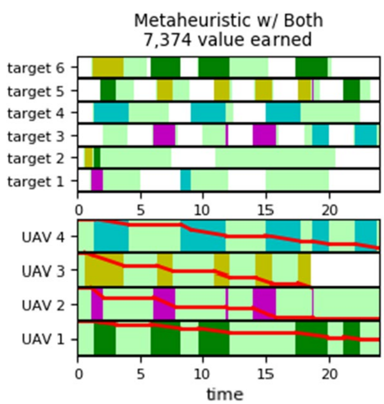

Fig. 5 The mapping produced by the Metaheuristic with preemption and filtering for our example scenario in Tables 1 and 2. This mapping earns a total of 7,374 value during the 24 h we consider. This mapping combines the benefits of preemption to quickly switch to high priority targets and the benefits of filtering by conserving energy until later in the day compared to Fig. 3. Colors in this figure have the same meaning as in Fig. 3

With filtering, there will be some mapping events where a UAV will have all of its options removed. This is intentional, as sometimes there is benefit to doing nothing when none of the targets available for a UAV to surveil are efficient options. In those cases, it is usually better to wait for an efficient option to appear and to leave the available targets for other UAVs that may have sensors that are a better fit for the sensor affinities of the targets. The effects of combining preemption and filtering are demonstrated in Fig. 5. When compared to Fig. 3, preemption and filtering allow the UAVs to spend most of the day surveilling the most efficient targets. For example, UAV 1 is not significantly constrained by energy and spends almost all of its time surveilling the high priority target 6; and UAV 3 is significant constrained by energy and spends most of the day surveilling target 5, which is efficient in terms of value per unit of energy consumed.

4 Simulation setup

4.1 Generation of baseline set of randomized scenarios

4.1.1 Effect of energy consumption rate

Each scenario that we use to evaluate the heuristics is defined by a set of UAV characteristics and a set of target characteristics. To compare and evaluate the heuristics, we consider a wide variety of scenarios to understand the kinds of scenarios for which each heuristic is most effective.

We generate 10,000 baseline scenarios by sampling from probability distributions for the number of UAVs and targets in a scenario in addition to the value for each characteristic of the UAVs and targets. In each case, distributions are selected to attempt to model distributions of parameters that may occur in real-world environments. The details of these distributions are as follows.

4.1.2 Generating UAVs

The number of UAVs available during the 24-h period of a scenario is sampled from a Poisson distribution with the Poisson parameter $\lambda = 9$. The characteristics of each UAV are then generated. The total energy available to the UAV is sampled from a beta distribution with a mean of 0.8 and a standard deviation of 15% of the mean. We use beta distributions for many parameters in this work because many of our UAV and target characteristics are fixed between a minimum and a maximum value. The energy consumption rate for each sensor is sampled from a beta distribution with a mean of 0.05 and a standard deviation of 50% of the mean. The total energy and energy consumption rates are sampled in this way so that UAVs can be expected to operate for an average of 16 h.

The number of sensors available to each UAV is generated by using a Rayleigh distribution with a scale parameter of 2. Any values below 1 are increased to 1 and any values above 4 are decreased to 4. The sensor type for each sensor is selected using probabilities of 0.5, 0.2, 0.2, and 0.1 for the VIS, SAR, IR, and LIDAR sensor types, respectively. Because each UAV can only have one sensor of each type, a sensor type that has been selected for a UAV is no longer a candidate for that UAV and the next sensor is chosen among the remaining sensor types after normalizing their probabilities so that the sum is 1.0. The quality of each sensor is found using a beta distribution with a mean of 0.6 and a standard deviation of 40% of the mean. This value is then truncated to an integer and is clamped between 1 and 10, inclusive, such that all UAVs have a sensor quality between 1 (worst) and 10 (best).

4.1.3 Generating targets

The number of targets available to surveil during the 24-h period is obtained using a Poisson distribution with $\lambda = 14$. Because the number of UAVs was generated with $\lambda = 9$, these scenarios in general will be oversubscribed. The priority of each target is first sampled from a gamma distribution with a mean of 4 and a standard deviation of 60% of the mean. The same method described above for sensor qualities is then used to get integers between 1 (worst) and 10 (best). A gamma distribution was used here because the positive skew results in a slightly larger number of high priority targets instead of the highest priority having the smallest number of occurrences. To obtain the required surveillance time for each target, we use a uniform distribution ranging from 1 to 3 h.

Differing from what was used for UAVs, we obtain the number of allowed sensor types for each target by adding 1 to the value obtained from a binomial distribution with $p = 0.5$ and $n = 3$. The allowed sensor types selected to match this number are uniformly selected from VIS, SAR, IR, and LIDAR. To get the sensor affinity for each sensor type, we use a beta distribution with a mean of 0.7 and a standard deviation of 30% of the mean and use the same method described above for sensor qualities to get integers between 1 (worst) and 10 (best).

Surveillance intervals are arranged such that the average duration of an interval is three hours and the average time between two intervals is one hour. Starting from time 0, the start time of the first interval is found by sampling an exponential distribution with a mean of one hour and the end of the interval is found by sampling from a gamma distribution with a mean of three hours and a standard deviation of 20% of the mean. This same process is then repeated from the end of the first interval and continues until the next interval would start after the end of the day (24 h). Note that although these intervals are generated statically in advance, they are dynamic in our system model as described in Sect 2.3 and the heuristics are not aware of where the future intervals will be.

4.1.4 Generating dynamic events

We utilize Poisson processes to generate the dynamic events for our scenarios. Poisson processes are commonly used to model the occurrences of independent events with a known mean rate. In our baseline set of scenarios, the expected rate of each

Table 3 Dynamic event rates

Event type	Rate (events per day)
Add a UAV	1
Remove a UAV	1
Remove a sensor from a UAV	0.5
Modify sensor qualities of a UAV	0.5
Add a target	2
Remove a target	2
Change the priority of a target	4
Change the surveillance time of a target	6
Add allowed sensor types to a target	6
Remove allowed sensor types from a target	6
Modify sensor affinities of a target	2

type of dynamic event is shown in Table 3. For each of the dynamic event types, we model an independent Poisson process where the time between each occurrence of an event of that type is sampled from an exponential distribution with λ equal to the expected rate of events of that type from Table 3.

Each dynamic event that occurs uses methods similar to those described earlier in this section to determine the specific dynamic changes that will occur. If a UAV or target is to be added to the scenario, then a new UAV or target is generated as described in Sects. 4.1.2 or 4.1.3, respectively. The other event types will affect existing UAVs or targets. For these event types, the UAV or target that is affected is selected randomly (using a uniform distribution).

If a sensor type would be removed from a UAV, that sensor type is selected randomly from the set of sensor types available on the UAV (using a uniform distribution). If the UAV only has one sensor type available, then this is equivalent to removing the UAV from the scenario. When the sensor qualities of a UAV are dynamically changed, they are resampled as described in Sect. 4.1.2 as if a new UAV were being created.

The process for changing the characteristics of a target is similar. If the event would affect the sensor affinity or allowed sensor types of a target, these are handled in the same way as sensor quality and sensor types for a UAV, except that new allowed sensor types may be added to a target if it does not already allow all four sensor types that we model. When a new type is added, it is randomly selected from the sensor types that are currently not allowed on the target (using a uniform distribution). When this is done, a sensor affinity for that type is also sampled for the type as described in Sect. 4.1.3. Finally, events that dynamically change either the priority or surveillance time of a target simply resample the quantities as described in in Sect. 4.1.3.

4.2 Generation of additional scenarios for parameter sweeps

Because the baseline set of 10,000 scenarios in Sect 4.1 may have characteristics that are favorable to the performance of individual heuristics, we use parameter sweeps to evaluate the heuristics in a diverse set of environments. We generate 20 sets of 10,000 scenarios each for the parameter sweeps of six characteristics of the environment. The characteristics we vary are the mean number of targets in a scenario (two sets in addition to the baseline), the mean number of UAVs in a scenario (two sets in addition to the baseline), the mean rate of energy consumption for sensors (three sets in addition to the baseline), the mean rate at which dynamic events occur (three sets in addition to the baseline), the standard deviation of the priority of targets (four sets in addition to the baseline), and the fixed interval at which mapping events occur (five sets in addition to the baseline). The number of sets for each characteristic was selected such that the impact of each characteristic on the performance of the heuristics is clearly demonstrated.

We examine the effect of varying the number of targets and number of UAVs by generating scenarios for the cases with λ values of 10, 14, and 18 for the number of targets, and 5, 9, and 13 for the number of UAVs. We vary the mean energy consumption of sensors with scenarios where the mean is 0.05, 0.1, 0.15, and 0.2. The rate of dynamic events is varied by multiplying the rates given in Table 3 by 0, 1, 4, and 16. The coefficient of variation of target priority is varied between 0.2, 0.4, 0.6, 0.8, and 1.0. Finally, to analyze the effect of the mapping interval, we consider mapping intervals of 1, 5, 10, 30, 60, and 120 min.

4.3 Generation of large-scale scenarios

We also generate a baseline set of 500 scenarios in the same way as the 10,000 described in Sect. 4.1, except that the scenario is ten times as large on average. Specifically, all characteristics of the scenario are generated as described in Sect. 4.1, except that there is an average of 90 UAVs and 140 targets. In addition, the dynamic events are generated as described in Sect. 4.1.3, but with expected rates of events that are ten times as high as listed in Table 3. This baseline set of scenarios is then expanded following the same process described in Sect. 4.2 with the number of targets and UAVs scaled up by ten times. This set of scenarios is used to explore how effectively the heuristics scale when large scenarios are considered.

4.4 Simulation platform

We implement the system described in Sect. 2 in an event-based simulator using Python. The simulations were run in parallel using an AMD Ryzen 9 3900X processor.

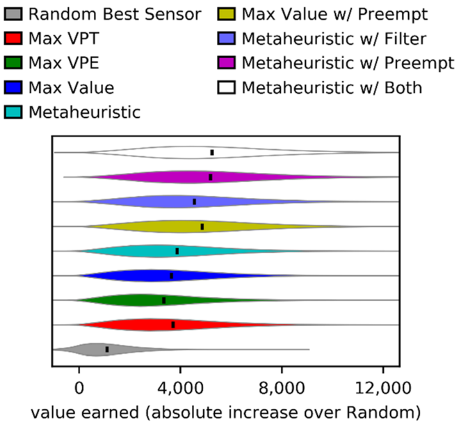

Fig. 6 A violin plot showing the difference between the surveillance value earned by each heuristic and the Random heuristic for the set of 200,000 small scenarios described in Sects. 4.1 and 4.2. The mean difference for each heuristic is indicated by the black marker in each distribution

5 Simulation results

5.1 Results for randomized set of small scenarios

5.1.1 Overview

As described in Sect. 4.2, the results shown in this section consist of parameter sweeps where the means of the distributions described in Sect. 4.1 are varied. Figure 6 is a violin plot, which shows the overall performance of each heuristic in all 200,000 of the scenarios we generated in Sects. 4.1 and 4.2. This overview of our results indicates that the Metaheuristic is among the best value-based heuristics without preemption or filtering modifications. Additionally, when modified with preemption or filtering, the average performance of the Metaheuristic improves significantly.

5.1.2 Effect of energy consumption rate

In Fig. 7, the subset of results where we vary the rate of energy consumption is shown. When the rate of energy consumption is low, the preemption modification results in heuristics that perform significantly better than the others. This is because in scenarios where UAVs can surveil targets for the entire day without running out of energy, it is most important to ensure that surveils on high priority targets begin as soon as possible with the UAVs that have high quality sensors with the best affinity for those targets. With preemption, heuristics can immediately start surveilling those targets with the optimal UAVs with no delay. When the rate of energy consumption increases, preemption still brings significant benefits, but the filtering modification becomes more effective than preemption because it is also important to ensure that the energy of the UAVs that are best for high priority targets later in the day is conserved.

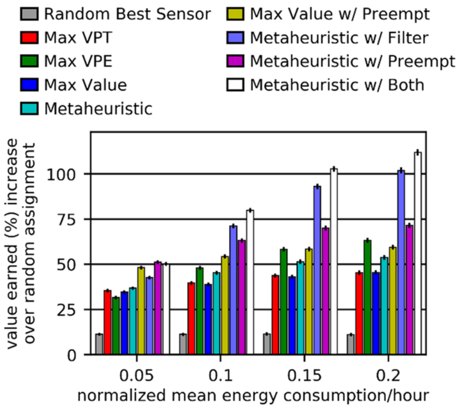

Fig. 7 A comparison of the percentage increase in surveillance value earned when compared to the Random heuristic in 10,000 small randomized scenarios. The mean rate of energy consumption per hour is varied from the baseline set of scenarios with a rate of 0.05 normalized units of energy per hour. Except for the rate of energy consumption, the other characteristics of the scenario use the values from the baseline case described in Sect. 4.1. The 95% mean confidence intervals are shown for each bar

When the rate of energy consumption is low (e.g., in the case with a mean energy consumption rate of 0.05 units of energy per hour), the filtering technique is less effective than preemption. A reason for this is because filtering causes the UAVs to sometimes not surveil any target to conserve energy and when the rate of energy consumption is very low, this can be counterproductive.

These results appear very different from Fig. 6 because Fig. 6 displays the distribution of results from all 200,000 of the small scenarios, while the scenarios shown in Fig. 7 are only 40,000 of those scenarios. The remaining 160,000 scenarios use a mean rate of energy consumption of 0.05 units of energy per hour, which matches the first set of bars in Fig. 7. When comparing this first set of bars to the overall results in Fig. 6, the pairwise relative performance of all heuristics appears similar.

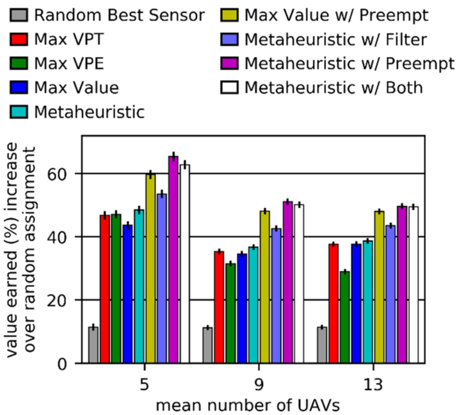

Fig. 8 A comparison of the percentage increase in surveillance value earned when compared to the Random heuristic in 10,000 small randomized scenarios. The mean number of UAVs is varied from the baseline set of scenarios with a mean of 9 UAVs. Except for the number of UAVs, the other characteristics of the scenario use the values from the baseline case described in Sect. 4.1. The 95% mean confidence intervals are shown for each bar

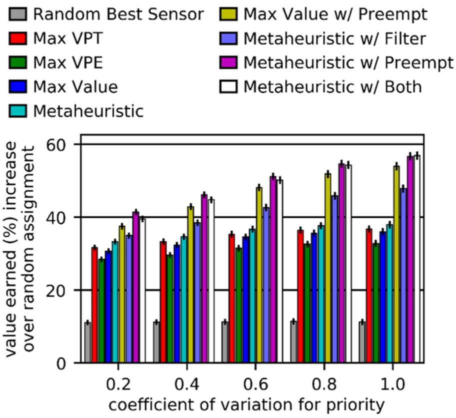

Fig. 9 A comparison of the percentage increase in surveillance value earned when compared to the Random heuristic in 10,000 small randomized scenarios. The coefficient of variation is varied from the baseline set of scenarios with coefficient of variation of 0.6. Except for this coefficient of variation, the other characteristics of the scenario use the values from the baseline case described in Sect. 4.1. The 95% mean confidence intervals are shown for each bar

5.1.3 Effect of number of UAVs and targets

When the number of UAVs is varied as shown in Fig. 8, our results show that the Metaheuristic with preemption performs among the best heuristics in all cases. When there are many UAVs, energy is not a significant constraint because there will be unused UAVs still available when some UAVs start running out of energy, which results in similar relative performance of heuristics to the cases with a low rate of energy consumption shown in Fig. 7.

When the number of targets in each scenario is varied, the results are very similar to varying the number of UAVs, except that trends seen when increasing the number of UAVs are instead seen when decreasing the number of targets. This is logical, because varying the number of UAVs and the number of targets are two different ways to vary the number of UAVs per target in a scenario.

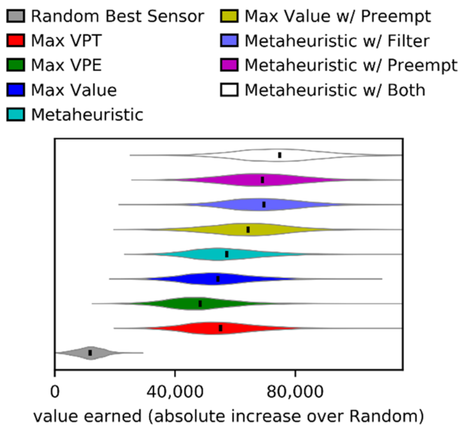

Fig. 10 A violin plot showing the difference between the surveillance value earned by each heuristic and the Random heuristic for the set of 10,000 large scenarios described in Sect. 4.3. The mean difference for each heuristic is indicated by the black marker in each distribution

5.1.4 Effect of coefficient of variation of target priorities

As the coefficient of variation of target priorities increases for the results in Fig. 9, there are three notable characteristics of the surveillance value earned by the heuristics. First, the Random and Random Best Sensor heuristics perform equally for all five values of the coefficient of variation that we considered (0.2, 0.4, 0.6, 0.8, and 1.0). This is expected because these heuristics select targets without considering priority. All value-based heuristics perform better as the coefficient of variation increases with the preemptive heuristics performing the best overall in all cases. The preemption and filtering modifications improve the most in performance when the variance in priority increases because they help to ensure that all high priority targets are surveilled. However, the Metaheuristic with preemption is still among the best overall heuristics because it is more likely that high priority targets will become available while low priority targets are being surveilled in the scenarios with a higher coefficient of variation.

5.2 Results for randomized set of large scenarios

The results in this subsection were generated using the set of large scenarios detailed in Sect 4.3. As shown in Fig. 10, the most significant differences between the overall results when compared to the smaller scenarios shown in Fig. 6 are that: (a) the Metaheuristic with both preemption and filtering is the best overall heuristic, (b) the filtering modification performs equal to preemption on its own when applied to the Metaheuristic, and (c) the performance of all value-based heuristics relative to Random has increased. Preemption performs well for the same reasons given in Sect. 5.1. The filtering modification performs significantly better in the larger scenarios because there is a larger number of UAVs and targets. This increases the chance that there are targets that have characteristics that are highly efficient for each UAV, which means it is more important to save energy to surveil those targets.

5.3 Discussion of results

The results in Sects. 5.1 and 5.2 indicate that depending on the scenario, either Max Value or Max Value Per Energy is an effective real-time heuristic for maximizing surveillance value. The Metaheuristic combines the strengths of both heuristics and is effective in all scenarios. When considering scenarios where the energy of UAVs will not be fully consumed during the day, Max Value Per Energy is ineffective. Based on these results, our proposed Metaheuristic is the best option to use in all cases where the characteristics of the scenario may change unexpectedly.

Additionally, the results demonstrate that both of our proposed modifications for the heuristics, preemption and filtering, improve performance of the Metaheuristic on average. The preemption modification always results in significant improvement in average surveillance value earned because it allows immediate response when options that would result in higher value surveils become available. For example, UAVs can immediately begin surveilling high priority targets without completing their currently active surveil first. The filtering modification performs extremely well in scenarios where a significant portion of the UAVs have options in the scenario that are significantly more efficient than others in terms of value earned per unit of energy consumed. This is because filtering conserves some of the energy for each UAV until later in the day when these efficient options that result in high surveillance value may be available.

In our simulations, mapping events for the slowest heuristic, the Metaheuristic with both preemption and filtering, took an average of less than 10 ms for each mapping to complete for the small set of scenarios and an average of less than 1 s each for the large set of scenarios. This demonstrates that any of these heuristics can be used to find mission schedules in real-time.

6 Related work

Developing a complete mission schedule for UAVs involves solving multiple problems, many of which have been studied in the past such as planning the specific routes used by UAVs, which we do not consider in this study, and assigning specific tasks to UAVs. Some studies solve these problems through time-consuming optimization techniques such as mixed integer linear programming (MILP), while others use techniques ranging from time-intensive metaheuristics like genetic algorithms (GAs) to fast and efficient greedy heuristics to find effective solutions.

In [5], mission planning is divided into two subproblems: task scheduling and route planning. The task scheduling problem is the one we consider in our study. The task scheduling problem is solved using an MILP approach to minimize the completion time of all tasks as opposed to our work that aims to maximize surveillance value. Similarly, in [6] a swarm of UAVs is also optimized to perform tasks while minimizing total completion time using an MILP approach. In [7], UAVs are assigned to attack and attempt to destroy clusters of targets through expressing three objective measures into one weighted measure (the success probability of the attack, the cost of the attack as a function of fuel consumption and risk to the UAV, and how well the timing of the attack will match a desired window), which is used to apply integer programming methods to find a solution. No-fly zones are considered in [8], which compares MILP and heuristic techniques to solve a task assignment problem where UAVs must complete a sequence of tasks. A solution here is represented by a directed acyclic graph (DAG). In [9], possible solutions for mapping a UAV to any combination of targets is represented by a decision tree, where moving from the root to a node represents assigning the UAV to the target corresponding to that node. A best first search (BFS) method is used to find solutions to this problem. In [10], a fleet of UAVs must be used to provide continuous 5G network coverage to the region of interest. The goal of this work is to determine both the required number of UAVs to guarantee coverage and to create a mission schedule for these UAVs. This is accomplished through a brute-force combinatorial technique to find the optimal solution, which is applicable in this case due to the size of problems considered. The most significant difference between these studies and our work is that we consider scenarios where decisions must be made in real-time throughout the day in a dynamic environment based on a mathematical model of a performance metric.

Mission planning for UAVs is sometimes studied as an orienteering problem [11, 12]. For example, in [11] the authors utilize a model where UAVs originate from a depot and gain profit from traveling a path through nodes and back to the depot. Robust optimization techniques are used to maximize profit while taking uncertainty into account to avoid running out of fuel early. This work differs significantly from ours because distance between targets and UAVs is considered and the focus is on optimization of UAV movement instead of sensing. In [12], UAVs again depart from a depot, but before departure they can select a specific set of sensors, which will impact their weight and the information they can gather. The authors solve this problem using both an MILP approach when ample time is available for finding solutions and several heuristic techniques for larger problem sizes that cannot be solved in a reasonable amount of time using the MILP approach. Our work differs significantly from these studies in part because a full mission plan is generated by the MILP approach and it is not modified during the day. In our work, the heuristics dynamically schedule UAVs to assign targets many times throughout the day and these decisions depend on the current state of the scenario. Additionally, our work considers the energy consumption of each sensor, which can greatly impact scheduling decisions.

Some studies have a greater focus on motion planning of the UAVs [13–22]. For example, the homogenous UAVs in [21] form a swarm that must search for and destroy heterogenous mobile targets. A hybrid artificial potential field and ant colony optimization (HAPF-ACO) technique is developed and implemented, which allows the swarm to efficiently search the grid for targets while dynamically avoiding collisions and threats. The pheromone levels for the ant colony optimization prevent UAVs from searching specific areas excessively and the forces of the artificial potential field attract UAVs towards targets to attack them and repel the UAVs from hazards. In the simulated results, the HAPF-ACO technique performs significantly better in all metrics compared to the comparison techniques. The UAVs in [22] form search teams that explore an area and must periodically return to their base to refuel. To ensure continuous exploration of the target area, other UAVs in backup teams will be scheduled to arrive and switch with an active search team when a search team is scheduled to return to base to refuel. The goal of [22] is to minimize total fuel consumption, which is achieved through a brute force examination of many different time intervals for switching searching UAVs with backup UAVs. The focus on motion planning in these studies is outside the scope of our contribution in this work. Our focus is on characterizing UAVs, sensors, and targets to develop a mathematical model that can be used as a system-wide performance measure that quantifies the success of surveils. We use this model as a basis for designing, evaluating, and comparing various dynamic mission scheduling heuristics through extensive simulation studies. Furthermore, in our study, the set of UAVs have heterogenous properties rather than being homogeneous.

When the environment can change dynamically, it is important to make use of mission scheduling strategies that can efficiently react to these changes. In [23], a decentralized strategy called the consensus-based bundle algorithm (CBBA) is used to assign tasks to UAVs in a conflict-free manor. To adapt this algorithm to a dynamic environment where new tasks become available, the authors propose CBBA with local replanning (CBBA-LR). The CBBA-LR technique is shown to achieve equal performance in score to fully recalculating the CBBA solution from scratch for the environment after dynamic changes despite performing significantly less work. Some key differences between our work and [23] include that targets in our model can be partially surveilled to get partial value, the amount of surveilling a UAV can do is determined by the energy consumption of its sensors, and our mission schedules are determined by a centralized scheduler. In [23], tasks assigned to a UAV are always fully completed, the UAVs are limited by a maximum number of tasks they are allowed to complete, and the scheduler is decentralized.

7 Conclusions and future work

We created a novel metric to quantity system-wide performance of a distributed system of UAVs surveilling targets in dynamic scenarios. To effectively compare mission schedules using simulations, we also detailed a model for generating randomized scenarios where a heterogeneous set of UAVs surveils a set of heterogenous targets.

We designed a set of value-based heuristics used to conduct mission planning for UAV surveillance of a set of targets (Max Value, Max Value Per Time, Max Value Per Energy, and the Metaheuristic). We conducted a simulation study to evaluate, analyze, and compare these heuristics in a variety of scenarios. We found that while Max Value and Max Value Per Energy are each good heuristics for a subset of the scenarios considered, the Metaheuristic found solutions with among the highest surveillance value for all scenarios.

In addition, we developed two modifications to these heuristics to improve their performance (preemption and filtering). Our simulations demonstrate that both preemption and filtering can significantly improve the performance of our heuristics and can be combined to take advantage of the benefits provided by both modifications. We found that preemption is effective in all environments we considered and performs better than filtering on average; however, in environments where there is little energy available to the UAVs, the filtering modification performs much better than the preemption modification. These results make it clear that it is an effective choice to always employ the preemption modification and that filtering should also be included in energy-constrained environments.

A major topic for future consideration is to also model the position of UAVs and targets, including taking into account the potential for UAVs to fly to and refuel at a base station to replenish their energy. New heuristics could be designed to take this option into account, such as a heuristic that compares the expected total value that the UAV could earn based on the currently known targets with and without refueling to determine if spending time to refuel would be worth it. Additionally, when considering the position of UAVs and targets, the model for surveillance value could be extended to take the distance between the UAV sensor and target into account, because this can have an impact on the quality of information obtained by the sensor.

Acknowledgements Preliminary portions of this material appeared in conference papers [24–26]. The differences between this work and those preliminary versions include: (a) this work models surveillance intervals for each target instead of allowing targets to be surveilled at any time, (b) the characteristics of the UAVs and targets are not static in this work and can dynamically change during the day, (c) we design, analyze, and evaluate two new modifications to our heuristics called preemption and filtering, (d) we improve our simulation model by using more realistic distributions for the characteristics of the randomized scenarios, and (e) we simulated the heuristics in both small-scale and large-scale scenarios to demonstrate their scalability for use in real-time environments. The authors thank Patrick J. Burns and John N. Carbone for their comments on this research.

Author's contribution All authors contributed to the conception and design of the study. DM ran the simulations for the study and prepared the results. All authors met and discussed the simulations and results throughout the process. DM wrote the first draft of the manuscript and all authors reviewed and commented on the draft. All authors read, reviewed, and approved the final manuscript.

Funding No funding was received for conducting this study.

Data availability The datasets generated during and/or analyzed during the current study are available from the corresponding author on reasonable request. However, the process used to randomly generate all input data used in this study is detailed in Sect. 4.

Declarations

Conflict of interest The authors have no relevant competing interests to declare.

Ethical approval This declaration is not applicable to this study.

References

1. Crowder J and Carbone J (2018) "Autonomous Mission Planner and Supervisor (AMPS) for UAVs." 20th International Conference on Artificial Intelligence (ICAI '18), 195–201,
2. Garey MR, Johnson DS (1979) Computers and intractability: a guide to the theory of np-completeness. W. H. Freeman and Co.
3. Machovec D, Khemka B, Kumbhare N, Pasricha S, Maciejewski AA, Siegel HJ, Akoglu A, Koenig GA, Hariri S, Tunc C, Wright M, Hilton M, Rambharos R, Blandin C, Fargo F, Louri A, Imam N (2019) Utility-based resource management in an oversubscribed energy-constrained heterogeneous environment executing parallel applications. Parallel Comput 83:48–72
4. Khemka B, Friese R, Pasricha S, Maciejewski AA, Siegel HJ, Koenig GA, Powers S, Hilton M, Rambharos R, Poole S (2015) Utility maximizing dynamic resource management in an oversubscribed energy-constrained heterogeneous computing system. Sustain Comput: Inf Syst 5:14–30
5. Wang JJ, Zhang YF, Geng L, Fuh JYH, and Teo SH (2014)"Mission planning for heterogeneous tasks with heterogeneous UAVs." 13th International Conference on Control Automation Robotics & Vision (ICARCV), 1484–1489, Mar. 2014.
6. Schumacher C, Chandler P, Pachter M, and Pachter L (2003) "UAV Task assignment with timing constraints." AIAA Guidance, Navigation, and Control Conference and Exhibit, 9 pp.
7. Zeng J, Yang X, Yang L, Shen G (2010) Modeling for UAV resource scheduling under mission synchronization. J Syst Eng Electron 21(5):821–826
8. Leary S, Deittert M, and Bookless J (2011) "Constrained UAV mission planning: a comparison of approaches." 2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops), pp. 2002–2009,

9. Faied M, Mostafa A, and Girard A (2010) "Vehicle routing problem instances: application to multi-UAV mission planning." AIAA Guidance, Navigation, and Control Conference, 12 pp.,
10. Tipantuña C, Hesselbach X, Sánchez-Aguero V, Valera F, Vidal I, Nogales B (2019) An NFV-based energy scheduling algorithm for a 5G enabled fleet of programmable unmanned aerial vehicles. Wireless Commun Mobile Comput 15:20
11. Evers L, Dollevoet T, Barros AI, Monsuur H (2012) Robust UAV Mission planning. Ann Oper Res 222:293–315
12. Mufalli F, Batta R, Nagi R (2012) Simultaneous sensor selection and routing of unmanned aerial vehicles for complex mission plans. Comput Oper Res 39:2787–2799
13. Chung W, Crespi V, Cybenko G, and Jordan A (2005) "Distributed sensing and UAV scheduling for surveillance and tracking of unidentifiable targets." Proceedings of SPIE 5778, Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense IV, pp. 226–235
14. Pascarella D, Venticinque S, and Aversa R (2013) "Agent-based design for UAV mission planning." 8th International Conference on P2P, Parallel, Grid, Cloud and Internet Computing (3PGCIC), pp. 76–83
15. Kim J, Song BD, Morrison JR (2013) On the scheduling of systems of UAVs and fuel service stations for long-term mission fulfillment. J Intell Rob Syst 70:347–359
16. Atencia CR, Ser JD, Camacho D (2019) Weighted strategies to guide a multi-objective evolutionary algorithm for multi-UAV mission planning. Swarm Evol Comput 44:480–495
17. Zhen Z, Chen Y, Wen L, Han B (2020) An intelligent cooperative mission planning scheme of UAV swarm in uncertain dynamic environment. Aerospace Sci Technol 100:16
18. Li GQ, Zhou XG, Yin J, Xiao QY (2014) An UAV scheduling and planning method for post-disaster survey. ISPRS – Int Arch Photogrammetry, Remote Sens Spatial Inf Sci 10:169–172
19. Yang W, Lei L, and Deng J (2014) "Optimization and improvement for multi-UAV cooperative reconnaissance mission planning problem." 11th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP), pp. 10–15,
20. Thibbotuwawa A, Bocewicz G, Radzki G, Nielsen P, Banaszak Z (2020) UAV mission planning resistant to weather uncertainty. Sensors 20:24
21. Zhen Z, Chen Y, Wen L, Han B (2020) An Intelligent cooperative mission planning scheme of UAV swarm in uncertain dynamic environment. Aerospace Sci Technol 100:16
22. Choi M, Contreras P, Shon PVST, Choi HH (2016) Exploring an area by groups of UAVs in the presence of a refueling base. J Supercomput 72:3409–3427
23. Chen J, Qing X, Ye F, Xiao K, You K, Sun Q (2022) Consensus-based bundle algorithm with local replanning for heterogeneous Multi-UAV system in the time-sensitive and dynamic environment. J Supercomput 78:1712–1740
24. Friese RD, Crowder JA, Siegel HJ, and Carbone JN (2018) "Bi-objective study for the assignment of unmanned aerial vehicles to targets." 20th International Conference on Artificial Intelligence (ICAI '18), pp. 207–213
25. Friese RD, Crowder JA, Siegel HJ, and Carbone JN (2019), "Surveillance mission planning: model, performance measure, bi-objective analysis, partial surveils." 21st International Conference on Artificial Intelligence (ICAI '19), 7 pp.
26. Machovec D, Crowder JA, Siegel HJ, Pasricha S, and Maciejewski AA (2020) "Dynamic heuristics for surveillance mission scheduling with unmanned aerial vehicles in heterogeneous environments." 22nd International Conference on Artificial Intelligence (ICAI '20), 21 pp.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

Authors and Affiliations

Dylan Machovec[1] · Howard Jay Siegel[1,2] · James A. Crowder[3] · Sudeep Pasricha[1,2] · Anthony A. Maciejewski[1] · Ryan D. Friese[4]

Howard Jay Siegel
hj@colostate.edu

James A. Crowder
jim.crowder@caes.com

Sudeep Pasricha
sudeep@colostate.edu

Anthony A. Maciejewski
aam@colostate.edu

Ryan D. Friese
ryan.friese@pnnl.gov

1 Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO, USA

2 Department of Computer Science, Colorado State University, Fort Collins, CO, USA

3 CAES Advanced Technology and Engineering (AT&E), Colorado Springs, CO, USA

4 Pacific Northwest National Laboratory, , Richland, WA, USA