Robust Planning For Coupled Cooperative UAV Missions

L. F. Bertuccelli, M. Alighanbari and J. P. How

Aerospace Controls Laboratory
Massachusetts Institute of Technology
{lucab, mehdi_a, jhow}@mit.edu

***Abstract*— This paper presents a new formulation for the UAV task assignment problem with uncertainty in the environment. The problem is posed as a weapons task assignment with uncertainty in the cost data, and we apply a modified robustness technique that allows the operator to tune the level of robustness in the optimization. This robust formulation is then used to solve the assignment problem for a heterogeneous fleet of vehicles operating in an environment with target identities and locations that are uncertain. The key aspect of this formulation is that it directly addresses the coupling inherent in deciding how to assign vehicles to perform reconnaissance tasks that provide the most benefit to the strike part of the missions. We demonstrate that the robust solution to this coupled problem can be solved as single mixed-integer linear problem. The paper presents simulations and discusses hardware testbeds that will be used for future experiments.**

I. Introduction

Future UAV missions will require more autonomous high-level planning capabilities onboard the vehicles using information acquired through sensing or communicating with other UAVs in the group. This information will include battlefield parameters such as target identities/locations, but will be inherently uncertain due to real-world disturbances such as noisy sensors or even deceptive adversarial strategies. This paper presents a new approach to the high-level planning (i.e., task assignment) that accounts for uncertainty in the situational awareness of the environment.

Except for a few recent results [1], [8], [10], the controls community has largely treated the UAV task assignment problem as a deterministic optimization problem with perfectly known parameters. However, the Operations Research and finance communities have made significant progress in incorporating this uncertainty in the high-level planning and have generated techniques that make the optimization *robust* to the uncertainty [2], [3], [9], [11]. While these results have mainly been made available for Linear Programs (LPs) [2], robust optimization for Integer Programs (IPs) has only recently been provided with elegant and computationally tractable results [4], [9]. The latter formulation allows the operator to tune the level of robustness included by selecting how many parameters in the optimization are allowed to achieve their worst case values. The result is a robust design that reflects the level of risk-aversion (or acceptance) of the operator. This is by no means a unique method to tune the robustness, as the operator could want to restrict the worst case deviation of the parameters in the optimization, instead of allowing only a few to go to their worst case. This paper investigates including robustness to the uncertainty in the environment in the task assignment process to ensure that we obtain designs that are less sensitive to the errors in our situational awareness.

Environmental uncertainty also creates an inherent coupling between the missions of the heterogeneous vehicles in the team. Future UAV mission packages will include both strike and reconnaissance vehicles (possibly mixed), with each type of vehicle providing unique capabilities to the mission. For example, strike vehicles will have the critical firepower to eliminate a target, but may have to rely on reconnaissance vehicle capabilities in order to obtain valuable target information. Including this coupling will be critical in truly understanding the cooperative nature of missions with heterogeneous vehicles.

This paper investigates the impact of uncertain target identity by formulating a *weapon task* assignment problem with uncertain data. We assume that sensing errors cause uncertainty in the classification of a target. In the presence of this uncertainty, we robustly assign a set of vehicles to a subset of these targets in order to maximize a performance criterion. We then extend this robustness formulation to solve a mission with heterogeneous vehicles (namely, reconnaissance and strike) with coupled actions operating in an uncertain environment.

II. Robust Formulation

Consider a weapon-target assignment problem – given a set of N_T targets and a set of N_V vehicles, we seek to assign the vehicles to the targets to maximize the score of the mission. We assume that each target has a score associated with it based on the current classification, and that the vehicle accrues that score if it is assigned to that target. If a vehicle is not assigned to a target, it receives a score of 0. The mission score is the sum of the individual scores accrued by the vehicles; in order for the vehicles to visit the "best" targets, we assume that $N_V < N_T$. Due to sensing errors, deceptive adversarial strategies, or even poor intelligence, these scores will be uncertain, and we need to incorporate this lack of perfect information in our planning. We extend the stochastic and robust formulations that have been introduced in OR literature to deal with this uncertainty.

The basic stochastic programming formulation of this problem replaces the deterministic target scores with expected target scores [5], and mathematically, the goal is to maximize the following objective function at time k

$$\begin{aligned}\max_x J_k &= \sum_{i=1}^{N_T} \bar{c}_{k,i} x_{k,i} \qquad (1)\\ \text{subject to:} & \quad \sum_{i=1}^{N_T} x_{k,i} = N_V,\ x_i \in \{0,1\}\end{aligned}$$

(We henceforth summarize the constraints as $x \in X$.) The binary variable $x_{k,i}$ is 1 if a vehicle is assigned to target i and zero if it is not, and $\bar{c}_{k,i}$ represent the expected score of the i^{th} target at time k. We assume that any vehicle can be assigned to any target and (for now) all the vehicles are homogeneous. This formulation however only incorporates first moment information, which could be misleading; for example, two targets with equal expected scores but different variances in these scores would appear identically attractive to a strike vehicle with this formulation. In reality however, we would certainly prefer to assign a strike vehicle to the target that has the lower variance, and thus we require a formulation that incorporates this higher moment information.

Robust formulations have been developed to account for uncertainty in the data by incorporating uncertainty sets for the data [2]. These uncertainty sets can be modeled in various ways. One way is to generate a set of realizations (or scenarios) based on statistical information of the data, and using them explicitly in the optimization; another way is by using the values of the moments (mean and standard deviation) directly. Using either method, the *robust* formulation of the weapon task assignment is posed as

$$\begin{aligned}\max_x \min_c J_k &= \sum_{i=1}^{N_T} c_{k,i} x_{k,i}\\ \text{subject to:} & \quad x \in X \qquad (2)\\ & \quad c_{k,i} \in C_k\end{aligned}$$

The optimization becomes to obtain the "best" worst-case score when each $c_{k,i}$ is assumed to lie in the uncertainty set C_k. Characterization of this uncertainty set depends on any *a priori* knowledge we have of the uncertainty. The choice of this uncertainty set will generally result in different robust formulations that are either computationally intensive (many are NP-hard [7]) or extremely conservative.

One formulation that falls in the latter case is the Soyster formulation [11]. The appeal of the Soyster formulation however is its simplicity, as we will subsequently show. Here we investigate a modification of the Soyster formulation applied to integer programs. It allows a designer to solve a robust formulation in the same manner as an integer program while allowing a designer to tune the level of robustness desired in the solution. Here, the expected target scores, $\bar{c}_{k,i}$, are assumed to lie in the interval $[\bar{c}_{k,i} - \sigma_{k,i}, \bar{c}_{k,i} + \sigma_{k,i}]$, where $\sigma_{k,i}$ indicates the standard deviation of target i at time k. In this case the Soyster formulation solves the following problem

$$\begin{aligned}\max_x J_k &= \sum_{i=1}^{N_T} (\bar{c}_{k,i} - \sigma_{k,i}) x_{k,i}\\ \text{subject to:} & \quad x \in X \qquad (3)\end{aligned}$$

This formulation assigns vehicles to the targets that exhibit the highest "worst-case" score. Note that the use of expected scores and standard deviations is not restrictive; quite the opposite, they are rather general, providing sufficient statistics for the unknown true target scores. In general, solving the Soyster formulation results in an extremely conservative policy, since it is unlikely that each target will indeed achieve its worst case score; furthermore, it is unlikely that each target will achieve this score *at the same time*. We therefore introduce a straightforward modification to the cost function allowing the operator to accept or reject the uncertainty, by introducing a parameter (μ) that can vary the degree of uncertainty introduced in the problem. The modified robust formulation then takes the form

$$\begin{aligned}\max_x J_k &= \sum_{i=1}^{N_T} (\bar{c}_{k,i} - \mu\sigma_{k,i}) x_{k,i}\\ \text{subject to:} & \quad x \in X \qquad (4)\end{aligned}$$

μ restricts the $\mu\sigma$ deviation that the mission designer expects and serves as a tuning parameter to adjust the robustness of the solution. Note that $\mu = 0$ corresponds to the basic stochastic formulation (which relies on expected scores, and ignores second moment information), while $\mu = 1$ recovers the Soyster formulation. Higher values of μ may be used if we desire to be robust to larger variations of the data. Furthermore, we need not restrict μ to a positive scalar; μ could actually be a vector with elements μ_i which penalize each target score differently. This would certainly be useful if the operator desires to accept more uncertainty in one target than another.

III. SIMULATION RESULTS

We now demonstrate numerical results of this robust optimization for the case of an assignment with uncertain data, and compare them to the stochastic programming formulation (where we replace the target scores with the expected target scores). We took the case of 10 targets having random score $c_{k,i}$ and standard deviation σ_i, and evaluated the assignments generated from the *robust* and stochastic formulation, when the scores were allowed to vary in the interval $[c_{k,i} - \sigma_i, c_{k,i} + \sigma_i]$. We then compared the expected mission score, standard deviation, minimum, and maximum scores attained in 1000 numerical simulations, and the results may be seen in Table I. The simulations confirm the expectation that the robust optimization results in a lower *but more certain* mission score; while the robust mission score is 2.8% lower than the stochastic programming score, there is a 65% reduction in the standard deviation of this resulting score.

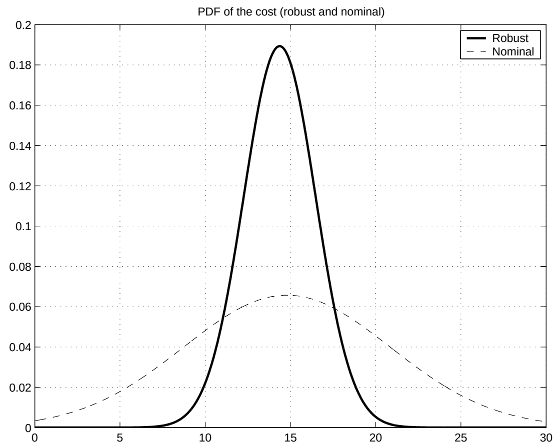

Fig. 1. Probability Density Functions

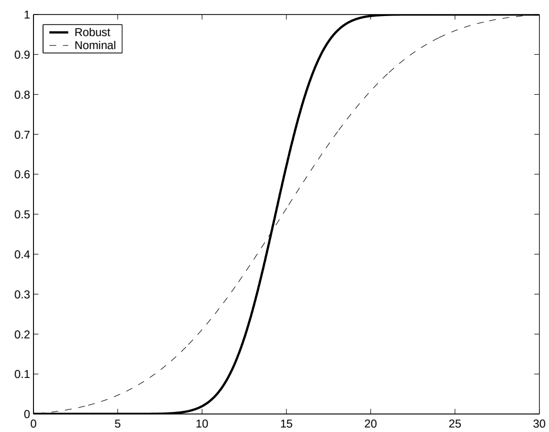

Fig. 2. Probability Distribution Functions

This results in less variability in the resulting mission scores, seen by considering a 2σ range for the mission scores: for the stochastic formulation this is $[2.65, 26.93]$ while for the robust one this is $[10.15, 18.59]$. Although the expected mission score is indeed lower, there is much more of a guarantee for this score. Furthermore, we note that the robust optimization has a higher minimum score in the simulations of 11.43 compared to 6.30 of the stochastic optimization, indicating that with the given bounds on the cost data, the robust optimization has a better guarantee of "worst-case" performance. This can also be seen in the probability density functions shown in Figure 1 (and the associated probability distribution functions in Figure 2). We can see that as the numerical results indicate, the stochastic formulation results in larger spread in the mission scores than the robust formulation, which restricts the range of possible missions scores. Thus, while the maximum achievable mission score is lower in the robust formulation than that obtained by the stochastic one, the missions scores in the range of the mean occur with much higher probability.

We also compared this robust formulation to the Conditional Value at Risk (CVaR) formulation in [8] in another

TABLE I

COMPARISON OF STOCHASTIC AND MODIFIED SOYSTER

Optimization	$\bar{J}$	σ_J	$\max$	$\min$
Stochastic	14.79	6.07	23.50	6.30
Robust	14.37	2.11	17.20	11.43

TABLE II

COMPARISON OF CVAR WITH MODIFIED SOYSTER

Number of Scenarios	$\bar{J}$	σ_J	$\max$	$\min$
10	18.01	2.41	22.80	13.20
20	17.33	1.87	20.98	13.62
50	17.39	1.99	20.98	13.62
100	16.51	1.18	18.90	14.10
200	16.58	1.31	19.16	14.04
500	16.51	1.18	18.90	14.10
Robust	16.51	1.18	18.90	14.10

series of experiments, where we had to assign 5 strike vehicles and 10 targets. CVaR is a modified version of the VaR optimization, which allows the operator to choose the level of "protection" in a probabilistic sense, based on given number of scenarios (N_{scen}) of the data. These scenarios are generated from realizations of the data in the range of $[c_{k,i} - \sigma_i, c_{k,i} + \sigma_i]$. This optimization can be expressed as

$$\begin{aligned} \max_x J_{VaR,k} \quad &= \quad \gamma + \frac{1}{N_{scen}(1-\beta)} \sum_{m=1}^{N_{scen}} [\gamma - c_m^T x]^+ \\ \text{subject to:} \quad & \quad \gamma \leq \sum_{i=1}^{N} c_{k,i} x_{k,i} \\ & \quad \sum_{i=1}^{N_T} x_{k,i} = N_V \\ & \quad x_{k,i} \in \{0,1\} \end{aligned} \tag{5}$$

Here x is the assignment vector, $x = [x_1, x_2, \ldots, x_{N_V}]^T$, c_m is the m^{th} realization of the target score vector, and $[g]^+ \equiv \max(g, 0)$. In our simulations we chose a value of $\beta = 0.01$, allowing for a 1% probability of exceeding our "loss function". The target scores were varied in same interval as before, $[c_{k,i} - \sigma_i, c_{k,i} + \sigma_i]$. We compared this to the modified Soyster ($Robust$ entry in the table) formulation using a value of $\mu = 3$. The numerical results can be seen in Table II.

Note that the CVaR approach depends crucially on the number of scenarios; for lower number of scenarios, the robust assignment generated by CVaR results in higher expected mission score, but also higher standard deviation. With 50 scenarios, the CVaR approach results in a higher mission score than the robust formulation, but also has a higher standard deviation. As the number of scenarios is increased to 100, the CVaR approach results match with the modified Soyster results; note that at 200 scenarios, a different assignment is generated, and the mission score is increased (as well as standard deviation). Beyond 500 scenarios, the two approaches generated the same assignments, and thus resulted in the same performance. In the next section, we extend the modified Soyster formulation to account for the coupling between the reconnaissance and strike vehicles.

IV. Modification for Cooperative Reconnaissance/Strike

As stated previously, future UAV missions will involve heterogeneous vehicles with coupled mission objectives. For example, the mission of reconnaissance vehicles is to reduce uncertainty in the environment and is coupled with the objective of the strike vehicles (namely, destroying targets in the presence of this uncertainty). First, we introduce the uncertainty and estimation models used in our work; we then use the robust formulation to pose and solve a mission with coupled reconnaissance and strike objectives.

A. Estimator model

For our estimator model, the target's state at time k is represented by its target type (i.e. its score). The output of a classification task is assumed to be a measurement of the target type, corrupted by some sensor noise ν_k

$$z_k = Hc_k + \nu_k \tag{6}$$

where c_k represents the true target state (assumed constant); ν_k represents the (assumed zero-mean, Gaussian distributed) sensor noise, with covariance $E[\nu_k^2] = R$. The estimator equations for the updated expected score and covariance that result from this model are [6]

$$\bar{c}_{k+1} = \bar{c}_k + L_{k+1}(z_{k+1} - \hat{z}_{k+1|k}) \tag{7}$$
$$P_{k+1}^{-1} = P_k^{-1} + HR^{-1}H^T \tag{8}$$

Here, $\bar{c}_k$ represents the estimate of the target score at time k; L_{k+1} represents an estimator gain on the innovations; the covariance $P_k = \sigma_k^2$; and $\hat{z}_{k+1|k} = H\bar{c}_k$. Note that here, $H = 1$ since we are directly observing the state of the target.

It is clear from Eq. 8 that the updated estimate relies on a new observation. However, this observation will *only* become available once the reconnaissance vehicle has actually visited the target. As such, at time k, our best estimate of the future observation (e.g. at time $k+1$) is

$$\tilde{z}_{k+1|k} = E\left[Hc_{k+1|k} + \nu_{k+1}\right] = H\bar{c}_k \tag{9}$$

We can use this expected observation in the estimator equations to update our predictions of the target classification.

$$\begin{aligned}\bar{c}_{k+1|k} &= \bar{c}_{k|k} + L_{k+1}(\tilde{z}_{k+1|k} - \hat{z}_{k+1|k}) \\ &= \bar{c}_k + L_{k+1}(H\bar{c}_k - H\bar{c}_k) = \bar{c}_k \qquad (10) \\ P_{k+1|k}^{-1} &= P_k^{-1} + HR^{-1}H^T \qquad (11)\end{aligned}$$

This update is the key component of the coupled reconnaissance/strike problem discussed in this paper. By rearranging Eq. 11 for the scalar case ($H = 1$), the modification to the uncertainty in the target classification as a result of assigning a future reconnaissance task can be rewritten as

$$\sigma_{k+1|k} = \sqrt{\frac{\sigma_k^2 R}{R + \sigma_k^2}} \tag{12}$$

or equivalently as the difference

$$\sigma_{k+1|k} - \sigma_k = \sigma_k \left\{ \sqrt{\frac{R}{R + \sigma_k^2}} - 1 \right\} \tag{13}$$

Note that in the limiting cases $R \rightarrow \infty$ (i.e., a very poor sensor), then $\sigma_{k+1|k} = \sigma_k$, and the uncertainty does not change. In the case of $R = 0$ (i.e., a perfect sensor) then $\sigma_{k+1|k} = \sigma_k = 0$ and the uncertainty in the target classification will be eliminated by the measurement. In summary, these equations present a means to analyze the expected reduction in the uncertainty of the target type by a future reconnaissance prior to visiting the target.

B. Preliminary reconnaissance/Strike formulation

Reconnaissance and strike vehicles have inherently different mission goals – the objective of the former is to reduce the uncertainty of the information about the environment, the objective of the latter is to recover the maximum score of the mission by destroying the most valuable targets. Thus, it would be desirable for a reconnaissance vehicle to be assigned to higher variance targets (equivalently, targets with higher standard deviations), while a strike vehicle would likely be assigned to targets exhibiting the best "worst-case" score. One could then derive an optimization criterion for these mission objectives as

$$\begin{aligned}&\max_{x,y} J_k = \sum_{i=1}^{N_T} (\bar{c}_{k,i} - \mu\sigma_{k,i})x_{k,i} + \mu\sigma_{k,i}y_{k,i} \\ &\text{subject to:} \quad \sum_{i=1}^{N_T} y_{k,i} = N_{VR}\ , \ \sum_{i=1}^{N_T} x_{k,i} = N_{VS} \qquad (14) \\ &\qquad x_{k,i}, y_{k,i} \in \{0,1\}\end{aligned}$$

Here, $x_{k,i}$ and $y_{k,i}$ represent the assignments for the strike and reconnaissance vehicles respectively, and the maximization is taken over these assignments. N_{VS} and N_{VR} represent the total number of strike and reconnaissance vehicles respectively. Note that this optimization can be solved separately for x and y, as there is no coupling in the objective function.

With this decoupled objective function, the resulting optimization is straightforward. However, this approach does not capture the cooperative behavior that is required between the two types of vehicles. For example, it would be beneficial for the reconnaissance vehicle to do more than just update the knowledge of the environment by visiting the most uncertain targets. Since the ultimate goal is to achieve the best possible mission score, we would like to modify the reconnaissance mission to account for the strike mission, and vice versa. We can achieve this by coupling the mission objectives and using the estimator results on the reduction of uncertainty due to reconnaissance.

We can solve for this cooperation by considering an objective function that couples the individual mission objectives. As mentioned previously, the target's score will remain the same if a reconnaissance vehicle is assigned to it (since an observation has not yet arrived to update its score), but its uncertainty (given by σ) will decrease from σ_k to $\sigma_{k+1|k}$. We can use this reduction in the uncertainty into the assignment problem for the strike vehicle. The result would exhibit *truly cooperative behavior* in the sense that the reconnaissance vehicle will be assigned to observe the target whose reduction in uncertainty will prove most beneficial for the strike vehicles, thereby creating this coupled behavior between the vehicle missions. The optimization for the coupled mission can be written as

$$\max_{x,y} J_k = \sum_{i=1}^{N_T} \left(\bar{c}_{k,i} - \mu\sigma_{k,i}(1-y_{k,i}) - \mu\sigma_{k+1|k,i}\ y_{k,i}\right) x_{k,i}$$
$$\text{subject to: } \sum_{i=1}^{N_T} y_{k,i} = N_{VR}\ ,\ \sum_{i=1}^{N_T} x_{k,i} = N_{VS} \quad (15)$$
$$x_{k,i}, y_{k,i} \in \{0,1\}$$

This objective function implies that if a target is assigned to be visited by a reconnaissance vehicle, then $y_{k,i} = 1$, and thus the uncertainty in target score i decreases from $\sigma_{k,i}$ to $\sigma_{k+1|k,i}$. Similarly, if a reconnaissance vehicle is not assigned to target i, the uncertainty does not change. Note that by coupling the assignment, if both a strike and reconnaissance vehicle are assigned to target i, the strike vehicle recovers an improved score.

We can simplify the objective function by combining similar terms to give

$$\max_{x,y} J_k = \sum_{i=1}^{N_T} (\bar{c}_{k,i} - \mu\sigma_{k,i})x_{k,i} + \mu(\sigma_{k,i} - \sigma_{k+1|k,i})x_{k,i}y_{k,i}$$

Note that this is a nonlinear objective function that cannot be solved as a Mixed-Integer Linear Program (MILP), but we can define $v_{k,i} \equiv x_{k,i}y_{k,i}$ as an additional optimization variable, and constrain it as follows

$$\begin{aligned} v_{k,i} &\leq x_{k,i} \\ v_{k,i} &\leq y_{k,i} \\ v_{k,i} &\geq x_{k,i} + y_{k,i} - 1 \\ v_{k,i} &\in \{0,1\} \end{aligned} \quad (16)$$

This change of variables enables the problem to be posed and solved as a MILP of the form

Algorithm #1

$$\max_{x,y} J_k = \sum_{i=1}^{N_T} (\bar{c}_{k,i} - \mu\sigma_{k,i})x_{k,i} + \mu(\sigma_{k,i} - \sigma_{k+1|k,i})v_{k,i}$$
$$\text{subject to: } \sum_{i=1}^{N_T} y_{k,i} = N_{VR}\ ,\ \sum_{i=1}^{N_T} x_{k,i} = N_{VS} \quad (17)$$
$$x_{k,i}, y_{k,i}, v_{k,i} \in \{0,1\}$$

TABLE III

TARGET PARAMETERS

Target	$\bar{c}$	σ_k	σ_{k+1}
1	20	4	0.3152
2	22	7	0.3159

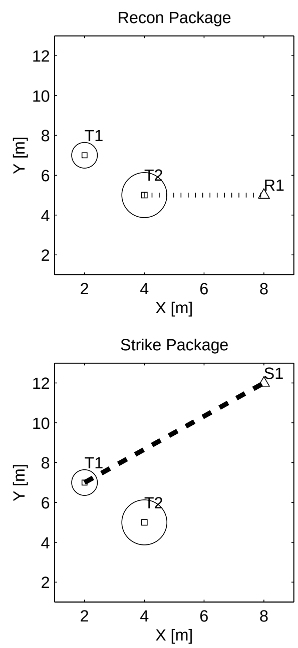

Fig. 3. Decoupled mission

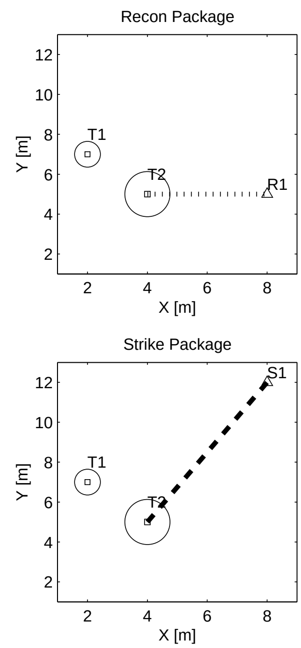

Fig. 4. Coupled mission

$$\begin{aligned} v_{k,i} &\leq x_{k,i} \\ v_{k,i} &\leq y_{k,i} \\ v_{k,i} &\geq x_{k,i} + y_{k,i} - 1 \end{aligned} \quad (18)$$

The key point with this formulation is that it captures the coupling in the cooperative heterogeneous mission by assigning the reconnaissance and strike vehicles *together*, taking into account the individual missions.

As a straightforward example, we consider a 2 target case with one strike and reconnaissance vehicle to be assigned (Figure 3). This problem is simple enough to visualize and be used as a demonstration of the effectiveness of this approach. The reconnaissance (R_1) and strike (S_1) vehicles are represented by $\star$ and Δ, respectively, and the i^{th} target, T_i, is represented by $\square$. The expected score of each target is proportional to the size of the box, and the uncertainty in the target score is proportional to the radius of the surrounding circle . The target parameters for this experiment are given in Table III (μ=1).

Figures 3 and 4 compare the assignments of the reconnaissance and strike vehicle for the decoupled and coupled cases. In the decoupled case, strike vehicle S_1 is assigned to T_1, while reconnaissance vehicle R_1 is assigned to T_2. Here the optimization is completely decoupled in that the strike vehicle and reconnaissance vehicle assignments are found independently. In the coupled case, both strike vehicle S_1

and reconnaissance vehicle R_1 are assigned to T_2. We can see that without reconnaissance to T_1, the expected worst case score is higher in T_1; however, with reconnaissance to that target, uncertainty is reduced for both targets, and T_1 then has a higher expected worst score. Note that with the two formulations, the strike vehicles are assigned to *different* targets. This serves to demonstrate that solving the optimization in Eq. 15 does not result in the same assignment as the coupled formulation. This is key: if we were able to solve the decoupled formulation for the strike vehicle assignments, we could then simply assign reconnaissance vehicles to those targets and obtain our reconnaissance/strike mission. As these results show, that is not the case.

To demonstrate these results numerically, we conducted a straightforward two-stage mission analysis. In the first stage, the above two optimizations were solved with the target parameters; after this first stage, the vehicles progressed toward their intended targets. At the second stage, it was assumed that the reconnaissance vehicle had actually *reached* the target to which it was assigned, and thus, there was *no* uncertainty in the target score. The optimization in Eq. 6 was then solved for the strike vehicle, with the updated target scores (from the reconnaissance vehicle's observation) and standard deviations. Note that this target score could have actually been *worse* than predicted, as the observation was made only at time of the reconnaissance UAV arrival; the target that was not visited by the reconnaissance vehicle maintained its original expected score and uncertainty. In order to compare the two approaches, we tabulated the scores accrued by the strike vehicles at the second stage and discounted them by their current distance to the (possibly new) target to visit. Both vehicles incurred this score penalty, but since the targets were en route to their previously intended targets, a re-assignment to a different target incurred a *greater* score penalty, and hence reduction in score.

Of interest in this experiment is the time delay between the assignment of the reconnaissance vehicle to a target, and its observation of that target. Clearly, if a reconnaissance vehicle had a high enough speed such that it could update the "true" state (*i.e.*, score) of the target almost immediately, then the effects of a coupled reconnaissance and strike vehicle would likely be identical to those obtained in a decoupled mission, since the strike vehicles would be immediately reassigned. This time delay however is present in these typical reconnaissance/strike missions; our time discount "penalty" for a change in reassignment does reflect that a reassignment as a result of improved information will result in a lower accrued score for the mission.

The numerical results of 1000 simulations are given in Table IV, where $\bar{J}$ indicates the average mission score of each approach, and σ_J indicates the standard deviation of this score. Note that the score accrued by the coupled approach has a much improved performance over the decoupled approach. Furthermore, we note that the variation of this mean performance is almost equivalent for the two approaches (though we note that this is troubling for the decoupled approach due to its lower mean). From this simple example, we can thus see the importance of coupling the missions of the two types of vehicles.

TABLE IV
NUMERICAL COMPARISONS OF DECOUPLED AND COUPLED RECONNAISSANCE/STRIKE

Reconnaissance/Strike	$\bar{J}$	σ_J
Coupled	61.19	26.56
Decoupled	41.50	23.12

C. Improved reconnaissance/Strike formulation

While the above example shows that the coupled approach performs better than a decoupled one, using Eq. 19 for more complex missions can result in an incomplete use of resources if there are more reconnaissance vehicles than strike vehicles, or if we seek to also reward reconnaissance as a mission objective in its own right. The cost function mainly rewards the strike vehicles, by improving their score if a reconnaissance vehicle is assigned to that target. However, it does not fully capture the reward for the reconnaissance vehicles that are, for example, not assigned to strike vehicle targets. With the previous algorithm, these unassigned vehicles could be assigned anywhere, but we would like them to explore the remaining targets based on a certain criteria. Such a criterion could be to assign them to the targets with the highest standard deviation, or to targets that exhibit the "best-case" score ($c_{k,i}+\sigma_{k,i}$) so as to incorporate the notion of cost in the optimization. Either of these options can be included by adding a an extra term to the cost function

Algorithm #2

$$\max_{x,y} J_k = \sum_{i=1}^{N_T} (\bar{c}_{k,i} - \mu\sigma_{k,i})x_{k,i} + \mu(\sigma_{k,i} - \sigma_{k+1,i})v_{k,i} + K\sigma_{k,i}(1 - x_{k,i})y_{k,i} \quad (19)$$

For small K this cost function keeps the strike objective as the principal objective of the mission, while the weighting on the latter part of the cost function assigns the remaining reconnaissance vehicles to highly uncertain targets. Again, this is a behavior that is intuitive to capture. Since the coupling between reconnaissance vehicles and strike vehicles is captured in the first part of the cost function, it seems appropriate to assign the remaining reconnaissance vehicles to targets that have the highest uncertainty. The term $(1 - x_{k,i})y_{k,i}$ captures the fact that these extra reconnaissance vehicles will be assigned to targets that have not been assigned (recall when the targets are unassigned, $x_{k,i} = 0$). Note that this approach is quite general, since the $K\sigma_{k,i}$ term can be replaced by any expression that captures an alternative objective function for the reconnaissance vehicle.

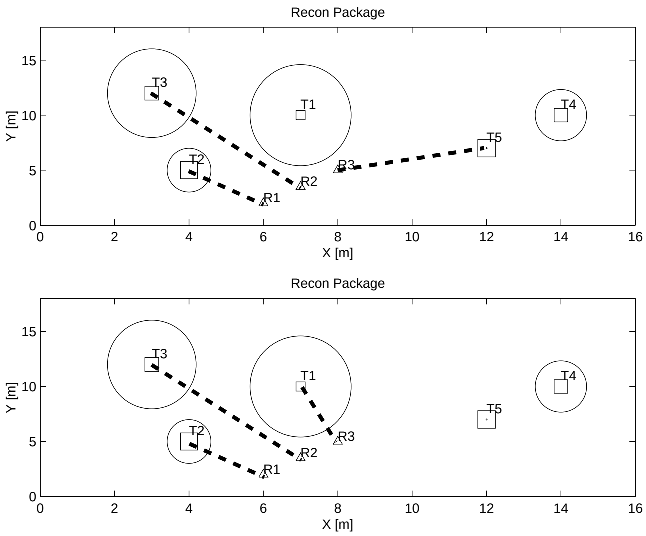

Fig. 5. Comparison of Algorithm 1 (top) and Algorithm 2 (bottom) formulations

We demonstrate this change in the objective function in Figure 5. In this example, we consider the assignment of 3 reconnaissance and 2 strike vehicles (strike assignments remained identical in both cases), and $K = .01$. In the earlier formulation, R_3 is assigned to T_5, a target with virtually no uncertainty (note that the target score is virtually certain since it has such a low uncertainty), since in this instance there was no reward for decreasing the uncertainty in the environment. The extra reconnaissance vehicle was assigned randomly, as assignment to any target did not improve the cost function. Note that there is benefit in the extra reconnaissance vehicle going to T_3 instead of T_5 since it will inherently decrease the uncertainty in the environment, and in fact this is what happens in the modified formulation. Thus, the modified formulation captures more intuitive results via a better allocation of resources.

V. Hardware testbed implementation

We are currently developing a heterogeneous testbed that consists of blimps (taking the role of reconnaissance vehicles) and autonomous trucks (taking the role of strike vehicles). The blimps are 7-ft spheres and carry a Sony VAIO to do onboard control. Currently, assignment- and waypoint-generation are done off-board. Position and heading information are obtained with the use of a ArcSecond 3Di Constellation system, which includes four transmitters and a receiver (one receiver per vehicle). We have a four transmitter setup that uniquely determines a set of inertial reference axes for the environment in which we do our experiments. Each vehicle is equipped with its own sensor to obtain unique position solution accurate to within millimeters; velocity is estimated via a Kalman filter based on truck and blimp dynamics, and is very accurate for navigational purposes. Heading information for the blimps is obtained with a second onboard receiver, whose position is numerically differenced from that of the primary sensor to determine a relative position vector with respect to the transmitter axes. The intent is to transition these newly developed algorithms on these testbeds, thereby demonstrating the applicability of these algorithms in the presence of real-world disturbances. Furthermore, we can begin demonstrating cooperative assignments based on unique vehicle capabilities (i.e., the ability of blimps to overfly obstacles). These capabilities can be captured in more complex formulations of the coupled robust approach.

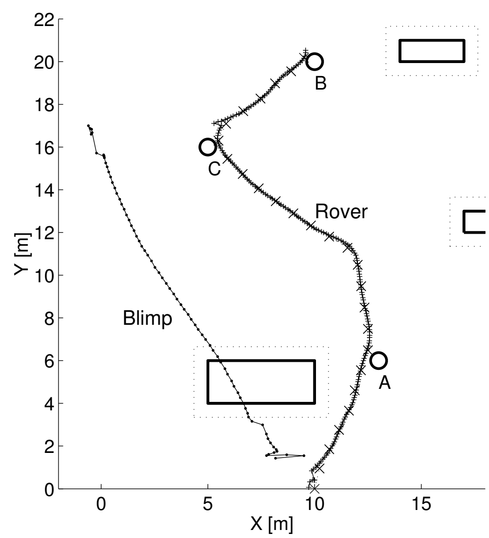

Fig. 6. Truck and blimp scenario

Figure 6 shows a mission with one of the blimps and a truck – here the blimp simulated doing reconnaissance for the benefit of the truck. Both vehicles operated autonomously in closed loop control. The truck began its mission unaware of the existence of target C, while the blimp was sent out to explore the environment. When it discovers target C, this target was added to the mission of the truck, and it can be seen that the truck corrects its trajectory to visit target C. While this mission did not incorporate the aforementioned algorithms, these will be implemented on our testbeds soon.

VI. Conclusions

This paper has presented a novel approach to the problem of mission planning for a team of heterogeneous vehicles with uncertainty in the environment. We have presented a simple modification of a robustness approach that allows for a direct tuning of the level of robustness in the solution. This robust formulation was then extended to account for the coupling between the reconnaissance (tasks that reduce uncertainty) and strike (tasks that directly increase the score) parts of the combined mission. Although nonlinear, we show that this coupled problem can be solved as a single

MILP. Future work will investigate the use of time discounting explicitly in the cost function, thereby incorporating the notion of distance in the assignment, as well as different vehicle capabilities and performance (speed). We are also investigating alternative representations of the uncertainty in the information of the environment.

ACKNOWLEDGMENTS

The research was funded by AFOSR Grant # F49620-01-1-0453 and funding for the testbed was provided by the DURIP Grant # F49620-02-1-0216.

REFERENCES

[1] J. S. Bellingham, M. J. Tillerson, A. G. Richards, J. P. How, "Multi-Task Assignment and Path Planning for Cooperating UAVs," *Conference on Cooperative Control and Optimization*, Nov. 2001.

[2] A. Ben-Tal and A. Nemirovski, "Robust Solutions of Uncertain Linear Programs" *Operations Research Letters*, Vol. 25, pp. 1-13.

[3] D. Bertsimas and M. Sim, "Robust Discrete Optimization and Network Flows," submitted to *Operations Research Letters*, 2002.

[4] D. Bertsimas and M. Sim, "Price of Robustness," submitted to *Mathematical Programming*, 2002.

[5] J. R. Birge and F. Louveaux, "Introduction to Stochastic Programming", *Springer-Verlag*, 1997.

[6] A. Gelb, "Applied Optimal Estimation", *MIT Press*, 1974.

[7] P. Kouvelis and G. Yu, "Robust Discrete Optimization and Its Applications," *Kluwer Academic Publishers*, 1996.

[8] P. Krokhmal, Murphey, R., Pardalos, P., Uryasev, S., and G. Zrazhevsky (2003) "Robust Decision Making: Addressing Uncertainties in Distributions", In: S. Butenko et al. (Eds.) "Cooperative Control: Models, Applications and Algorithms", Kluwer Academic Publishers, 165-185.

[9] P. Krokhmal, J. Palmquist, and S. Uryasev. "Portfolio Optimization with Conditional Value-At-Risk Objective and Constraints." *The Journal of Risk*, Vol. 4, No. 2, 2002.

[10] W. M. McEneaney and B. Fitzpatrick, "Control for UAV Operations Under Imperfect Information," *AIAA*, 2002-3418.

[11] A. L. Soyster, "Convex Programming with Set-Inclusive Constraints and Applications to Inexact Linear Programming," *Operations Research*, 1973, pp. 1154–1157.