Use of sUASs for Infrastructure Monitoring in Communication-Constrained Environments

René Andrés García Franceschini (ragarcia@mit.edu)
Saurabh Amin (amins@mit.edu)

***Abstract*—The United States Federal Government considers the timely detection and repair of public infrastructure a top priority of national and financial security. Because of their ease of maneuverability and low personnel risk, small Unmanned Aerial Systems (sUASs) are an increasingly popular choice for assessing infrastructure. As these systems become less reliant on human operators and more reliant on GPS and 4G LTE signal, it is becoming imperative to know how to route these drones in communications-constrained environments, such as rural areas. This paper poses the problem of routing of agents in networks with uncertain linkages, in the context of infrastructure inspections by sUASs. Our solution for this problem combines heuristic algorithms for the Team Orienteering Problem with kriging to simultaneously route drones and make inferences on the layout of the map. This approach results in solutions that are comparable in score to a non-kriging approach, but significantly better in terms of computational time and the number of drones that make it back to base.**

I. Introduction

The proper maintenance and care for the United States public infrastructure has been a growing concern for the federal government and for the Departments of Energy and Transportation. Public infrastructure, such as pipes, bridges, electrical grids, etc., can be damaged through corrosion, operator carelessness, intentional theft or attacks (either physical or cyber-attacks), and natural disasters such as earthquakes and hurricanes, among others. In the case of pipelines, the recent Protecting our Infrastructure of Pipelines and Enhancing Safety (PIPES) Act of 2016 is evidence of the increasing attention that Congress is putting towards this issue [1].

A fair amount of critical infrastructure (pipelines, power grids, etc.) is in remote, isolated areas that have wildly inconsistent GPS or communications coverage, which hampers coordination between multiple agents. *Figure 1* shows the 4G LTE coverage in the city of Houston: note the decreased coverage in areas more distant to the city center. Even regions that are close to the urban caucus might suffer from poor communications coverage during disaster conditions. The aforementioned city of Houston, cities along the southern coast of Florida and the entirety of Puerto Rico have had their communications infrastructure down in certain areas for extended periods of time. Lack of proper planning for some of these extreme scenarios has seriously hampered relief efforts, leading to extraordinary financial, environmental, and humanitarian losses [2].

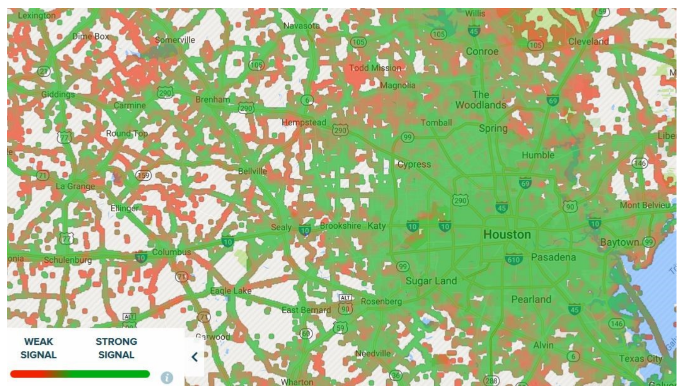

Fig. 1. 4G LTE signal strength in and around Houston. Note the variability in signal strength, especially in rural areas [3].

These narratives provide the motivation for our work in dynamic routing of small Unmanned Aerial Systems (sUASs) under coverage constraints. sUASs are readily and commercially available, and can reach inaccessible areas at a fraction of the time compared to ground-based survey teams. sUASs can also be equipped with highly specialized sensors that can more rapidly detect leaks. Their growing popularity has prompted the Federal Aviation Agency (FAA) to more readily allow users to use sUASs in their respective fields of work [4]. Perhaps most crucially, although many sUASs are currently controlled through RC frequencies, companies such as Qualcomm are currently developing drones that can be controlled using the more powerful 4G LTE infrastructure. This makes it likely that LTE-controlled sUASs could replace our current RC ones in the coming years, underscoring the importance of being able to effectively route them through communications coverage constraints [5].

With this project, we consider the following problem statement: In the face of restricted or denied communications coverage, how can we improve the routing of sUASs through a network of nodes that need to be inspected? This problem can be decomposed into two well-known ones. First is the Team Orienteering Problem, where a group of agents (in this case, sUASs) traverse a network obtaining bounties (a measure of how important inspecting a particular piece of infrastructure is) along the way, given a time constraint (which is indicative of battery life). Second is the prediction of spatially correlated quantities, where you have to estimate the value of an attribute of a certain node given the observed quantities of the nodes you have observed (which, in our case, represents the communications signal strength at each node, given that they are spatially correlated). This project proposes an adaptive routing policy that traces routes for sUASs using all of the known information about the communications space, and continually makes predictions about the communications signal strength to improve these routes.

II. Related Work

Literature for the orienteering problem (and its many particular cases) is abundant. Some of the first to work on routing of sUASs were Sundar et al, who proposed a multi-vehicle, fuel-constrained routing problem for inspecting infrastructure. Their approach used a mixed-integer program (MIP) formulation that was able to solve the routing problem to optimality. Their work was limited by the fact that their model was completely static and did not account for uncertainty in any way, as a justification for the longer solving time of the MIP [6]. Sundar et al also presented the problem of routing sUASs in areas that had poor GPS connectivity. Their work built up on previous localization literature that determined that sUASs needed to be within line of sight of two landmarks to achieve localization [7]. This work is immediately connected to ours, as it considers routing in areas where the communications space is not known a priori. However, practical limitations include the need to set these landmarks prior to inspection, which may not be a luxury that is available in an emergency situation.

Given the speed needed to adjust routes on the fly, heuristics are much more appealing than exact MIP solutions. Souffriau et al combined previous work on vehicle routing using a Greedy Randomized Adaptive Search Procedure (GRASP) approach with a Path Relinking step to arrive at good solutions to the team orienteering problem. In short, GRASP constructs several independent routes, improves upon each of them using a local search procedure, and returns the best one. Path Relinking, on the other hand, takes two solutions that have already been made and attempts to connect them in order to improve the total score. The authors implemented a slow and fast version of the algorithm, which simply dictates how many iterations of improvements the algorithm will go through. The performance was impressive: the fast version could solve in 5 seconds instances that competitive algorithms took upwards of 3 minutes to solve, with just a .39% average gap from the optimal solution [8]. However, as before the model assumes a known configuration of the network and does not account for uncertainty of any kind. Our problem improves upon this by adding uncertainty and adaptability into the setup.

Compared to the orienteering problem, the prediction of spatially correlated quantities has much more literature surrounding it. Most of the literature concerning spatial interpolation of communications coverage uses some variant of a kriging, a well-studied technique for spatial interpolation that uses a Gaussian process with prior covariances. Krishnakumar et al used a multivariate Gaussian model depict received signal strength in order to characterize the uncertainty in signal strength-based location estimation techniques [9]. Fink et al used a Gaussian process model that allowed for prediction of radio signals using mobile robots, and introduced a policy for collecting samples of signal strength [10]. Wagle and Frew also used spatio-temporal Gaussian processes to characterize the signal strength of radio frequencies in airborne environments [11]. Molinari et al saw that ordinary kriging with a squared exponential kernel is the most effective technique at predicting cellular coverage from crowdsourced data [12]. Although these and many others have done estimation of spatially correlated quantities, none of their work incorporates any sort of routing element, which is incorporated in our project.

III. Technical Approach

A. Model

To tackle this problem, we need an apt model that describes where the sUAS can travel, what is the reward associated with visiting a certain node, and how poor coverage prevents the sUAS from visiting a certain node. The model we decided to employ is illustrated in *Figure 2*. Here, a real map is discretized into a 2D grid of nodes. The orange node in the middle represents where the operator is located, and where all the sUASs commence their operation. Black nodes represent those that the sUAS can travel to, but do not have an associated reward. Blue nodes are those that the sUAS can travel to and have some reward associated with it. In practice, some infrastructure is more important than others. This is reflected in the value of the reward associated with the node. Finally, red nodes are those that have poor connectivity, and thus cannot be visited. Therefore, any path that goes through them must be rerouted.

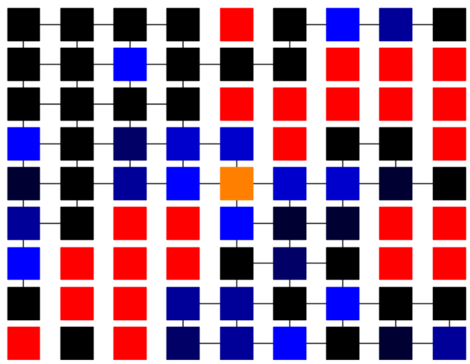

Fig. 2. Visual representation of the model. Red nodes have poor coverage, black nodes have good coverage and no reward, blue nodes have good coverage and some reward, and the orange node is the start node.

There are key features that make this model appropriate for this context. First, any real field can be discretized in this form. Second, the equal weights of the edges mean that at each timestep, an sUAS will be on a node, never in transit between two nodes. This makes the rerouting process much easier. Third, the regularity allows us to assign connection coverage values that spatially correlated, as real coverage maps are. Finally, most coverage maps are measured in a similar fashion (crowdsourced data points of signal strength), meaning that it would be easy to compare our models with real life measurements.

For this problem, we assume that the coverage map is time invariant; that is, a node that is accessible at a certain time is accessible at any other time. Thus, the map of the field is realized at the beginning of the mission. Although in reality the coverage map does change with time, drone battery lives are around 20 minutes at the very high-end, making the approximation appropriate. Crucially, *the operator does not know the layout of the map at the beginning of the mission*. They only know the bounty associated with each node (in a realistic scenario, they would know how important it is to inspect a certain piece of infrastructure), the connectivity of the map assuming all nodes have good coverage, and a prior probability distribution of the connection strength at each node. The operator can only know with absolute certainty if a node has poor coverage by visiting a node next to it.

B. Kriging

The kriging step of our algorithm was modelled after the work by Molinari et al [12]. We used a squared exponential kernel:

$$k_{SE}(i,j) = e^{\frac{||\vec{p}_i - \vec{p}_j||_2}{2}} \tag{1}$$

where i and j are two arbitrary nodes on the map, and $\vec{p}_i$ and $\vec{p}_j$ are their corresponding position vectors. From a set of nodes that have been observed V, the expected value of the coverage at an unknown node u is given by:

$$\hat{y}_u = \vec{k}_{SE}(u,V)^T * K(V)^{-1} * \vec{y}_V \tag{2}$$

where $\vec{k}_{SE}(u,V)$ is the vector of squared exponential kernels between u and every node in V, $K(V)$ is the Gram matrix of the nodes in V, and $\vec{y}_V$ is the vector of observed signal strengths at nodes in V.

C. Algorithm

A mission is be carried out in the following way:

1) At the start of the mission, we use a heuristic to determine a good solution to the Team Orienteering Problem presented *assuming that all nodes have good coverage*. Although any good heuristic would work here, we used the Team Orienteering algorithm designed by Souffriau et al for its flexibility in balancing speed and optimality gap [8]. For this initial stage, we used 10 iterations.
2) The sUASs then commence their exploration of the graph, collecting the rewards of the nodes they visit and updating the map based on the coverage of the nodes they visit. The maps are updated in either one of two ways:
   - Only updating based on what the sUASs observe directly, which is just the nodes neighboring their position (no predictive algorithms).
   - Kriging as described above based on the observed coverages.
3) At every time step, the exploration will stop and the drones are rerouted using the same algorithm with 5 iterations.
4) Once the time limit for the mission has been reached, the simulation ends. We then keep track of three key metrics:
   - The total score of the mission (the sum of the bounties of all nodes visited by sUASs *that made it back to the start node at the end of the simulation*).
   - The computational time of the simulation
   - The *downed sUASs*, or the number of sUASs that were unable to make it to the start node by the end of the simulation.

D. Simulation

100 randomly generated maps were tested for each combination of the parameters listed below. The goal was to observe how performance varied as the different parameters were changed. These parameters are:

1) Size of the map: *9x9* and *11x11* grids were considered. These sizes were chosen for their relative ease of computation and because it was sufficiently large to observe differences in signal strength.
2) Fuel amount: 15 and 25 timestep horizons were considered. Timesteps were chosen for their relative ease of computation.
3) Minimum bounty for nodes: for one set of tests, the minimum bounty was 0. For the other, it was 0.5. We wanted to see whether having an arbitrary small minimum bounty would encourage further exploration of the map.

Thus, eight total configurations were run. 100 maps were randomly generated for each configuration, and each was run with the kriging and non-kriging update procedure. The following parameters were kept fixed among all tests:

1) Probabilistic instantiation: the field signal strength values were generated by sampling a multivariate normal distribution with mean 0, and covariance matrix derived from the squared exponential kernel function.
2) Cutoff for poor connection: nodes whose randomly generated sampling of the multivariate normal field had a *cdf* less than .3 were considered as having poor connection. This value was selected arbitrarily.
3) Bounty probability: each node had a .5 probability of holding a bounty.
4) Bounty amounts: of the nodes that held a bounty, the bounty amount was a random uniformly distributed integer from 1 to 5.

The results of these simulations and the comparison between kriging and non-kriging method for the three metrics mentioned above are presented and discussed in the following section.

IV. Results

The full results for all 1600 simulations are shown in *Figures 3, 4* and *5*, and summarized in *Table I*. In each of the subsections below, the results in terms of the three metrics described above are discussed in more detail. For the remainder of this section, a particular combination of experimental parameters will be referred to as a scenario.

A. Score

Figure 3 shows the difference in score for the same map between the kriging and the non-kriging methods for all eight scenarios. Because the maps were all randomly generated in terms of signal strength and bounty, it is unlikely that two maps have the same amount of bounty. Even if they do, the differences in which nodes the sUASs can access make some maps' bounties more accessible than others. Thus, we saw it appropriate to examine the increase or decrease in bounty within the same map between the kriging and non-kriging methods, as opposed to the absolute magnitude of the total bounty.

From our simulations, it is difficult to attribute increases or decreases in score to anything other than the fact that we used a stochastic heuristic. All eight scenarios show score differences with means close to 0, and standard errors that are close to, and oftentimes larger than, the mean. Since the motivation for using the .5 minimum bounty was to encourage greater exploration in the kriging method, it is a bit disappointing to see that ultimately the effect is either negligible or non-existent. Ultimately, it seems that in terms of score, using kriging is about as effective as avoiding it.

B. Computational Time

Figure 4 shows the difference in computational time for the same map between the kriging and non-kriging methods for the eight different scenarios. The computational time of the algorithm we chose is roughly a function of three different factors: the number of iterations the algorithm went through (which was kept constant between the two methods), the number of nodes that are available to traverse, and the number of nodes that have some sort of bounty. The latter two factors depend on the map that is generated, again making it difficult to compare across maps. Thus, we display our results in terms of the increase or decrease in computational time per map between the kriging and non-kriging methods.

Here we start to see major differences in the two methods. In all eight scenarios, we observe a statistically significant decrease in the computational time at the 5% level, and in all but one scenario it is significant at the 1% level. For the scenarios with a fuel amount of 25 and a grid of *11x11*, many of the simulations show decreases of more than two minutes.

We believe that this decrease has an intuitive explanation. By using kriging, the sUASs can disregard nodes that are assumed to have poor connectivity, thus having less nodes to consider in the routing algorithm. This is especially pronounced when the minimum bounty is 0.5.

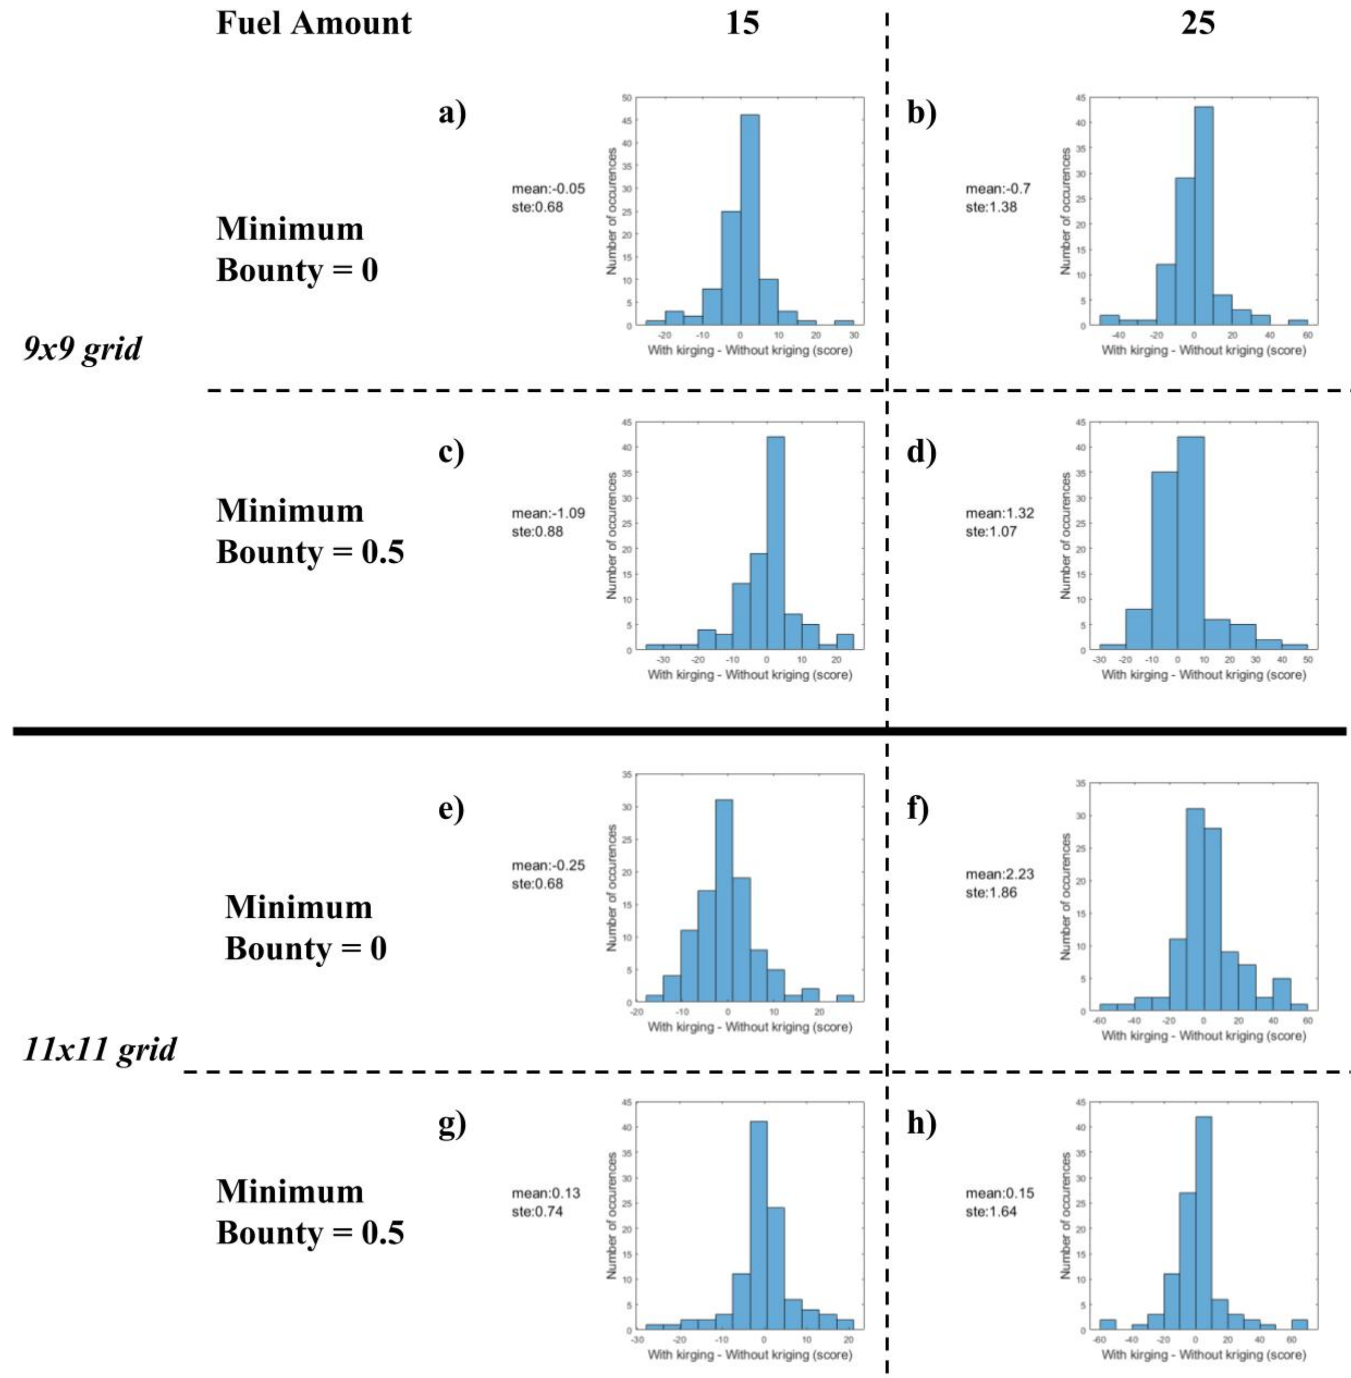

Fig. 3. Histogram of scores for all combinations of parameters. The horizontal axis represents the difference between the total score at the end of the simulation with kriging and without kriging for the same (randomly generated) map.

C. Downed sUASs

Figure 5 shows the number of instances where one and two sUASs were unable to make it back to the start node at the end of the simulation. A great majority of simulations had 0 sUASs that were downed; these were excluded from the histograms. To obtain the number of instances where 0 sUASs were downed, simply subtract the total number of instances that are depicted in the histogram from 100, per method.

By this metric, the kriging method is also an improvement over the non-kriging method. We can see that across all scenarios, there are fewer instances where at least one drone is downed. We also see that there are no instances where two sUASs are downed for the kriging

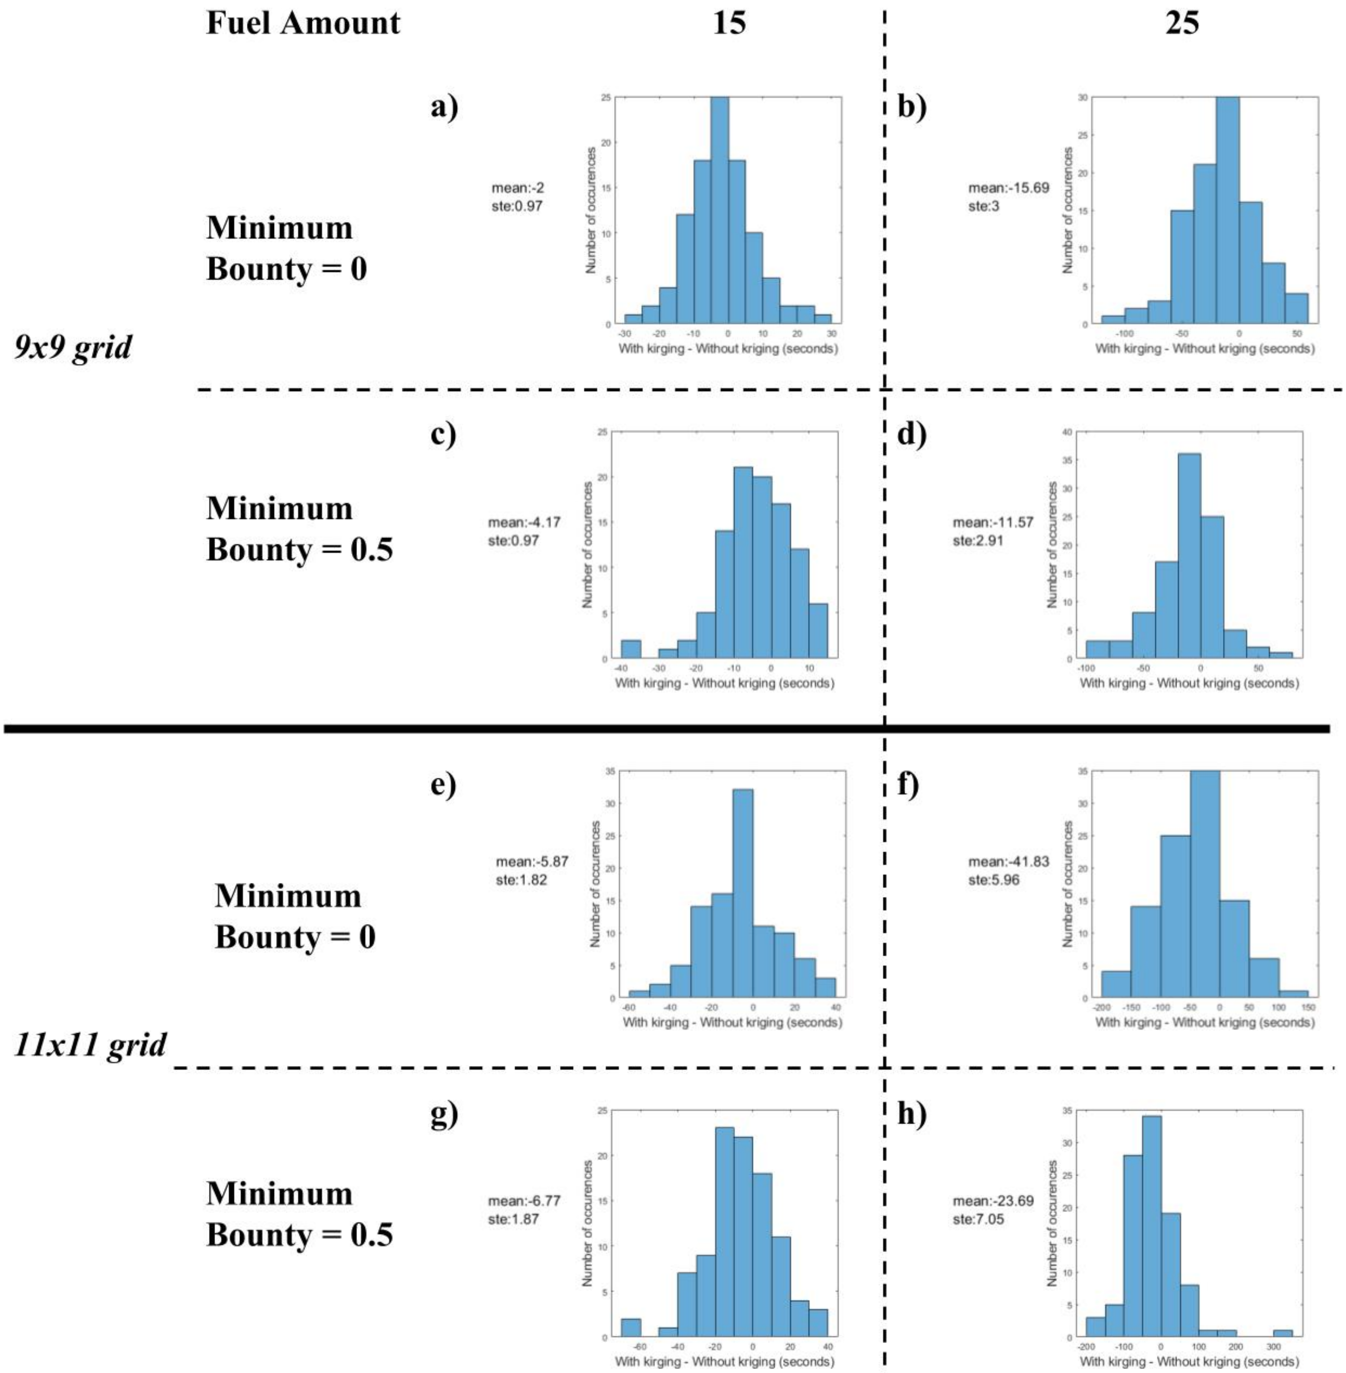

Fig. 4. Histogram of computational times for all combinations of parameters. The horizontal axis represents the difference between the computational time in seconds with kriging and without kriging for the same (randomly generated) map.

method, as opposed to three instances for the non-kriging method. There is one more instance of a downed sUAS for the kriging method in the scenario (h), but it should be seen in light of the fact that the non-kriging method had two instances of two drones being downed.

We believe that using kriging minimizes risk of not making it back by making sUASs aware of potential roadblocks on the way back to the start node. Thus, they are less likely to find themselves in a scenario where they do not have enough fuel to make it back.

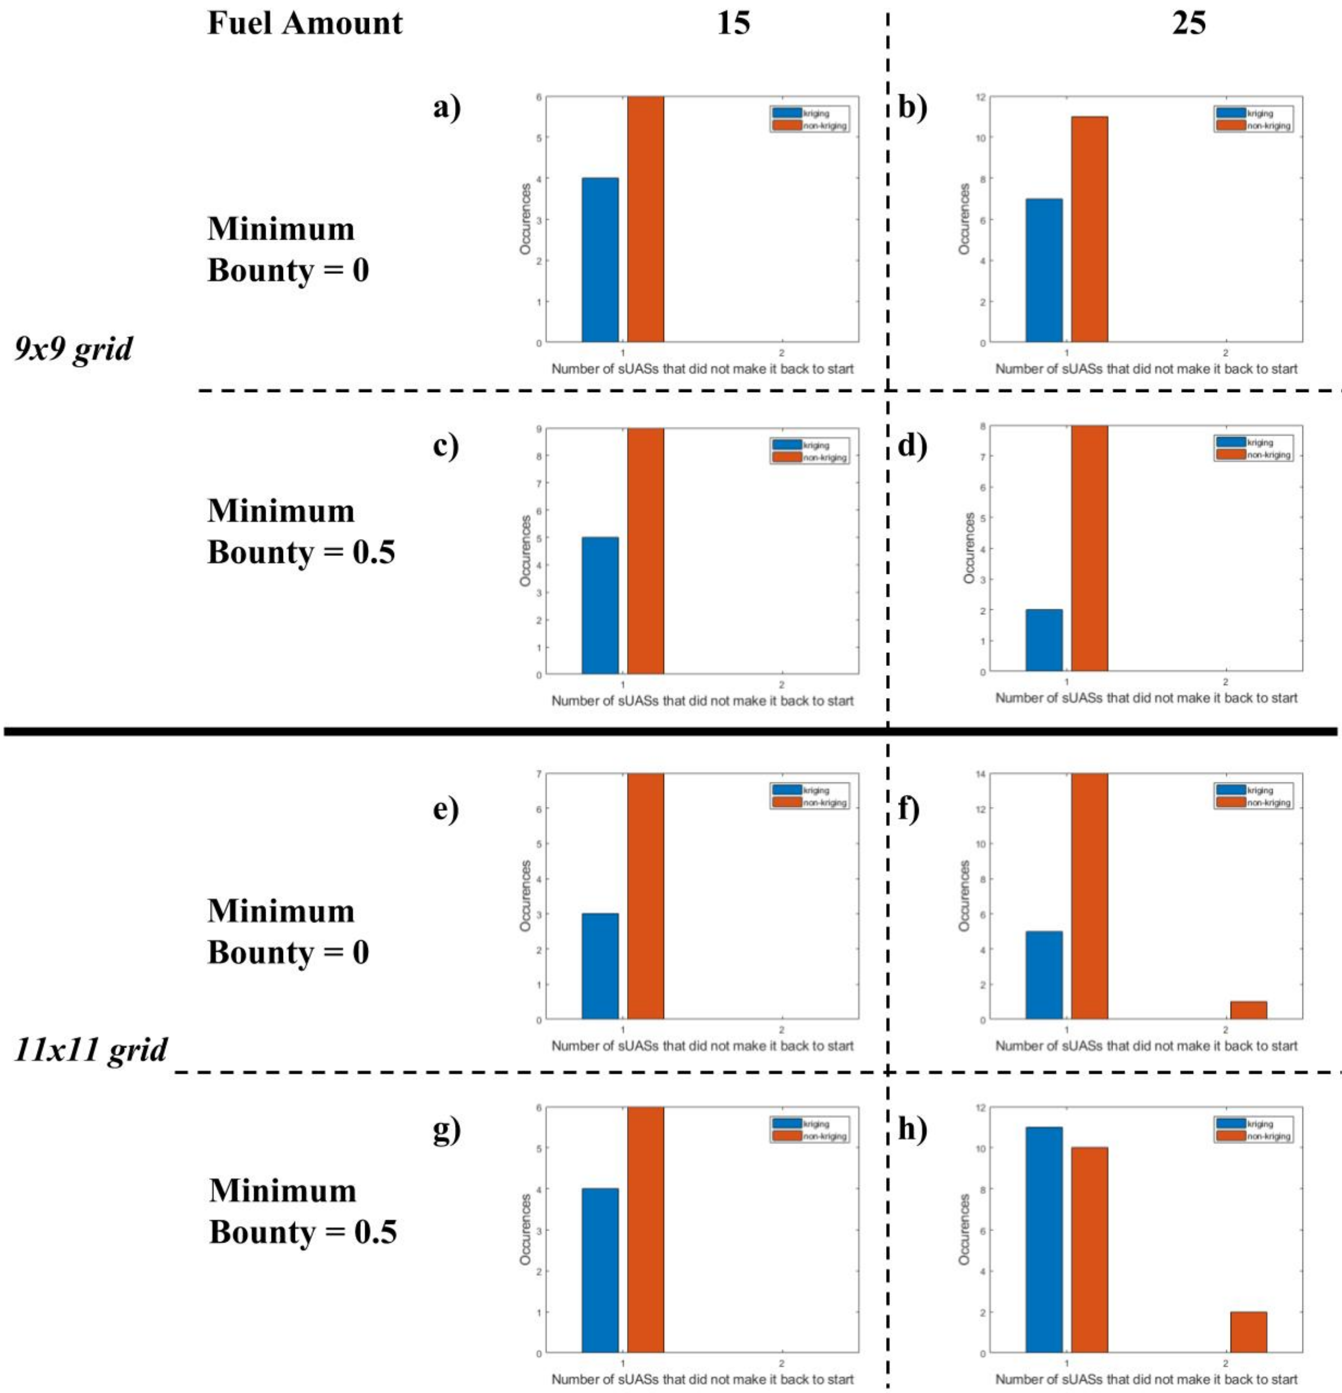

Fig. 5. Histogram of the number of instances where sUASs did not make it back to the start node at the end of the simulation, separated by how many sUASs did not make it back in the same mission, with kriging and without kriging.

V. Conclusion

In this paper, a new Team Orienteering problem with uncertain network linkages was presented, in the context of routing of sUASs through communications constrained environments. This real-world context was modelled based on previous research in communication fields. To solve this problem, an algorithm that combines previous Orienteering work with kriging was employed and discussed. This algorithm was tested against one without kriging in a variety of configurations, and it outperformed in terms of computational time and probability of a downed sUAS, while maintaining comparable results in terms of score.

Future work falls roughly into three categories.

SCENARIO	a	b	c	d	e	f	g	h
SCORE DIFF.	-0.05	-0.7	-1.09	1.32	-0.25	2.23	0.13	0.15
	(0.68)	(1.38)	(0.58)	(1.07)	(0.68)	(1.86)	(0.74)	(1.64)
COMP. TIME DIFF.	-2.00*	-15.69**	-4.17**	-11.57**	-5.87**	-41.83**	-6.77**	-23.69**
	(0.97)	(3.00)	(0.97)	(2.91)	(1.82)	(5.96)	(1.87)	(7.05)

TABLE I

SUMMARY OF RESULTS. SHOWN ARE THE AVERAGE DIFFERENCES IN SCORES AND COMPUTATIONAL TIMES FOR NON-KRIGING VS. KRIGING APPROACHES. * SIGNIFICANT AT 5%; ** SIGNIFICANT AT 1%

First, although computational time is lowered, it is still not good enough for an actual drone to reroute on the fly. Thus, non-adaptive routing policies that incorporate kriging will need to be developed to further decrease computational time. Second, many of the parameters that do affect performance were kept constant because of time constraints. A more thorough review of how these different parameters make kriging more or less effective will need to be conducted. Finally, our method of generating maps of communications coverage could be validated by crowdsourced data to make our simulations more realistic.

REFERENCES

[1] P. W. Parfomak, “DOTs Federal Pipeline Safety Program: Background and Key Issues for Congress,” p. 37.

[2] J. Kirby, “How Puerto Rico Struggled to Communicate After Hurricane Maria,” Oct. 2017. [Online]. Available: http://nymag.com/daily/intelligencer/2017/10/how-puerto-rico-struggled-to-communicate-after-maria.html

[3] Open Signal, “3g and 4g LTE Cell Coverage Map - OpenSignal.” [Online]. Available: https://opensignal.com/

[4] “Fact Sheet Small Unmanned Aircraft Regulations (Part 107).”

[5] “Qualcomm Technologies releases LTE drone trial results,” May 2017. [Online]. Available: https://www.qualcomm.com/news/onq/2017/05/03/qualcomm-technologies-releases-lte-drone-trial-results

[6] K. Sundar, S. Venkatachalam, and S. Rathinam, “Formulations and algorithms for the multiple depot, fuel-constrained, multiple vehicle routing problem,” *arXiv:1508.05968 [cs]*, Aug. 2015, arXiv: 1508.05968. [Online]. Available: http://arxiv.org/abs/1508.05968

[7] K. Sundar, S. Misra, S. Rathinam, and R. Sharma, “Routing Unmanned Vehicles in GPS-Denied Environments,” *arXiv:1708.03269 [cs]*, pp. 62–71, Jun. 2017, arXiv: 1708.03269. [Online]. Available: http://arxiv.org/abs/1708.03269

[8] W. Souffriau, P. Vansteenwegen, G. Vanden Berghe, and D. Van Oudheusden, “A Path Relinking approach for the Team Orienteering Problem,” *Computers & Operations Research*, vol. 37, no. 11, pp. 1853–1859, Nov. 2010. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0305054809001464

[9] A. S. Krishnakumar and P. Krishnan, “On the accuracy of signal strength-based estimation techniques,” in *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, vol. 1, Mar. 2005, pp. 642–650 vol. 1.

[10] J. Fink and V. Kumar, “Online methods for radio signal mapping with mobile robots.” IEEE, May 2010, pp. 1940–1945. [Online]. Available: http://ieeexplore.ieee.org/document/5509574/

[11] N. Wagle and E. Frew, “Spatio-temporal characterization of airborne radio frequency environments,” in *2011 IEEE GLOBECOM Workshops (GC Wkshps)*, Dec. 2011, pp. 1269–1273.

[12] M. Molinari, M.-R. Fida, M. K. Marina, and A. Pescape, “Spatial Interpolation based Cellular Coverage Prediction with Crowdsourced Measurements.” ACM Press, 2015, pp. 33–38. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2787394.2787395