Proceedings of Australasian Conference on Robotics and Automation, 2-4 Dec 2013, University of New South Wales, Sydney Australia

# Large-Scale Near-Optimal Decentralised Information Gathering with Multiple Mobile Robots

**Timothy Patten, Robert Fitch** and **Salah Sukkarieh**
Australian Centre for Field Robotics
The University of Sydney, NSW 2006, Australia
{t.patten,rfitch,salah}@acfr.usyd.edu.au

## Abstract

Information gathering at large spatial scales can be addressed with teams of decentralised robots. Many existing methods search over a limited time horizon and do not provide strong performance guarantees. Near-optimal methods that exploit submodular objective functions have been proposed, given a fixed time budget. We propose a revised problem formulation that seeks to near-optimally maximise information gain quickly. We present a novel, near-optimal polynomial-time decentralised algorithm for multiple robots and analyse the expected path length with respect to the number of robots, the size of the area, and the number of observations. Our approach is based on area partitioning and is practically beneficial in that it allows for superlinear speedup in the time required to maximise the submodular objective function, is decentralised, and is easy to implement. We show extensive simulation results that compare the performance of our algorithm to existing sequential allocation methods.

## 1 Introduction

Decentralised information gathering with multiple mobile robots has broad and established significance in environmental monitoring applications [Singh *et al.*, 2010]. Because such applications often comprise time-varying phenomena over large geographic areas, it is important to consider methods appropriate for large spatial and small time scales. A team of mobile robots is advantageous in this case because multiple robots can perform observations in parallel at many locations. In order to harness this parallelism and ensure accurate global estimates, it is critical to understand the theoretical performance guarantees of decentralised information gathering algorithms.

One class of decentralised coordination algorithms is *myopic* in that robots act to maximise information gain over a limited time horizon [Gan *et al.*, 2012]. Unfortunately, myopic methods offer no general performance guarantees and can produce arbitrarily poor information utility in the general case. Recently, analysis of submodular functions has shown that a near-optimal set of observation locations can be computed greedily in polynomial time [Golovin and Krause, 2011]. These approximation results have led to considerable interest in their application to information gathering with mobile agents.

Although greedy approximation of submodular objective functions is powerful theoretically, critical challenges remain in the practical exploitation of these results for robotics. In the pioneering work of Nemhauser et. al. [Nemhauser *et al.*, 1978], observations are not constrained by physical travel time. Since this "omniscient" sensor model is not feasible for robots with limited sensing range, recent work has investigated the related problem of maximising information utility within a fixed resource budget. This budget can model the travel time required to make an observation. Singh et al. [Singh *et al.*, 2009a; 2009b] propose an algorithm to solve this problem in the multi-robot case based on the idea of *sequential allocation*, where the path of each robot is chosen in sequential fashion. Sequential allocation is a naturally *centralised* approach, since robots do not plan their paths simultaneously [Stranders *et al.*, 2010].

In this paper, we propose an alternative approach that is *decentralised*. We propose an algorithm and analysis for decentralised multi-robot information gathering with near-optimal information quality guarantees, polynomial computation time, and known travel time bounds.

Instead of specifying a fixed budget as input, we instead specify the number of desired observations as input. This problem variation is useful for several reasons. By removing the fixed time budget, it is possible to partition the workspace and efficiently distribute workload. This reduces the search space of individual robots and thus leads to a superlinear speedup in the computation time required to maximise the objective function. Even with lazy evaluation [Krause *et al.*, 2008], this maximisa-

tion operation can be prohibitively expensive given the large number of possible observation locations involved in gathering information over wide geographical areas. Practically, our approach provides a mechanism to reduce the overall problem size by increasing the number of robots.

Our approach also considers two properties not provided by the existing sequential allocation algorithms. First, we can consider non-zero "service time" in making an observation, which models the time spent in processing a sensor observation. Second, partitioning obviates the need for inter-robot collision avoidance during information gathering. Robots are spatially segregated via partitioning and thus are not subject to inter-robot collisions as they make observations.

In this work, we consider a class of submodular objective functions called *(r,γ)-local* [Singh *et al.*, 2009b]. This class is known to include useful objectives such as mutual information [Guestrin *et al.*, 2005]. After reviewing known performance guarantees, we present a new decentralised algorithm based on area partitioning. This algorithm is inspired by work in Dynamic Vehicle Routing (DVR) [Bullo *et al.*, 2011], which leads to our analysis of time bounds. Instead of fixing a time budget, we provide both *quality* and *time* guarantees.

Finally, we present extensive evaluation in simulation in two information gathering scenarios. This evaluation 1) empirically validates our theoretical claims, 2) shows the practical benefit of using multiple robots in large environments with up to 1000 observed locations, and 3) compares the performance of our method to sequential allocation.

## 2 Related Work

The problem of selecting a near-optimal subset of observations has been the studied in [Guestrin *et al.*, 2005] and [Krause *et al.*, 2006] for the purpose of monitoring a spatial phenomenon with wireless sensors. This work shows that selecting sensing locations based on maximising mutual information has strong theoretical guarantees due to the submodularity property [Nemhauser *et al.*, 1978]. These ideas were extended to informative path planning for mobile robots [Binney *et al.*, 2013] and multiple mobile robots [Singh *et al.*, 2007; 2009a; 2009b] where a greedy selection of observations combined with sequential allocation was proven to be near-optimal. In particular, an adaptive algorithm was presented in [Singh *et al.*, 2009b] where robots re-plan after every observation. Recent work by Hollinger et al. [Hollinger *et al.*, 2012] implemented these ideas for an autonomous underwater vehicle inspecting an underwater surface. Collectively, this results give guarantees in terms of the information gain due to their maximisation of a submodular function. However, they do not provide explicit time bounds for a given number of observations.

In the area of combinatorial optimisation, the study of the Vehicle Routing Problem has led to the development of time-optimal routing policies. The static vehicle routing problem is to determine the optimal set of routes to be performed by a fleet of vehicles to serve a given set of customers [Toth and Vigo, 2001]. The Dynamic Vehicle Routing Problem is an extension to this where the problem is to plan routes through demands that arrive during mission execution [Bullo *et al.*, 2011]. Early work by Bertsimas and van Ryzin [Bertsimas and Van Ryzin, 1991; 1993a] introduced what is now known as the Dynamic Traveling Repairman Problem (DTRP) or the Multiple-Vehicle Dynamic Traveling Repairman Problem (m-DTRP) [Bertsimas and Van Ryzin, 1993b] for the case of multiple servicing vehicles. For this problem, theoretical lower bounds are derived for the average time that any demand must wait in the system before being serviced [Bertsimas and Van Ryzin, 1993b; Xu, 1995]. Policies that achieve near optimal performance were presented in [Pavone *et al.*, 2011], along with several extensions to the general problem such as including time windows and customer constraints [Pavone *et al.*, 2009], team forming [Smith and Bullo, 2009] and priority classes [Smith *et al.*, 2010]. Our work applies analysis inspired by this area of literature in deriving time bounds for near-optimal information gathering.

## 3 Problem Statement

We consider the task of choosing informative observations for a team of $m$ mobile robots that cooperatively reduces the estimate uncertainty with respect to a particular environment state. Consider a bounded convex region $Q$ that is gridded into $N$ cells $\mathcal{C} \subset \mathbb{R}^2$ where each cell is a possible sensing location for a robot. We are interested in the situation where only a limited number of observations can be made, thus the task at hand is to select the most informative set of sensing locations of size $K < N$. The uncertainty at each cell location $c \in \mathcal{C}$ is represented by a random variable. When a robot visits a location it makes a noisy observation $z$ which updates the belief about the state of the world $X$. A prior probability distribution, $\phi(x)$, represents the uncertainty about the environment and we assume that $\int_Q \phi(x)dx = 1$. This could be generated from prior knowledge about the task at hand, for example, when searching for a lost target, the last known location could serve as an initial estimate. If no such information is available, then a uniform probability distribution is sufficient. Mutual information is used to quantify the reduction of uncertainty with each observation. Given $X$, the information gained by making a sequence of obser-

vations $Z_K = \{z_1, z_2, ..., z_K\}$ is defined as

$$F_X(Z_K) = I(X; Z_K) = H(X) - H(X|Z_K), \quad (1)$$

where $H$ is the conditional entropy. Therefore the problem of minimising the uncertainty about the state of the world is to determine the set of locations $Z_K$ which maximises the mutual information

$$\underset{Z_K \in \mathcal{Z}}{\operatorname{argmax}} F_X(Z_K), \quad (2)$$

where $\mathcal{Z}$ is the set of all possible collections of $K$ observations. Once the locations are known the problem is to then coordinate the group of robots along optimal routes such that the total time to visit and make the observations is minimised.

## 4 Information Gathering

The challenge of computing the optimal subset of observations is intractable for large problems because it requires an exhaustive search through all $\binom{n}{k}$ combinations. This is clearly infeasible for any real world application where there are constraints on processing resources and it is not ideal for on-line applications which need quick solutions. The problem that we investigate has been shown to satisfy the intuitive property of diminishing returns, also known as submodularity [Krause *et al.*, 2006], which states that adding an element $e$ to a smaller set $S$ has more advantage than adding that element to a larger set $T$. Formally, a set function $f$ is submodular if for any two subsets $S$ and $T$ and element $e$:

$$f(e \cup S) - f(S) \geq f(e \cup T) - f(T) \qquad \forall S \subseteq T. \quad (3)$$

In addition to this, a function $f$ is monotonic if

$$f(e \cup S) \geq f(S), \quad (4)$$

which is the "information never hurts" property. The result by Nemhauser et al. [Nemhauser *et al.*, 1978] states that greedily selecting elements to maximise a submodular function is within a factor $(1 - 1/e)$ of the optimal set of elements. That is, for a submodular function $f$, greedily choosing a set of elements $Z$ is guaranteed quality

$$f(Z) \geq \left(1 - \frac{1}{e}\right) f(Z^*), \quad (5)$$

where $Z^*$ is the optimal set of elements. Guestrin et al. [Guestrin *et al.*, 2005] showed that mutual information is a submodular and monotonic function. Thus, solving the problem of (2) by choosing each sensing location $z_i$ greedily

$$z_i = \underset{c \in \mathcal{C} \setminus Z_{i-1}}{\operatorname{argmax}} F_X(\{c\} \cup Z_{i-1}), \quad (6)$$

is guaranteed to achieve a near-optimal solution.

## 5 Decentralised Algorithm

Our decentralised algorithm for information gathering is presented in Alg. 1. The robots initially partition the environment into $m$ regions, each having a subset of locations to observe $\mathcal{C}_r$ for $r = 1, ..., m$ which can be done using a decentralised *CoverageControl* algorithm [Schwager *et al.*, 2007] that takes as input parameters the locations of the robots $L$ and the state of the environment $X$. An important modification is that the partitions are equalised according to the square root of the underlying density function as denoted by $X_{1/2}$. The motivation for this weighted partitioning is to balance the expected distance that each robot travels between observation locations, which has been shown to be proportional to the square root of the distribution of points along the path [Beardwood *et al.*, 1959]. Once the robots are in their partitions they can work simultaneously and independently. They each greedily select observations within their own region and compute a TSP tour through the locations as described in Alg. 2. TSP tours can be computed optimally for small problem sizes, and efficient approximation algorithms exist. After all robots have made their observations, standard decentralised data fusion techniques can be used to build a synchronised global estimate.

**Algorithm 1** Decentralised Information Gathering

**INPUT:** $L$, $K$, $X$

Robots configure into equal partitions
  $\mathcal{C}_{1:m} \leftarrow CoverageControl(X_{1/2}, L)$
Then each robot performs simultaneously:
  $P_r \leftarrow FindPath\left(\frac{K}{|L|}, \mathcal{C}_r, X\right)$
  $Traverse(P_r)$

**Algorithm 2** FindPath

**INPUT:** $K$, $\mathcal{C}$, $X$
**OUTPUT:** $P$ the path of way points for the robot

$Z_K = \emptyset$
**for** $i = 1 \rightarrow K$ **do**
  $e \leftarrow \operatorname{argmax}_{c \in \mathcal{C}} F_X(\{c\} \cup Z_K(i-1))$
  $Z_K(i) \leftarrow Z_K(i-1) \cup e$
  $\mathcal{C} \leftarrow \mathcal{C} \setminus e$
**end for**
$P \leftarrow TSP(Z_K)$
**return** $P$

The algorithm we have proposed has the following performance guarantees.

**Theorem 1** (Performance of Decentralised Algorithm)**.** *Suppose we have a bounded convex region $Q$ with area $A$ that is discretised to a resolution of $\delta$. Given a team*

*of $m$ robots moving within it at constant velocity $v$ and spending an average time of $s$ when making an observation, then the set of $K$ observations, $Z_K$, made by the robots implementing the decentralised algorithm will reduce the uncertainty of the environment state $X$ by an amount*

$$\frac{F_X(Z_K)}{F_X(Z_K^*)} \geq \frac{1}{m}\left(1-\frac{1}{e}\right), \tag{7}$$

*within a time of*

$$T \leq \frac{diam(Q)}{v} + \frac{K}{m}\left(\frac{\beta}{v\delta}\sqrt{\frac{A}{Km}} + s\right), \tag{8}$$

*where $Z_K^*$ is the optimal set of $K$ observations that could have been selected.*

*Proof.* The procedure for selecting observations in line 3 of Alg. 2 maximises the mutual information with every selection. Guestrin et al. [Guestrin *et al.*, 2005] state that the information that will be achieved by the greedy selection is guaranteed to be at least $(1-1/e)$ as informative as the optimal selection. That is,

$$F_X(Z_{Kglobal}) \geq \left(1-\frac{1}{e}\right)F_X(Z_K^*), \tag{9}$$

where $Z_{Kglobal}$ is the subset of locations selected from the entire environment $Q$. Now consider that the robots partition the environment into $m$ areas $q_1, q_2, ..., q_m$ where each area has a particular subset of sensing locations $\mathcal{C} = \{\mathcal{C}_1, \mathcal{C}_2, ..., \mathcal{C}_m\}$ according to line 2 in Alg. 1. Each robot makes $\gamma = K/m$ observations such that the set $Z_r = \{Z_1, Z_2, ..., Z_m\}$ is the collection of all the sensing locations that are chosen individually. Assuming that these are selected independently then the worst case performance is when the globally greedy best subset, $Z_{Kglobal}$, happens to be located in a single partition. In this partition the value of the $\gamma$ observations will be

$$F_X(Z_\gamma) = \frac{1}{m}F_X(Z_{Kglobal}). \tag{10}$$

Given that this partition has the highest value, the total information gain is bounded by

$$\begin{aligned} F_X(Z_K) &\geq \frac{1}{m}F_X(Z_{Kglobal}) \\ &\geq \frac{1}{m}\left(1-\frac{1}{e}\right)F_X(Z_K^*), \end{aligned} \tag{11}$$

which proves (7).

The time for a single robot to make a set of observations is

$$T_r = \text{travel time} + \text{observation time}. \tag{12}$$

Given a set of $n$ points, it was shown in [Beardwood *et al.*, 1959] that the expected length of the shortest closed path through the points, $\mathcal{L}(n)$, is

$$\lim_{n\to+\infty} n^{-(d-1)/d}\mathcal{L}(n) = \beta_{TSP,d}\int_Q \psi^{(d-1)/d}(x)dx, \tag{13}$$

where $d$ is the dimension of the environment and $\psi$ is the density of the point distribution. We assume the robots only travel in the $2D$ Euclidean plane, thus $d = 2$ . $\beta_{TSP,2}$ is a constant that has been estimated numerically as $0.7120 \pm 0.0002$ [Percus and Martin, 1996] and we let $\beta = \beta_{TSP,2}$ for simplicity. Each robot is in a partition $q_r$ and computes a TSP through its $\gamma$ observations. The length of this tour is bounded by (13) since it gives the shortest path and so the total time for a single robot to make $\gamma$ observations is

$$T_r \leq \frac{\text{diam}(Q)}{v} + \sqrt{\gamma}\beta\frac{\int_{q_r}\phi^{1/2}(x)dx}{v} + \gamma s. \tag{14}$$

We have upper bounded the travel time during the partitioning by $\text{diam}(Q)/v$ where $\text{diam}(\cdot)$ denotes the longest distance within a given bounded environment. The *CoverageControl* algorithm takes as input the square root of the density function $\phi^{1/2}(x)$ and equalises the partitions according to this such that $\int_{q_1}\phi^{1/2}(x)dx = \int_{q_2}\phi^{1/2}(x)dx = ... = \int_{q_m}\phi^{1/2}(x)dx$. This means that $\int_{q_r}\phi^{1/2}(x)dx = \int_Q \phi^{1/2}(x))dx/m$ and so the time can be re-written as

$$T_r \leq \frac{\text{diam}(Q)}{v} + \sqrt{\gamma}\beta\frac{\int_Q\phi^{1/2}(x)dx}{mv} + \gamma s. \tag{15}$$

From the theory of $L^p$ spaces and *p-norms* we have that

$$\| u \|_1 \leq \sqrt{n} \| u \|_2, \tag{16}$$

where $\| u \|_p = \left(\int_S | u |^p d\mu\right)^{1/p} < \infty$ for $1 \leq p < \infty$ and $u$ is a measurable function over a measurable space $(S, \Sigma, \mu)$. Using this property with $u = \phi^{1/2}(x)$ we have

$$\begin{aligned} \int_Q \phi^{1/2}(x)dx &\leq \sqrt{n}\left(\int_Q \left(\phi^{1/2}(x)\right)^2 dx\right) \\ &= \sqrt{n}\left(\int_Q \phi(x)dx\right) \\ &= \sqrt{n}, \end{aligned} \tag{17}$$

since $\int_Q \phi(x)dx = 1$ for the probability density function. The total time of operation for a single robot can be bounded by

$$T_r \leq \frac{\text{diam}(Q)}{v} + \sqrt{\gamma}\beta\frac{\sqrt{n}}{mv} + \gamma s, \tag{18}$$

where $n$ represents the number of potential sensing locations $| \mathcal{C}_r |$ in its region $q_r$. If the environment has an

area of $A$ and is discretised into cells with resolution $\delta$ we have

$$\begin{aligned} T_r &\leq \frac{\mathrm{diam}(Q)}{v} + \sqrt{\gamma}\beta\frac{\sqrt{A/\delta^2}}{mv} + \gamma s \\ &= \frac{\mathrm{diam}(Q)}{v} + \sqrt{\gamma}\beta\frac{\sqrt{A}}{mv\delta} + \gamma s. \end{aligned} \quad (19)$$

The robots work in parallel so the time for one robot is the same as the total time for $m$ robots. Thus the total time for all the robots to collect $K$ observations is

$$\begin{aligned} T &= T_r \\ &\leq \frac{\mathrm{diam}(Q)}{v} + \sqrt{\gamma}\beta\frac{\sqrt{A}}{mv\delta} + \gamma s \\ &= \frac{\mathrm{diam}(Q)}{v} + \frac{K}{m}\left(\frac{\beta}{v\delta}\sqrt{\frac{A}{Km}} + s\right). \end{aligned} \quad (20)$$

□

For the case where the partitioning and observation times are negligible, the mission time is approximately $T \lesssim \frac{1}{m^{3/2}}\sqrt{K}\frac{\beta\sqrt{A}}{v\delta}$, which is a reduction by a factor of $m^{3/2}$ of the time to travel through $K$ points in an environment of area $A$. The lower bound given for the information gain is considered near-optimal since it is a constant factor approximation of the optimal solution in terms of the number of observations. Although it is clear that the performance is reduced as the number of robots is increased, we show empirically in our results that the performance of the algorithm often exceeds this lower bound.

### 5.1 Analysis

In this section we analyse the computational complexity of our algorithm. The dominant computational step in the *CoverageControl* algorithm is the formation of the *Voronoi* cells. We use the implementation of Cortes et al. [Cortes *et al.*, 2004], which is a modified version of Lloyd's algorithm. This algorithm is known to run very fast in practice, and the worst case bound has recently been shown to be polynomial in the number of Voronoi cells (in our case, the number of robots $m$) [Arthur *et al.*, 2009]. In *FindPath*, the operation for solving (6) in line 3 searches through all $N$ cells and is done $K$ times, thus the for loop runs in $\mathcal{O}(NK)$ time. Line 7 invokes a TSP solver, but although the TSP problem is NP-hard, many polynomial-time algorithms that approximate it exist. For example, Christophide's algorithm [Christofides, 1976] runs in $\mathcal{O}(n^3)$ time, which for $K$ observations as in our case is $\mathcal{O}(K^3)$. Thus, the total time complexity is polynomial in $N$, $m$, and $K$.

We also show that increasing the number of robots leads to a decrease in total (parallel) computation time. First, the search space for each robot is reduced by a factor of $m$ since each robot only needs to perform $N/m$ operations within each iteration of the for loop in *FindPath*. Second, the work is performed in parallel, so instead of solving one problem for $K$ observations we solve $m$ simultaneous subproblems of size $K/m$. This leads to a (superlinear) speedup of $m^2$ time. In contrast, the running time of sequential allocation increases linearly in the number of robots [Singh *et al.*, 2009b].

### 5.2 Relation to Dynamic Vehicle Routing

One of the strategies for dealing with the challenges of the DVR problem is to divide the environment between the vehicles such that each agent is responsible for the work in one localised region. In [Pavone *et al.*, 2011] it was proved that if an optimal single vehicle routing policy is performed by each vehicle in each region then an equitable partition of the environment yields an optimal multi-vehicle routing policy. They also extended this to show that if the single vehicle policy is within a constant factor $\pi$ of optimal then an equitable partitioning yields a multi-vehicle routing policy that is also within $\pi$ of optimal. The problem we have presented here is similar in nature to DVR and it is an open question whether our method of partitioning can be extended to show the same time-optimal performance.

## 6 Results

Simulations were performed to verify the performance of our proposed algorithm. For an initial experiment we measured the running time for computing the best 1000 sensing locations for problems of different sizes. We compared the time to compute a centralised solution for the entire environment with the time taken by multiple robots to compute the solution in parallel. As shown in Fig. 1 the computation time for the centralised algorithm grows rapidly as the size of the world increases. We performed the simulation with 2 and 10 robots and even with 2 robots the performance is dramatically improved. This supports our motivation to divide the work between multiple distributed robots because it massively reduces the computation effort for any single robot.

### 6.1 Search

The first application to which we applied our algorithm is searching for a static target. We performed the search within a $100m \times 100m$ square grid for different probability density functions. An example is shown in Fig. 2. The sensor model in the search problem is a Gaussian function

$$S(x_0) = 1 - e^{-\frac{(x-x_0)^2}{2\sigma^2}}, \quad (21)$$

where $x_0$ is the location of the robot making the observation and we used $\sigma = 0.5$. The probability of

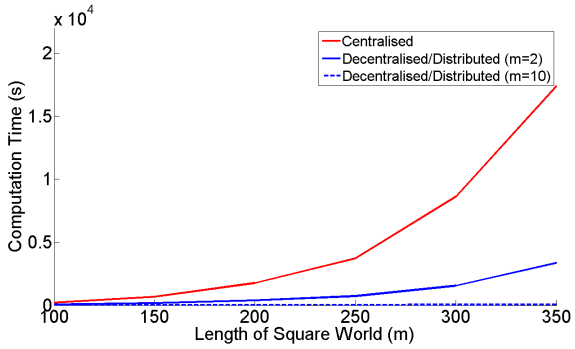


Figure 1: Running time of complete and distributed computation of 1000 sensing locations.

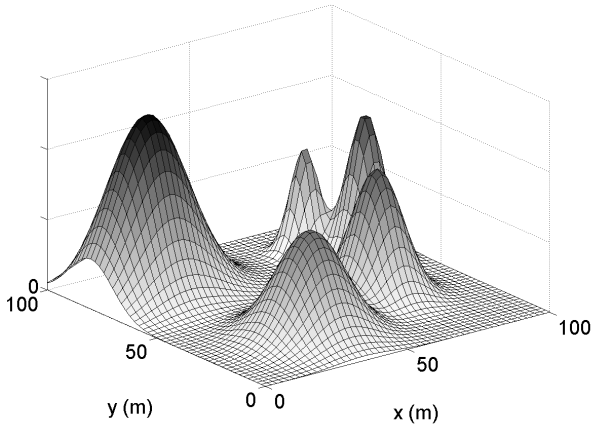


Figure 2: An example of a prior probability distribution of the target used in the search task.

detection of the target is a useful metric to quantify the performance of search. Initially, the uncertainty is at its maximum and so the probability of detection is 0. The relationship between the probability of detection and information gain is $ProbabilityOfDetection = 1 - InformationGain$. For all simulations the robots are homogeneous and move with a constant velocity of $0.5m/s$, and the time it takes for each observation is $1s$. An illustrative example of a search mission with three robots is shown in Fig. 3.

Figure 4 shows the total time for 100 observations to be made by robot teams of different sizes and the results are averaged over 10 different randomly generated probability density functions. For comparison we have plotted the time for sequential allocation [Singh *et al.*, 2009b] and a centralised algorithm which was implemented by choosing sensing locations from the entire environment, solving a $k$-TSP and then allocating tours to robots with a Hungarian algorithm. The plot shows that the decentralised solution obtains the fastest result as compared

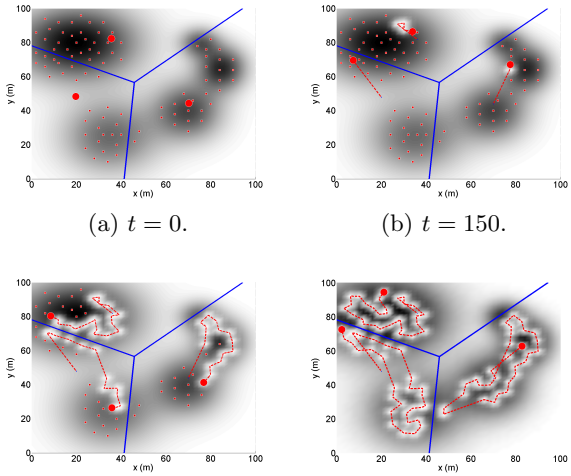


(a) $t = 0$. (b) $t = 150$.

(c) $t = 300$. (d) $t = 444$.

Figure 3: Simulation of the algorithm shown at four different time instances. The squares are the selected observations and the circles are the robots. The density function is shown in the background and the partitions are highlighted by the thick blue lines.

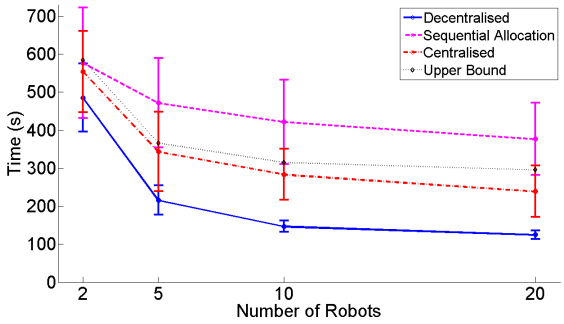


Figure 4: Time taken to make 100 observations, averaged over 10 probability density functions in the search task. The thin black line is the theoretical upper bound of our algorithm.

to the other two methods for all team sizes.

Figure 5a compares the performance of our decentralised algorithm with its theoretical lower bound on information gain, which in this case is probability of detection. The dashed line represents the lower bound claimed in Theorem 7. We report observed information gain normalised with respect to this bound. The observed results clearly respect this bound. Figure 5b plots the average rate of increase in information gain (probability of detection), computed as the final probability of detection divided by the time taken to make 100 observations.

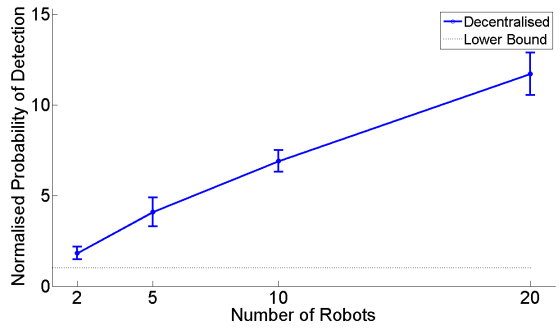


(a) Probability of detection, normalised with respect to our claimed lower bound.

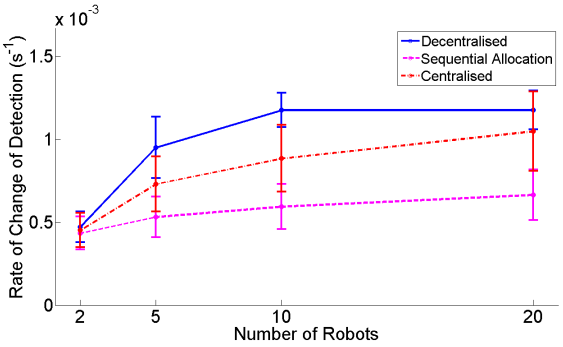


(b) Average rate of increase in information gain (probability of detection).

Figure 5: The average normalised probability of detection and average rate of increase in the probability of detection (probability/time) for 100 observations. The average is performed over 10 different probability density functions in the search task.

## 6.2 Estimating a spatial phenomenon

The second task to which we applied our algorithm is estimating a spatial phenomenon. This phenomenon is arbitrary and could represent tasks in environmental monitoring where samples are taken in order to estimate an environmental process. We used a Gaussian process to estimate the field to be sampled as well as to represent the uncertainty. An example of a spatial phenomenon is shown in Fig. 6.

Mutual information is computed from the covariance matrix using methods from [Singh *et al.*, 2010] and we quantify the performance by the error between the estimate and the true field. The simulations were performed for 100 observations with teams of different sizes, averaged over 10 randomly generated spatial fields. Figure 7a shows the reduction in error as compared to the theoretical lower bound (similar to Fig. 5a), and Fig. 7b shows the average rate of error reduction for our algorithm in comparison to the centralised and sequential allocation implementations. Larger error reduction indicates better

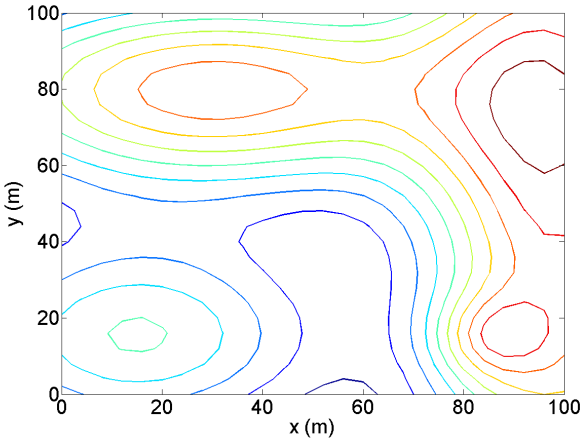


Figure 6: An example contour plot of a spatial phenomenon to be estimated by multiple robots.

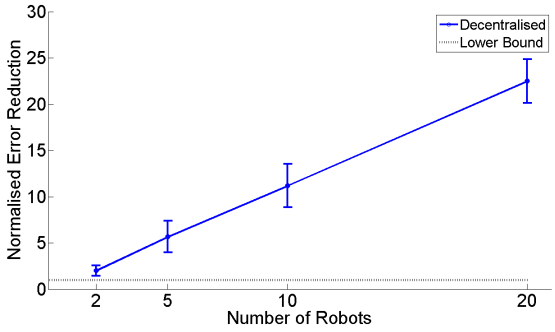


(a) Error reduction, normalised with respect to our claimed lower bound.

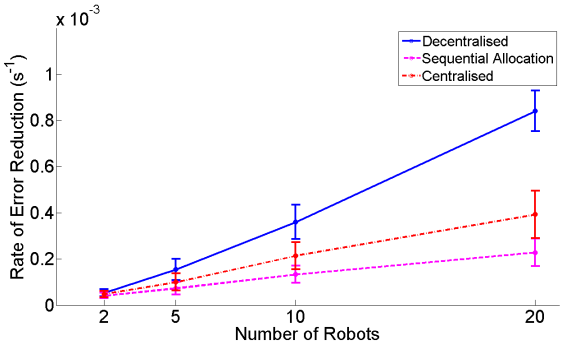


(b) Average rate of error reduction.

Figure 7: The average normalised error reduction and average rate of error reduction (error reduction/time) for 100 observations. The average is performed over 10 different spatial fields in the estimation task.

performance in terms of information gain. The results demonstrate that our decentralised algorithm respects its lower bound on information gain.

## 7 Conclusion

The problem of information gathering is ubiquitous in robotics and it has a variety of applications. As the size of the problem grows, however, the computation becomes intensive and current implementations which rely on a central processor cannot cope with the demands. In this paper we have presented a fully decentralised algorithm for information gathering by multiple robots. The decentralised nature allows the workload to be shared which dramatically reduces the load for any one agent. We perform a greedy search to select observations based on maximising the mutual information. This technique has been shown to satisfy the conditions for submodularity and so we proved that our method achieves near-optimal information gain. Furthermore, we analysed the time performance and gave an upper bound on the mission time. As a consequence, our implementation achieves near-optimal quality within a guaranteed amount of time.

Our technique was tested in simulation for the problems of searching for a static target and estimating a spatial phenomenon with a Gaussian process. The results showed empirically that our algorithm satisfies the claimed guarantees on mission time and information gain. We also compared our algorithm to a fully centralised implementation and a sequential allocation implementation. These experiments showed that our algorithm performed better in terms of information gain rate. One possible interpretation is that our algorithm results in shorter paths between observations, and thus makes more observations in a given time interval as compared to the other algorithms. The main advantage, however, is that for all instances of the problem our approach distributes the workload and so it can be applied to problems which may not be feasibly solved by existing methods.

Avenues for future work include extending our time bounds to consider robots moving in a 3D workspace. Another direction is to investigate information gathering tasks for non-submodular functions. In [Golovin and Krause, 2011] it was shown that there are still performance guarantees for problems maximising non-submodular functions at the expense of making a slightly larger selection of elements. We are currently investigating ways to incorporate this.

## Acknowledgements

This work was supported in part by the Australian Centre for Field Robotics and the New South Wales State Government.

## References

[Arthur *et al.*, 2009] D. Arthur, B. Manthey, and H. Röglin. k-Means has polynomial smoothed complexity. In *Proc. of IEEE FOCS*, pages 405–414, 2009.

[Beardwood *et al.*, 1959] J. Beardwood, J. H. Halton, and J. M. Hammersley. The shortest path through many points. *Math. Proc. Cambridge*, 55(4):299–327, 1959.

[Bertsimas and Van Ryzin, 1991] D. J. Bertsimas and G. Van Ryzin. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Oper. Res.*, 39(4):601–615, 1991.

[Bertsimas and Van Ryzin, 1993a] D. J. Bertsimas and G. Van Ryzin. Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles. *Oper. Res.*, 41(1):60–76, 1993.

[Bertsimas and Van Ryzin, 1993b] D. J. Bertsimas and G. Van Ryzin. Stochastic and dynamic vehicle routing with general demand and interarrival time distributions. *Adv. Appl. Probab.*, 25(4):947–978, 1993.

[Binney *et al.*, 2013] J. Binney, A. Krause, and G. S. Sukhatme. Optimizing waypoints for monitoring spatiotemporal phenomena. *Int. J. Rob. Res.*, 32(8):873–888, 2013.

[Bullo *et al.*, 2011] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith. Dynamic vehicle routing for robotic systems. *Proc. IEEE*, 99(9):1482–1504, 2011.

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

[Cortes *et al.*, 2004] J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *Trans. Robot. Automat.*, 20(2):243–255, 2004.

[Gan *et al.*, 2012] S. Gan, R. Fitch, and S. Sukkarieh. Real-time decentralized search with inter-agent collision avoidance. In *Proc. of IEEE ICRA*, pages 504–510, 2012.

[Golovin and Krause, 2011] D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *J. Artif. Intell. Res.*, 42(1):427–486, 2011.

[Guestrin *et al.*, 2005] C. Guestrin, A. Krause, and A. P. Singh. Near-optimal sensor placements in Gaussian processes. In *Proc. of ICML*, pages 265–272, 2005.

[Hollinger *et al.*, 2012] G. Hollinger, B. Englot, F. Hover, U. Mitra, and G. Sukhatme. Uncertainty-driven view planning for underwater inspection. In *Proc. of IEEE ICRA*, pages 4884–4891, 2012.

[Krause *et al.*, 2006] A. Krause, C. Guestrin, A. Gupta, and J. Kleinberg. Near-optimal sensor placements: maximizing information while minimizing communication cost. In *Proc. of ACM/IEEE IPSN*, pages 2–10, 2006.

[Krause *et al.*, 2008] A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *J. Mach. Learn. Res.*, 9:235–284, 2008.

[Nemhauser *et al.*, 1978] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions–I. *Math. Program.*, 14(1):265–294, 1978.

[Pavone *et al.*, 2009] M. Pavone, N. Bisnik, E. Frazzoli, and V. Isler. A stochastic and dynamic vehicle routing problem with time windows and customer impatience. *Mob. Netw. Appl.*, 14(3):350–364, 2009.

[Pavone *et al.*, 2011] M. Pavone, E. Frazzoli, and F. Bullo. Adaptive and distributed algorithms for vehicle routing in a stochastic and dynamic environment. *IEEE Trans. Automat. Contr.*, 56(6):1259–1274, 2011.

[Percus and Martin, 1996] G. Percus and O. C. Martin. Finite size and dimensional dependence in the Euclidean traveling salesman problem. *Phys. Rev. Lett.*, 76(8):1188–1191, 1996.

[Schwager *et al.*, 2007] M. Schwager, J. J. Slotine, and D. Rus. Decentralized, adaptive control for coverage with networked robots. In *Proc. of IEEE ICRA*, pages 3289–3294, 2007.

[Singh *et al.*, 2007] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batalin. Efficient planning of informative paths for multiple robots. In *Proc. of IJCAI*, pages 2204–2211, 2007.

[Singh *et al.*, 2009a] A. Singh, A. Krause, C. Guestrin, and W. Kaiser. Efficient informative sensing using multiple robots. *J. Artif. Int. Res.*, 34(1):707–755, 2009.

[Singh *et al.*, 2009b] A. Singh, A. Krause, and W. Kaiser. Nonmyopic adaptive informative path planning for multiple robots. In *Proc. of IJCAI*, pages 1843–1850, 2009.

[Singh *et al.*, 2010] A. Singh, F. Ramos, H. Durrant-Whyte, and W. Kaiser. Modeling and decision making in spatio-temporal processes for environmental surveillance. In *Proc. of IEEE ICRA*, pages 5490–5497, 2010.

[Smith and Bullo, 2009] S. L. Smith and F. Bullo. The dynamic team forming problem: Throughput and delay for unbiased policies. *Syst. Control Lett.*, 58(10-11):709–715, 2009.

[Smith *et al.*, 2010] S. L. Smith, M. Pavone, F. Bullo, and E. Frazzoli. Dynamic vehicle routing with priority classes of stochastic demands. *Siam J. Control Optim.*, 48(5):3224–3245, 2010.

[Stranders *et al.*, 2010] R. Stranders, F. M. Delle Fave, A. Rogers, and N. R. Jennings. A decentralised coordination algorithm for mobile sensors. In *Proc. of AAAI Conference on Artificial Intelligence*, pages 874–880, 2010.

[Toth and Vigo, 2001] P. Toth and D. Vigo, editors. *The vehicle routing problem*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2001.

[Xu, 1995] H. Xu. *Optimal Policies for Stochastic and Dynamic Vehicle Routing Problems*. PhD thesis, Massachusetts Institute of Technology, Department of Civil and Environmental Engineering, 1995.