Multi-robot Persistent Coverage with Stochastic Task Costs

Derek Mitchell, Nilanjan Chakraborty, Katia Sycara, and Nathan Michael

***Abstract*— We propose the Stochastic Multi-Robot Persistent Coverage Problem (SMRPCP) and correspondant methodology to compute an optimal schedule that enables a fleet of energy-constrained unmanned aerial vehicles to repeatedly perform a set of tasks while maximizing the frequency of task completion and preserving energy reserves via recharging depots. The approach enables online modeling of uncertain task costs and yields a schedule that adapts according to an evolving energy expenditure model. A fast heuristic method is formulated that enables online generation of a schedule that concurrently maximizes task completion frequency and avoids the risk of individual robot energy-depletion and consequential platform failure. Failure mitigation is introduced through a recourse strategy that routes robots based on acceptable levels of risk. Simulation and experimental results evaluate the efficacy of the proposed methodology and demonstrate online system-level adaptation due to increasingly certain costs models acquired during the deployment execution.**

I. Introduction

We propose adaptive deployment algorithms that consider the Stochastic Multi-Robot Persistent Coverage Problem (SMRPCP) and enable energy-constrained robot teams operating over long durations to execute missions given uncertain energy expenditure models. Representative systems exhibiting these characteristics include the aerial delivery programs proposed by Amazon [1], DHL [2], and Matternet [3]. The aerial robot systems employed in these applications are equipped with finite energy reserves and approximate energy expenditure models that vary between platforms due to differences in actuators, batteries, and other perishable system characteristics. We propose a deployment algorithm that adapts the mission- and robot-level plans via a centralized formulation given the current energy expenditure model and potential risk of severe energy-depletion and consequential platform failure.

The persistent coverage problem may be addressed through strategies that seek to ensure sufficient energy reserves through rendezvous at charging stations with the assumption that expenditure models are known a priori [4, 5]. Uncertain energy expenditure characterization is addressed through health management and awareness approaches that seek to identify robot-level performance degradation (mechanical wear and sensor failure) through stochastic measurements and online modeling via Markov Decision Processes

Derek Mitchell, Katia Sycara, and Nathan Michael are affiliated with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA. {derekm,katia,nmichael}@cmu.edu.

Nilanjan Chakraborty is affiliated with the Department of Mechanical Engineering, Stony Brook University, Stony Brook, NY 11790, USA. nilanjan.chakraborty@stonybrook.edu.

We gratefully acknowledge the support of ONR grants N00014-13-1-0821 and N00014-09-1-0680 as well as the NSF award 1218542.

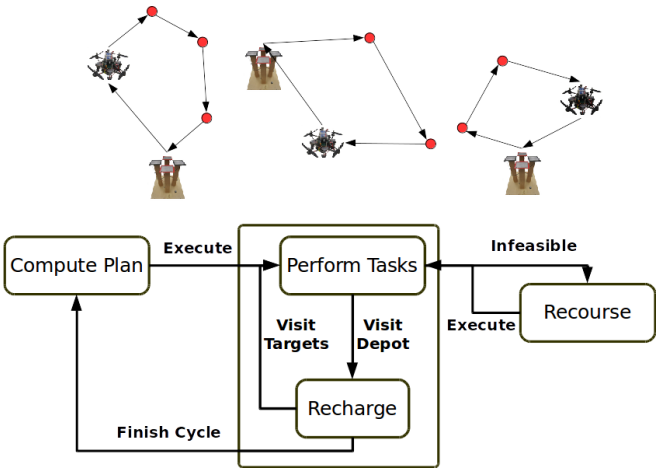


Fig. 1. Depiction of persistent deployment system with quadrotors and charging stations (top). Diagram of algorithm progression.

and demonstrated in the context of single robot, single package delivery [6] . We propose a modified interpretation of persistent coverage as a combinatorial optimization problem that constrains the robots to maintain a sufficient energy level to ensure continued operation given an initially unknown expenditure model. The primary contribution in this work is a methodology that admits failure mitigation strategies through recourse and the incorporation of an acceptable risk associated with critical levels of energy depletion that manifest in platform failure.

The proposed methodology builds on a wealth of prior related literature on the topic of optimal constrained deployments formulated as a Mixed Integer Linear Program (MILP) [7]. In this context, the SMRPCP is expressed as a multi-depot, multi-vehicle routing problem with energy constraints and an objective that prioritizes cost distribution through flow constraints. MILP formulations are limited due to computational complexity that grows with the size of the solution space with significant solution times that preclude online evaluation for problem sizes as detailed in this work. Consequently, heuristic methods that leverage iterative local search techniques are proposed that address the deterministic fuel constrained multi-robot routing problem yielding sub-optimal feasible solutions with approximate bounds [8]. Following prior work [9], we formulate the Multi-Robot Persistent Coverage Problem (MRPCP) as an optimization problem with an energy-aware objective formulation (Sect. II-A), integer and degree constraints (Sect. II-B), and sub-tour elimination (Sect. II-C). In this work, we extend this approach to incorporate energy depletion constraints through *chance*

Symbol	Usage
$\mathcal{K}$	set of robot fleet indices
$\mathcal{D}$	set of all charging depots
$\mathcal{B}$	set of initial charging depots ($\mathcal{B} \subseteq \mathcal{D}$)
$\mathcal{T}$	set of visitation targets
$\mathcal{S}$	set of Monte-Carlo sample indices
$\mathcal{L}$	combined set of all location indices corresponding to target and depot indices ($\mathcal{T} \cup \mathcal{D}$)
F	total energy capacity of each robot
C	modified cost matrix
l_i	location associated with index $i \in \mathcal{L}$
b_k	starting depot index of robot k
r_i	energy reserve of robot performing task i
t_i^e	execution time of task i
f_i^e	execution energy cost of task i
$t_{i,j}^l$	time duration to travel from l_i to l_j
$f_{i,j}^l$	energy cost to travel from l_i to l_j
$\hat{f}_+$	rate of energy gain
$\hat{f}_-$	rate of energy loss
q	ratio of rate of energy loss to energy gain
γ	risk factor associated with the likelihood of energy depletion
$x_{i,j}^k$	binary variables indicating if robot k traverses from l_i to l_j
z_s	binary variable indicating if a sample violates energy constraints

constraints [10] that model the initially uncertain (stochastic) energy expenditure associated with task execution (Sect. II-D). The proposed problem is therefore a stochastic variant of the MRPCP (or SMRPCP). Heuristic methods that yield solutions with bounded sub-optimality are proposed to enable the real-time performance toward enabling online system-level deployment adaptation (Sect. III). As chance constraints allow for the violation of energy reserve constraints, failure mitigation is introduced through recourse and an associated acceptable risk of platform failure.

II. Problem Definition

The proposed methodology seeks to generate optimal deployments for energy constrained robots that maximize the expected frequency of target visitation while ensuring, with high probability, that individual robots do not fully deplete energy reserves. Consider a fleet of $|\mathcal{K}|$ robots with finite energy reserves, F, tasked to repeatedly visit a set of target locations, $\mathcal{T}$, and recharging depots, $\mathcal{D}$, with infinite energy storage capacity distributed throughout the environment. The visitation targets and recharging depots represent the set of locations, $\mathcal{L} = \mathcal{T} \cup \mathcal{D}$, that the system visits persistently, with each location denoted as l_i, $i \in \mathcal{L}$. Assuming that robot k begins at depot b_k with the set of all initial depots defined by $\mathcal{B}$, the energy reserves of each robot ($k \in \mathcal{K}$) is assumed to evolve as follows:

- energy expenditure, $f_{i,j}^l$, due to traversal from l_i and l_j;
- energy expenditure, f_i^e, while operating at l_i for a finite duration t_i^e; and
- energy gain while recharging at a depot modeled as the amount of energy required to achieve full capacity.

We assume that the traversal energy expenditure, $f_{i,j}^l$, is known a priori and that the task execution energy expenditure, f_i^e, may be initially unknown but can be modeled as an independent stationary process. For ease of presentation, we assume that energy expenditure is modeled as a Gaussian distribution but note that the proposed techniques generalize to a broader class of probability distributions. We assume that all target locations are accessible from a depot and that it is always possible to transit from a depot to a another depot or target location and return to a potentially different depot without fully depleting energy reserves.

A. Objective Formulation

We define the optimal solution to the deterministic MRPCP problem as the deployment that minimizes the time required to perform a set of tasks over a single cycle [9]. Therefore, given binary variables that indicate if robot k traverses from l_i to l_j:

$$x_{i,j}^k \in \{0,\, 1\} \quad \forall \{i,\, j\} \in \mathcal{L},\, k \in \mathcal{K}$$

the optimal deployment minimizes the objective:

$$J^* = \min_{x_{i,j}^k} \max_{k \in \mathcal{K}} \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{L}} (1+q) \left(t_{i,j}^l + \frac{t_i^e + t_j^e}{2} \right) x_{i,j}^k$$

In this work, we introduce energy loss and gain through task execution and recharging. As energy expenditure is modeled as a stochastic process, the objective formulation is adapted accordingly:

$$J^* = \min_{x_{i,j}^k} \max_{k \in \mathcal{K}} \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{L}} \mathbb{E} \left[(1+q) \left(t_{i,j}^l + \frac{t_i^e + t_j^e}{2} \right) x_{i,j}^k \right]$$

with the time spent expending energy in between nodes i and j defined as $t_{i,j}^l + (t_i^e + t_j^e)/2$. Denoting $\hat{t}_{i,j}^l$ as the time required to replenish $f_{i,j}^l$ energy units:

$$f_{i,j}^l = t_{i,j}^l \hat{f}_-$$
$$\hat{t}_{i,j}^l = \frac{f_{i,j}^l}{\hat{f}_+} = t_{i,j}^l \frac{\hat{f}_-}{\hat{f}_+}$$

and define $q = \hat{f}_- / \hat{f}_+$, with a similar approach applied to task cost. Thus charging time is represented by the q factor multiplied by the total time spent expending energy and we can assume that there is no added time cost to dwell at depots (e.g., $t_j^e = 0$ for depot j).

B. Integer and Degree Constraints

Integer constraints enforce that the decision variable representing edge traversal, $x_{i,j}^k$, is integer for all edges and binary for edges that contain a task.

$$x_{i,j}^k \in \{0,\, 1\} \quad \forall i \in \mathcal{T} \text{ or } j \in \mathcal{T}, k \in \mathcal{K}$$
$$x_{i,j}^k \in \{0,\, 1,\, 2, \ldots,\, |\mathcal{T}|\} \quad \forall \{i,\, j\} \in \mathcal{D}, k \in \mathcal{K}$$

Denoting the starting location of robot k as b_k and the set of all starting locations, $\mathcal{B}$, with indices correspondant to robots $k \in \mathcal{K}$, the degree constraints enforce that every task

is performed once and that robots approaching a node must depart from the same node.

$$\begin{aligned}
\textstyle\sum_{k\in\mathcal{K}}\sum_{j\in\mathcal{L}} x^k_{i,j} = 1 \qquad & \forall i \in \mathcal{T} \\
\textstyle\sum_{k\in\mathcal{K}}\sum_{j\in\mathcal{L}} x^k_{j,i} = 1 \qquad & \forall i \in \mathcal{T} \\
\textstyle\sum_{j\in\mathcal{L}} x^k_{b_k,j} \leq 1 \qquad & \forall k \in \mathcal{K} \\
\textstyle\sum_{j\in\mathcal{L}} x^k_{j,b_k} \leq 1 \qquad & \forall k \in \mathcal{K} \\
\textstyle\sum_{i\in\mathcal{L}} \left(x^k_{i,j} - x^k_{j,i}\right) = 0 \quad & \forall j \in \mathcal{L}, \forall k \in \mathcal{K}
\end{aligned}$$

C. Subtour Elimination Constraints

Subtour elimination constraints ensure that each robot plan consists of a single connected cycle and correspond to flow constraints where robot k is given a finite number of units of a commodity and drops off one unit whenever it performs a task [11]. For brevity, we only note that $p^k_{i,j}$ is a decision variable introduced to utilize a standard flow constraint formulation to enforce connected tours and defer to prior work for further clarification [9].

$$\begin{aligned}
\textstyle\sum_{i\in\mathcal{L}} \left(p^k_{b_k,i} - p^k_{i,b_k}\right) = \sum_{i\in\mathcal{T},\, j\in\mathcal{L}} x^k_{i,j} \qquad & \forall k \in \mathcal{K} \\
\textstyle\sum_{j\in\mathcal{L}\setminus\{i\}} \left(p^k_{j,i} - p^k_{i,j}\right) = \sum_{j\in\mathcal{L}} x^k_{i,j} \qquad & \forall i \in \mathcal{T},\, k \in \mathcal{K} \\
\textstyle\sum_{j\in\mathcal{L}\setminus\{i\}} \left(p^k_{j,i} - p^k_{i,j}\right) = 0 \qquad & \forall i \in \mathcal{D},\, k \in \mathcal{K} \\
0 \leq p^k_{i,j} \leq |\,\mathcal{T}\,|\, x^k_{i,j} \qquad & \forall \{i,j\} \in \mathcal{L},\, k \in \mathcal{K}
\end{aligned}$$

D. Energy Depletion Constraints

The proposed SMRPCP arises from the introduction of non-deterministic costs representing uncertain energy expenditure models and corresponding energy depletion constraints. A chance constraint formulation introduces uncertain constraints that encode a *risk factor* associated with the probability of constraint violation, denoted as γ. Feasible solutions are those deployments that yield constraint satisfaction with probability $1-\gamma$ given a set of costs generated through Monte Carlo sampling and drawn from task cost distributions. Therefore, the sum of all binary variables corresponding to the validity of the constraints, denoted by z_s, is compared to the fraction of samples that must be valid for the constraint to be held true:

$$\sum_{s\in S} z_s \geq (1-\gamma)|\,\mathcal{S}\,| \qquad z_s \in \{0,1\},\ s \in \mathcal{S}$$

The proposed formulation introduces both execution and transition costs that may be considered as node and edge costs, respectively. However, the proposed solution methodology is only applicable to systems with undirected edge costs. Therefore, to transform the proposed formulation into a weighted undirected graph with only edge costs, we define the modified edge cost associated with energy and time expenditure [12]:

$$\bar{f}_{i,j} = f^l_{i,j} + \frac{f^e_i + f^e_j}{2} \quad \text{and} \quad \bar{t}_{i,j} = t^l_{i,j} + \frac{t^e_i + t^e_j}{2} \qquad (1)$$

We denote the available energy reserve of the robot currently executing task i by r_i with capacity constraints:

$$0 \leq r_i \leq F \qquad \forall i \in \mathcal{T}$$

A feasible deployment ensures that the energy reserves of individual robots are sufficient to enable task execution and the transition between target locations. Following Sundar and Rathinam [11], we introduce energy reserve constraints that ensure that the modified expenditure cost is lower than the available energy reserve costs for the sampled deployment:

$$\begin{aligned}
r_j - r_i + \bar{f}_{i,j}(s) &\leq \mathcal{M}\left(2 - z_s - x^k_{i,j}\right) \\
r_i - r_j - \bar{f}_{i,j}(s) &\leq \mathcal{M}\left(2 - z_s - x^k_{i,j}\right) \\
&\qquad \forall \{i,\, j\} \in \mathcal{T},\, k \in \mathcal{K},\, s \in \mathcal{S}
\end{aligned}$$

$$\begin{aligned}
r_j - F + \bar{f}_{i,j}(s) &\leq \mathcal{M}\left(2 - z_s - x^k_{i,j}\right) \\
F - r_j - \bar{f}_{i,j}(s) &\leq \mathcal{M}\left(2 - z_s - x^k_{i,j}\right) \\
&\qquad \forall i \in \mathcal{D} \text{ and } j \in \mathcal{T},\, k \in \mathcal{K},\, s \in \mathcal{S}
\end{aligned}$$

$$\begin{aligned}
-r_i + \bar{f}_{i,j}(s) &\leq \mathcal{M}\left(2 - z_s - x^k_{i,j}\right) \\
&\qquad \forall i \in \mathcal{T} \text{ and } j \in \mathcal{D},\, k \in \mathcal{K},\, s \in \mathcal{S}
\end{aligned}$$

where $\bar{f}_{i,j}(s)$ represents the modified expenditure cost associated with the graph defined by the sample s and $\mathcal{M}$ represents an arbitrarily large value [13].

We now detail a multi-step heuristic approach that solves the stochastic MRPCP and generates plans that ensure that the robots maintain sufficient energy reserves with probability $1-\gamma$. The approach updates an expenditure model online and recomputes deployments based on the current expenditure cost model (Sect. III-D). A recourse strategy mitigates failures that may arise due to significant differences between the current and true expenditure models (Sect. IV).

III. Heuristic Approach

In this section, we present a Monte-Carlo approach to solve the stochastic MRPCP detailed in Sect. II. The algorithm consists of three steps:

1) Formulate an expenditure cost matrix that ensures sufficient energy reserves through the introduction of detours to charging depots.
2) Compute feasible robot tours via the expenditure cost matrix.
3) Refine the computed robot tours to enforce energy constraints.

The first and third steps differ from the previously proposed MRPCP heuristic methodology and capture the stochastic nature of the current problem. We abbreviate the presentation of the second step and defer to prior work for a detailed discussion [9].

A. Cost Adjustment

Algorithm 1 formulates a modified cost matrix, $C \in \mathbb{R}^{|\,\mathcal{L}\,|\times|\,\mathcal{L}\,|}$, with entries that encode the transition costs between tasks to enable the computation of an optimal deployment. Given the undirected graph formulation with the modified edge costs detailed in Sect. II-D, we wish to ensure that the transition between any two tasks i and j is feasible. The worst case scenario arises when the maximum available energy reserves, F, is lower than the cost to transit from the closest depot to i, execute i, transit to j, execute j, and return to the closest depot to j. Therefore, the following condition must hold true to ensure edge feasibility between

Algorithm 1: Cost Adjustment

Input: set of targets $\mathcal{T}$, set of depots $\mathcal{D}$, fuel capacity F, sample index set $\mathcal{S}$, modified energy expenditure cost $\bar{f}_{i,j}$, modified time expenditure cost $\bar{t}_{i,j}$, risk factor γ

Output: modified cost matrix C

begin

  for $\{i,j\} \in \mathcal{T}, i \neq j$ **do**

    if $|\{\forall s \in \mathcal{S} : \bar{f}_{i,j} \leq F - \min_{d\in\mathcal{D}} \bar{f}_{d,i} - \min_{d\in\mathcal{D}} \bar{f}_{j,d}\}| > (1-\gamma)|\mathcal{S}|$ **then**

      $c_{i,j} = \bar{t}_{i,j}$

    else

      $P_{i,j} = \text{minimumTimePath}(i, \mathcal{D}, j)$

      $c_{i,j} = \text{pathTimeCost}(P_{i,j})$

any two tasks:

$$\bar{f}_{i,j} \leq F - \min_{d\in\mathcal{D}} \bar{f}_{d,i} - \min_{d\in\mathcal{D}} \bar{f}_{j,d} \tag{2}$$

given the modified energy expenditure cost (1) and the assumption that there is no time execution cost at depots.

We require that condition (2) hold true with probability $1-\gamma$ for $s \in \mathcal{S}$ and define feasible costs in C as $c_{i,j} = \bar{t}_{i,j}$. If the probability of excessive energy expenditure exceeds the acceptable risk level, a detour is introduced that represents the minimum time expenditure path from l_i to l_j via one or multiple depots. The resulting costs associated with the minimum time path are introduced into C accordingly.

B. Overview of Tour Computation

Given the modified cost matrix determined by Alg. 1, we compute a tour for each robot using the following steps: (a) compute a minimum length Hamiltonian path using Christofides algorithm [14], (b) split the path into n subpaths (with n less than or equal to the number of robots), and (c) assign the robots to the subpaths so that the maximum cost of a subpath is minimized (by formulating and solving a bottleneck assignment problem). For details, we defer to the prior work except to note that the modified cost matrix obtained from Sect. III-A may not satisfy the triangle inequality but the shortcutting step of Christofides algorithm can be modified to ensure that there is no increase in cost during shortcutting [9].

C. Strand Repair

The tours generated from the splitting algorithm form the basis of our schedule. However, they do not yet guarantee feasibility. To ensure that the probability of failure falls below the risk factor, γ, we insert refueling detours into each tour utilizing a method that decomposes a tour into strands and makes each strand feasible [15]. A *strand* refers to a sequence of nodes that begins and ends with a depot, but otherwise consists of only of tasks. The approach entails adding recharging detours to each node and greedily remove them where feasible. The only change necessary in our approach is that the feasibility check must be valid with probability $1-\gamma$. This condition is enforced by checking the sum of all costs for the segment in question and ensuring that the percentage of samples that exceed the remaining energy reserve is less than γ.

D. Cost Update Model

We model the stochastic expenditure costs as Gaussian with the task execution expenditure process model from time $t-1$ until t:

$$f_i^e(t) = f_i^e(t-1) - \hat{f}^-$$

and assume we can measure the current energy reserve level, $\tilde{f}_i^e$:

$$\tilde{f}_i^e(t) = f_i^e(t) + v$$

where v is zero mean Gaussian noise with variance R_v.

As a robot visits a target, the current energy reserve level is monitored and the amount of energy expended is introduced as a measurement into a Kalman-filter [16] that updates the current estimate of the task expenditure cost, f_i^e. If the robot is unable to finish a task, the measurement is limited to the amount of energy expended at that moment. The true measurement is extrapolated from the partial input into the Kalman-filter. In this manner, we ensure that the system is always progressing toward a more accurate representation of the environment.

IV. Recourse

While the proposed methodology seeks to ensure feasibility, it is possible that energy constraints may become violated during a deployment due to unanticipated energy expenditure. We propose a greedy search for the optimal configuration of tasks that satisfy the energy constraints that is sufficiently fast to execute online assuming that the number of targets a robot has yet to visit in its current strand is sufficiently small. The proposed recourse strategy computes the maximum number of targets that can be kept while maintaining a sufficiently low risk of full energy depletion and minimizing the time to complete the strand.

Algorithm 2 initially prioritizes the number of targets in the sequence and increments the number of potential dropped targets until a valid configuration sequence is determined by comparing the percentage of sample graphs for which the new strand is feasible with probability $1-\gamma$. To prevent a robot from running out of energy while executing a task or traveling, we set the minimum energy threshold to the capacity required to visit the closest depot. If the energy reserve level reaches this threshold, we abandon the remainder of the plan and travel to that depot.

V. Simulation Evaluation

We evaluate the efficacy of the proposed methodology via realistic simulations developed from expenditure models established through experiments (Fig. 2) that accurately models the system dynamics and energy expenditure. The simulation studies evaluate: (1) the convergence rate of target dropping

Algorithm 2: Recourse

Input: current position $i \in \mathcal{T}$, next depot in sequence $d \in \mathcal{D}$, sequence of targets, $\mathcal{T}' \subseteq \mathcal{T}$, between i and depot, d, remaining energy reserve $F' \leq F$, modified edge costs $\bar{t}_{i,j}$ and $\bar{f}_{i,j}(s) \quad \forall s \in \mathcal{S}$, risk factor γ

Output: adjusted sequence of targets

begin

  for $n = 0 : |\mathcal{T}'|$ **do**

    $Q \leftarrow$ set of all subsets of $\mathcal{T}'$ with n targets removed

    for $q \in Q$ **do**

      $A = \{i, q, d\}$ $C_q = \sum_{m=1}^{|A|-1} \bar{t}_{A_m, A_{m+1}}$

      $V_q = \left| \{\forall s \in \mathcal{S} : \sum_{m=1}^{|A|-1} \bar{f}_{A_m, A_{m+1}}(s) > F'\} \right|$

    if $\exists q \in Q$ *such that* $V_q \leq \gamma |\mathcal{S}|$ **then**

      return $\arg\min_{q \in \{Q : V_q \leq \gamma |\mathcal{S}|\}} C_q$

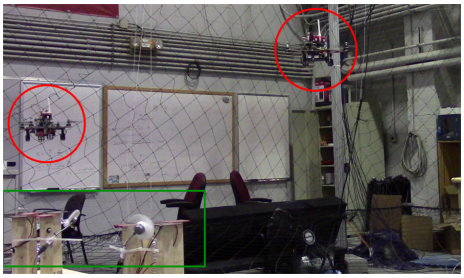


Fig. 2. Experimental setup with two quadrotors and depots inside a motion capture arena. Red circles indicate the quadrotors and a green square highlights the charging stations. The energy expenditure models employed during the simulation trials were developed through experimental characterization.

relative to the risk, and (2) the efficacy of recourse. For (1), we expect that higher acceptable levels of risk associated with larger γ values will yield faster convergence to a lower number of targets dropped as the system approaches steady state. Conversely, operations with a lower γ value will tend toward preserving energy reserves. For (2), we expect that the energy level will always remain adequate for operation with recourse and yield platform failures otherwise.

A. Experimental Design

Depots were distributed randomly such that a neighboring depot exists that is reachable given $0.2F$ reserve capacity. Targets were distributed such that a neighboring depot exists that is reachable given $0.1F$. For each run, we chose a random distribution of depots and tasks with execution expenditure costs, f_i^e, of $0.25F$. To learn these costs over time, we implemented a Kalman filter that updates the predicted state with measurement noise $R_v = \frac{1}{12}F$.

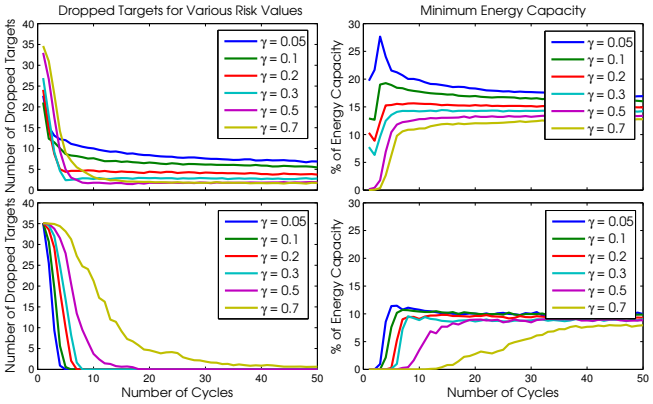


Fig. 3. Simulation results averaged over fifty runs. Dropped targets per cycle for various levels of risk (top left). Minimum energy reserve level experienced during each cycle for all robots (top right). Results computed without recourse (bottom).

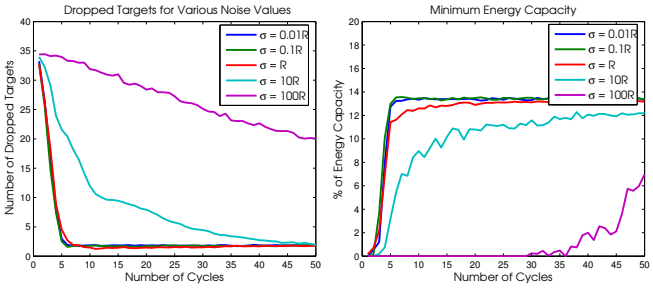


Fig. 4. Simulation Results averaged over fifty runs with risk $\gamma = 0.5$, varying over the amount of measurement noise σ relative to $R = \frac{1}{12}F$.

B. Impact of Acceptable Risk and Recourse

Figure 3 confirms the expected algorithmic response to changes in risk with an average of 29 cycles to reach 10% convergence when $\gamma = 0.05$, 25 cycles for $\gamma = 0.1$, and 5 cycles for $\gamma = 0.2$. Note that the convergence rate begins to increase again as the risk factor gets larger than $\gamma = 0.3$, suggesting a point of diminishing return when it is no longer beneficial to increase the level of acceptable risk.

Figure 3 also depicts the impact of risk on the available energy reserves. As the risk factor increases the reserve decreases while lower risk results in a larger reserve. Note that without recourse, the energy reserve is significantly reduced and well below the designed environmental bound of $0.1F$ required to transit from task locations to depots.

C. Impact of Measurement Uncertainty

Figure 4 depicts the rate of convergence of the cost model given varying levels of measurement uncertainty. As expected, increasingly uncertain measurements greatly increase the dropped target convergence rate and more accurate measurements offer limited improvement after a point. We also observe a similar degradation in minimum energy capacity as noise increases.

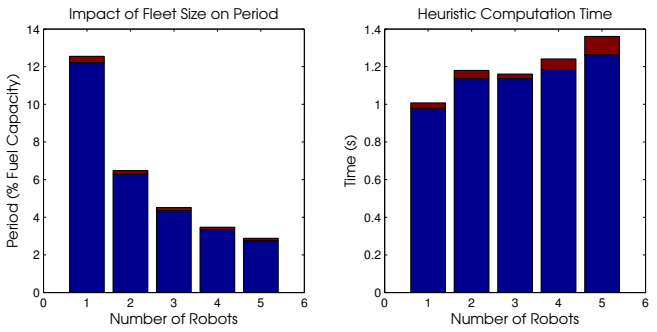


Fig. 5. Graphs depicting the maximum energy expended during a cycle (left) and heuristic computation time (right) relative to the fleet size averaged over fifty runs.

D. Fleet Size Evaluation

Figure 5 shows the impact of fleet size on period. As more robots are introduced, the energy expenditure becomes increasingly distributed across the system. This distribution is limited to at least one robot per task and thus the maximum period will be determined by the largest travel and task cost.

VI. Summary and Future Work

In this paper, we present an algorithm to generate plans for performing a set of tasks persistently over long durations where the cost of performing a task is non-deterministic. Coupling this algorithm with another capable of learning an accurate model of the cost probability distribution allows the user to generate plans that restrict the risk of full energy depletion to an acceptable level. We also present a recourse algorithm capable of updating plans as targets must be dropped to maintain an energy reserve level to support emergency detours to charging depots. We evaluate the performance through simulation results and analyze the impact of risk, measurement uncertainty, and fleet size on algorithm performance. The results of the evaluation suggest that the risk factor adequately captures a measure of robustness or convergence rate that an operator can adapt based on application requirements.

This work presents preliminary steps toward integrating stochastic cost into the deployment approach detailed in prior work [9]. We plan to pursue approximation bounds established through our prior work by exploring the bounds that can be established when using Monte-Carlo sampling with chance constraints. We will then extend the approach to handle directed graphs, collision avoidance, and depot conflict. The end result will be a fully realized system that can drive a team of robots to persistently perform mission tasks in an unknown or dynamic environment.

References

[1] (2015, July) Prime air. Amazon. [Online]. Available: http://www.amazon.com/b?node=8037720011

[2] (2015, July) Unmanned aerial vehicles – ready for take-off? DHL. [Online]. Available: http://www.dhl.com/en/about_us/logistics_insights/dhl_trend_research/uav.html

[3] (2015, July) Matternet. Matternet. [Online]. Available: http://matternet.us/

[4] J. Barnes, V. Wiley, J. Moore, and D. Ryer, "Solving the aerial fleet refueling problem using group theoretic tabu search," *Mathematical and Computer Modelling*, vol. 39, no. 6–8, pp. 617 – 640, Mar. 2004.

[5] N. Mathew, S. Smith, and S. Waslander, "Multirobot rendezvous planning for recharging in persistent tasks," *IEEE Trans. Robot.*, vol. 31, no. 1, pp. 128–142, Feb. 2015.

[6] A.-A. Agha-mohammadi, N. Ure, J. How, and J. Vian, "Health aware stochastic planning for persistent package delivery missions using quadrotors," in *Proc. of the IEEE/RSJ Int. Conf. on Intelli. Robots and Systs.*, Sept. 2014, pp. 3389–3396.

[7] C. Prins, P. Lacomme, and C. Prodhon, "Order-first split-second methods for vehicle routing problems: A review," *Transportation Research Part C: Emerging Technologies*, vol. 40, pp. 179–200, Mar. 2014.

[8] D. Levy, K. Sundar, and S. Rathinam, "Heuristics for routing heterogeneous unmanned vehicles with fuel constraints," *Mathematical Problems in Engineering*, vol. 2014, p. 12 pages, Apr. 2014.

[9] D. Mitchell, M. Corah, N. Chakraborty, K. Sycara, and N. Michael, "Multi-robot long-term persistent coverage with fuel constrained robots," in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, May 2015, pp. 1093–1099.

[10] S. Ahmed and A. Shapiro, "Solving chance-constrained stochastic programs via sampling and integer programming," in *State-of-the-Art Decision-Making Tools in the Information-Intensive Age*, 2015, ch. 13, pp. 261–269.

[11] K. Sundar and S. Rathinam, "Algorithms for routing an unmanned aerial vehicle in the presence of refueling depots," *IEEE Trans. Autom. Sci. Eng.*, vol. 11, no. 1, pp. 287–294, Jan. 2014.

[12] N. Chakraborty, S. Akella, and J. Wen, "Coverage of a planar point set with multiple robots subject to geometric constraints," *IEEE Trans. Autom. Sci. Eng.*, vol. 7, no. 1, pp. 111–122, Jan. 2010.

[13] M. Kojima, S. Mizuno, and A. Yoshise, "A little theorem of the big $\mathcal{M}$ in interior point algorithms," *Mathematical Programming*, vol. 59, no. 1-3, pp. 361–375, Mar. 1993.

[14] N. Christofides, "Worst-case analysis of a new heuristic for the travelling salesman problem," Graduate School of Industrial Administration, Carnegie Mellon University, Technical Report 388, Feb. 1976.

[15] S. Khuller, A. Malekian, and J. Mestre, "To fill or not to fill: The gas station problem," in *Algorithms – ESA 2007*, ser. Lecture Notes in Computer Science, Oct. 2007, vol. 4698, pp. 534–545.

[16] R. Kalman, "A new approach to linear filtering and prediction problems," *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, Mar. 1960.