Rational Swarms for Distributed On-line Search

Alfredo Garcia, Chenyang Li and Fernan Pedraza

University of Virginia

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 likelihood that a target is present. These trails are stored in a geographically distributed array of stationary motes and are reminiscent of the pheromone trail used by ant colonies to find the shortest path between their nest and a food source. Each mobile 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 locate all targets) through active manipulation of the trails. This feature guarantees the proposed scheme is both reconfigurable and scalable.

Key words: Distributed control, sensor networks, on-line 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, [1] and [2]). *Unlike* the established paradigms of swarm intelligence that evolve from the idea of flocking and following (see [2] and [13]), 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 m 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 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, each agent individually 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 converge to a one-to-one allocation) 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 probabilities regarding target locations. 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

Preprint submitted to Automatica *30 December 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 search is formalized. In section 3, the convergence with probability one to a one-to-one allocation of sensors to targets is proven. Numerical tests are conducted in Section 4 to test the scheme's scalability, speed of convergence and the ability to track moving targets. A physical testbed is developed in Section 5 to implement the proposed the search scheme. Finally, in the paper's last section, we offer some concluding remarks with an emphasis on future research tasks.

1 Literature Review

There has been an increasing interest over recent years in the use of multi-robot (multi-agent) systems in applications such as humanitarian demining, search and rescue operations, and planetary exploration, where cooperation amongst robots and to some extent their coordination, promise to improve mission effectiveness and robustness. Chan et. al [5] for instance, investigated engineering trade-offs for some of these coordination approaches: *i*) Contract-Net (C-net) [20], [17], inspired on the way managers engage contractors to perform tasks for them, *ii*) Stigmergy ([21], [18] and [23]) borrowed concepts from social insects such as ants and bees, and *iii*) Signaling, that mimics scouts in a treasure-hunt. Chan et. al reported that in order for decentralized schemes like Stigmergy and Signaling to be viable, their corresponding implementation platforms should cost at most 40% of that of centralized schemes (e.g. C-net). However, other factors such as scalability and sensor technology, limitations not accounted for in [5], may support a stronger case for decentralized schemes.

Roughly speaking, approaches based on stigmergy (see Gaudiano [10], Parunak [14] and Brueckner [3]) rely on the agents' ability to share data via pheromone-like trails. Other swarming-based search schemes (Pugh and Martinoli's [15]) are based on Particle Swarm Optimization (PSO). There, the idea is to avoid employing artificial pheromone trails by only making use of sensor readings and agent-to-agent wireless communication to accomplish desired coordinated behaviors (e.g. fish schooling [11]). Approaches like the one presented in [15] are well suited for searching a relatively low number of targets, as agents tend to cluster around some targets, thus potentially obstructing others agents' discovery efforts. In more recent work, Pugh and Martinoli [16] have addressed this issue, by introducing a "random-repulse" search scheme. Our scheme resembles "random-repulse" in the usage of an inter-agent repulsive element, but differ in its conceptualization: agents in our approach are driven away from certain areas by other agents' exploration profile not by their actual presence.

Probabilistic detection measurement considering target mobility and imprecise sensing technology was explored by Furukawa et. al [9] and [22]. There, a heterogeneous team of agents (UAVs) traverse the objective region searching and tracking multiple targets using recursive Bayesian filtering. Although in this paper we present a somewhat similar updating strategy as that of [9], the objective function for each agent is constructed and used differently.

2 Distributed Search by a Swarm of Sensors

Consider a situation where n mobile sensors are tasked to find m targets ($m \geq n$) in a search region $\mathcal{X} = (X, A)$ modeled as a connected graph with locations as vertices in the set X and arcs $A \subset X \times X$. We shall denote by $X^* \subset X$ the set of feasible locations for the m targets. Let $\mu^0(x)$ be the *a priori* probability that a target is located at $x \in X$ at stage $t = 0$. Let α and β denote the conditional probability of obtaining "false positives" and "false negatives", respectively. The formula for updating μ is:

$$\mu^{t+1}(x) = (1 - \rho)\, \mu^t(x) + \rho 1_x \qquad (1)$$

where

$$1_x = \begin{cases} 1 \text{ with probability } 1 - \beta \\ 0 \text{ with probability } \beta \end{cases}$$

whenever $x \in X^*$ and

$$1_x = \begin{cases} 1 \text{ with probability } \alpha \\ 0 \text{ with probability } 1 - \alpha \end{cases}$$

when $x \notin X^*$ and $\max\{\alpha, \beta\} < \frac{1}{2}$.

Let $\lambda^t(x) = \left(\lambda_i^t(x), \lambda_{-i}^t(x)\right)$ denote a measure of probing activity in location $x \in X$, where $\lambda_i^t(x)$ is a measure of probing activity in location x by sensor i and $\lambda_{-i}^t(x)$ is the corresponding measure for sensors other than i. 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 $N(x)$ denote "neighborhood" of x, i.e. the set of locations in X that are reachable from x within one iteration. In other words, $y \in N(x)$ if and only if $(x, y) \in A$. We assume $x \in N(x)$, meaning a sensor may remain at its current location. Given a current location s_i^t, let

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

Agent i's next location s_i^{t+1} is chosen as follows

$$s_i^{t+1} \in \begin{cases} B_i^t(s_i^t) \text{ if } \min\{\mu^t(x), \lambda_i^t(x) - \lambda_{-i}^t(x)\} \geq \theta \\ \\ N(s_i^t) \text{ otherwise} \end{cases}$$

where selections from $B_i^t(s_i^t)$ and $N(s_i^t)$ are made randomly and $\theta \in (0,1)$. Note that $\lambda_{-i}^t(x)$ can be interpreted as the "pheromone trail" left by other sensors, i.e. a simple form of *stigmergy* (see [19]). One of the most cited examples of stigmergy is the mechanism used by ant colonies to find the shortest path between their nest and a food source. In our setup for distributed search, agents' collective intelligence or *swarm intelligence* is encapsulated in the local optimization rule in (2) according to which an agent's next choice of location maximizes the chance of finding a target *without duplication.* In choosing locations according to the rule in (2), agents will tend to stay away from each other (hence the term, $-\lambda_{-i}^t(x)$) while reinforcing good choices (hence, $\lambda_i^t(x)$). After the choice of location is made, agents 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 $\rho \in (0,1)$, $\mathbf{1}_{\{s_i^{t+1}=x\}}$, $\mathbf{1}_{\{s_{-i}^{t+1}=x\}}$ are the indicator random variables of events $\{s_i^{t+1} = x\}$ and $\{s_j^{t+1} = x, j \neq i\}$ respecitively.

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. $\lambda^t(x)$, $x \in X$) *regardless of the identity of the sensor(s) that executed the probes*, and the updated probability that a target is present at location x, i.e. $\mu^t(x)$, $x \in X$. 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), thus 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. The deployment of these stationary motes could be accomplished in various ways. For instance, these motes could be scattered randomly over the search region before the search begins. Alternatively, the network of stationary motes could be deployed by the mobile agents themselves (i.e. each agent carrying a number of motes and dropping one in locations where no signal from any other stationary mote is detected).

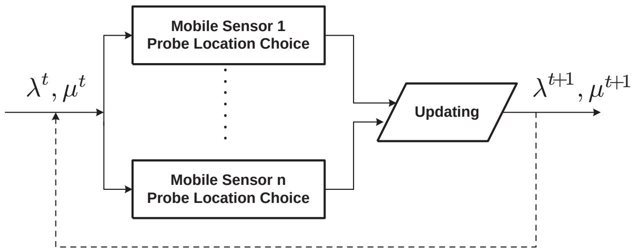

Fig. 1. Schematic for Distributed Search.

3 Convergence

In this section we study the convergence of the scheme where n mobile sensors are tasked to find m targets ($m \geq n$). We assume that an agent i can only move within his/her neighborhood $N(s_i^t)$ at every iteration. Let us denote by $S \subset X^n$ the set of joint locations of all agents and S^* the set of different one-to-one allocations of agents to targets, i.e. $s^* = (s_1^*, s_2^*, ..., s_n^*) \in S^*$ if and only if $s_i^* \in X^*$ and $s_i^* \neq s_j^*$.

Lemma 1: $s^t = (s_1^t, s_2^t, ..., s_n^t) \in S^*$ *infinitely often.*

Proof: By contradiction, assume $s^t \notin S^*$ eventually. Thus, there exist a subset of agents, say $\mathcal{I} \subset \{1, ..., n\}$ such that for $i \in \mathcal{I}$ either $s_i^t \notin X^*$ or $s_i^t = s_j^t \in X^*$ eventually. This implies $\min\{\mu(s_i^t), \lambda_i^t(s_i^t) - \lambda_{-i}^t(s_i^t)\} \to 0$. Thus, there exist $\tau > 0$ such that for all $t > \tau$, all agents in $\mathcal{I}$ randomly choose probing locations. Thus, with probability 1, there exists $t > \tau$ such that $s^t \in S^*$, a contradiction. ■

Theorem 1: $s^t \to s^* \in S^*$ *with probability one.*

Proof: Let $s^{t+1} = s^*$. That is, $s_i^{t+1} = x \in X^*$ and $s_{-i}^{t+1} \neq x$, for all $i \in \{1, \ldots, N\}$. We shall show that $s_i^{t+\tau} = x$ and $s_{-i}^{t+\tau} \neq x$, for $\tau \geq 1$ with positive probability $p > 0$. To show this, suppose $s_i^{t+1} = x \in B_i^t(s_i^t)$ for all $i \in \{1, \ldots, N\}$. Note that this is equivalent to assuming $\min\{\mu^t(x), \lambda_i^t(x) - \lambda_{-i}^t(x)\} \geq \theta$ for all $i \in \{1, \ldots, N\}$. Now

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

Since $x \in B_i^t(s_i^t)$ we have that for all $y \in N(s_i^t)$,

$$\mu^t(x) + \lambda_i^t(x) - \lambda_{-i}^t(x) \geq \mu^t(y) + \lambda_i^t(y) - \lambda_{-i}^t(y)$$

and

$$\mu^{t+1}(y)+\lambda_i^{t+1}(y)-\lambda_{-i}^{t+1}(y) = (1-\rho)(\mu^t(y)+\lambda_i^t(y)-\lambda_{-i}^t(y))$$

thus we conclude

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

or equivalently, $B_i^t(s_i^{t+1}) = \{x\}$ and it it follows that $s_i^{t+\tau} = x$ and $s_{-i}^{t+\tau} \neq x$, for $\tau \geq 1$.

Now, suppose that $\min\{\mu^t(x), \lambda_i^t(x)-\lambda_{-i}^t(x)\} = \lambda_i^t(x)-\lambda_{-i}^t(x) < \theta$ for at least one agent $i \in \{1,\ldots,N\}$. With probability $\left(\frac{1}{|N(x)|}\right)^\tau$ the agent remains at the same location x for τ times in a row. When this happens

$$\lambda_i^{t+\tau}(x)-\lambda_{-i}^{t+\tau}(x) = (1-\rho)^\tau[\lambda_i^t(x)-\lambda_{-i}^t(x)] + \sum_{0\leq k<\tau}(1-\rho)^k\rho$$

Thus, it follows that there exist an upper bound on the number of iterations τ (say τ_i) needed to guarantee $\lambda_i^{t+\tau}(x) - \lambda_{-i}^{t+\tau}(x) \geq \theta$. Thus, with probability $p_i^* = \left(\frac{1}{|N(x)|}\right)^{\tau_i}$ we have that $s_i^{t+\tau} = x$ and $s_{-i}^{t+\tau} \neq x$, for $\tau \geq 1$.

Now suppose $\min\{\mu^t(x), \lambda_i^t(x) - \lambda_{-i}^t(x)\} = \mu^t(x) < \theta$ for at least one agent $i \in \{1,\ldots,N\}$. With probability $\left((1-\beta)\frac{1}{|N(x)|}\right)^\tau$ the agent remains at the same location x for τ times in a row and

$$\mu^{t+\tau}(x) = (1-\rho)^\tau\mu^t(x) + \sum_{0\leq k<\tau}(1-\rho)^k\rho$$

Thus, it follows that there exists an upper bound on the number of iterations τ (say $\tilde{\tau}_i$) needed to guarantee $\mu^{t+\tau}(x) \geq \theta$. Thus, with probability $\tilde{p}_i = \left(\frac{1-\beta}{|N(x)|}\right)^{\tilde{\tau}_i}$ we have that $s_i^{t+\tau} = x$ and $s_{-i}^{t+\tau} \neq x$, for $\tau \geq 1$. It follows that whenever $s^{t+1} = s^*$, we have $s^{t+\tau} = s^*$,$\tau \geq 1$ with probability

$$p^* = \min_i\{p_i^*, \tilde{p}_i\} > 0$$

To finalize the proof, by Lemma 1 we know $s^t \in S^*$ *infinitely often*. Hence, $s^t = s^*$ *eventually.* ■

We have proved convergence to one-to-one-agent-target allocation when $m > n$. Note however that some targets might remain uncovered. To address this, a small modification to the algorithm can be introduced: an agent is "released" when the likehood measure μ of its current location reaches a certain threshold enabling it to continue to search. Alternatively, seach can continue after the target is successfully removed or tagged by an agent.

In the case where we have more agents than targets, i.e. $m > n$, duplicate search efforts are inevitable and convergence to a one-to-one allocation is clearly not possible. The oversupply of agents implies that the system as a whole does not incur in any inefficiency when two or more agents settle for the same target.

3.1 Convergence with Bayesian updating rule

The updating rule for μ presented in (1) could also be replaced with a Bayesian scheme. Let $Z(x)$ be the probing 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:

$$\mu^{t+1}(x) = \begin{cases} \frac{(1-\beta)\mu^t(x)}{\alpha(1-\mu^t(x))+(1-\beta)\mu^t(x)} & Z(x) = 1 \\ \\ \frac{\beta\mu^t(x)}{(1-\alpha)(1-\mu^t(x))+\beta\mu^t(x)} & Z(x) = 0 \end{cases} \quad (4)$$

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 [8]) that $\mu^t(x) \to 1$(respectively, $\mu^t(x) \to 0$) with probability one. Therefore, with the Bayesian update, Lemma 1 will still hold. Let us set a threshold $\delta \in (0,1)$ for $\mu^t(x)$, such that once we have $\mu^t(x) > \delta$, x is believed to have a target in it and no further probing is needed. It is easy to verify that with $1-\beta > \alpha$, and $x \in X^*$, $\mu^{t+1}(x) > \mu^t(x)$ if a true positive is obtained. This means whenever $s^t \in S^*$ happens, there exists an upper bound on the number of iterations τ (say $\tilde{\tau}$), such that with at least probability $\tilde{p} = \left(\frac{1-\beta}{|N(x)|}\right)^{n\tilde{\tau}}$, we have $\mu^{t+\tilde{\tau}}(s) > \delta$. Then by a similar argument as in the proof above, we have the same result obtained in Theorem 1.

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 cur-

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

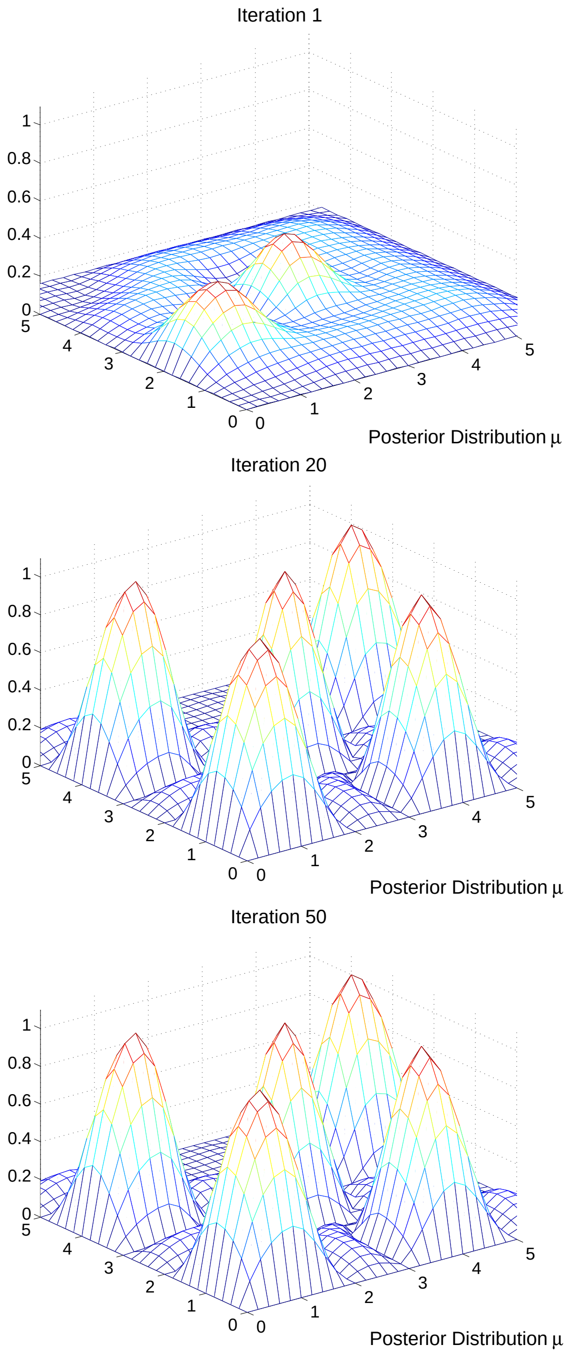

Fig. 2. Evolution of μ^t.

rent location.

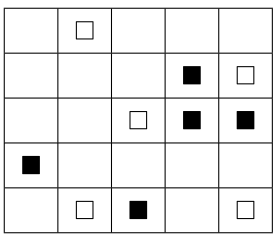

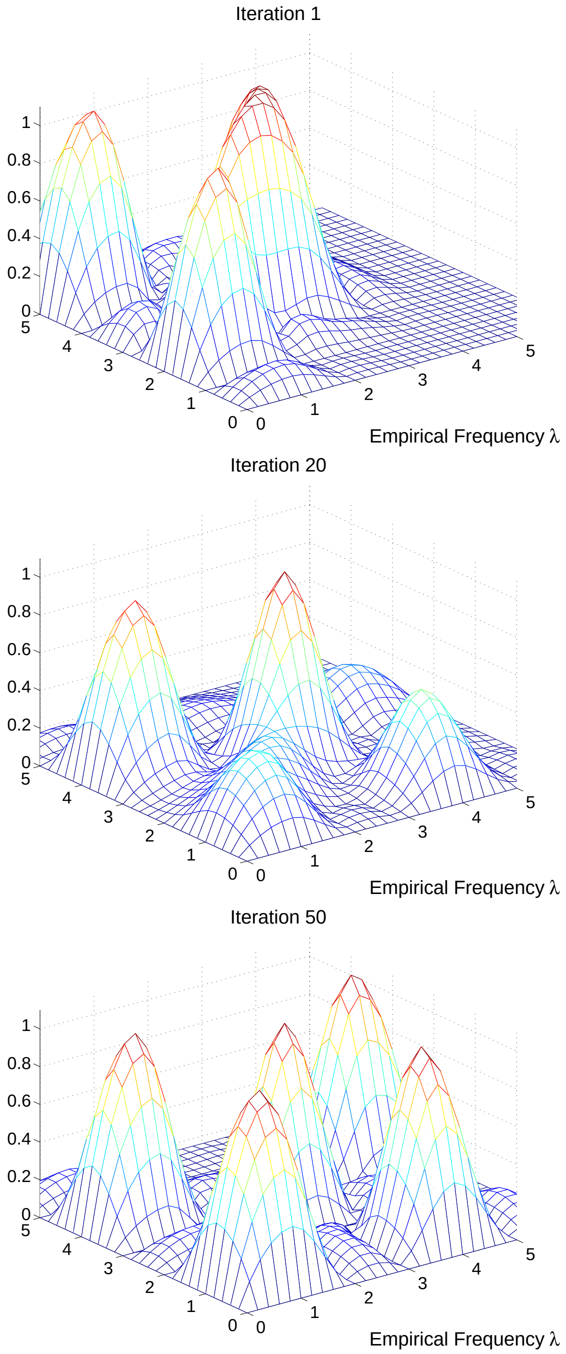

Fig. 3. Evolution of λ^t.

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. This is a very stringent evaluation as one could conceivably "release" agents that have already located targets (up to a level of precision in $\mu(x)$, say 0.95) in order to help with the search of the remaining unidentified target locations. Otherwise, the tasks of locating the last few targets becomes increasingly difficult for the last few uncommitted agents.

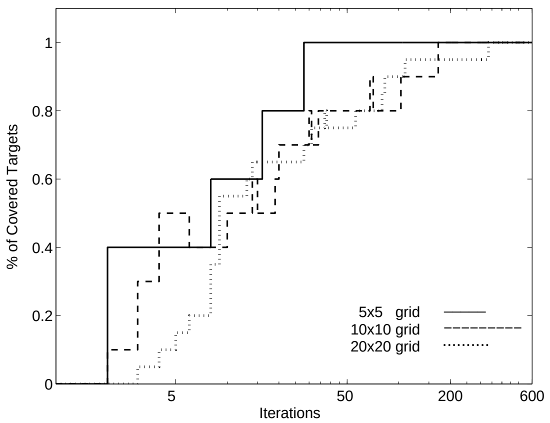

Fig. 4. Test for scalability.

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.

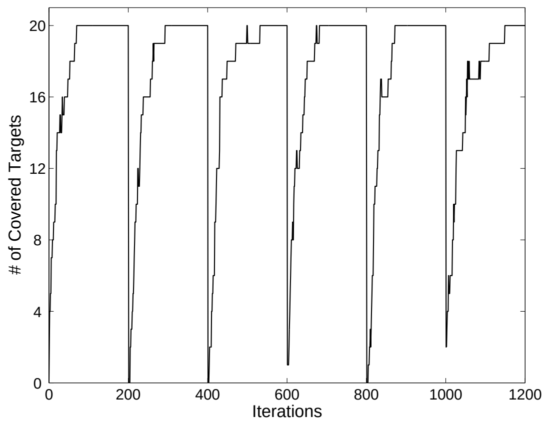

Fig. 5. Adaptation to Target Location Changes.

Finally, we compare the performance of our distributed search algorithm with a *centralized* search scheme and a *randomized* search scheme. In the centralized search scheme, there is a "virtual" base station that instructs each agent on the best next location to probe by selecting the closest location in the agent's neighborhood with the highest value of μ. Centralization avoids any duplication in search efforts. In a distributed random search scheme, each agent randomly makes a decision on the next location to probe and stops at location x, when $\mu(x) > 0.95$ and no other agent is currently at x. We tested the three algorithms with a total of $5, 10, 20$ targets, $6, 12, 24$ mobile agents, and area size 5^2, 10^2, 20^2 respectively, given $\alpha = 0.15$ and $\beta = 0.1$. For each scenario, we repeated the experiment 200 times with randomized targets locations each time. The results are shown in Table 2.

Table 2: Comparison

Mean (STD)	**# Iterations to Full Detection**		
# Targets	**Distributed**	**Centralized**	**Randomized**
5	20.02 (10.83)	12.54 (6.39)	39.55(16.43)
10	49.05 (20.18)	44.29 (17.37)	139.31(48.45)
20	139.80 (49.43)	137.39 (50.02)	401.90(121.25)

This evidence suggests the performance of our distributed search scheme is very close to the centralized scheme and is much more efficient than the randomized search. Furthermore, its performance improves when the scale is increased.

Even though in the centralized search scheme there is no need to keep track of λ^t, this scheme is not as robust as the distributed scheme since communication link failures may leave agents idle and it may also require significant networking capabilities (e.g. multi-hopping to communicate to the central base station).

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 individual sensing took place.

The testbed can be described as the integration of the following main components:

Agents Four Lego NXT Mindstorm robots using a three-wheel configuration were used as mobile platforms.

Positioning System An elevated webcam together with a pair of LEDs mounted on each Agent were used to determine agent position and direction within the searching field; emulating a GPS-like localization system.

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 to each agent's sensing unit.

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

Restricted Movement Agents were only allowed 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

For the movement in step 6 a close loop control was developed, using positioning information to correct direction errors. On the first set of experiments we forced the agents to coordinate their movement so each of these was in position to measure target presence at the beginning of each cycle, but this constrain can be easily relaxed for future application examples.

In Figure 6 we can see an output from the physical demo [2].

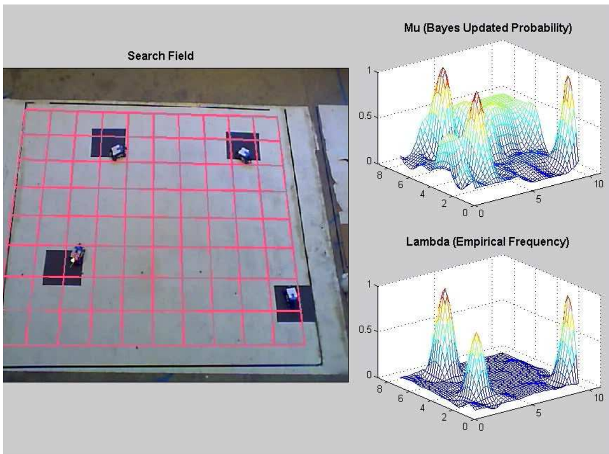

Fig. 6. GUI Output from the Physical Testbed

6 Conclusions

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

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 locate all 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 target-present 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.

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

References

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

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

[3] Brueckner S. and Parunak H. V. D. 2000. "Multiple Pheromones for Improved Guidance," In Proceedings of the 2^nd DARPA-JFACC Symposium on Advances in Enterprise Control, Minneapolis, MN, USA.

[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] Chan Y. M., Wong S., Foo M. C.and Teo R. 2004. "Engineering Intuition for Designing Multi-Robot Search and Rescue Solutions," In Proceedings of the 2004 IEEE Conference on Cybernetics and Intelligent Systems, Singapore.

[6] Colorni A., Dorigo M. and Maniezzo V. 1991. "Distributed Optimization by Ant Colonies," In Proceedings of ECAL91 European Conference on Artificial Life, pp. 134-142, Paris, France.

[7] 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.

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

[9] Furukawa T., Bourgault F., Lavis B. and Durrant-Whyte H. F. 2006. "Recursive Bayesian Search-and-Tracking Using Coordinated UAVs for Lost Targets," In Proceedings of the 2006 IEEE International Conference on Robotics and Automation, Orlando, FL.

[10] Gaudiano P., Shargel B., Bonabeau E. and Clough B.T. 2003. "Swarm Intelligence: a new C2 Paradigm with an Application to Control of Swarms of UAVs," 8^th ICCRTS Command and Control Rearch and Technology Symposium.

[11] Han X., Shen C. C. and Rossi L. 2007. "Autonomous Navigation of Wireless Robot Swarms with Covert Leaders," In Proceedings of ROBOCOMM 2007, Athens, Greece.

[12] 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

[13] Kennedy J., Eberhart R., 1995. "Particle swarm optimization," IEEE Proceedings of the International Conference on Neural Networks, Perth, Australia, pp. 1942–1948.

[14] Parunak H. V. D., Purcell M. and Connell R. O. 2002. "Digital Pheromones for Autonomous Coordination of Swarming UAVs," AIAA First Technical Conference and Workshop on Unmanned Aerospace Vehicles, Systems and Operations, No. 3446.

[15] Pugh J. and Martinoli A. 2007. "Inspiring and Modeling Multi-Robot Search with Particle Swarm Optimization," In Proceedings of the 2007 IEEE Swarm Intelligence Symposium, Hawaii.

[16] Pugh J. and Martinoli A. 2007. "The Cost of Reality: Effects of Real-World Factors on Multi-Robot Search," In Proceedings of the 2007 IEEE International Conference on Robotics and Automation, Roma, Italy.

[17] Sondholin T. 1993. "An Implementation of the Contract-net Protocol based on Marginal Cost Calculations", In Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI-93), pp. 256-262.

[18] Svennebring J. and Koenig S. 2004. "Building Terrain-Covering Ant Robots: Feasibility Study," Autonomous Robots, Vol 16, No 3, pp. 313-332.

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

[20] Excelente-Toledo C.B. and Rlennings N. 2004. "The Dynamic Selection of Coordination Mechanisims," Journal of Autonomous Agents and Multi-Agent Systems, Vol. 9, No. 1-2, pp. 55-85.

[21] Wagner I.A., Lindenbaum M. and Bruckstein A.M.1999 "Distributed Covering by Ant-Robots using Evaporating Traces," IEEE Robotics and Automation, vol. 15, no. 5, pp. 918-933.

[22] Wong E., Bourgault F., Furukawa T. 2005. "Multi-vehicle Bayesian Search for Multiple Lost Targets," In Proceedings of the 2005 IEEE International Conference on Robotics and Automation, Barcelona, Spain.

[23] Zohdi T.I.2003. "Computational Design of Swarms." International Journal for Numerical Methods in Engineering, Vol. 57, No. 15, pp. 2205-2219.