Rational Swarms for Distributed On-line Bayesian Search

Alfredo Garcia, Chenyang Li and Fernan Pedraza

Department of System and Information Engineering, University of Virginia
{agarcia,chenyangli,fapedraza}@virginia.edu

Abstract

We present a novel scheme for distributed search in mobile sensors networks that is inspired by collective forms of intelligence present in many biological systems. Unlike the established paradigms of swarm intelligence, we posit a form of individual rationality governing each agent's decision. In the scheme proposed, a network of mobile sensors is tasked to find several targets over a search area. The sensing technology is imperfect so there are non-negligible probabilities for false positives and false negatives. Mobile sensors leave two data 'trails' across potential target locations that have been explored. One trail is associated with the frequency with which a given location has been probed while the other relates to the Bayes updated likelihood that a target is present. These trails are stored in a geographically distributed array of stationary motes. Each sensor processes the implicit information encapsulated in the two trails and chooses a decision that is aimed at maximizing the chance of detecting a target without unnecessary duplication in probing. By endowing mobile sensors with this simple optimization rule, we show that a form of 'rational swarm' intelligence emerges as sensors successfully coordinate indirectly (i.e. they achieve a one-to-one allocation of agents and targets) through active manipulation of the trails. This feature guarantees the proposed scheme is both reconfigurable and scalable.

Key words: Distributed control, sensor networks, Bayesian search, swarm intelligence.

Introduction

In this paper, we are interested in peer-to-peer network applications involving the search, identification, and location of targets across a search area that is often vast and possibly dangerous to navigate. Mobile robots and/or sensors employed for such applications provide promise to more rapidly and safely locate targets in different military, homeland security and/or disaster recovery scenarios. Advantages of these type of applications include reduced risk for human search and/or rescue teams and significantly enhanced search capabilities.

Specifically, we present a novel scheme for distributed search in mobile sensors networks that is inspired by collective forms of intelligence present in many biological systems typically referred to as "swarm intelligence" (see for instance, [2] and [3]). *Unlike* the established paradigms of swarm intelligence, we posit a form of individual rationality governing each agent's decision. Hence the term "rational" swarm. Under the proposed scheme a network of n mobile sensors is tasked to find n targets. The sensing technology is imperfect so there are non-negligible probabilities for false positives and false negatives. Mobile sensors leave two 'trails' across potential target locations that have been explored. One trail is associated with the frequency with which a given location has been probed while the other relates to the Bayes updated likelihood that a target is present. These trails are reminiscent of the pheromone trail used by ant colonies to find the shortest path between their nest and a food source. *Unlike* the established paradigms of swarm intelligence, agents process the implicit information encapsulated in the two trails and choose a decision that is aimed at maximizing the chance of detecting a target without unnecessary duplication in probing. By endowing mobile sensors with this simple optimization rule, we show that a form of 'rational swarm' intelligence emerges as sensors successfully coordinate indirectly (i.e. they achieve a one-to-one allocation of agents and targets) through active manipulation of the trails. This feature guarantees the proposed scheme is both *reconfigurable* and *scalable*. Reconfigurability follows from the fact that each sensor only need to know how often a given location has been probed by itself and by sensor(s) other than itself in the past (regardless of the identity of the other sensor(s) that executed the probes) and the updated Bayesian probabilities. Thus, sensors do not need to know the makeup of the group so new sensors can enter the network and others can exit. Scalability follows from the fact that bilateral communication amongst sensors is not required. Instead, agents must be able to access the values of the

ROBOCOMM 2007, 14th – 16th Oct 2007, Athens, Greece.

ISBN 978-963-9799-08-0
DOI 10.4108/ICST.ROBOCOMM2007.2181

21 August 2007

two trails. This can be achieved by having a geographically distributed array of stationary motes in charge of keeping track of these values.

The structure of this paper is as follows: in section 1 we provide a brief literature review on distributed control schemes for multi-agent systems that are inspired by natural phenomena. In section 2, our proposed scheme for on-line distributed Bayesian search is formalized. In section 3, the convergence with probability one to an efficient allocation of sensors to targets is proven. Numerical tests are conducted in Section 4 to test scalability, the speed of convergence vis-a-vis a centralized search scheme and the ability to track moving targets. In section 5 an implementation of the algorithm on a physical testbed is demonstrated. Finally, in the paper's last section, we offer some concluding remarks with an emphasis on future research tasks.

1 Literature Review

The design and implementation of distributed control schemes for multi-agent systems that are inspired by natural phenomena (e.g. swarming, flocking, herding etc.) has been the subject of intense activity by various research communities including artificial intelligence, control theory and operations research (see for instance, [2] and [3]). Coordination in multiple robot systems poses new fundamental challenges not present in single-robot applications (see [4] and the references therein). In this section, we shall not attempt to provide a comprehensive literature review. Instead, we briefly discuss some of the relevant research publications relating to relevant challenges in the context of this paper:

(1) *Communication*: Sharing information in multi-robot systems is particularly critical for ensuring successful coordination. Nonetheless, difficulties associated with limited bandwidth, rapidly changing topologies and link failures, call for more sophisticated routing and flow control protocols (see [15] and [17]). In more recent work (see [8] and [13]), ant colony based routing algorithms for time-varying operating conditions have been proposed and tested with encouraging results.

(2) *Localization, mapping and exploration*: Multi-robot systems have enabled new capabilities in this area. For example, a multi robot system can take measurements from distant robots to increase positioning precision and/or allow sensor calibration. Teams of robots that cooperate can speed up field exploration by exchanging captured data and/or optimize sensing coverage. Notable recent contributions in this area include [7], [18] and [5].

(3) *Motion coordination*: Path planning in the context of multi-robot systems must address issues such as collision avoidance (see [9] and [16]), the generation and maintenance of robot formations and coordinated target tracking, where the optimization of each robot's motion path is aimed at reducing the search time (see [6], [1], [11] and [20]).

(4) *Architecture, task allocation and control*: The platform architecture provides the structure upon which collective behavior is implemented. As such, it determines most of the system capabilities and limitations. Task allocation and control relates to delegation of responsibilities and goal distribution within both homogeneous or heterogeneous multi-robot teams, and how these affect performance on the entire system (see [4], [14] and [12]).

2 Distributed Search by a Swarm of Sensors

Consider a situation where n mobile sensors are tasked to find n targets in a grid (without obstacles) which we shall denote by X. We shall denote by X^* the set of feasible locations for the n targets, i.e. $X^* = \{x_1^*, x_2^*, \ldots, x_n^*\}$, $x_i^* \in X$. Let $\mu^0(x)$ be the *a priori* probability that a target is located at $x \in X$. Let α and β denote the conditional probability of obtaining "false positives" and "false negatives", respectively. Let $Z(x)$ be the probe result for a given location $x \in X$, where $Z(x) = 1$ if a "positive" result is obtained and $Z(x) = 0$, otherwise. The Bayesian update [1] is:

$$\mu^1(x) = \begin{cases} \frac{(1-\beta)\mu^0(x)}{\alpha(1-\mu^0(x))+(1-\beta)\mu^0(x)} & Z(x) = 1 \\ \\ \frac{\beta\mu^0(x)}{(1-\alpha)(1-\mu^0(x))+\beta\mu^0(x)} & Z(x) = 0 \end{cases} \tag{1}$$

Let $\lambda^t(x)$ denote a measure of probing activity in location $x \in X$. With knowledge of $\lambda^t(x)$ and $\mu^t(x)$ each sensor i (currently located at $s_i^t \in X$) is to decide where to probe next. Let $N(s_i^t)$ denote the set of locations in X that are reachable from s_i^t within one iteration. We assume $s_i^t \in N(s_i^t)$, meaning a sensor may remain at its current location. Also, $\lambda_{-i}^t(x)$ is a measure of probing activity in location $x \in X$ by sensors other than i. In determining his/her next probe location s_i^{t+1}, robot i solves the following simple optimization problem

$$s_i^{t+1} \in B_i^t = \arg\max_{x \in N(s_i^t)} [\mu^t(x)(\lambda_i^t(x) - \lambda_{-i}^t(x))] \tag{2}$$

where the choice of s_i^{t+1} is random in the case B_i^t is not a singleton. Note that $\lambda_{-i}^t(x)$ can be interpreted as the "pheromone trail" left by other sensors, i.e. a simple form of *s tigmergy* (see [19]). In this context, robots' collective intelligence or *swarm intelligence* is encapsulated

[1] The updating formula for more than one probe for a given location at the same time is ommitted for space considerations.

in the local optimization rule in (2) according to which a robot's next choice of location maximizes the chance of finding a target *without duplication.* After the choice of location is made, robots proceed to sample the location and the updated probability distribution μ^{t+1} is constructed as described in (1) above and the frequency trail $\lambda^t = (\lambda_i^t, \lambda_{-i}^t)$ is updated as follows:

$$\begin{aligned}\lambda_i^{t+1}(x) &= \lambda_i^t(x) + \rho(\mathbf{1}_{\{s_i^{t+1}=x\}} - \lambda_i^t(x)) \\ \lambda_{-i}^{t+1}(x) &= \lambda_{-i}^t(x) + \rho(\mathbf{1}_{\{s_{-i}^{t+1}=x\}} - \lambda_{-i}^t(x))\end{aligned} \tag{3}$$

where $\rho \in (0,1)$ and $\mathbf{1}_{\{s_i^{t+1}=x\}}$ and $\mathbf{1}_{\{s_{-i}^{t+1}=x\}}$ are the indicator random variables of the event $\{s_i^{t+1} = x\}$ and $\{s_j^{t+1} = x, j \neq i\}$.

The basic structure of the algorithm sketched is illustrated in Figure 1. The network architecture implicitly represented in Figure 1 is reconfigurable because agents only need to know how often a given location has been probed in the past (i.e. $\lambda^t(x)$, $x \in X$) and the updated Bayesian beliefs (i.e. $\mu^t(x)$, $x \in X$) *regardless of the identity of the sensor(s) that executed the probes.* Thus, sensors do not need to know the makeup of the group so new sensors can enter the network and others can exit (e.g., in a low fuel state), providing a flexible and adaptable network for dynamic mission objectives. Note also that while bilateral communication is not required, agents must be able to access at all times the repository of the "state" of the system (i.e. the values of $\lambda^t(x)$ and $\mu^t(x)$, $x \in X$). This can be achieved by having a geographically distributed array of stationary motes in charge of keeping track of $\lambda^t(x)$ and $\mu^t(x)$, for locations $x \in X$, within their communication radius.

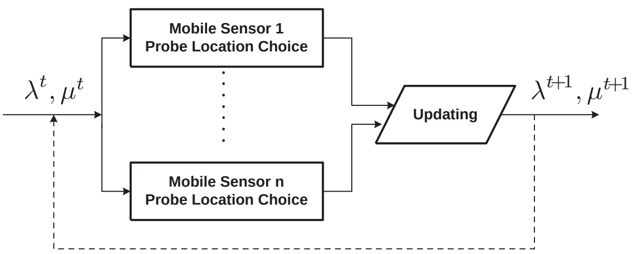

Fig. 1. Schematic for Distributed Search.

3 Convergence

3.1 Complete Information

In this section we study the convergence of the scheme when target locations are known (i.e. $\mu^t(x) = 1$ if $x \in X^*$ and $\mu^t(x) = 0$, otherwise). We shall also assume that robots can move to any location within one time period, i.e. $N(s_i^t) = X$ for all locations $s_i^t, t > 0$. Our first Lemma characterizes the dynamics of the set B_i^t.

Lemma 1: $B_i^{t+1} \subseteq B_i^t$

Proof: Let $x \notin B_i^t$. It follows that

$$\begin{aligned}\lambda_i^{t+1}(x) - \lambda_{-i}^{t+1}(x) &= (1-\rho)[\lambda_i^t(x) - \lambda_{-i}^t(x)] - \rho\mathbf{1}_{\{s_{-i}^{t+1}=x\}} \\ &\leq (1-\rho)[\lambda_i^t(x) - \lambda_{-i}^t(x)]\end{aligned} \tag{4}$$

For $y = s_i^{t+1} \in B_i^t$ we have

$$\lambda_i^{t+1}(y) - \lambda_{-i}^{t+1}(y) = (1-\rho)[\lambda_i^t(y) - \lambda_{-i}^t(y)] + \rho(1 - \mathbf{1}_{\{s_{-i}^{t+1}=y\}})$$

Hence,

$$\lambda_i^{t+1}(y) - \lambda_{-i}^{t+1}(y) \geq (1-\rho)[\lambda_i^t(y) - \lambda_{-i}^t(y)] \tag{5}$$

From (4) and (5) it follows that

$$\begin{aligned}&\lambda_i^{t+1}(y) - \lambda_{-i}^{t+1}(y) - [\lambda_i^{t+1}(x) - \lambda_{-i}^{t+1}(x)] \\ &\geq (1-\rho)[\lambda_i^t(y) - \lambda_{-i}^t(y)] - (1-\rho)[\lambda_i^t(x) - \lambda_{-i}^t(x)] \\ &> 0\end{aligned}$$

This implies $x \notin B_i^{t+1}$. ■

Although Lemma 1 clearly suggests each agent's best reply set converges to a singleton, it doesn't necessarily guarantee the outcome to be a one-to-one allocation of agents to targets. Actually, an example of inefficient allocation (more than 1 agent end up with the same target) can be found without much difficulty.

To deal with this problem, we now introduce a slight variation to the scheme where agents make small "errors" when choosing their best location to probe. In words, at any point in time an agent has a "small" probability $p > 0$ of choosing a location that is ϵ-optimal (i.e. the value of the objective function is within ϵ of the optimal value, for $\epsilon \in (0,1)$). This variation is formalized by introducing the following notation:

$$\tilde{B}_i^t = \begin{cases} B_i^t & \text{with probability } 1-p \\ \\ B_i^t(\epsilon) & \text{with probability } p \end{cases} \tag{6}$$

In what follows, we study the convergence of the sequence of choice locations induced by the random rule formalized in (6). In the lemma below, we show that an agent converges upon a target location whenever he/she has been "almost" the only one to visit it (i.e. other agents have rarely visited that same location).

Lemma 2: *If* $\lambda_i^t(x) - \lambda_{-i}^t(x) > 1 - \delta$ *for some* $t < \infty$ *and* $x \in X^*, \delta < \frac{1-\epsilon}{2}$ *then*

$$\lim_{t\to\infty}[\lambda_i^t(x) - \lambda_{-i}^t(x)] = 1$$

Proof: By hypothesis, $\lambda_i^t(x) > 1-\delta$ which implies $\lambda_i^t(y) < 1-(1-\delta) = \delta$ for $y \in X\setminus\{x\}$. Thus,

$$\lambda_i^t(x) - \lambda_{-i}^t(x) - [\lambda_i^t(y) - \lambda_{-i}^t(y)] > 1-2\delta > \epsilon$$

Hence, $B_i^t(\epsilon) = B_i^t = \{x\}$. Also, $\lambda_{-i}^t(x) < \delta$ and $\lambda_i^t(x) > 1-\delta$ respectively imply that $\lambda_j^t(x) < \delta$ and $\lambda_{-j}^t(x) > 1-\delta$ for all $j \in \{1,2,\ldots,n\}\setminus\{i\}$. Moreover, there exists a choice of location, say $y \in X\setminus\{x\}$, for which $\lambda_j^t(y) - \lambda_{-j}^t(y) \geq 0$. Hence,

$$\lambda_j^t(x) - \lambda_{-j}^t(x) < 2\delta - 1 < -\epsilon < 0 \leq \lambda_j^t(y) - \lambda_{-j}^t(y)$$

Or equivalently, $x \notin B_j^t(\epsilon) \cup B_j^t$. Thus,

$$\lambda_i^{t+1}(x) - \lambda_{-i}^{t+1}(x) > \lambda_i^t(x) - \lambda_{-i}^t(x) > 1-\delta$$

and by a similar argument as above, $s_i^{t+2} = x$ and $s_j^{t+2} \neq x$ for $j \in \{1,2,\ldots,n\}\setminus\{i\}$. The argument follows by induction. ■

Before we proceed with the main result in this section (i.e. Theorem 1), we need a final preliminary result:

Lemma 3: *Assume $\rho > \frac{\epsilon}{1+\epsilon}$. If $s_i^{t+1} = x \in B_i^t(\epsilon)$ and $s_{-i}^{t+1} \neq x$ then $B_i^{t+1} = \{x\}$ and $x \notin B_j^{t+1}$ for $j \neq i$.*

Proof: Note that

$$\begin{aligned}\lambda_i^{t+1}(x) - \lambda_{-i}^{t+1}(x) &= (1-\rho)(\lambda_i^t(x) - \lambda_{-i}^t(x)) + \rho \\ &\geq (1-\rho)(\lambda_i^t(x^*) - \lambda_{-i}^t(x^*) - \epsilon) + \rho\end{aligned}$$

where $x^* \in B_i^t$. Since $\rho > (1-\rho)\epsilon$, it follows that

$$\begin{aligned}\lambda_i^{t+1}(x) - \lambda_{-i}^{t+1}(x) &> (1-\rho)(\lambda_i^t(x^*) - \lambda_{-i}^t(x^*)) \\ &\geq (1-\rho)(\lambda_i^t(y) - \lambda_{-i}^t(y)) \\ &\geq \lambda_i^{t+1}(y) - \lambda_{-i}^{t+1}(y)\end{aligned}$$

for all $y \in X\setminus\{x\}$. Let $y \in B_j^t$. It follows that

$$\begin{aligned}\lambda_j^{t+1}(x) - \lambda_{-j}^{t+1}(x) &= (1-\rho)(\lambda_j^t(x) - \lambda_{-j}^t(x)) - \rho \\ &< (1-\rho)(\lambda_j^t(y) - \lambda_{-j}^t(y)) - \rho \\ &\leq \lambda_j^{t+1}(y) - \lambda_{-j}^{t+1}(y)\end{aligned}$$

Thus, $x \notin B_j^{t+1}$. ■

In the main result in this section, we show how non-convergence or even limit cycles can not occur as the probability of reaching the "absorbing" condition in Lemma 2 has a uniform positive lower bound.

Theorem 1: *Assume $\rho > \frac{\epsilon}{1+\epsilon}$. For every $x \in X^*$ there exists $i \in \{1,2,\ldots,n\}$ such that*

$$\lim_{t\to\infty}[\lambda_i^t(x) - \lambda_{-i}^t(x)] = 1 \text{ w.p. } 1$$

Proof: Let $x \in X^*$ and suppose there exists $i \in \{1,2,\ldots,n\}$ such that

$$s_i^{t+1} = x \text{ and } s_{-i}^{t+1} \neq x \;\; \textit{infinitely often.}$$

It follows that there exists a subsequence (indexed by $t(k)$) along which $s_i^{t(k)+1} = x$ and $s_{-i}^{t(k)+1} \neq x$. By Lemma 3, $B_i^{t(k)+1} = \{x\}$ and $x \notin B_j^{t(k)+1}$, $j \neq i$. For $\tau \geq 1$ with probability $[(1-p)^n]^\tau$:

$$\begin{aligned}&\lambda_i^{t(k)+\ell+1}(x) - \lambda_{-i}^{t(k)+\ell+1}(x) \\ &> \lambda_i^{t(k)+\ell}(x) - \lambda_{-i}^{t(k)+\ell}(x)\end{aligned}$$

for $\ell \in \{0,1,\ldots,\tau\}$. It follows that with probability $[(1-p)^n]^{\tau^*}$ where $\tau^* = \left\lceil \frac{1-\delta}{\rho} \right\rceil$, we would have

$$\lambda_i^{t+\tau^*+1}(x) - \lambda_{-i}^{t+\tau^*+1}(x) > 1-\delta$$

Thus, by Lemma 2, $\lambda_i^t(x) - \lambda_{-i}^t(x) \to 1$ with probability 1. Suppose now that for all $i \in \{1,2,\ldots,n\}$,

$$s_i^{t+1} \neq x \text{ or } s_i^{t+1} = s_{-i}^{t+1} = x \;\; \textit{eventually.}$$

Note that if $x \notin B_i^t(\epsilon)$ then $x \notin B_i^{t+1}(\epsilon)$. Thus, either $s_i^{t+1} \neq x$ *eventually* (in which case, $\lambda_i^t(x) \to 0$) or $s_i^{t+1} = s_{-i}^{t+1} = x$ *eventually* (in which case, $\lambda_i^t(x) \to 1$ and $\lambda_i^t(x) - \lambda_{-i}^t(x) \to 0$). This implies that all agents converge to a certain location. If $\lim_{t\to\infty}\lambda_i^t(x) = \lim_{t\to\infty}\lambda_{-i}^t(x) = 1$ we have that at most $n-2$ agents are converging on $n-1$ remaining locations, thus there exists a choice of location $y \in X^*$ such that $\lim_{t\to\infty}\lambda_{-i}^t(y) = 0$ and $T(\epsilon)$ such that for all $t > T(\epsilon)$

$$\max\{\lambda_i^t(x) - \lambda_{-i}^t(x); \lambda_i^t(y) - \lambda_{-i}^t(y)\} < \frac{\epsilon}{2}$$

This implies $y \in B_i^t(\epsilon)$ for $t > T(\epsilon)$. Choosing x over y infinitely often is a event with zero probability. ■

3.2 Incomplete Information

In this section we study the convergence of the scheme when target locations are not known. The different probes by mobile sensors results in an updated probability distribution μ^t. Suppose location $x \in X^*$ (respectively, $x \in X\setminus X^*$) is probed infinitely often. Provided

$1-\beta \neq \alpha$, it is well known (see [10]) that $\mu^t(x) \to 1$ with probability one (respectively, $\mu^t(x) \to 0$).

Lemma 4: *Assuming incomplete information,*

a. $P\left(s_i^{t+1} = s_{-i}^{t+1} = x \;\; \textit{eventually}\right) = 0$ *for* $x \in X^*$ *and* $i \in \{1, 2, ..., n\}$.
b. $P\left(\{s_i^{t+1} = x \; \textit{eventually}\}\right) = 0$ *for* $x \in X \backslash X^*$ *and* $i \in \{1, 2, ..., n\}$.

Proof: Part (a). $s_i^{t+1} = s_{-i}^{t+1} = x$ *eventually* implies $\lim_{t\to\infty} \lambda_i^t(x) = \lim_{t\to\infty} \lambda_{-i}^t(x) = 1$ and

$$\lim_{t\to\infty} [\mu^t(x)\left(\lambda_i^t(x) - \lambda_{-i}^t(x)\right)] = 0.$$

Let $y \in X \backslash X^*$. If y is probed infinitely often then $\mu^t(y) \to 0$. If y is not probed infinitely then $\lambda^t(y) \to 0$. Thus,

$$\lim_{t\to\infty} [\mu^t(y)\left(\lambda_i^t(y) - \lambda_{-i}^t(y)\right)] = 0$$

It follows that there exists $T(\epsilon)$ such that for all $t > T(\epsilon)$, $y \in B_i^t(\epsilon)$. Therefore,

$$\begin{aligned} 0 &\le P\left(s_i^{t+1} = s_{-i}^{t+1} = x \; \textit{eventually}\right) \\ &\le P\left(s_i^{t+1} = x \; \textit{infinitely often}\right) \\ &= 0 \end{aligned}$$

Part (b). Let $y \in X \backslash X^*$ and $y \neq x$. As argued above $\lim_{t\to\infty}[\mu^t(y)\left(\lambda_i^t(y) - \lambda_{-i}^t(y)\right)] = 0$, thus there exists $T(\epsilon)$ such that for all $t > T(\epsilon)$, $y \in B_i^t(\epsilon)$. Therefore,

$$\begin{aligned} 0 &\le P\left(s_i^{t+1} = x \; \textit{eventually}\right) \\ &\le P\left(s_i^{t+1} = x \; \textit{infinitely often}\right) \\ &= 0. \end{aligned}$$ ■

The convergence proof for the scheme under incomplete information builds up on the proof of Theorem 1.

Theorem 2: *Assume* $\rho > \frac{\varepsilon}{1+\varepsilon}$ *and incomplete information. For every* $x \in X^*$, *there exists* $i \in \{1, 2, ..., n\}$ *such that*

$$\lim_{t\to\infty} [\mu^t(x)(\lambda_i^t(x) - \lambda_{-i}^t(x))] = 1 \;\; \text{w.p. } 1$$

Proof: Let $x \in X^*$. By Lemma 4(a) and 4(b), x will be probed infinitely often and $\mu^t(x) \to 1$ with probability one. By Lemma 4(a) there exists $i \in \{1, 2, \dots, n\}$ such that

$$s_i^{t+1} = x \text{ and } s_{-i}^{t+1} \neq x \;\; \textit{infinitely often.}$$

It follows that there exists a subsequence (indexed by $t(k)$) along which $s_i^{t(k)+1} = x$ and $s_{-i}^{t(k)+1} \neq x$. By a similar argument to Lemma 3 above, $B_i^{t(k)+1} = \{x\}$ and $x \notin B_j^{t(k)+1}$, $j \neq i$

with probability $1-\beta$. For $\tau \ge 1$ with probability $[(1-\beta)(1-p)]^{N\tau}$:

$$\begin{aligned} &\mu^{t(k)+\ell+1}(x)(\lambda_i^{t(k)+\ell+1}(x) - \lambda_{-i}^{t(k)+\ell+1}(x)) \\ &> \mu^{t(k)+\ell}(x)(\lambda_i^{t(k)+\ell}(x) - \lambda_{-i}^{t(k)+\ell}(x)) \end{aligned}$$

for $\ell \in \{0, 1, \dots, \tau\}$. It follows that for some finite τ^* we would have

$$\mu^{t+\tau^*+1}(x)((\lambda_i^{t+\tau^*+1}(x) - \lambda_{-i}^{t+\tau^*+1}(x)) > 1-\delta$$

with probability $[(1-\beta)(1-p)]^{N\tau^*} > 0$. Also, $\mu^t(x) \to 1$ implies the existence of $T(\delta)$ such for fixed t and $t' > \max\{T(\delta), t+\tau^*+1\}$

$$\mu^{t'}(x)((\lambda_i^{t+\tau^*+1}(x) - \lambda_{-i}^{t+\tau^*+1}(x)) > 1-\delta$$

By a similar argument to Lemma 2, we can show that $B_i^{t'}(\epsilon) = B_i^{t'} = \{x\}$ and $x \notin B_{-i}^{t'}(\epsilon) \cup B_i^{t'}$. ■

Now we have proved in both complete and incomplete scenarios, our proposed rational swarm searching scheme will achieve a one-to-one-agent-target allocation with probability one given equally many agents and targets. Obviously, the results of Theorem 1 and Theorem 2 still hold when we have more agents than targets. In the case where there are less agents, some targets will remain uncovered. However this issue can be sorted out by giving some small modification to the algorithm. For example, an agent can be released to search other area when the posterior probability μ of the current spot reaches a certain threshold or after the target is successfully removed by the agent.

4 Numerical Experiments

To illustrate, consider a situation where $n = 5$ mobile sensors (whose initial location is represented by the symbol □) are tasked to find n targets (whose stationary location is represented by the symbol ■) in a $n \times n$ area depicted below. We further assume there are no obstacles in the area and that within one iteration a sensor may move to any *adjacent* location or remain at its current location.

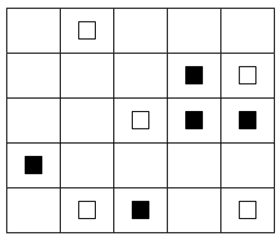

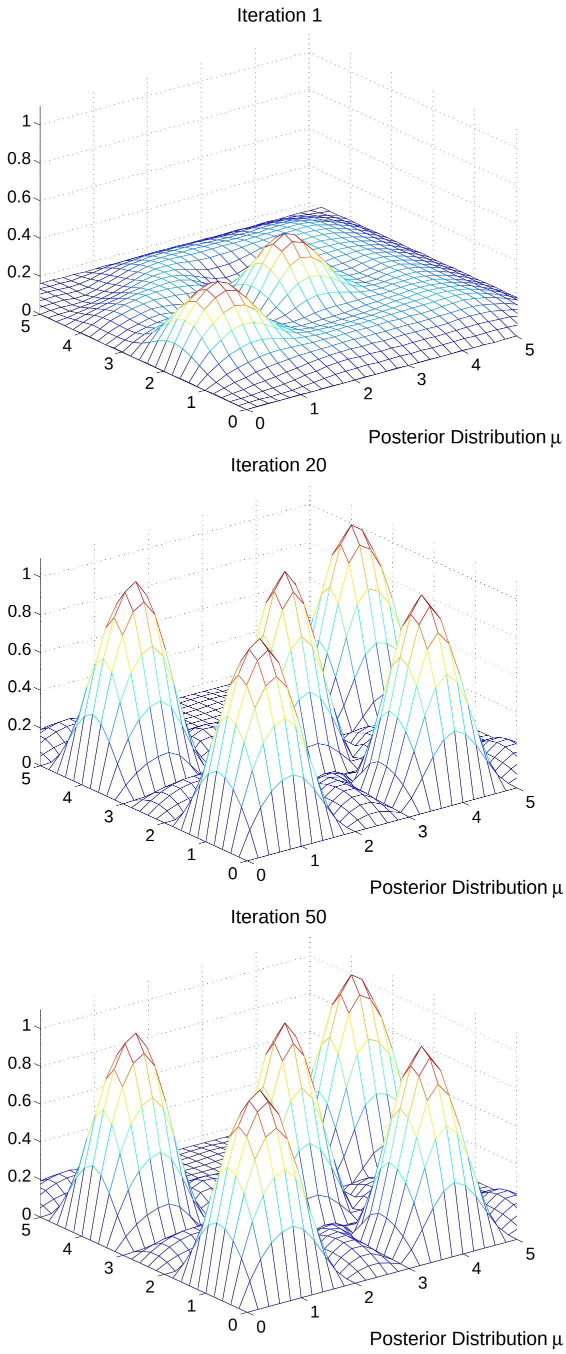

Fig. 2. Evolution of μ^t.

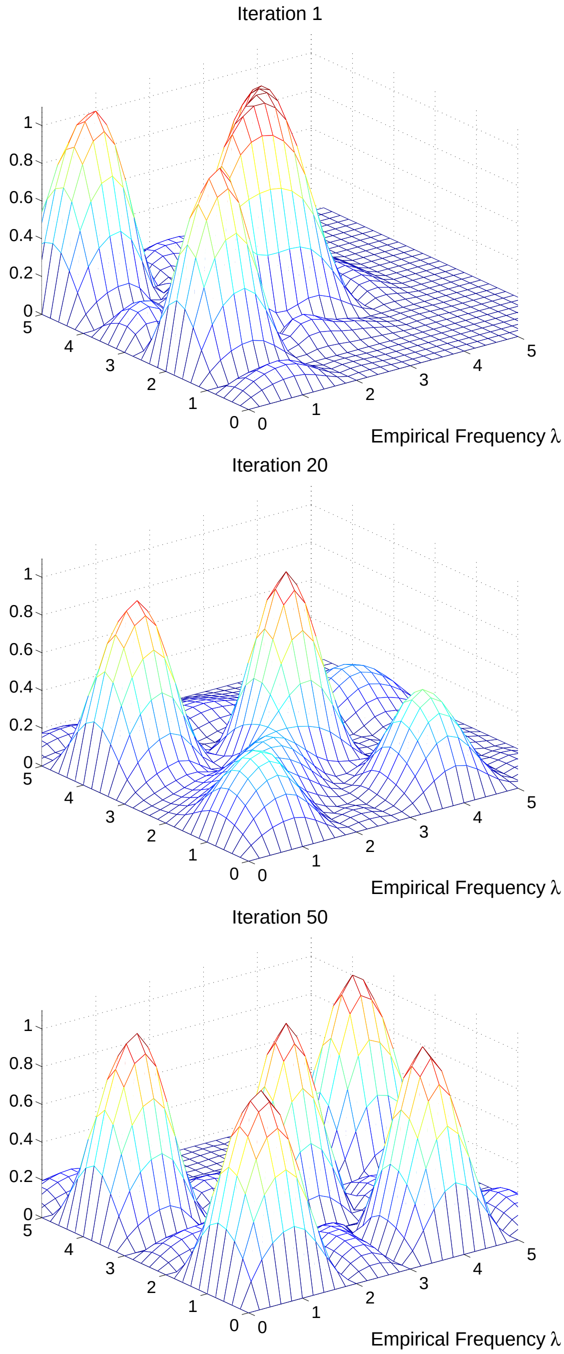

Fig. 3. Evolution of λ^t.

Figure 2 and 3 show the evolution of the posterior distribution μ and the empirical frequency λ as the agents are probing the area and $\alpha = 0.15$ and $\beta = 0.10$. Note how "peaks" on probing activity emerge right in the exact target locations.

We tested how well the scheme scales up. In Figure 4, we show the results of applying such modification to a set of problems for an increasing number of targets/agents, $N = \{5, 10, 20\}$, and where the size of the grid increases squarely in the number of agents, i.e. $G = \{5^2, 10^2, 20^2\}$. Note the time to full coverage degrades linearly in scale.

Suppose now that the target location changes in an unpredictable fashion. In Figure 5, we present the results where we have 20 targets in a 20×20 area and *target locations change randomly every 200 iterations* and the number of agents exceeds the number of targets by 20 percent.

Finally, we compare the performance of our distributed search algorithm with a *centralized* search scheme, where

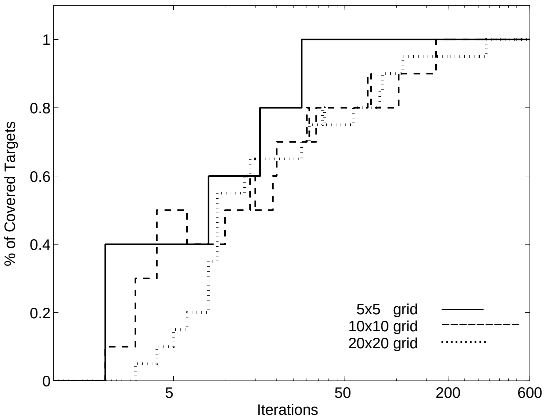

Fig. 4. Test for scalability.

for each iteration t, there is a "virtual" base station that instructs each agent on the best next location to probe. Note that such in such a centralized scheme, unnecessary duplication in search efforts is avoided. We tested the two algorithms with a total of $5, 10, 20$ targets, $6, 12, 24$ mobile agents, and area size 5^2, 10^2, 20^2 respectively. For each scenario, we repeated the experiment 200 times with randomized targets locations each time. The results are shown in Table 2.

Mean (Std Dev)	# Iterations to Full Detection	
# Targets	Distributed	Centralized
5	19.33 (8.87)	12.95 (6.61)
10	48.71 (22.60)	45.56 (18.82)
20	138.20 (47.59)	132.15 (49.13)

Table 2: Comparison

This evidence suggests the performance of our distributed search scheme is very close to the centralized scheme and it improves when the scale is up. Although in the centralized search scheme there is no need to

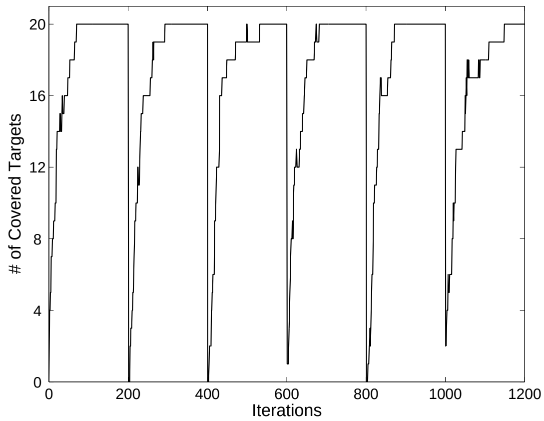

Fig. 5. Adaptation to Target Location Changes.

keep track of λ^t, this scheme is not as robust as the distributed scheme since communication link failures leave agents idle and it may also require significant communication capabilities (e.g. multi-hopping).

5 Physical Testbed

For the purpose of testing the proposed searching scheme we built a testbed consisting of four agents, and four targets distributed in an 8 by 10 feet square searching field. A virtual grid was generated to divide the field into 1 foot square cells, where the individual sensing took place. The testbed can be described as the integration of the following main components:

Agents A set of four Lego Mindstorm NXT robots using a three-wheel configuration where used as mobile platforms

Positioning System An elevated webcam together with onboard LED's were used to determine agent positions and facing direction within the searching field

Target Detection A light sensor pointing downwards was placed in each agent to detect darker sectors (targets) that contrasted with the white field

Measurement Error False positive and false negative outputs were introduced within each agent posterior target sensing

Data Handling Bluetooth links were established between a data repository (Laptop) and each agent independently to share the desired information

Restricted Movement Agents are only allow to move to non-diagonal adjacent cells

The searching scheme runs under some simple steps that are described below,

(1) While $\mu_i < .99$, $\forall i = 1, .., n$
(2) read agent i's sensor output and determine its position and facing direction
(3) update λ and μ
(4) retrieve position, facing direction and adjacent cells' information to agent i
(5) agent i computes its best next reply x_i
(6) agent i moves to x_i
(7) set i to be the next available agent
(8) end

Figure 6 is an output from the physical demo [2].

6 Conclusions

The scheme proposed in this paper has direct implications for applications involving the search, identification,

[2] The full version of the physical demo can be watched at http://people.virginia.edu/~ag7s/papers/Lego_video.avi

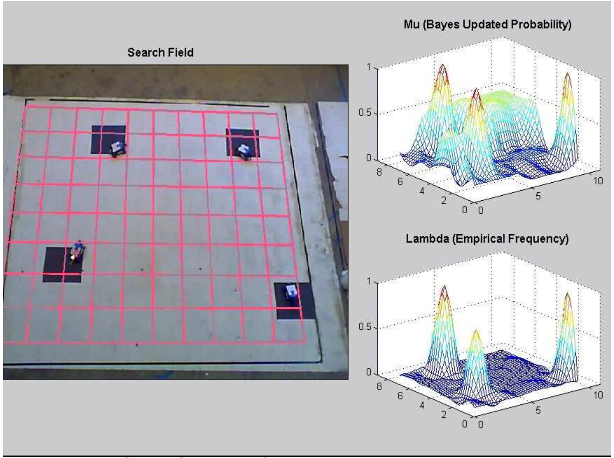

Fig. 6. GUI Output from the Physical Testbed

and location of targets across a search area that maybe vast and dangerous to navigate. Mobile robots employed for such applications provide promise to more rapidly and safely locate human beings and other targets. Such applications reduce risk for human rescue teams. Recent tragedies such as the Hurricane Katrina disaster management efforts and World Trade Center recovery efforts have highlighted some of the most extreme examples where such applications could be brought to benefit.

Unlike the established paradigms of swarm intelligence, in our scheme, agents process the implicit information encapsulated in two "trails" and choose a decision that is aimed at maximizing the chance of detecting a target without unnecessary duplication in probing. By endowing mobile sensors with this simple optimization rule, we have shown that a form of 'rational swarm' intelligence emerges as sensors successfully coordinate indirectly (i.e. they achieve a one-to-one allocation of agents and targets) through active manipulation of the trails. This feature guarantees the proposed scheme is both *reconfigurable* and *scalable*. Reconfigurability follows from the fact that agents only need to know how often a given location has been probed in the past (regardless of the identity of the sensor(s) that executed the probes) and the updated Bayesian probabilities. By enabling more rapid and safer target location discovery, the scheme we propose could be widely applicable in different military, homeland security and/or disaster recovery scenarios.

References

[1] Belta C. and Kumar V. 2004. "Abstraction and Control for Groups of Robots", IEEE Transactions on Robotics, Vol. 20, No. 5, pp.865-875.

[2] Beni G. and Wang U. "Swarm Intelligence in Cellular Robotic Systems". 1989. In NATO Advanced Workshop on Robots and Biological Systems, Il Ciocco, Tuscany, Italy.

[3] Bonabeau E., Dorigo M. and Theraulaz G. 1999. *Swarm Intelligence: From Natural to Artificial Systems*, Santa Fe Institute Studies on the Sciences of Complexity.

[4] Bullo F. and Cortes J. 2005. "Adaptive and Distributed Coordination Algorithms for Mobile Sensing Networks" in *Cooperative Control*, Lecture Notes in Control and Information Sciences, V. Kumar, N. Leonard and A. Morse, Eds. Springer Verlag, New York, pp. 43-62.

[5] Campos-Nañez E., Garcia A. and Li C. 2006. "A Game Theoretic Approach to Efficient Power Management in Sensor Networks", *forthcoming* Operations Research (available at *http://www.sys.virginia.edu/techreps/*)

[6] Clark C. and Barfoot T. 2004. "Motion Planning for Formations of Mobile Robots". Journal of Robotics and Autonomous Systems, Vol. 46 pp. 65–78.

[7] Cortes J., Martinez S. Karatas T., and Bullo F. 2004. "Coverage Control for Mobile Sensing Networks" IEEE Transactions on Robotics and Automation, Vol 20 No. 2 pp. 243-255.

[8] Di Caro, G., Ducatelle F. and Gambardella L. 2005. "Swarm Intelligence for Routing in Mobile Ad Hoc Networks", In Proceedings of the 2005 IEEE Swarm Intelligence Symposium.

[9] Dimarogonas D. and Kyriakopoulos K. 2005. "Formation Control and Collision Avoidance for Multi-Agent Systems and a Connection between Formation Infeasibility and Flocking Behavior" Proceedings of the 2005 IEEE International Conference on Robotics and Automation, Barcelona, Spain, pp. 84-92.

[10] De Groot M. 1970. *Optimal Statistical Decisions*. McGraw-Hill.

[11] Desai J., Ostrowski J., and Kumar V. 1998. "Controlling Formations of Multiple Mobile Robots" In Proceedings of the International Conference on Robotics and Automation, pp 2864-2869.

[12] Gazi V. and Passino K. M. 2003. "Stability Analysis of Swarms," IEEE Transaction on Automatic Control, vol. 48, no. 4, pp. 692–697, April 2003.

[13] Gunes M., Sorges U., Bouazizi I, 2002. "ARA – The Ant-Colony Based Routing Algorithm for MANETs", in Proceedings 2002 International Conference on Parallel Processing

[14] Jadbabaie A. , Lin J., and Morse A. S., 2003. "Coordination of Groups of Mobile Autonomous Agents using Nearest Neighbor Rules," IEEE Transactions on Automatic Control, vol. 48, no. 6, pp. 988–1001

[15] Johnson D. and Maltz D., 1996. "Dynamic Source Routing in Ad Hoc Wireless Networks," Mobile Computing, T. Imielinski and H. Korth, Eds., Ch. 5, Kluwer, pp. 153–81

[16] Lindhe M., Ogren P and Johansson K. 2005 "Flocking with Obstacle Avoidance: A New Distributed Coordination Algorithm Based on Voronoi Partitions" Proceedings of the 2005 IEEE International Conference on Robotics and Automation, Barcelona, Spain, pp. 1785-1791.

[17] Rybski P. E. , Stoeter S. A. , Gini M. , Hougen D. F. , and Papanikolopoulos N. P. 2002. "Performance of a Distributed Robotic System using Shared Communications Channels". IEEE Transactions on Robotics and Automation. pp. 713—727.

[18] Schwger M., McLurkin J. and Rus D. 2006. "Distributed Coverage Control with Sensory Feedback for Networked Robots" in Proceedings of Robotics: Science and Systems, Philadelphia, PA.

[19] Threaulaz G. and Bonabeau E. 1999. "A Brief History of Stigmergy". Artificial Life, Special Issue on Stigmergy Vol. 5. pp. 97-116.

[20] Xi W., Tan X. and Baras J.S., 2006. "Gibbs sampler-based self-organization of autonomous swarms", Automatica Vol. 42, 7, pp. 1107-1119