Adaptive Sampling for Multi-Robot Wide-Area Exploration

Kian Hsiang Low, Geoffrey J. Gordon, John M. Dolan, and Pradeep Khosla

***Abstract*— The exploration problem is a central issue in mobile robotics. A complete coverage is not practical if the environment is large with a few small hotspots, and the sampling cost is high. So, it is desirable to build robot teams that can coordinate to maximize sampling at these hotspots while minimizing resource costs, and consequently learn more accurately about properties of such environmental phenomena. An important issue in designing such teams is the *exploration strategy*. The contribution of this paper is in the evaluation of an adaptive exploration strategy called *Adaptive Cluster Sampling* (ACS), which is demonstrated to reduce the resource costs (i.e., mission time and energy consumption) of a robot team, and yield more information about the environment by directing robot exploration towards hotspots. Due to the adaptive nature of the strategy, it is not obvious how the sampled data can be used to provide unbiased, low-variance estimates of the properties. This paper therefore discusses how estimators that are *Rao-Blackwellized* can be used to achieve low error. This paper also presents the first analysis of the characteristics of the environmental phenomena that favor the ACS strategy and estimators. Quantitative experimental results in a mineral prospecting task simulation show that our approach is more efficient in exploration by yielding more minerals and information with fewer resources and providing more precise mineral density estimates than previous methods.**

I. Introduction

The problem of exploring an unknown environment is a central issue in mobile robotics. Typically, it requires sampling the entire terrain [1]. However, a complete coverage is not practical in terms of resource costs if the environment is large with only a few small-scale features of interest or "hotspots", and the sampling cost is high. This arises in applications like planetary exploration (e.g., antarctic meteorite search, and prospecting for mineral deposits or localized methane sources on Mars), and environment and ecological monitoring (e.g., monitoring of ocean phenomena (plankton bloom), rare species, pollution, or contamination).

In this paper, we consider the above exploration problem with a team of robots, which can potentially complete the task faster than a single robot. A robot team is also more robust to failures by providing redundancy, but its performance may be adversely affected by physical interference between robots. Our goal is to design and build robot teams that can coordinate to (1) explore intelligently by maximizing sampling at these hotspots while minimizing resource costs, and consequently (2) learn more accurately about properties of such environmental phenomena. In particular, we focus on building a robot team for surface prospecting of *in situ* mineral resources on Mars, which is crucial to planning and establishing large, self-sufficient planetary settlements (e.g., site selection, processing equipment, and manufactured products). Without loss of generality, the work in this paper will be discussed in the context of this prospecting task.

The first aspect of our goal pertains to the *exploration strategy*: how do the robots decide where to explore next? Traditionally, conventional sampling methods [2] such as *Raster Scanning* (RS), *Simple Random Sampling* (SRS), and stratified random sampling have been used in single-robot exploration. The first approach acquires measurements at uniform intervals, thus incurring high sampling and travel costs to achieve adequate sampling density. The second approach selects a random sample of locations and makes measurements at each of the selected locations. However, it ignores the fact that hotspots such as mineral deposits are clustered and sometimes rare. This results in an imprecise (i.e., large variance) mineral density estimate of the explored region. Stratified random sampling requires prior knowledge of the mineral distribution for allocating the appropriate sampling effort among strata. Without such information, its efficiency degrades to that of SRS. There is one other conventional sampling scheme called *Systematic Sampling* (SS) [3], which spaces out the selected locations in a systematic manner. Though it has not been utilized in robot exploration, it will be used as a method of comparison in our paper.

This paper presents a multi-robot wide-area exploration strategy that is based on adaptive sampling. Assume that the explored region is discretized into a grid of N sampling units. In contrast to conventional sampling, *adaptive sampling* refers to sampling strategies in which the procedure for selecting units to be included in the sample depends on the sampling data observed during exploration. To satisfy the second aspect of our goal, adaptive sampling can exploit the characteristics of the environmental phenomena (i.e., spatial clustering of mineral deposits) to obtain more precise estimates of the properties (e.g., mineral density of explored region) than conventional strategies for a given sample size.

In this paper, we describe and evaluate a specific exploration strategy known as *Adaptive Cluster Sampling* (ACS) (Section III), which has desirable benefits: it (1) yields more minerals and information about the explored region by directing robot exploration towards hotspots (i.e., areas of high mineral density), thus providing detailed maps of the boundaries of such areas, and (2) reduces the resource costs of the robot team (Section VI).

The adaptive nature of this scheme incurs a considerable bias in conventional estimators due to a large proportion of sampling data observed in the hotspots. Consequently, two unbiased estimators are proposed in [3] for the ACS strategy

The authors are with Carnegie Mellon University, Pittsburgh, PA 15213, USA (bryanlow@cs.cmu.edu, ggordon@cs.cmu.edu, jmd@cs.cmu.edu, pkk@ece.cmu.edu).

(Section IV). This paper examines how the error of these estimators can be reduced through *Rao-Blackwellization* (Section V), in which the outputs of the estimators are averaged over several different ordered samples that are constructed by permuting the original sampled data. We have also presented the first analysis of the characteristics and distribution of the environmental phenomena that favor the ACS strategy (Section III-A) and estimators (Section IV-C). Before discussing the ACS strategy and estimators, an overview of the multi-robot architecture will be presented.

II. Multi-Robot Architecture

The multi-robot architecture comprises the teleoperation base and robot prospectors. To facilitate the teleoperator's analysis and monitoring, the base maintains a plan of the robot tours to visit the selected units to be sampled, and a list of sampled units and their corresponding mineral content. Each robot maintains an individual tour of its assigned units to be sampled, and shares spectrometric data of its sampled units with the base and robot team. The base continuously receives sampling data from the robots, selects new sampling units based on the ACS strategy described in Section III, and replans the robot tours to visit the new and current sampling units. After all selected units have been sampled, it determines the mineral density estimates of the explored region. These estimates can also be computed in a distributed manner among the robots as discussed below.

Our planning problem is an instance of the k-traveling salesman problem where k is the number of robots. The selected sampling units can be considered as cities to be visited. We consider two different optimality criteria: minimizing (1) total energy consumption of all robots, and (2) maximum mission time of any robot. In general, this problem is NP-hard. So, our centralized planner at the base uses a modified Minimum Spanning Tree (MST) heuristic proposed in [4] to obtain 2- and $2k$-competitive tour allocation for the first and second criterion respectively. Alternatively, the centralized planner can be easily replaced by a distributed auction-based planner [4] in every robot to eliminate central point of failure. In this case, each robot uses the ACS strategy directly for exploration; whenever it encounters new sampling units, it initiates an auction with the other robots to allocate the new and current sampling units. The bids are constructed according to the chosen criterion discussed above. This process is elaborated in [4].

In terms of computational complexity, the centralized and distributed planners require $O((k+n)n\log_2 n)$ and $O(n^2\log_2 n)$ time respectively where n is the number of selected units to be sampled. Note that n is usually much greater than k and this results in the same polynomial time complexity for both planners. In terms of communication complexity, the centralized and distributed planners require, respectively, $O(k)$ and $O(kn)$ messages to replan whenever new sampling units are selected.

III. Adaptive Cluster Sampling

In a large environment with only a few small hotspots, it is often useful after locating a sampling unit in a hotspot to continue exploring its neighborhood. One way of doing so is by the ACS strategy [3], which proceeds as follows: an initial sample of size n_1 is taken using SRS without replacement by the base or robot auctioneer. If the observed mineral content of an initially sampled unit satisfies a certain condition C (e.g., mineral content $\geq$ predefined threshold), the unit's neighborhood is added to the sample. For every unit, its neighborhood consists of the unit and a set of "neighboring" units (e.g., top, bottom, left, and right units). If any other units in that neighborhood satisfy C, their neighborhoods are also included in the sample. This process is repeated until no more units that satisfy C are encountered.

At this stage, clusters of units are obtained. Each *cluster* contains units that satisfy C and a boundary of *edge units*. An *edge unit* is a unit that does not satisfy C but is in the neighborhood of a unit that does. The final sample of size ν consists of up to n_1 clusters. There can be fewer than n_1 distinct clusters, since two units in the initial sample that satisfy C could have been selected from the same cluster. If a unit in the initial sample does not satisfy C, it is considered to be a cluster of size one.

Let the *network* $\mathcal{A}_i$ that is generated by unit i be defined as a cluster generated by that unit with its edge units removed. A selection of any unit in $\mathcal{A}_i$ leads to the selection of all units in $\mathcal{A}_i$. Any unit that does not satisfy C is a network of size one since its selection does not lead to the inclusion of any other units. This implies that any edge unit is also a network of size one. Hence, any cluster of size larger than 1 can be decomposed into a network with units that satisfy C, and also networks (edge units) of size one that do not satisfy C. Clusters may overlap on their edge units. In contrast, networks are disjoint and form a partition of the entire population of units.

Fig. 1b illustrates the adaptive cluster sample technique. The values in this table are obtained in a simulation test run on the prospecting region in Fig. 1a, which is discretized into a 28×20 grid of square sampling units (thus, the total number of units $N = 560$). The condition for sampling a unit's neighborhood is defined as $C = (y \geq 1.0$ wt$\%)$ where y is the observed mineral content of a sampling unit. The boxed values correspond to units from the initial sample. The lightly and darkly shaded units correspond, respectively, to the network and edge units of a cluster. Note that the network within the cluster is intersected twice by the initial sample.

A. Cost Analysis of ACS

This section analyzes the cost of ACS over SRS for a given final sample size ν. If the cost of a multi-robot exploration strategy is attributed primarily to sampling and motion (e.g., see Section VI), only the motion costs of ACS and SRS can differ due to the same ν. In particular, we compare the worst-case motion costs of ACS and SRS; let the ratio of these costs be ρ. Then, using the MST heuristic (Section II), ρ is the ratio of the MST size on the n_1 randomly selected initial sample units together with the cost of $(\nu - n_1)$ adaptively added units to the MST size on the ν randomly selected units. So, if $\rho < 1$, the worst-case tour allocation cost of ACS is less than that of SRS for either of the optimality criteria. The results below were processed from the MST sizes for

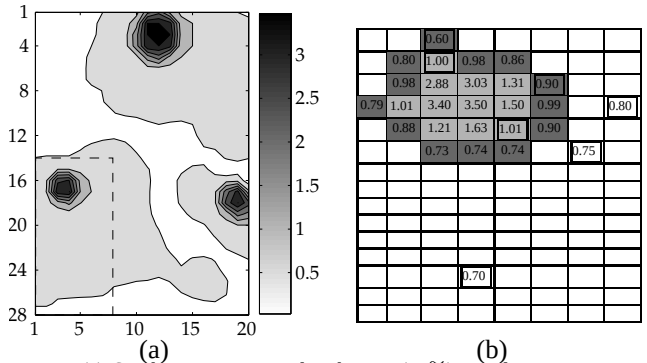

Fig. 1. (a) Synthetic zirconium distribution (wt%) in a lunar prospecting region with population mean $\mu = 0.648$ [5]. (b) Partial ACS example corresponding to boxed area (dashed) in (a).

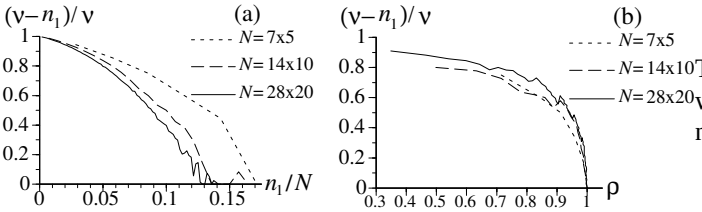

Fig. 2. Graphs of proportion of adaptively added units in the final sample vs. (a) initial sample size, and (b) cost ratio with varying grid resolutions.

$n_1 = 1, \ldots, N$, each of which was obtained by averaging over 1000 test runs.

Fig. 2a shows the largest proportion of units that can be adaptively added (i.e., largest value of $(\nu - n_1)/\nu$) when ACS does not cost more than SRS (i.e., $\rho \leq 1$). This is illustrated with varying grid resolutions (i.e., $7\times5, 14\times10, 28\times20$ grids of square sampling units); the cost of each adaptively added unit is thus a unit's width. Note that when no units can be adaptively added, the final sample size ν is equal to the initial sample size n_1. As a result, ACS degrades to SRS and they have the same costs (i.e., $\rho = 1$). The results show that the maximum proportion of units that can be adaptively added decreases with increasing initial sample size. This implies that for sampling large hotspots, a smaller initial sample ($n_1/N \leq 0.05$) has to be used in order for ACS to cost less than SRS: this allows more units to be adaptively added from the large hotspots located by the initial sample. Note that a smaller initial sample is also sufficient for locating large hotspots. On the other hand, a larger initial sample ($0.1 \leq n_1/N \leq 0.13$) can be used for small hotspots, as they require fewer adaptively added units and are harder to find. Also, though a higher grid resolution decreases the proportion of adaptively added units, it provides a more detailed mapping of the hotspot boundaries.

Fig. 2b shows the proportion of units that can be adaptively added (i.e., $(\nu - n_1)/\nu$) where the cost ratio ρ is minimized for varying n_1. The results show that the proportion of units that can be adaptively added decreases with increasing cost ratio; since an adaptively added unit does not cost more than an initial sample unit, decreasing the proportion of adaptively added units increases the cost of ACS.

IV. Unbiased ACS Estimators

Since the ACS scheme results in a large proportion of high mineral content data sampled from the hotspots, it will incur a considerable bias with the conventional sample mean estimator $\bar{\mu} = \nu^{-1} \sum_{i=1}^{\nu} y_i$ (Section VI). Kriging (or Gaussian process regression) [3] is a more sophisticated alternative but will be similarly biased. Hence, unbiased estimators are needed for the ACS scheme. Two of these are presented in this section.

A. Modified Horvitz-Thompson Estimator

The first ACS estimator is modified from the Horvitz-Thompson (HT) estimator [3]. Let $\mathcal{B}_i$ be the set of units in the ith network and m_i be the number of units in $\mathcal{B}_i$. Note that $\mathcal{B}_i$ is defined in the same way as network $\mathcal{A}_i$ in Section III except that its index i refers to the network label rather than the unit label. The probability that the initial sample intersects network $\mathcal{B}_i$ is

$$\pi_i \stackrel{\text{def}}{=} 1 - \binom{N - m_i}{n_1} \Big/ \binom{N}{n_1} . \tag{1}$$

The total mineral content of the explored region can be written as the sum of the mineral contents of the individual networks. So, the average mineral content is

$$\mu = \frac{1}{N} \sum_{i=1}^{K} y_i^*$$

where y_i^* is the total mineral content of the ith network and K is the total number of distinct networks in the population.

μ cannot be computed directly due to the unknown y_i^*'s for unsampled networks. So, to form an unbiased estimator of μ, each term in the sum can be multiplied by I_i/π_i, where I_i is an indicator variable of value 1 if the initial sample intersects $\mathcal{B}_i$, and 0 otherwise. The expected value of I_i/π_i is 1, so our estimator is unbiased; since I_i is 0 for unsampled networks, information about these networks are not needed to calculate our estimator. Applying this trick yields the *modified HT estimator* of μ:

$$\hat{\mu}_{HT} = \sum_{i=1}^{K} \frac{y_i^* I_i}{N\pi_i} = \sum_{i=1}^{\kappa} \frac{y_i^*}{N\pi_i} \tag{2}$$

where κ is the number of distinct networks intersected by the initial sample.

For practical use of the HT estimator, it is important to be able to estimate its variance from the sample. There is a simple closed-form formula which can be used for this purpose. π_i has been defined to be the probability that the initial sample intersects the ith network. Define π_{jk} to be the probability that the initial sample intersects both the jth and kth networks. If $j = k$, then $\pi_{jk} = \pi_j$. Otherwise, to compute π_{jk}, notice that the probability that the initial sample intersects neither network j nor network k is

$$\mathrm{P}(I_j \neq 1 \cap I_k \neq 1) = \binom{N - m_j - m_k}{n_1} \Big/ \binom{N}{n_1} .$$

So, the probability that the initial sample intersects either jth or kth network is $1 - \mathrm{P}(I_j \neq 1 \cap I_k \neq 1)$, and

$$\pi_{jk} = \pi_j + \pi_k - (1 - \mathrm{P}(I_j \neq 1 \cap I_k \neq 1)) .$$

Since $\hat{\mu}_{HT}$ is a sum of several terms, its variance can be derived by taking the sum of covariances between these terms:

$$\begin{aligned} \mathrm{var}[\hat{\mu}_{HT}] &= \sum_{j=1}^{K} \sum_{k=1}^{K} \mathrm{cov}[\frac{y_j^* I_j}{N\pi_j}, \frac{y_k^* I_k}{N\pi_k}] \\ &= \sum_{j=1}^{K} \sum_{k=1}^{K} \frac{y_j^*}{N\pi_j} \frac{y_k^*}{N\pi_k} \mathrm{cov}[I_j, I_k] . \end{aligned} \tag{3}$$

(3) cannot be computed from the sample data since not all the networks in the population are necessarily sampled. So, to obtain an unbiased estimator of the variance, we can use a similar trick as before: each term is multiplied by $I_j I_k/\pi_{jk}$ (which has an expected value of 1) to get

$$\begin{aligned}\widehat{\text{var}}[\hat{\mu}_{HT}] &= \sum_{j=1}^{K}\sum_{k=1}^{K}\frac{y_j^* I_j}{N\pi_j}\frac{y_k^* I_k}{N\pi_k}\frac{\text{cov}[I_j, I_k]}{\pi_{jk}} \\ &= \frac{1}{N^2}\left[\sum_{j=1}^{\kappa}\sum_{k=1}^{\kappa}\frac{y_j^* y_k^*}{\pi_{jk}}\left(\frac{\pi_{jk}}{\pi_j\pi_k}-1\right)\right] .\end{aligned} \tag{4}$$

The second equality follows because cov[I_j,I_k] is π_{jk}–$\pi_j\pi_k$.

The network formulation of these estimators allows their computations to be readily distributed among the robots if desired; the networks are allocated to the robots such that each robot is responsible for the computations within its assigned networks. The resulting network data can then be aggregated by a robot or a cyclic message-passing algorithm to obtain the estimates. This can be similarly achieved for the second estimator described next.

B. Modified Hansen-Hurwitz Estimator

The second ACS estimator is modified from the Hansen-Hurwitz (HH) estimator [3]. In Section IV-A, we mention that the total mineral content of the explored region is the sum of the mineral contents of the individual networks. The mineral content of each network can be written as the average mineral content of all units in this network summed over its number of network units. So, the average mineral content of the explored region can also be expressed as

$$\mu = \frac{1}{N}\sum_{i=1}^{N} w_i$$

where w_i is the average mineral content of the network $\mathcal{A}_i$ containing unit i.

μ cannot be computed directly due to the unknown w_i's for unsampled networks. Using the same trick as in Section IV-A, an unbiased estimator of μ can be formed by multiplying each term in the sum with NJ_i/n_1, where J_i is an indicator variable of value 1 if unit i is included in the initial sample, and 0 otherwise. The expected value of NJ_i/n_1 is 1, so our estimator is unbiased; since J_i is 0 for units not in the initial sample, information about these units is not needed to calculate our estimator. Applying this trick yields the *modified HH estimator* of μ:

$$\hat{\mu}_{HH} = \frac{1}{n_1}\sum_{i=1}^{N} w_i J_i = \frac{1}{n_1}\sum_{i=1}^{n_1} w_i \tag{5}$$

Note that $\hat{\mu}_{HH}$ can be interpreted as the conventional sample mean obtained using SRS of size n_1 from a population of w_i values rather than y_i values. So, using the theory of SRS [3],

$$\text{var}[\hat{\mu}_{HH}] = \frac{N-n_1}{Nn_1(N-1)}\sum_{i=1}^{N}(w_i-\mu)^2 \tag{6}$$

with unbiased estimator

$$\widehat{\text{var}}[\hat{\mu}_{HH}] = \frac{N-n_1}{Nn_1(n_1-1)}\sum_{i=1}^{n_1}(w_i-\hat{\mu}_{HH})^2 . \tag{7}$$

C. Efficiency Analysis of ACS Estimators

The estimator efficiency of ACS over SRS depends on the characteristics of the environmental phenomena (i.e., mineral distribution being sampled). In particular, $\hat{\mu}_{HH}$ is more efficient than the conventional sample mean $\hat{\mu}$ for SRS if var[$\hat{\mu}_{HH}$] < var[$\hat{\mu}$]. Using the theory of SRS [3],

$$\text{var}[\hat{\mu}] = \frac{N-\nu}{N\nu(N-1)}\sum_{i=1}^{N}(y_i-\mu)^2 \tag{8}$$

The total sum of squared difference between y_i and μ in (8) can be partitioned into within-network and between-network components:

$$\sum_{i=1}^{N}(y_i-\mu)^2 = \sum_{i=1}^{N}(y_i-w_i)^2 + \sum_{i=1}^{N}(w_i-\mu)^2 \tag{9}$$

Using (6), (8), and (9), var[$\hat{\mu}_{HH}$] < var[$\hat{\mu}$] if and only if

$$\left(1-\frac{n_1}{\nu}\right)\sum_{i=1}^{N}(y_i-\mu)^2 < \left(1-\frac{n_1}{N}\right)\sum_{i=1}^{N}(y_i-w_i)^2 \tag{10}$$

It can be observed from (10) that $\hat{\mu}_{HH}$ is more efficient than $\hat{\mu}$ if (1) the within-network variance of the population (rightmost term) is sufficiently high, (2) the final sample size ν is not much larger than the initial sample size n_1 for $\hat{\mu}_{HH}$ so that $1-n_1/\nu$ is small, and (3) $n_1 \ll N$ so that $1-n_1/N$ is large. However, conditions 2 and 3 can oppose condition 1 because a small difference between initial and final sample size, and a small initial sample size usually mean small within-network variance. So, ACS with $\hat{\mu}_{HH}$ performs better than SRS with $\hat{\mu}$ if the networks are small enough to restrict the final sample size but large enough for the within-network variance to represent the population variance reasonably. That is, it works better with environmental phenomena that are clustered into a few small hotspots. Even though drastically lowering the threshold for condition C can increase the within-network variance and improve condition 1, it increases the final sample size tremendously and violates condition 2 easily.

Although it is straightforward to compare var[$\hat{\mu}_{HT}$] (3) and var[$\hat{\mu}$], the result cannot be easily interpreted since var[$\hat{\mu}_{HT}$] involves the intersection probabilities. However, empirical results in Section VI show that $\hat{\mu}_{HT}$ is consistently more efficient than $\hat{\mu}$. In the next section, we will show how the variances of the ACS estimators can be reduced to be even more efficient.

V. Unbiased Rao-Blackwellized ACS Estimators

An estimator $t(D_o)$ of a population characteristic μ is a function t which maps our observed data D_o to an estimate of μ. Saying that μ is a population characteristic means there is a parameter vector $\boldsymbol{\theta}$ which completely describes the distribution of our population, and $\mu = \mu(\boldsymbol{\theta})$ is a function of $\boldsymbol{\theta}$. In our setting, D_o is an ordered list of pairs $\langle i_s, y_{i_s}\rangle$ where i_s is the unit sampled at step s and y_{i_s} is its mineral content. The population characteristic of interest μ is the average mineral content of the explored region. The population parameter is $\boldsymbol{\theta} = \langle y_1, \ldots, y_N\rangle$, which is the vector of true mineral contents for all units in the population. The estimators of μ that we are interested in are $\hat{\mu}_{HT}$ and $\hat{\mu}_{HH}$.

To evaluate an estimator $t(D_o)$, its distribution conditioned on a possible value of $\boldsymbol{\theta}$ can be examined. Good estimators have low Mean-Squared Errors (MSEs), i.e., the distribution $\mathrm{P}(t(D_0) - \mu|\boldsymbol{\theta})$ is concentrated around 0. We will now describe how to reduce the MSEs of $\hat{\mu}_{HT}$ and $\hat{\mu}_{HH}$.

Rao-Blackwellization is a procedure that can reduce the MSE of an arbitrary estimator $t(D_o)$ [3]. The improved estimator is $\mathrm{E}(t(D_o)|\mathcal{D})$, where $\mathcal{D}$ is a reduced description of our data that omits some redundant information. In particular, $\mathcal{D}$ is defined as a *statistic* if it is a function of our data D_o (i.e., $\mathcal{D} = g(D_o)$), and $\mathcal{D}$ is defined as a *sufficient statistic* if it contains all relevant information in D_o about $\boldsymbol{\theta}$, i.e., $\mathrm{P}(D_o|\mathcal{D}, \boldsymbol{\theta}) = \mathrm{P}(D_o|\mathcal{D})$. Given these definitions, Rao-Blackwellization is the process of computing $\mathrm{E}(t(D_o)|\mathcal{D})$ when $\mathcal{D}$ is a sufficient statistic. In our case, $\mathcal{D}$ is set to be the *unordered* set of distinct, labeled observations, i.e., $\mathcal{D} = \{\langle i, y_i\rangle |\ i \in \mathcal{S}\}$ where $\mathcal{S}$ is the set of distinct unit labels in our data sample.

The following theorem, adapted from the Rao-Blackwell theorem, justifies the use of Rao-Blackwellized estimator:

Theorem 1: Let $t = t(D_o)$ be a (not necessarily unbiased) estimator of μ. Define $t_{\mathcal{D}} = \mathrm{E}[t|\mathcal{D}]$. Then
(a) $t_{\mathcal{D}}$ is an estimator;
(b) $\mathrm{E}[t_{\mathcal{D}}] = \mathrm{E}[t]$;
(c) $\mathrm{MSE}[t_{\mathcal{D}}] \leq \mathrm{MSE}[t]$ with strict inequality for all $\boldsymbol{\theta}$ such that $\mathrm{P}_{\boldsymbol{\theta}}(t \neq t_{\mathcal{D}}) > 0$.

Corollary 1: If t is unbiased,

$$\begin{aligned} \mathrm{var}[t_{\mathcal{D}}] = &\ \mathrm{var}[t] - \mathrm{E}_{\mathcal{D}}\mathrm{E}[(t - t_{\mathcal{D}})^2|\mathcal{D}] \\ = &\ \mathrm{var}[t] - \mathrm{E}_{\mathcal{D}}\{\mathrm{var}[t|\mathcal{D}]\} \ . \end{aligned} \tag{11}$$

The proofs of Theorem 1 and Corollary 1 are provided in [6]. From (11), $\mathrm{var}[t_{\mathcal{D}}] \leq \mathrm{var}[t]$ since the variance reduction term $\mathrm{E}_{\mathcal{D}}\{\mathrm{var}[t|\mathcal{D}]\} \geq 0$.

Rao-Blackwellization does nothing if $g(D_o)$ is already a function of $\mathcal{D}$. On the other hand, it achieves the largest possible reduction in variance when $\mathcal{D}$ is a *minimal sufficient statistic*. A minimal sufficient statistic is one that reduces D_o as much as possible without losing information about $\boldsymbol{\theta}$:

Definition 1: A sufficient statistic $\mathcal{D} = g(D_o)$ is *minimal sufficient* for $\boldsymbol{\theta}$ if, for any other sufficient statistic $\mathcal{D}' = g'(D_o)$, $\mathcal{D}$ is a function of $\mathcal{D}'$.

In our case, $\mathcal{D}$ is minimal sufficient, and $\hat{\mu}_{HT}$ and $\hat{\mu}_{HH}$ are not functions of $\mathcal{D}$; they depend on the order of selection [6]. In order to Rao-Blackwellize $\hat{\mu}_{HT}$ and $\hat{\mu}_{HH}$, we will need several notations. Let $G = \binom{\nu}{n_1}$ be the number of combinations of n_1 distinct initial sample units from the ν units in the final sample and let these combinations be indexed by the label g where $g = 1, 2, \ldots, G$. Let τ_g be the value of an estimator t when the initial sample consists of combination g, I_g be an indicator variable of value 1 if the gth combination can result in $\mathcal{D}$ (i.e., is *compatible* with $\mathcal{D}$), and 0 otherwise. The number of compatible combinations is then $\xi = \sum_{g=1}^{G} I_g$. It follows that $\mathrm{P}(t = \tau_g|\mathcal{D}) = 1/\xi$ for all compatible g. So, the improved Rao-Blackwellized estimator is

$$t_{RB} = \mathrm{E}[t|\mathcal{D}] = \frac{1}{\xi}\sum_{g=1}^{G}\tau_g I_g = \frac{1}{\xi}\sum_{g=1}^{\xi}\tau_g \ . \tag{12}$$

The variance of t_{RB} is obtained using (11) where $t_{\mathcal{D}} = t_{RB}$. The unbiased estimator of $\mathrm{var}[t_{RB}]$ is then

$$\widehat{\mathrm{var}}[t_{RB}] = \widehat{\mathrm{var}}[t] - \mathrm{var}[t|\mathcal{D}] = \widehat{\mathrm{var}}[t] - \frac{1}{\xi}\sum_{g=1}^{\xi}(\tau_g - t_{RB})^2 \ . \tag{13}$$

Since (12) and (13) are based on samples compatible with $\mathcal{D}$, naively, the ξ compatible samples have to be identified from the G combinations and their corresponding ξ estimators have to be evaluated. ξ and G can be potentially large, which would render the Rao-Blackwellized method computationally infeasible. However, closed-form expressions exist for the Rao-Blackwellized HT (RBHT) and HH (RBHH) estimators, which are described in [6]. These expressions are computationally efficient if relatively few networks of size larger than 1 are intersected by the initial sample. This assumption is valid if the prospecting region contains only a few hotspots.

VI. Experiments and Discussion

This section presents quantitative evaluations of the ACS strategy and its estimators for wide-area exploration with a team of four robots. The experiments were performed using Webots, a mobile robot simulator, which incorporated 10% white noise in its sensors and actuators. 16 directed distance sensors with 0.3 m range were modelled around the 0.32 m (L) $\times$ 0.27 m (W) $\times$ 0.2 m (H) robot body. Each robot could sense its global position through GPS[1], and communicate spectrometric and tour data with the base. The robots used the potential fields method for navigation between sampling units and obstacle avoidance. Each robot could move at a maximum speed of 0.425 m/s and consumed about 28.2 J/m. It used the Alpha Particle X-Ray Spectrometer (APXS) (1.3 W) for sampling, which required about 2 hours to obtain a high-quality x-ray spectrum of the mineral content. So, sampling each unit would use about 9.5 kJ. The 6.46 km $\times$ 4.61 km prospecting region is discretized into a 28×20 grid of sampling units such that each unit's width is about 231 m (Fig. 1a). The robots were placed at a sampling unit in the center of the region and had to rendezvous at this same unit after all selected units were sampled.

To compare the performance of the estimators, the Root Mean-Squared Error (RMSE) criterion is used to measure their efficiency:

$$\mathrm{RMSE}[t] = \left[\frac{1}{R}\sum_{i=1}^{R}(\tau_i - \mu)^2\right]^{\frac{1}{2}}$$

where $R = 20$ is the number of test runs, τ_i is the mean mineral content estimate obtained in test run i.

Using this measure, a quantitative test was conducted to compare the estimators described above. For ACS, the initial sample size n_1 was 10, 20, or 40 sampling units. After 20 test runs for each n_1, it resulted in an average final sample size $\mathrm{E}[\nu]$ of approximately 41, 64, and 92 units, which corresponded to 7.4%, 11.5%, and 16.4% of the 560 total sampling units. The SRS and SS schemes were conducted using the same sample sizes as $\mathrm{E}[\nu]$.

[1]Deployment of space exploration infrastructure would ultimately result in GPS or similar localization capability on the Moon and Mars.

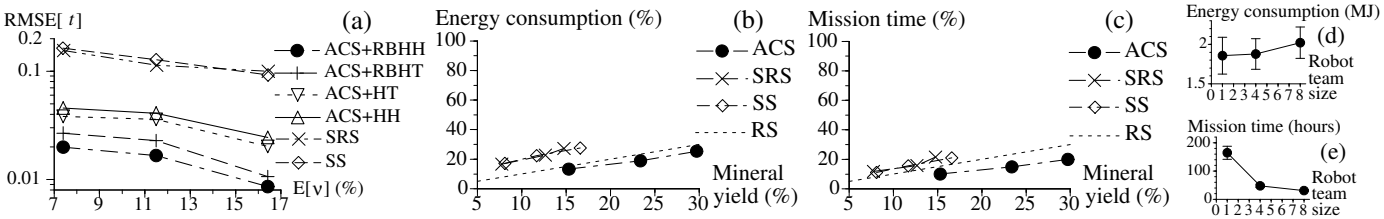

Fig. 3. Comparison of (a) RMSEs of different estimators and sampling strategies, (b) energy consumption of different sampling strategies, (c) mission time of different sampling strategies, (d) energy consumption, and (e) mission time of ACS strategy with different robot team sizes.

Test results (Fig. 3a) show, with statistical significance, that the ACS estimators perform better than the non-ACS estimators, and among the ACS estimators, the Rao-Blackwellized estimators achieved lower RMSE. This implies that the ACS estimators, especially the Rao-Blackwellized ones, are practically more appealing because more accurate mineral density estimates can be obtained with a reasonably small sample size. Using t-tests ($\alpha = 0.1$), the differences in RMSEs between the estimators have been verified to be statistically significant if these differences are more than 0.007, 0.013, and 0.008 for the sample sizes of 41, 64, and 92 units respectively. Note that the biased sample mean estimator $\bar{\mu}$ under the ACS scheme is not included in Fig. 3a; it has extremely large RMSEs of 0.682, 0.670, and 0.524 corresponding to 7.4%, 11.5%, and 16.4% of the total sampling units.

To compare the robot team performance between different sampling/exploration strategies, the previously mentioned optimality criteria are considered: minimizing (1) total energy consumption of all robots, and (2) maximum mission time of any robot. Figs. 3b and c show the results after 20 test runs for the first and second criterion respectively; the mineral yield, energy consumption, and mission time recorded for the various sampling strategies are given as a percentage of the corresponding values for RS (i.e., complete sampling of 560 units). As a result, varying the size of the prospecting region does not change our results. Note that each strategy (other than RS) has three different records in its plot, which correspond to E[ν] of 41, 64, and 92 units; a smaller sample size gives a smaller mineral yield. The line for RS shows a constant ratio of energy consumption or mission time to mineral yield. We observe that the ACS strategy yields more minerals than SRS and SS with less energy and mission time. The differences in mineral yield, energy consumption or mission time between ACS and the other two strategies have been verified using t-tests ($\alpha = 0.1$) to be statistically significant.

Furthermore, in contrast to SRS and SS, we observe that ACS falls below the dotted line of RS, which implies it achieves a lower ratio of energy consumption or mission time to mineral yield than RS. Hence, it is both energy- and time-efficient to utilize ACS for exploration in place of RS. We also expect the ACS strategy to be even more efficient in exploration when the cost of sampling/sensing increases. For example, the Mössbauer spectrometer runs at 2 W and needs 6 hours. In our experiments for ACS, the spectrometry incurs 35% of the total energy consumption and 73% of the overall mission time for a typical sample size of 92 units. These figures will increase substantially if the Mössbauer spectrometer is used instead.

Lastly, the robot team performance is compared with varying team sizes (i.e., 1, 4, and 8 robots) for the ACS strategy. Figs. 3d and e show the results after 20 test runs for the first and second criterion respectively. Fig. 3d shows that the team of 8 robots is less energy-efficient than 1 robot and the team of 4 robots; the larger team incurs a greater amount of physical interference during rendezvous. Fig. 3e shows that the reduction in mission time decreases with more robots; the teams of 4 and 8 robots achieve, respectively, 28.9% and 18.7% of the mission time taken by 1 robot, which are greater than the expected 25% and 12.5%. This is due to the competitive ratio of the tour allocation, which increases with the number of robots (Section II).

VII. Conclusion and Future Work

This paper describes the application of the ACS strategy and estimators to multi-robot wide-area exploration. They can exploit the clustering nature of the environmental phenomena (i.e., hotspots) and therefore perform better than SRS in such environments as shown in the analysis. Quantitative experimental results in the mineral prospecting task simulation show that the ACS strategy is most efficient in exploration by yielding more minerals and information with fewer resources, and the Rao-Blackwellized ACS estimators can provide more precise mineral density estimates than previous methods. In future work, we will apply these techniques on a larger robot team and real robots. Our planner will be improved using other minimum spanning tree heuristics or stochastic search strategies to reduce the tour lengths so that ACS can be even more efficient than RS. We will also consider the effect of noisy and multivariate mineral content data on the ACS strategy and estimators. Lastly, adaptive systematic and stratified sampling will be examined.

References

[1] W. Burgard, M. Moors, C. Stachniss, and F. E. Schneider, "Coordinated multi-robot exploration," *IEEE Trans. Robotics*, vol. 21, no. 3, pp. 376–386, 2005.

[2] M. Rahimi, R. Pon, W. J. Kaiser, G. S. Sukhatme, D. Estrin, and M. Srivastava, "Adaptive sampling for environmental robotics," in *Proc. IEEE ICRA*, 2004, pp. 3536–3544.

[3] S. K. Thompson, *Sampling*. NY: John Wiley & Sons, Inc., 2002.

[4] M. G. Lagoudakis, E. Markakis, D. Kempe, P. Keskinocak, A. Kleywegt, S. Koenig, C. Tovey, A. Meyerson, and S. Jain, "Auction-based multi-robot routing," in *Proc. Robotics: Science and Systems*, 2005, pp. 343–350.

[5] G. J. Taylor and L. M. V. Martel, "Lunar prospecting," *Adv. Space Res.*, vol. 31, no. 11, pp. 2403–2412, 2003.

[6] K. H. Low, G. J. Gordon, J. M. Dolan, and P. Khosla, "Adaptive sampling for multi-robot wide area prospecting," Robotics Institute, CMU, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-05-51, October 2005.