

Correlated Orienteering Problem and its Application to Informative Path Planning for Persistent Monitoring Tasks

Jingjin Yu Mac Schwager Daniela Rus

Abstract—We propose a novel non-linear extension to the Orienteering Problem (OP), called the Correlated Orienteering Problem (COP). We use COP to plan informative tours (cyclic paths) for persistent monitoring of an environment with spatial correlations, where the tours are constrained to a fixed length or time budget. The main feature of COP is a quadratic utility function that captures spatial correlations among points of interest that are close to each other. COP may be solved using mixed integer quadratic programming (MIQP) that can plan multiple disjoint tours that maximize the quadratic utility function. We perform extensive characterization of our method to verify its correctness, as well as its applicability to the estimation of a realistic, time-varying, and spatially correlated scalar field.

I. INTRODUCTION

We envision control algorithms for teams of robots where the robots use correlations established between different parts of the environment for specific events using historical data to determine more effectively the next action. For example, consider a scenario where a robot senses a specific event, say an intruder at a specific location in the environment. Using correlations and prediction a different robot can adjust its location to intercept the intruder. Toward that goal, in this paper, we consider how previously learned correlations can improve the performance and capabilities of teams of robots tasked with persistent surveillance and monitoring with limited travel budgets. Figure 1 provides a graphical example illustrating the problem addressed in this work.

In persistent surveillance and monitoring tasks using mobile robots, such as unmanned aerial vehicles (UAVs), the mobile robots usually have fixed base stations from which they must depart and return. Moreover, these robots have limited travel distance or time budget. Thus, when a large number of points of interest (nodes) must be surveyed, it may well be the case that only a subset of the nodes can be visited by the robots. Then, choices among the nodes must be made to accommodate two conflicting goals: 1. each robot must follow a tour (cyclic path) whose total cost does not exceed its travel budget, and 2. the robots must visit as many nodes as possible to maximize the amount of collected information, as measured by some utility (reward) function. When nodes have utilities

Jingjin Yu is with the Computer Science and Artificial Intelligence Lab at the Massachusetts Institute of Technology and the Mechanical Engineering Department at Boston University. E-mail: jingjin@csail.mit.edu.

Mac Schwager is with the Mechanical Engineering Department at Boston University. E-mail: schwager@bu.edu.

Daniela Rus is with the Computer Science and Artificial Intelligence Lab at the Massachusetts Institute of Technology. E-mail: rus@csail.mit.edu.

This work was supported in part by ONR projects N00014-12-1-1000 and N00014-09-1-1051.

We thank the anonymous reviewers for their constructive comments.

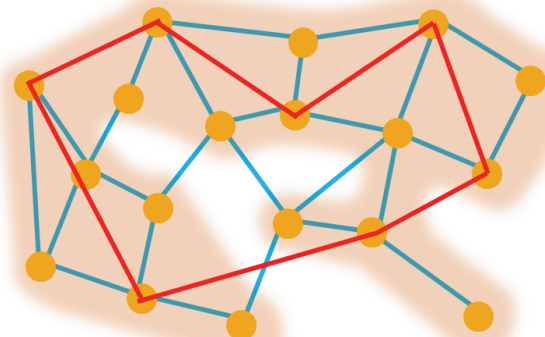
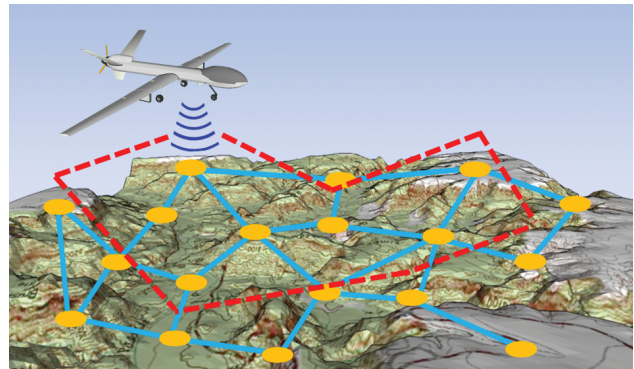


Fig. 1. [top] A surveillance scenario in which an UAV with limited range and sensing budgets is faced with the problem of covering a large number of nodes. [bottom] Using correlations among nodes in a network, it becomes possible to infer at least partial information for nodes that are not directly surveyed by the UAV. Following the red tour path, which is much shorter than a traveling salesman (TSP) tour, the UAV can provide at least partial information about every node in the network.

that are additive, an *Orienteering Problem* (OP) [1], a problem closely related to the well known *Traveling Salesman Problem* (TSP) [2], arises. Research on OP has yielded many effective algorithms for solving many versions of the problem, including *Team Orienteering Problem* (TOP)[3], in which multiple tours must be planned.

In practice, however, the information to be collected at the nodes are frequently correlated between nodes that are close to each other, rendering the total utility a non-linear combination of individual node utilities. That is, it is often the case that such information can be viewed as forming a spatially continuous field (that also varies over time). Therefore, surveying a given node will also yield partial information about its neighbors. For

example, the nodes may be cities, city blocks and locations in reservoirs with the associated quantities being population dynamics, criminal activities, and water pollutant concentration, respectively. In this paper, assuming that the spatial correlation among the nodes are intrinsic (*i.e.*, determined by local structures and mostly time-invariant) we propose a *quadratic* extension to OP and TOP, called the *Correlated Orienteering Problem* (COP), to incorporate such correlations in the informative path planning phase. After formulating COP, we provide mixed integer quadratic programming (MIQP) models for solving the problem for a single robot as well as for multiple robots. Our simulations indicate that COP and the associated MIQP capture spatial correlations among the nodes quite well. In a nutshell, COP can be viewed as a relaxation from the problem of planning informative tours for multiple robots under a spatially correlated field; the relaxed problem is then solved exactly. We also note that we do not assume convexity or submodularity over the underlying field.

Our work brings together ideas from two relatively disjoint branches of research: (discrete) OP and (mostly continuous) informative path planning for persistent monitoring tasks. OP, as indicated by its name, originates from orienteering games [3], [4]. In such a game, rewards of different sizes are spatially scattered. A player (or players in TOP) is faced with the task of maximizing the additive rewards under a time constraint that translates naturally to a distance constraint. Thus, OP can be viewed as a variation of both the Knapsack Problem (KP) [5] and the Traveling Salesman Problem (TSP) [2]. For a detailed account of OP, see [1].

The literature on informative path planning for persistent monitoring is fairly rich [6]–[19], covering both theories and applications. On work most closely related to ours, in [8], iterative TSP paths are planned to minimize the maximum latency across all nodes in a connected network. The authors show that the approach yields $O(\log n)$ approximation on optimality in which n is the number of nodes in the network. The problem of generating speed profiles for robots along predetermined cyclic (closed) paths for keeping bounded the uncertainty of a varying field is addressed in [16], in which the authors characterize appropriate policies for both single and multiple robots. In [17], decentralized adaptive controllers were designed to morph the initial closed paths of robots to focus on regions of high importance.

Sampling based planning methods (*e.g.* PRM, RRT, RRT* and their variations [20]–[22]) have also been applied to informative path planning problems. In [18], Rapidly-Exploring Random Graphs (RRG) are combined with *branch and bound* methods for planning most informative path for a mobile robot. In [13], the authors tackle the problem of planning a cyclic trajectory for best estimation of a time-varying Gaussian Random Field, using a variation of RRT called Rapidly-Expanding Random Cycles (RRC).

Lastly, our problem, and OP in general, also has a *coverage* element. Coverage of a two-dimensional region has been extensively studied in robotics [23]–[25], as well as in purely geometric settings, for example, in [26], where the proposed

algorithms compute the shortest closed routes for continuous coverage of polygonal interiors under an infinite visibility sensing model. Coverage with limited sensing range was also addressed later [27], [28].

This work brings two main contributions. First, we propose COP as a novel non-linear extension to OP to model the spatial correlations that are frequently present in informative path planning problems for persistent monitoring tasks. In particular, our formulation addresses the difficult problem of planning tours (cyclic paths) under limited travel budget. Second, we provide complete mixed integer quadratic programming (MIQP) models for solving COP. These models, with a MIQP solver, can effectively compute tours for multiple robots, each with a fixed base node. We then show that our models are applicable to estimate time-varying, spatially corrected scalar field, with computational experiments.

The rest of the paper is organized as follows. Section II formulates the *Correlated Orienteering Problem* (COP) that we study. MIQP models are then outlined for solving the single-robot and multi-robot cases in Section III. In Section IV, we perform extensive computational experiments to evaluate the applicability and effectiveness of COP and the associated MIQP models. Section V concludes the paper.

II. PROBLEM STATEMENT

Table I lists the symbols used in this paper. We study the problem of using mobile sensors to periodically monitor a set of nodes for quantities of interest that can be measured at these nodes. In general, these quantities can be represented as a continuous scalar or vector field that changes over space and time. One common characteristic of such fields is that they have spatial continuity. That is, at any fixed time instance, the values do not fluctuate much for nodes that are (spatially) close to each other. Such relationships can be represented as some form of correlations among the nodes¹.

This paper focuses on node networks in which the correlations determined by spatial relationships are known and do not vary much over time (*i.e.*, it is time-invariant) whereas the field itself may vary significantly over time. Then, using these correlations, it becomes possible to evaluate the field's value at one node using nearby nodes at any given time. In particular, we are interested in using mobile sensors with limited travel range to “sample” nodes of the network and then estimate the field's value at the rest of the nodes using correlations. We assume that the field changes at a slower pace in comparison to the time it takes a mobile sensor to complete a set of measurements at sampled nodes (*i.e.*, the field remains relatively static during a trip by the mobile sensor(s)). Below, this problem is formulated as a *Correlated Orienteering Problem* (COP). More precisely, our formulation is a quadratic extension to the linear *Orienteering Problem* (OP), in which one or multiple team members use only paths with fixed end points to visit sites and gain rewards. Our

¹Here, we use the broad meaning of correlation, which could be, but is not necessarily, the correlation of random variables.

TABLE I
LIST OF FREQUENTLY USED SYMBOLS AND THEIR INTERPRETATIONS.

$V = \{v_i\}$	Node or point of interest, $ V = n$
$G = (V, E)$	Node network
\mathbf{p}_i	The two dimensional coordinate of v_i
r_i	Utility of knowing complete information about v_i
$\psi(v_i, t)$	Time-varying scalar field
$A = \{a_k\}$	Mobile robot k , $ \{a_k\} = m$
v_{b_k}	Start and end node for robot a_k
c_k	Travel budget for robot a_k
Π	π_1, \dots, π_m , a set of robot tours
$J(\Pi)$	The cost function for a given set of robot tours
w_{ij}	The weight measuring v_i 's influence on v_j
x_i	Binary variable indicating whether v_i is on a tour
x_{ij}	Binary variable indicating whether v_j is visited immediately after v_i
x_{ijk}	Binary variable indicating whether v_j is visited immediately after v_i , by robot a_k
u_i	Integer variable, $2 \leq u_i \leq n$, the order of v_i in a tour path, if used
u_{ik}	Integer variable, $2 \leq u_{ik} \leq n$, the order of v_i in a tour path, if used, by robot a_k
d_{ij}	Travel cost from v_i to v_j , maybe non-symmetric
α_{ij}, β_{ij}	Linear regression coefficients

particular focus in this paper are tours - cyclic paths with the same starting and ending nodes.

Let $V = \{v_1, \dots, v_n\}$ be a set of spatially distributed nodes in some workspace $W \subset \mathbb{R}^2$. Each node $v_i \in V$ is associated with coordinates $\mathbf{p}_i \in \mathbb{R}^2$. Let

$$r : V \rightarrow \mathbb{R}^+, v_i \mapsto r_i,$$

represent the importance (utility) of the nodes. Let $\psi(\mathbf{p}, t)$ be a scalar field over W that changes over time. The values on the nodes of V , with a slightly abuse of notation, are written as $\psi(v_i, t)$, $1 \leq i \leq n$. We assume that the spatial relationship among the nodes of V , as determined by ψ , induces a directed graph G over V . More precisely, $G = (V, E)$ has an edge (v_i, v_j) if and only if $\psi(v_j, t)$ is dependent on $\psi(v_i, t)$. That is, let $N_i = \{v_{i_1}, \dots, v_{i_k}\}$ be the neighbors of v_i (in G , with edges pointing to v_i), then for some time-invariant f_i ,

$$\psi(v_i, t) = f_i(\psi(v_{i_1}, t), \dots, \psi(v_{i_k}, t)). \quad (1)$$

We point out that these requirements are quite mild as we do not assume that ψ has properties such as convexity.

Let there be m mobile robots (sensors), $A = \{a_1, \dots, a_m\}$. For a single robot a_k , let its base (*i.e.*, where it must start and end in a tour) be a node $v_{b_k} \in V$. Let

$$c : A \rightarrow \mathbb{R}^+, a_k \mapsto c_k,$$

represent the maximum distance budget the mobile robots can travel. In TOP, a team of mobile robots, working together, get to collect a reward (utility) r_i for visiting v_i the first time. The robots do not gain more utility for subsequent visits to v_i . The goal is to find tours for the robots satisfying the travel budget constraint (*i.e.*, tour distance for robot a_k is no more than c_k) so that the total utility is maximized. Given a set of tours (paths) $\Pi = \{\pi_1, \dots, \pi_m\}$ taken by the robots, let $\{x_1, \dots, x_n\}$ be n binary variables, with $x_i = 1$ if and only if v_i (*i.e.*, \mathbf{p}_i)

is on some tour $\pi_k \in \Pi$. We emphasize that the robots are not constrained to stay on the graph G , but can move between any two nodes with a cost proportional to the distance between them. The graph G only represents the spatial correlations in the field.

Our generalization of OP to COP aims at incorporating correlations among the nodes during the tour planning phase. To obtain a problem amenable to mathematical programming techniques, we relax Equation (1) in the following way. Let

$$w : E \rightarrow \mathbb{R}^+, (v_j, v_i) \mapsto w_{ji},$$

represent the effectiveness (a weighting) of using $\psi(v_j, t)$ to estimate $\psi(v_i, t)$. One may view w_{ji} as representing the amount of information that $\psi(v_j, t)$ has about $\psi(v_i, t)$, independent of other neighbors of v_i . The utility that can be collected over a vertex v_i is defined as

$$r_i(x_i + \sum_{v_j \in N_i} w_{ij} x_j (x_j - x_i)),$$

in which the quadratic term $x_j(x_j - x_i)$ is non-zero if and only if $x_j = 1$ and $x_i = 0$. Obviously, for each $0 \leq i \leq n$, $\sum_{v_j \in N_i} w_{ji} \leq 1$. Note that we do not assume that $w_{ij} = w_{ji}$, which may be the case if the field's values at v_i and v_j are Gaussian random variables. Over a set of tours for the robots, Π , the total utility to be maximized is then

$$J(\Pi) = \sum_{i=1}^n (r_i x_i + \sum_{v_j \in N_i} r_j w_{ij} x_i (x_i - x_j)). \quad (2)$$

The tour length constraints and the cost function given by Equation (2) defines a quadratic extension to TOP. In the next section, we propose a quadratic integer programming model with quadratic cost functions and linear constraints (often known as *mixed integer quadratic programming* or MIQP) for solving the problem. Then we will look at several practical estimation scenarios, involving one or multiple mobile robots, that can be tackled using this fairly general method.

Remark. We note that OP, TOP, and COP are all NP-hard problems. Restricting to OP (*i.e.*, a one-member team or a single mobile robot), for a given travel budget, an algorithm for OP must answer the question of whether the budget is enough for going through all nodes in V . Therefore, OP requires solving (potentially many) TSPs. For COP, making the weights $\{w_{ij}\}$ sufficiently small reduces it to an OP, because the quadratic cost (the second summation in Equation (2)) then becomes negligible.

III. MIXED INTEGER QUADRATIC PROGRAMMING MODELS FOR QUADRATIC TEAM ORIENTEERING PROBLEM

In this section, we outline MIQP models for COP, starting from the case of a single mobile robot and then move to the case of multiple robots. Then, we show how these models may be applied to potential applications.

A. MIQP Model for COP with a Single Tour

We start with the case of a single tour ($m = 1$). Without loss of generality, let the single robot start from v_1 . We adapt the constraints from [1]. Let x_{ij} be a binary variable with $x_{ij} = 1$ if and only if the robot visits v_j immediately after it visits v_i . Note that this does not depend on the existence of an edge between v_i and v_j in G . Because it is never beneficial to return to revisit a vertex, the robot should only enter and leave any vertex at most once, yielding the following constraints.

$$\sum_{i=2}^n x_{1i} = \sum_{i=2}^n x_{i1} = x_1 = 1, \quad (3)$$

$$\sum_{j=1, j \neq i}^n x_{ij} = \sum_{j=1, j \neq i}^n x_{ji} = x_i \leq 1, \quad \forall 2 \leq i \leq n. \quad (4)$$

The equality in Equation (3) is due to the requirement that the base node must be used.

Equations (3) and (4) ensure that the robot will take a tour starting and ending at v_1 . They do not, however, prevent additional disjoint tours being created. To prevent this from happening, let $2 \leq u_i \leq n$ be integer variables for $2 \leq i \leq n$. The constraints

$$u_i - u_j + 1 \leq (n-1)(1 - x_{ij}), \quad 2 \leq i, j \leq n, i \neq j \quad (5)$$

guarantees that no additional tours not containing v_i get created.

With the definition of the variables $\{x_{ij}\}$, the tour distance constraint can be easily enforced as

$$\sum_{i=1}^n \sum_{j=1, j \neq i}^n x_{ij} d_{ij} \leq c_1, \quad (6)$$

in which d_{ij} is the distance between v_i to v_j and c_1 is the tour distance constraint for the single (first) robot. Note that the distance d_{ij} needs not to be symmetric. Moreover, it is easy to incorporate sensing cost at a node v_i by absorbing that cost to d_{ij} for all $j \neq i$. Alternatively, if the sensing cost is not compatible with the travel cost, an additional cost constraint can be added as well.

The objective function Equation (2), subject to the constraint Equations (3), (4), (5), and (6), define a complete MIQP model that can be solved using an MIQP solver.

B. MIQP Model for COP with Multiple Tours

Once the MIQP model for a single robot is fully specified, extending it for multiple robots is rather straightforward. To accommodate m robots, the variables $\{x_{ij}\}$ and $\{u_i\}$ are extended to x_{ijk} and u_{ik} , with $1 \leq k \leq m$ representing the robots. Constrain Equations (3) and (4) become

$$\sum_{i=1, i \neq b_k}^n x_{b_k i k} = \sum_{i=1, i \neq b_k}^n x_{i b_k k} = x_{b_k k} = 1, \quad 1 \leq k \leq m \quad (7)$$

and

$$\sum_{j=1, j \neq i}^n x_{ijk} = \sum_{j=1, j \neq i}^n x_{jik} \leq 1, \quad \forall 1 \leq i \leq n, 1 \leq k \leq m \quad (8)$$

$$\sum_{k=1}^m \sum_{j=1, j \neq i}^n x_{ijk} = \sum_{k=1}^m \sum_{j=1, j \neq i}^n x_{jik} = x_i \leq 1, \quad \forall 1 \leq i \leq n. \quad (9)$$

Equation (4) splits into Equations (8) and (9) because each vertex should only be used at most once by each robot as well as by all robots.

The constraints on u_{ik} become (for all $1 \leq k \leq m$)

$$u_{ik} - u_{jk} + 1 \leq (n-1)(1 - x_{ijk}), \quad i, j \neq b_k, i \neq j, 1 \leq i, j \leq n. \quad (10)$$

The traveled distance constraint, Equation (6), becomes

$$\sum_{i=1}^n \sum_{j=1, j \neq i}^n x_{ijk} d_{ij} \leq c_k, \quad 1 \leq k \leq m. \quad (11)$$

Finally, the cost function Equation (2) remains the same.

Remark. The above MIQP model for COP does not allow two robots to start from the same base. We can easily accommodate such scenarios via modifying Equations (7), (8), and (9) accordingly. Also, if a fixed base is not required, the models can be easily modified to accommodate this.

C. Application to Persistent Monitoring Tasks

We project that our MIQP models will be useful in approaching OPs in which the objective is not simple linear summations of individual utility at the nodes. Here, we illustrate COP as a relaxed problem for multi-robot persistent monitoring tasks. As an example application, we look at planar networks. Such networks are fairly prevalent in persistent monitoring tasks.² To connect to COP, we assume that the f_i 's in Equation (1) take the form

$$\psi(v_i, t) = \alpha_{0i} + \sum_{v_j \in N_i} \alpha_{ji} \psi(v_j, t), \quad (12)$$

in which α_{0i} and α_{ji} 's are coefficients. Note that when not all values of ψ are available for nodes in N_i , Equation (12) can still be applied. To map these coefficients to our relaxed COP models, for each w_{ji} , we compute two sets of coefficients (using historical data). The first set of coefficients α'_{ji} 's are computed assuming all of N_i are available; the second set, α''_{ji} 's, are computed assuming v_j is v_i 's only neighbor. We then compute the weight w_{ji} via

$$w_{ji} = \frac{\alpha'_{ji} + \alpha''_{ji}}{\sum_{k \in N_i} (\alpha'_{ki} + \alpha''_{ki})}. \quad (13)$$

Equation (13) was chosen to balance the impact of single neighbors as well as the impact of the entire neighborhood in estimating the value of ψ of a node using its neighbors.

IV. SIMULATION EXPERIMENTS

In this section, we first evaluate MIQP models for COP over various benchmark examples. Then, we apply the models to a realistic synthetic scalar field that varies spatially and temporally. All computations were performed on a computer with Intel Core-i7 3930K CPU under an 8GB JavaVM. The MIQP solver used is Gurobi[29].

²For example, cities and connecting roads form such natural node networks.

A. Validity of the MIQP Models for COP

We first verify the validity of the MIQP models, *i.e.*, we check whether the models actually maximizes the objective function given by Equation (2). For this task, we begin with a single robot. Here, 3×3 and 4×4 grid networks are used with unit edge length as test node networks. For such grids, we set the weights w_{ji} for a vertex i simply as $1/|N_i|$. For example, if vertex i has three neighbors, then all w_{ji} 's are set to $1/3$.

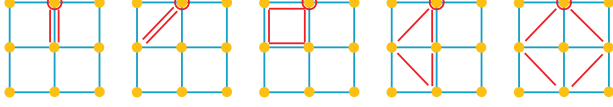


Fig. 2. Single robot tours for travel budgets 2, 3, 4, 5, and 6, in that order, from left to right.

For the 3×3 grid, we let the single robot start at the middle node on the top row (the circled node in Figure 2) and let the maximum allowed travel budget vary from 2 to 6 with unit increments. Each node has a unit utility. The computed tours for these budgets are illustrated in Figure 2, with utilities as 4.0, 4.5, 5.7, 7.3, and 9 (maximum possible), respectively. One can easily verify (manually) that these are consistent with the design of the MIQP model for a single robot.

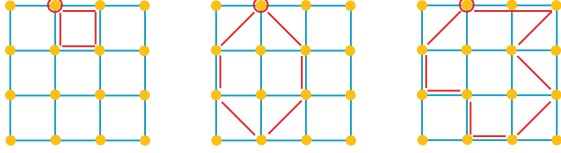


Fig. 3. Single robot tours for travel budgets 4, 8, and 12, in that order, from left to right.

For the 4×4 grid, under a similar setup, we get the tours as illustrated in Figure 3 for travel budgets 4, 8, and 12, respectively. The associated maximum utilities are 6.2, 11.5, and 16.0, respectively.

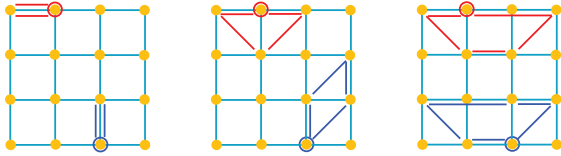


Fig. 4. Two-robot tours for travel budgets (per robot) 3, 5, and 7, in that order, from left to right.

Next, a two-robot setup is tested on the 4×4 grid, with the robots starting at opposite locations as indicated as the red and purple circled nodes in Figure 4, which illustrates the tours with individual travel budgets 3, 5, and 7, respectively. The associated maximum utilities are 7.2, 12.5, and 16.0, respectively. Each problem instance in this subsection took at most two seconds to solve.

B. Irregular Node Network

Our second experiment works with the irregular, realistic node network from Figure 1. The bounding rectangle of the network is roughly 13 units by 8 units. For this network, synthetic weights (w_{ij} 's) are again computed based on the number of neighbors. Up to three robots were attempted with the longest running time being about 100 seconds. The trial results and the associated parameters are given in Figure 5. The base nodes, indicated as colored circles, were hand picked (only once, *i.e.*, we did not try any other choices and then select the best one) to be roughly evenly distributed on the network.

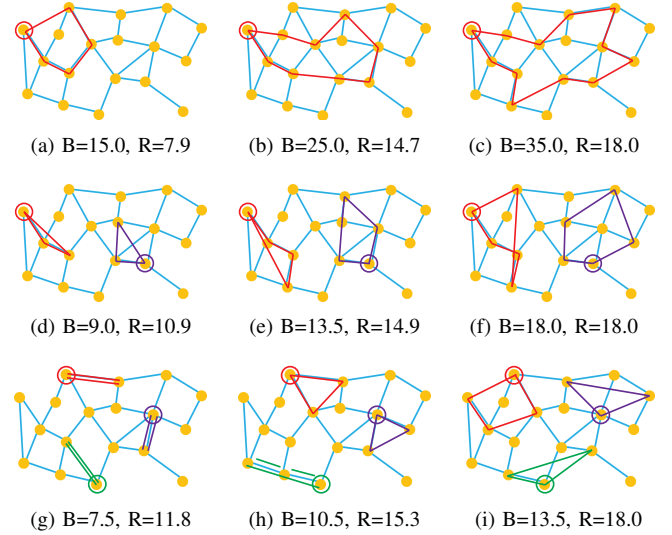


Fig. 5. Results from running MIQP models for COP on the irregular, realistic node network from Figure 1. The numbers under each picture indicate the budget (B) per tour/robot and the total utility (R), respectively.

From the result (Figure 5), we see that the MIQP model always select tours that do not have spatial overlaps, which is expected but nevertheless a nice feature to have. Also, regardless of the number of robots and tours, the total travel budgets (35.0 for one tour, 36.0 for two tours, and 40.5 for three) to ensure full coverage of the network appear to be similar. We note that for the cases with multiple tours, some of the individual budgets can be shortened. For example, the tours in Figure 5(f) can be (manually) updated to the tours in Figure 6, with reduced total (actual) travel budget but without reducing the collected utility. Varying individual budget is supported in the model by default.

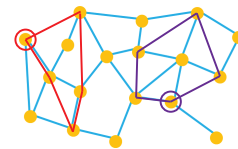


Fig. 6. Updated two-robot tours with the same utility as the tours in Figure 5(f).

C. Computational Performance

The third experiment seeks to establish the performance of our current MIQP models for solving COP. To evaluate the computational performance, we again start with a single robot and attempt grid networks with different sizes. Unlike the last experiment, the grid network is perturbed so that the edge lengths are no longer uniform; Figure 7 shows a typical 5×5 perturbed node network used in this experiment.

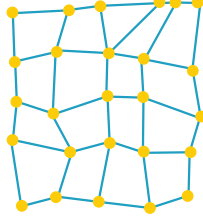


Fig. 7. A perturbed 5×5 grid network.

For a single robot, computation over a 7×7 grid network with 49 nodes for a full range of travel budgets can be completed within 20 minutes. We list the computational results in Table II. We notice from the experiment that, similar to TOP, the computational time can vary greatly depending on the travel budget for the same node network. The most computationally costly instances are these with utilities (for the given budget) close to the maximum possible utility. On the other hand, if the travel budget is too small to yield the maximum possible utility, then the computational cost is relatively low. Same holds when the travel budget is very large.

TABLE II
COMPUTATIONAL PERFORMANCE THE MIQP MODEL FOR A SINGLE ROBOT

Grid size		Trial #			
		1	2	3	4
3×3	budget	2.7	5.4	8.1	10.8
	time(s)	0.04	0.08	0.01	0.01
	utility	4.0	7.6	9.0	9.0
4×4	budget	3.6	7.2	10.8	14.4
	time(s)	0.02	0.92	0.16	0.01
	utility	5.3	11.2	15.4	16.0
5×5	budget	4.5	9.0	13.5	18.0
	time(s)	0.19	8.1	33.1	4.7
	utility	6.1	13.8	20.0	25.0
6×6	budget	10.8	16.2	21.6	27.0
	time(s)	31.6	48.3	23.6	0.89
	utility	16.25	24.9	32.6	36.0
7×7	budget	12.6	18.9	25.2	31.5
	time(s)	358	376	337	2387
	utility	20.2	31.2	39.8	46.8

We also tested the case of two and three robots using the same node networks; the model worked well with up to 6×6 grids. While all cases for 5×5 grids can be computed relatively quickly, hard instances for the 6×6 grid took over a day to compute, at which point we stopped the trial run. The results are listed in Tables III and IV for the two-robot and three-robot cases, respectively. It is interesting to see that

the computational time does not seem to vary much between those to cases.

TABLE III
COMPUTATIONAL PERFORMANCE FOR TWO ROBOTS

Grid size		Trial #			
		1	2	3	4
3×3	budget	2.4	3.6	4.8	6.0
	time(s)	0.04	0.04	0.04	0.04
	utility	6.5	9.0	9.0	9.0
4×4	budget	3.2	4.8	6.4	8.0
	time(s)	0.1	6.2	9.0	3.7
	utility	9.6	13.8	16.0	16.0
5×5	budget	4.0	6.0	8.0	10.0
	time(s)	0.4	166	1.5K	1.3K
	utility	10.5	16.2	20.8	25.0
6×6	budget	2.4	4.8	7.2	
	time(s)	0.2	18.2	27K	
	utility	7.3	14.3	20.5	

TABLE IV
COMPUTATIONAL PERFORMANCE FOR THREE ROBOTS

Grid size		Trial #			
		1	2	3	4
4×4	budget	2.1	3.2	4.3	5.3
	time(s)	0.5	0.2	1.8	3.2
	utility	9.8	13.5	14.6	16.0
5×5	budget	2.7	4.0	5.3	6.7
	time(s)	0.4	1.0	1.4K	20K
	utility	11.5	15.6	19.8	24.3
6×6	budget	3.2	4.8		
	time(s)	1.421	157		
	utility	14.7	20.3		

D. Measuring a Time-Varying Scalar Field

Lastly, we perform experiments to verify the effectiveness of Equation (13) in connecting actual scalar fields to our relaxed COP problem and models. We focus on the case of a single robot as the number of robots do not matter in measuring the quality of the tours produced by the MIQP model.

Our experiments are performed over a synthetic scalar field generated by three two-dimensional Gaussians. These Gaussians have fixed centers but varying magnitudes and covariance matrices over time; we fix the centers to ensure that the spatial correlations are relatively time-invariant. The field is simulated for 200 time steps; the snapshots of the field at time steps 0, 50, 100, 150, and 200 are provided in Figure 8. The node network used here is the same 5×5 randomized grids (see, e.g., Figure 7) scaled to the dimensions of the support of the scalar field. For each fixed travel budget, 100 random 5×5 node networks were generated. In each randomly generated network, the nodes of the network are given equal importance (*i.e.*, unit utility). To estimate α'_{ij} 's and α''_{ij} 's for computing the weights, data from the first fifty time steps were used. For running the models, the second diagonal node from the top-left corner was used as the base node.

After a tour is produced, the quality of the tour is estimated as follows. Visiting each node yields one *quality* point. For

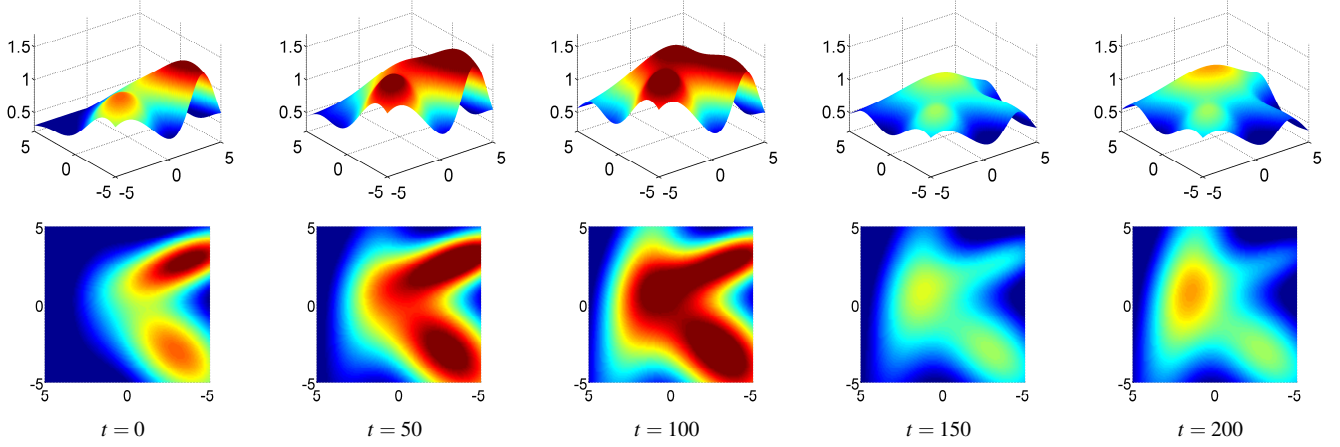


Fig. 8. Snapshots of a synthetic scalar field at time steps 0, 50, 100, 150, and 200, from left to right, respectively. [top row] Three-dimensional views. [bottom row] Two-dimensional heat-map views.

a node that is not visited but is a neighbor of one or more visited nodes, the values of the node at time steps 100, 150, and 200 are estimated in the following way. Let v_i be such a node and let N'_i be its neighbor set that is actually visited by the robot. The data for v_i and nodes in N'_i from $t \in [1, 50]$ were used to perform multiple linear regression according to

$$\psi(v_i, t) = \beta_{0i} + \sum_{v_j \in N'_i} \beta_{ji} \psi(v_j, t). \quad (14)$$

The resulting parameters were then used to estimate $\psi(v_i, t)$ for $t = 100, 150$, and 200 . Let the estimated value be $\psi'(v_i, t)$, we compute the quality of $\psi'(v_i, t)$ as

$$\frac{\sum_{t \in \{100, 150, 200\}} (\psi(v_i, t) - |\psi'(v_i, t) - \psi(v_i, t)|)}{\sum_{t \in \{100, 150, 200\}} \psi(v_i, t)} \quad (15)$$

To compare to our results, we also exhaustively search through the network for tours starting and ending at the same base node for the tour that minimizes the same quality defined by Equation (15) under the same travel budget. This experiment was limited to travel budgets 6 and 8. These budgets correspond to tours containing up to five nodes. While our model can produce tours with many more nodes, for comparing the result, we have to exhaustively search through all tours starting from the base node to find the best one, which becomes exponentially costly for tours with more than 5 nodes. The quality score obtained this way is denoted as “actual quality”.

TABLE V
MODEL FIDELITY OVER A SYNTHETIC SCALAR FIELD

Travel Budget	Model Quality	Actual quality	Relative error
6.0	7.16	7.64	0.48
8.0	8.46	9.38	0.92

The comparison result is given in Table V. Using the given metric, the average quality error was less than one, meaning that it was not more than the error incurred by omitting a single

node. In roughly 30% of the cases, the tour found using our method was identical to the one found using exhaustive tour search.

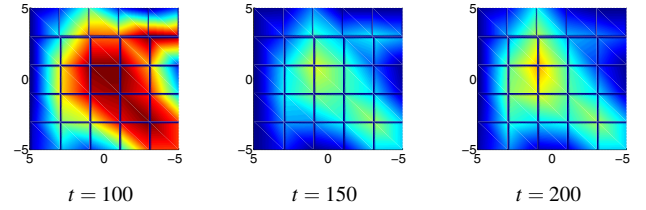


Fig. 9. Estimated scalar field based on the data obtained using the MIQP model.

As a secondary measure of the quality of our method, we put a regular 6×6 node network fitted over the same field (Figure 8) and run the MIQP model such that we just have enough budget to obtain a full utility of 36. We let the start node be the second left most node on the first row. From the output we can then estimate ψ for all nodes that are not visited on the tour. We then plot the much sparser survey data over the same space for time steps 100, 150, and 200 as shown in Figure 9. Comparing these figures with the corresponding ones from Figure 8, we observe that our models provide very reasonable estimation of the entire synthetic scalar field without the need to visit all the nodes.

V. CONCLUSION AND FUTURE WORK

In this paper, we introduce COP as an extension to OP with a non-linear cost function, to address the problem of planning tours for surveying a spatially correlated field that also varies over time. Our preliminary computational experiments show that the MIQP models for COP are effective in capturing the spatial correlation among nearby nodes, indicating that COP and the associated MIQP models are applicable to persistent monitoring tasks in which the mobile robots have limited travel range.

As a first attempt to introduce a new problem, our discussion of COP and its solution in the current paper is certainly limited in its scope. In the future, we expect to extend this study in several directions. Whereas our method is reasonably effective, it cannot yet handle networks with a large number of nodes quickly. We are working toward improving the current models to make them more efficient. To improve the computational efficiency, in addition to perfecting the MIQP model through techniques such as reachability analysis, we hope to understand from the solutions to the MIQP problems the general structures of these (optimal or near optimal) solutions. From there, we can then design heuristic algorithms that produce approximate solutions quickly. Currently, we are evaluating the performance of some preliminary heuristic algorithms; early results show that much larger problem instances involving hundreds of nodes can be solved quickly.

At the same time, we hope to improve our models to better capture real-world application scenarios. Toward this end, we note that the current weight selection criterion has lots of room for improvement. Through a more systematic approach, perhaps via statistical methods, we hope to derive more principled procedures for selecting the weights for the MIQP models.

Beyond these immediate steps, this paper only begins to address the problem of using correlations in informative path planning in a discrete fashion. The dual problem to this estimation problem is a learning problem: how can we learn the correlations among the nodes so as to apply the methods from this paper? Can we do learning and estimation simultaneously?

REFERENCES

- [1] P. Vansteenwegen, W. Souffriau, and D. V. Oudheusden, "The orienteering problem: A survey," *European Journal of Operational Research*, vol. 209, pp. 1–10, 2011.
- [2] G. Laporte, "The traveling salesman problem: An overview of exact and approximate algorithms," *European Journal of Operational Research*, vol. 59, pp. 231–247, 1992.
- [3] I. Chao, B. Golden, and E. Wasil, "Theory and methodology - the team orienteering problem," *European Journal of Operational Research*, vol. 88, pp. 464–474, 1996.
- [4] —, "Theory and methodology - a fast and effective heuristic for the orienteering problem," *European Journal of Operational Research*, vol. 88, pp. 475–489, 1996.
- [5] R. M. Karp, "Reducibility among combinatorial problems," in *Complexity of Computer Computations*, R. E. Miller and J. W. Thatcher, Eds. New York: Plenum, 1972, pp. 85–103.
- [6] N. Michael, E. Stump, and K. Mohta, "Persistent surveillance with a team of mavs," in *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 2708–2714.
- [7] R. N. Smith, M. Schwager, S. L. Smith, B. H. Jones, D. Rus, and G. S. Sukhatme, "Persistent ocean monitoring with underwater gliders: Adapting sampling resolution," *Journal of Field Robotics*, vol. 28, no. 5, pp. 714–741, September–October 2011.
- [8] S. Alamdari, E. Fata, and S. L. Smith, "Persistent monitoring in discrete environments: Minimizing the maximum weighted latency between observations," *International Journal of Robotics Research*, 2012, to appear.
- [9] E. Arvelo, E. Kim, and N. C. Martins, "Memoryless control design for persistent surveillance under safety constraints," September 2012, available at <http://arxiv.org/abs/1209.5805>.
- [10] C. G. Cassandras, X. Lin, and X. Ding, "An optimal control approach to the multi-agent persistent monitoring problem," *IEEE Transactions on Automatic Control*, vol. 58, no. 4, pp. 947–961, April 2013.
- [11] A. Girard, A. Howell, and J. Hedrick, "Border patrol and surveillance missions using multiple unmanned air vehicles," in *Proc. 43rd IEEE Conference on Decision and Control*, 2004, pp. 620–625.
- [12] B. Grocholsky, J. Keller, V. Kumar, and G. Pappas, "Cooperative air and ground surveillance," *IEEE Robotics and Automation Magazine*, vol. 13, no. 3, pp. 16–25, Sep 2006.
- [13] X. Lan and M. Schwager, "Planning periodic persistent monitoring trajectories for sensing robots in gaussian random fields," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA 13)*, May 2013, pp. 2407–2412.
- [14] N. Nigam and I. Kroo, "Persistent surveillance using multiple unmanned air vehicles," in *Proc. IEEE Aerospace Conference*, 2008, pp. 1–14.
- [15] J. L. Ny, M. A. Dahleh, E. Feron, and E. Frazzoli, "Continuous path planning for a data harvesting mobile server," in *47th IEEE Conference on Decision and Control*, 2008, pp. 1489–1494.
- [16] S. L. Smith, M. Schwager, and D. Rus, "Persistent robotic tasks: Monitoring and sweeping in changing environments," *IEEE Transactions on Robotics*, vol. 28, no. 2, pp. 410–426, April 2012.
- [17] D. E. Soltero, M. Schwager, and D. Rus, "Generating informative paths for persistent sensing in unknown environments," in *Proc. of the International Conference on Intelligent Robots and Systems (IROS 12)*, October 2012, pp. 2172–2179.
- [18] G. A. Hollinger and G. S. Sukhatme, "Sampling-based motion planning for robotic information gathering," in *Proc. Robotics: Science and Systems*, 2013.
- [19] J. Yu, S. Karaman, and D. Rus, "Persistent monitoring of events with stochastic arrivals at multiple stations," in *Proceedings IEEE International Conference on Robotics & Automation*, 2014.
- [20] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics & Automation*, vol. 12, no. 4, pp. 566–580, Jun. 1996.
- [21] S. M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," Oct 1998, tR 98-11, Computer Science Dept., Iowa State University.
- [22] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *Int. Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, June 2011. [Online]. Available: <http://ares.lids.mit.edu/papers/Karaman.Frazzoli.IJRR11.pdf>
- [23] H. Choset, "Coverage of known spaces: The boustrophedon cellular decomposition," *Autonomous Robots*, vol. 9, pp. 247–253, 2000.
- [24] —, "Coverage for robotics - a survey of recent results," *Annals of Mathematics and Artificial Intelligence*, vol. 31, pp. 113 – 126, 2001.
- [25] Y. Gabriely and E. Rimón, "Competitive on-line coverage of grid environments by a mobile robot," *Computational Geometry*, vol. 24, no. 3, pp. 197–224, 2003.
- [26] W.-P. Chin and S. Ntafos, "Optimum watchman routes," *Information Processing Letters*, vol. 28, pp. 39–44, 1988.
- [27] P. Hokayem, D. Stipanovic, and M. Spong, "On persistent coverage control," in *Proc. 46th IEEE Conference on Decision and Control*, 2008, pp. 6130–6135.
- [28] S. Ntafos, "Watchman routes under limited visibility," *Computational Geometry*, vol. 1, pp. 149–170, 1991.
- [29] I. Gurobi Optimization, "Gurobi optimizer reference manual," 2014. [Online]. Available: <http://www.gurobi.com>