Team Distribution between Foraging Tasks with Environmental Aids to Increase Autonomy

Juan M. Nogales and Gina Maira Barbosa de Oliveira
Department of Computer Science, Federal University of Uberlandia, Santa Monica, Uberlandia, Brazil

Keywords: Task Allocation, Foraging, Adaptability, Cooperation, Autonomy.

Abstract: In this paper, robots have to distribute themselves across a set of regions where they will serve in foraging tasks, transporting objects repetitively. Each region stores information about the performance of the subgroup of robots serving that region. Robots can also share information between them and identify which region is offering better conditions to forage. In particular, each region has a different rate to recover recently removed objects, which demands a different number of robot foragers. We explore the effects of the network structure in robot distribution and their performance. Results indicate a small dependence of robot-robot connections and a great dependence of robot-environment interaction. Since cooperative robots are going after a global goal, the proposed distribution rules combined with environmental aids allowed them to make better decisions autonomously, increasing the number of transported objects and reducing the number of travels.

1 INTRODUCTION

Social insects live in colonies, but they divide themselves into subgroups to complete different tasks, for instance, honeybees searching for nest sites (Seeley et al., 2006), wasps storing wood (Jeanne and Nordheim, 1996), or ants collecting food (Anderson and Bartholdi, 2000). These insects tend to distribute themselves between regions to perform specific jobs. However, they can switch between jobs whenever the colony needs it (Zahadat et al., 2015). Scientists attribute their success and recovery skills to the coordination within and across the subgroups of insects (Bonabeau et al., 1999).

In (Schmickl et al., 2012), the authors showed, with bees in a simulated environment, that their adaptability depends on regulated communication. The swarm has only a few receivers in the entrance of the hive. They get information about source quality, which foragers share them through a short-range communication (trophallaxis). Then, receivers spread this information through a long-range communication (waggle-dance) to help in the recruitment of other bees. Thus, bees could distribute themselves to forage for nectar by combining short and long-range communication.

Insects also exhibit a highly decentralized control; it seems they have no leader, known as divisional autonomy. On the other hand, they follow collective rules, known as distributed control (Anderson and Bartholdi, 2000). For instance, in the bee colony investigated in (De Marco and Farina, 2001), individuals can decide which recruiter to follow, exhibiting divisional autonomy. However, they have to forage for nectar for the colony, following distributed control rules. In other words, individuals are autonomous, but some group rules or objectives restrain their autonomy.

Due to insect success, robotics researchers brought forth the concept of swarm robotics. It is a novel concept inspired by insect strategies to solve complex tasks, which began to grow at the beginning of the 2000s. In particular, such solutions offer a far better alternative by employing simpler units. Designing simple robots seems easier than creating a big, expensive, and heavy robot. In (Şahin, 2005), the authors considered pertinent to describe the desirable properties of swarm robotics before new works blurred this concept through time, which continue until our days:

- **Robustness:** redundancy and decentralization should foster the swarm to continue operating despite failures or disturbances in the environment, although at a lower performance.
- **Flexibility:** requires the swarm to be able to generate modularized solutions to different tasks.
- **Scalability:** considers that the coordination

Nogales, J. and Oliveira, G.
Team Distribution between Foraging Tasks with Environmental Aids to Increase Autonomy.
DOI: 10.5220/0006588700250036
In *Proceedings of the 10th International Conference on Agents and Artificial Intelligence (ICAART 2018) - Volume 1*, pages 25-36
ISBN: 978-989-758-275-2

mechanism would be able to deal with a large number of relatively simple robots.

Modern applications with a swarm of robots range from navigation to surveillance problems, for instance: wildfires containment (Phan and Liu, 2008), intruders detection (Raty, 2010; Khan et al., 2016), and area exploration (Antoun et al., 2016). In those scenarios, robots have to coordinate among themselves to achieve multiple objectives. Therefore, it is necessary to allow them to divide into subgroups. As insects, robots can make their own decisions, but their behaviors need an orientation toward the group goal.

A common approach to design these decision-making strategies is to model the group as a Multi-Agent System (MAS), define a local utility function for each objective, and establish the common goal as the sum of all utility functions (Krause and Guestrin, 2007). Thus, if agents optimize the sum of all local utilities, they also achieve the common goal. Despite MAS solutions are generic and often abstract inherent complexities of Multi-Robot Systems (MRS), the decision-making rules stimulate individuals to cooperate among subgroups such that they may maximize the group utility. Some authors have reached optimal solutions by adapting multi-agent decision-making rules to robots as described in the survey presented in (Yan et al., 2013).

For instance, in (Krause and Guestrin, 2007), the authors employ the law of diminishing returns to design utility functions. They incorporate the benefits of assigning an extra robot to serve a particular location. In particular, robots have to surveil a water distribution system divided into regions. Each region has a probability of intrusions and a potential detriment if a robot is not serving there (e.g., the affected population by an intrusion). Robots make decisions based on the marginal contribution they yield when selecting a region. Thus, when robots optimize the utility functions, they reach an effective allocation for the surveillance. Unlike some agent models, nature of robots restrains them to be in one place at a time.

Here, we adapted the MAS decision-making strategies described in (Nogales and Finke, 2013) to reach a near-optimal distribution with foraging robots. Each task is associated with a region where robots have to forage for virtual objects. In particular, we mimic the receivers at the entrance of a hive of bees with some environmental aids to help robots to share information. Robot's decision-making strategies depend on the law of diminishing returns, which allow individuals to show both divisional autonomy and distributed control. Each individual can decide which task to serve, while it seeks to optimize the group utility. We tested three different decision-making models varying their information-sharing structure and robot autonomy. The proposed decision-making models stimulate robots to cooperate among subgroups such that they maximize the number of foraged objects.

The rest of the paper is organized as follows: Section 2 describes previous works and their strategies. Our proposal is detailed in Section 4, while the experiments and results are in Section 5. Finally, Section 6 provides a short discussion of the results and suggests future work.

2 RELATED WORKS

In this section, we focus on Multi-Robot Task Allocation (MRTA) for foraging. It is important to mention that Search and Rescue missions are related to foraging task (Ahuja et al., 2002; Liu and Nejat, 2013), but they are out of the scope of this work. The objective of MRTA is to assign M jobs to N robots. Unlike MAS, robots have to interact with a physical world and with one another (Cao et al., 1997). Thus, we first reviewed some works in MRTA solving foraging tasks, then, we described a pair of papers in MAS, which followed a similar way to the one we designed the local utilities for the decision-making rules. In particular, because they work with the law of diminishing returns to reach an optimal distribution of agents.

Task allocation has brought significant improvements in foraging tasks. Foraging robots are transporting objects from one place (e.g., a source) to another (e.g., a nest). Notwithstanding, the problem continues in terms of jobs and workers that must maximize the overall performance (Dasgupta, 2011). Theories from operational research and combinatorial optimizations underlie several approaches of task allocation. Such solutions employed concepts like utility functions (Sung et al., 2013), auctions (Viguria et al., 2007), market-based processes (Akbarimajd and Simzan, 2014), game theory payoffs (Marden et al., 2009), and the like to help robots in coordination. In those works, the impact of robot decision appears only after a robot takes action according to its selection.

One of the former and most cited works with robots transporting objects while employing task allocation strategies is (Kube and Bonabeau, 2000). The authors tested MURDOCH, an auction-based task allocation system, in a box-pushing experiment with heterogeneous robots. Their results show auctions as a promising strategy for foraging with task allocation. Later, several variations of auctions strategies appeared, e.g., repetitive auction processes, decentralized auctions, and role assignment strategy. In particular, the idea of role assignment came from robot soccer domain, where each robot calculates its utility for each role and periodically broadcasts these values to coordinate with its teammates (Stone and Veloso, 1999). The authors of (Chaimowicz et al., 2002) introduced roles, as functions robots must perform to transport an object. Thus, when a robot finds an object, it shares information about the utility of the available role and the need of helpers. Robots would listen this alternative, and some of them could consider it a better option. Then, similar to the bidding process of an auction, they would offer their help to the leader, which would choose the best-qualified helper to transport that object.

In other transportation tasks, the robots have previous knowledge of the object positions and can employ path-planning strategies to avoid collisions. For instance, in (Yan et al., 2012), the robots could minimize the total transportation time while keeping a low energy consumption on each robot. The environment includes a place of constant production of goods, but their rate of production is unknown. The authors implemented a heuristic that helps robots to estimate the rate of production and define their idle periods to reduce their energy consumption. They compared their heuristic solution to a centralized *replanner* described in (Wawerla and Vaughan, 2010). The results show that their strategy was faster in the implemented environments, and it only required a few more energy than the *replanner*.

In (Lerman and Galstyan, 2002), the authors examine a scenario for foraging objects where experiments show a decreasing average return effect, which is known as the law of diminishing returns. Loosely speaking, each additional robot working on a task would increase the performance, but the size of its improvement is gradually lower until the group reaches a size with which its performance declines (Färe, 1980). Several works in MRTA exhibit this phenomenon (Bonabeau et al., 1997; Schmickl et al., 2012; Akbarimajd and Simzan, 2014). However, few works exploit it to help the group to find an optimal distribution of agents (Nogales and Finke, 2013) or to reach a near-optimal distribution of robots (Krause and Guestrin, 2007).

In particular, this phenomenon appears due to both limited resources and space where robots interact with one another while exploring or going after an object. As the group grows, more interference appears lengthening the delivery of the items to nests. In some occasions, robots begin to focus on avoiding collisions, which holds them back from delivering objects. Finally, in (Lerman and Galstyan, 2002), the authors concluded that there is an optimal quantity of robots, and beyond that number, the benefits of parallelism begin to disappear.

Finally, the most accepted taxonomy of the classification of MRTA problems was found in (Gerkey and Matari, 2004), which divides problems as follows:

- **Single (ST) vs Multiple Tasks (MT):** refers to a number of tasks a robot can carry out simultaneously
- **Single (SR) vs Multiple Robots (MR):** refers to a number of robots needed to fulfill a task
- **Instantaneous (IA) vs. Time-extended Assignment (TA):** refers to the available information for planning future allocations

Although our robots reallocate themselves dynamically, we can consider that our proposal belongs to the group of ST-MR-IA, because several robots have to forage for objects, and each object requires one robot for its transport. They cannot predict the rates of object production of the environment as in (Yan et al., 2012), i.e., they are working with current (and possibly outdated) information. Different levels of communication are also explored. Three decision-making strategies must help robots to find a balance between inhibition and stimulation of autonomy through communication. For instance, if a robot informs that its region is providing objects more quickly, most robots could decide to arrive at that region and increase the congestion, which would hold them back, as a group. Therefore, we need to find a threshold between divisional autonomy and distributed control.

Besides that, we associate each subgroup to mathematical functions that satisfy the law of diminishing returns to regulate robots autonomy. We opted for the MAS strategy described in (Nogales and Finke, 2013), which we adapted for a MRTA system. In particular, because those decision-making rules consider: *i*) delays in both communication and movement of heterogeneous agents, *ii*) agents can serve in only one task at a time, and *iii*) since agents are of the same type, they generate the same contribution to the utility being indistinguishable from each other. Thus, our robots could employ utility functions based on the law of diminishing returns to distribute themselves and reach an (near-)optimal number of foraged objects when following those rules.

3 PROBLEM

Following the objective of MRTA, we have to define the tasks for our foraging robots and an environment that favors task allocation strategies. In particular, we adapted the decentralized task allocation strategy for multi-agent systems found in (Nogales and Finke, 2013) to work with a multi-robot system. Although in (Nogales and Finke, 2013), agents can be of various types, here, we worked with homogeneous robots to follow one of the conditions of swarm robotics (Şahin, 2005). However, another type of robots could work in parallel tasks.

We employed graph theory to model the environment and its tasks, besides complex network substrates for robot-robot and robot-environment information sharing. Thus, let nodes represent the regions, i.e., distributed locations in which robots perform foraging tasks. Nodes belong to a set N, indexed from 1 to n. Figure 1 illustrates how a graph models the rooms of a floor in a warehouse.

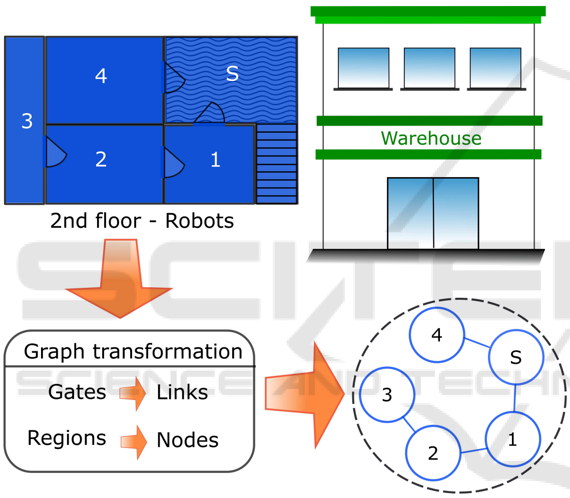

Figure 1: Transforming a warehouse scenario into a graph model. Small circles represent the nodes, while the dashed circle represent the floor.

In this building, robots have to move objects within a room (node), but each room requires a different number of robots. Therefore, robots should distribute themselves between the rooms. Note that the graph can be as complex as the designer needs to model the environment.

Challenge

As illustrated in the warehouse of Figure 1, we model the regions of the environment as nodes into a graph. Note that within each node, a robot should perform several subtasks, e.g., repetitive object transportation. For mathematical reasons, we assume that the number of robots is large enough to be appropriately represented by a continuous variable. However, as our experiments show, this is not a critical assumption for a practical implementation.

Thus, for a node i, the number of robots is defined by r_i. Let $\Delta_q \subset \mathbb{R}^n$ denote the $(n-1)$ dimensional simplex defined by the equality constraint $\sum_{i=1}^{n} r_i = q$, where q denotes the quantity of robots available. The benefit of having an amount of robots r_i foraging within node i is given by the utility function $f_i : \mathbb{R} \to [0,\infty)$. The total utility function is defined by $f : \mathbb{R}^n \to [0,\infty)$, $f(r) = \sum_{i=1}^{n} f_i(r_i)$, where $r = [r_1, \ldots, r_n]^\top$ represents the state of the system. Under the assumption of local information-sharing and decentralized decision-making, the objective is to identify conditions that allow us to solve the following optimization problem

$$\text{maximize } f(r), \text{ subject to } r \in \Delta_q. \tag{1}$$

In other words, we want to find the optimal allocation of all robots that maximizes the utility associated to each node. The following section details the proposed environment, the mathematical notation, and the decision-making mechanisms.

4 PROPOSAL

In this work, the information sharing structure has two layers that include robot-robot and robot-environment connections. Since TAMs are available in the environment, we can consider them as an additional aid for communication and coordination between robots. TAM devices emulate picking or dropping activities of virtual objects through color codes. Besides that, TAMs, as proposed in (Brutschy et al., 2015), can connect one another through a Zigbee network increasing the possibility of spreading information between regions. The second layer allows different network structures to underlie the communication between robots.

4.1 Environment

Since we want to evaluate the team performance in foraging through a task allocation strategy, we employed an environment consisting of several regions. Commonly, the environment for foraging includes at least two regions: one where they deposit objects and other where robots explore (Mataric, 1994). Notwithstanding, more regions can appear when: robots divide the environment into regions (Pini et al., 2014; Buchanan et al., 2016) or the researchers provide a predefined division (Bobadilla et al., 2012; Pitonakova et al., 2016). We opted for a predefined division of the environment by combining ground color, smart color-changing landmarks (Brutschy et al., 2015), and gates (Bobadilla et al., 2012) to separate the regions. Figure 2 shows the environment dimensions, gates, landmarks and their distribution.

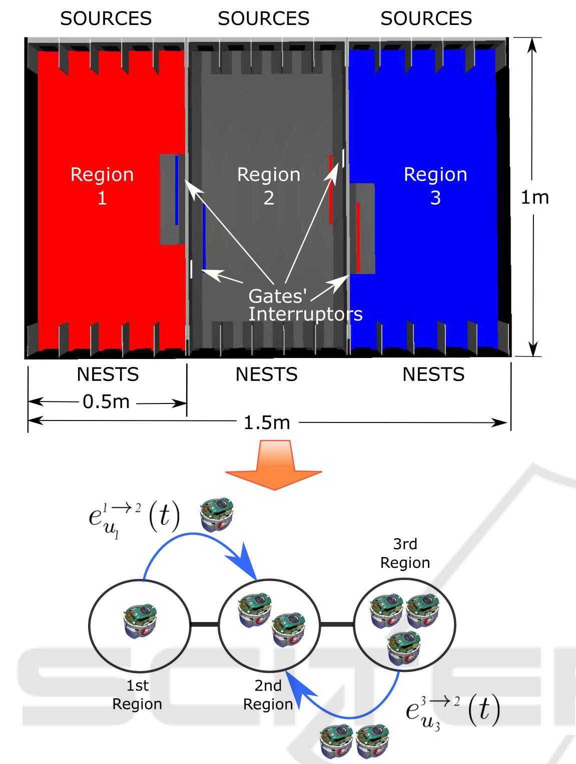

Figure 2: Environment for the foraging task. Upper image considers all dimensions of the environment. Below it is the transformation to a graph considering robot movements.

Note that these environmental aids are a cheap solution to distinguish regions. The more complex aids we used are TAMs, which are smart landmarks proposed in (Brutschy et al., 2015). TAMs' dimensions are $10 \times 10\ cm^2$ in order that an e-puck robot can enter them. Every region has five TAMs working as nests and five working as sources. In each region, TAMs working as nests store performance and utility information of that region and help the subgroup of robots working in the same region to update this information. On the other hand, TAMs working as sources have a rate for the delivery of objects, which is different for each region (as production rates of goods in companies). However, robots and TAMs do not know about these rates.

Note that the environment consists of three regions that are linked through hallways with automatic gates. Since transitions between regions last a finite time in real environments, we added a predefined cost (in time) in the gates. Once a robot reaches one of the switches, the gate opens and waits enough time for the robot to pass through. Next, when robots are transitioning, they are not foraging in any region, consequently, more travels leads to less objects foraged from TAMs.

Although we compute the optimal quantity of robots for the environment, in any moment, a region can hold more robots than its optimal number. Therefore, robots would need distributed rules inside their decision-making mechanism in order to be able to decide whether to abandon or remain there and to improve their performance in each region.

4.2 Notation and Model

We adapted the MAS strategy proposed in (Nogales and Finke, 2013) such that we could employ it in a MRTA scenario. We implemented two decision-making strategies: *i)* The deterministic model (D), which guarantees an optimal distribution and works as a reference for the other strategies and *ii)* A semi-stochastic version (SS), which allows validating the effect of connections with the environment as (SS-TAM). These strategies tried to solve the problem described in Section 3. However, we needed to add some assumptions to make it possible for robots to use the MAS solution. In particular, each utility f_i satisfies the following three assumptions (common in economic theory (Färe, 1980)):

A1 Each function f_i is continuously differentiable on $\mathbb{R}$.

A2 An increase in utility satisfies

$$\frac{f_i(r_i+u_i)-f_i(r_i)}{u_i} > \frac{f_i(r_i+w_i)-f_i(r_i)}{w_i} \quad (2)$$

where $r_i \in \mathbb{R}$, $w_i > u_i > 0$ represent a finite number of robots entering node i.

A3 An increase in the number of robots within a region increases the utility of that region, bounded by

$$0 < \frac{f_i(r_i+u_i)-f_i(r_i)}{u_i} < \infty \quad (3)$$

Assumption A2 represents the law of diminishing returns and implies that increasing the number of robots in a node will always yield decreasing average returns. Assumption A3 indicates that any additional robot should increase the utility moderately.

Using Eq. (2) and according to Assumption A3, the partial derivative of f_i with respect to r_i, denoted by s_i, satisfies

$$-a \leq \frac{s_i(x_i)-s_i(y_i)}{x_i-y_i} \leq -b \quad (4)$$

for any $x_i, y_i \in \mathbb{R}$, $x_i \neq y_i$, and constants $0 < b \leq a$. It can be shown that if Assumptions A1-A3 are satisfied, the marginal utility functions $s_i(\cdot)$ are continuous on $\mathbb{R}$, strictly decreasing, and non-negative, while $f_i(\cdot)$ is strictly concave (see (Nogales and Finke, 2013) for details). Note that each additional robot must yield a lower average return. We have to find the marginal utility functions for the environment of this proposal described in Figure 2.

Next, a connection between two nodes i and j (regions) means that robots can move back and forth between them, and robots at node i can obtain information about node j. By moving across nodes, robots may join or leave them at time indexes $t = 0, 1, 2, \ldots$ according to the asynchronous occurrence of discrete events. Let $e_{u_i}^{i \to k}(t)$ denote the decision of a number u_i of robots to leave node $i \in N$ to join a neighboring node $k \in N_i$ at time t. Let $e_{u_i}^{i \to N_i}(t)$ denote the set of all possible simultaneous decisions from node i to its neighboring nodes N_i. The set of events $\mathcal{E} = \mathcal{P}(\{e_{u_i}^{i \to N_i}(t)\}) - \{\emptyset\}$ represents all possible simultaneous decisions from all nodes. A single event $e(t) \in \mathcal{E}$ is defined as a set where each element represents a decision of a number of robots to abandon the nodes (see robot movements in the graph model between connected nodes in Figure 2).

If an event $e(t) \in \mathcal{E}$ occurs at time t, the update of the state of the system is given by $r(t+1) = g(r(t))$. For the robots belonging to node $i \in N$, $g(r(t))$ is defined as

$$r_i(t+1) = r_i(t) - \sum_{\{k: e_{u_i}^{i \to k}(t) \in e(t)\}} u_i(t) + \sum_{\{j: e_{u_j}^{j \to i}(t) \in e(t)\}} u_j(t) \quad (5)$$

In other words, the current amount minus those leaving plus some arriving. Solving (1) requires that the model satisfy the following assumptions on the network and its robots:

A4 The nodes are in a connected graph G_n.

A5 There is a large enough number of robots, q, such that there can be at least a robot within each node providing a positive utility when they reach the optimal distribution $r \in \Delta_q^{\star}$.

Assumption A4 implies that there is a path across all locations of the graph, placing minimum conditions on the sensing and possible decisions of robots across them. Assumption A5 requires a minimum number of robots, which, in general, depends on the nature of the utility functions. Moreover, under Assumptions A4 and A5, the optimal solution takes the form

$$\Delta_q^{\star} = \{r \in \Delta_q | \, \forall i \in N, \, \forall k \in N_i, \, s_i(r_i) = s_k(r_k)\} \quad (6)$$

Thus, for any finite number of robots, the optimal distribution $r \in \Delta_q^{\star}$ is unique (Bertsekas, 1999). The distribution $r \in \Delta_q^{\star}$ captures the optimal division of subgroups when all of them have the same marginal utility because no robot has incentives to abandon its node. How robots decide which node to serve is the focus of the next section.

4.3 Deterministic Decision-making

In this section, we detail how environmental aids regulate robot movements and a variation where robots recover their autonomy to decide which region to serve by using these aids information. Both decision-making models consider marginal utilities as the average number of packages delivered within a region's period. After a fixed period or interval of time, TAMs update robots information such that they can restore their performance metrics and update their information to make their decisions.

4.3.1 Deterministic Decision-making (D)

Movements may be stochastic, but any $e_{u_i}^{i \to N_i}(t) \in e(t)$ must satisfy the following rules in this model.

D-R1 If $s_i(r_i(t)) \geq s_j(r_j(t))$ for all $i \in N_i$, then $u_i(t) = 0$, i.e., robots remain in a node where they have a better marginal utility.

D-R2 If there exists a node $j \in N_i$ such that $s_i(r_i(t)) < s_j(r_j(t))$, then some robots could decide to abandon i to serve in the neighboring node j, which has the highest marginal utility among the neighbors of node i, N_i. In particular, the number of robots leaving node i is bounded by

$$0 < u_i(t) \leq \frac{1}{2}\phi\,[s_k(r_k(t)) - s_i(r_i(t))] \quad (7)$$

where $\phi \in (0, 1/a]$ represents the level of cooperation between robots, and $\forall j, k \in N_i$ $s_{k\ell}(r_k(t)) \geq s_j(r_j(t))$.

Rules D-R1 and D-R2 restrain the allowable events in the network. In particular, D-R2 captures the tendency of robots to join a node that has a higher marginal utility value than all other neighboring nodes. Within a particular node, robots show divisional autonomy in the sense that they are unconstrained in their decisions to serve that node. However, in this deterministic solution, nodes regulate robot transitions; they choose which robot must depart bereaving them of node-to-node movements. It would be as if in a warehouse scenario, robots serving within a room could not move between rooms unless the room indicates to do so.

4.3.2 Semi-stochastic Decision-making (SS)

We allowed robots to choose the node where they want to serve, that is, nodes have no authority upon them. In this case, we needed to add the following assumptions on communication:

A6 Each node offers information of its local utility to robots working on it.

A7 If there is a connection between two nodes, that is, $j \in N_i$, then there exists at least one link communicating a robot from node i with some robot in node j.

Assumptions A6 and A7 are guaranteeing communication, which is a critical requirement in our robotic task allocation. Robots need information about neighboring nodes to decide which node is the best for them. Note that Assumption A7 is a local version of Assumption A4; it allows robots to receive information about performances and local utilities in neighboring nodes. A robot serving in the neighboring nodes shares this information. Thus, robots can compare the options and make a decision.

Since robots following the semi-stochastic decision-making have no node regulation, they are free to move across nodes. Note that when robots detect another node with a better utility, they could depart massively and leave a node empty. Therefore, we have to design local rules to avoid such massive movements, i.e., to regulate robot movements such that some robots remain in a node, even when there is another node with a better utility. Thus, robots would need to know or at least estimate how many of them are working in that node. However, knowing this information means they would have a kind of global knowledge, which is commonly unavailable in swarm robotics (Şahin, 2005). However, they can estimate or guess how many are in the same node. Each robot keeps track of its performance at node i in $s_i^\ell(t)$, which relates to its marginal contribution. Since we are working with homogeneous robots, a robot ℓ can estimate how many of them are foraging in the same region by using the marginal utility of its region, $s_i(r_i(t))$, and its own performance, $s_i^\ell(t)$. Despite robots have a similar performance, it is not exactly the same. Thus, for robot ℓ, let $\hat{r}_i^\ell(t) = s_i(r_i(t))/s_i^\ell(t)$ represent its estimation of the number of robots working in the same node.

Next, let $p_{i\to j}^\ell(t)$ be the probability of robot ℓ departing from node i toward node j, which offers a better utility. When robots follow the proposed decision-making strategy, any movement $e_{u_i}^{i\to N_i}(t) \in e(t)$ must also satisfy the following rules.

S-R1 If $s_i(r_i(t)) = s_i^\ell(t)$ (i.e., $\hat{r}_i^\ell(t) = 1$), then $p_{i\to j}^\ell(t) = 0$, i.e., that robot is the only one serving there and must remain in it even when there is a node with a better utility.

S-R2 If $s_i(r_i(t)) > s_i^\ell(t)$ (i.e., $\hat{r}_i^\ell(t) > 1$), then $p_{i\to j}^\ell(t) > 0$, i.e., that robot should compute its probability of departing toward node k with a better utility. If robot ℓ is considering departing, it means that node i has a neighboring node $k \in N_i$ such that $s_k(r_k(t)) > s_i(r_i(t))$. Then, some robots could decide to abandon node i at the same time. The probability of robot ℓ departing is given by

$$p_{i\to j}^\ell(t) = \frac{1}{2}\phi_\ell \frac{[s_k(r_k(t)) - s_i(r_i(t))]}{\hat{r}_i^\ell(t)} \qquad (8)$$

where ϕ_ℓ is its level of cooperation.

Note that, unlike the deterministic solution, robots use probability functions, and therefore, there is no longer a warranty of achieving the optimal distribution. However, robots can reach a near-optimal solution as in (Krause and Guestrin, 2007).

4.4 Learning Mechanism

In this work, robots are learning about their performance in each region. They keep a historical estimation of the number of objects they transported. Recall that if they did not work in a region, information from other robots working there could spread up to them. TAMs are helping in these information-sharing processes by keeping track of the historical performance of the region where they are and by indicating robots the end of a period when they update local variables. This idea resembles the way enterprises pay their employees: they are asynchronous and deliver the payments at the end of a period of work (e.g., 15 days or monthly). Once TAMs indicate the end of a period in their region, robots will consider how many objects they transported as their marginal contribution.

Next, a robot ℓ uses the following exponential moving average function for learning updates with its own measurements and with incoming messages

$$s_i^\ell(t) = \alpha * s_i^\ell(t-1) + (1-\alpha) * M_i^\ell \qquad (9)$$

where $\alpha \in (0,1]$ is the rate of learning, M_i^ℓ is the score the robot measured (listened) during that period, and $s_i^\ell(t)$ defines the performance estimation that robot ℓ has of working at node i. Robots are sharing this estimation with neighboring robots, that is, those robots with which they have a connection. Therefore, changes in the communication structure may affect the task allocation process because they are affecting the learning process. Note that if robots have no path of connections with other regions, their only possible reality is that their current region is the best. They would never got information about other regions.

The following section describes the preliminary analysis to configure the environment parameters, simulations, and the results of both decision-making strategies that performed the foraging task in the proposed environment.

5 EXPERIMENTS

Initially, we needed to find the marginal utility functions of the environment. Recall that to find the optimal distribution point, $\Delta_q^\star$ in Eq. (6), it is necessary to have the marginal utility functions, because the optimal distribution is reached when all marginal utilities have the same value. We measured the marginal utility as the number of foraged objects in each region within a period. Since robots are transitioning between regions, we established a periodical evaluation of the marginal utilities every 1,000 steps. We ran our simulations in Netlogo and fixed cooperation for all robots to 1.

Thus, if the system follows the law of diminishing returns, the marginal utility of a region should decrease by adding more robots into that region (node). Since each region has the same dimensions, we chose one of them and variated the rate of recovery in that region. This rate indicates the capacity of sources to restore objects as soon as robots remove them. However, a source evaluates if it can recover the object every period (1,000 steps). The rate of recovery of the objects removed by robots can be 20%, 40%, 60%, 80%, and 100%. Figure 3 shows the marginal utility functions with each rate of recovery (30 simulations for each rate of recovery).

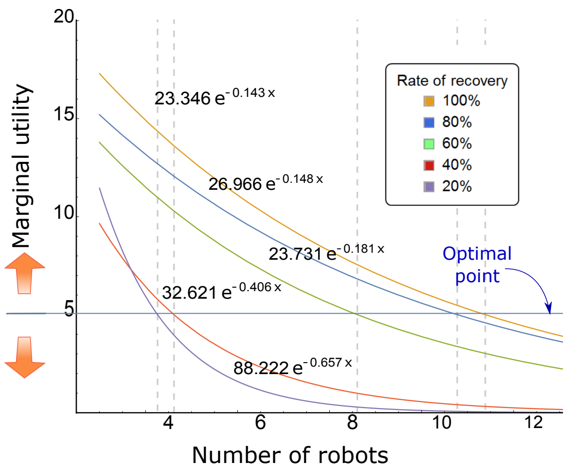

Figure 3: Marginal utility functions for the different rates of recovery of the sources.

By moving the point of optimal distribution and using the functions, we could get an estimation of the quantity of robots the environment requires. For instance, working with the optimal point of Figure 3 and defining the rates of the regions as follows: the first region (red ground) with 20% needs 4 robots, the second (gray ground) with 40% needs 4 robots, and the third (blue ground) with 80% needs 11. We would need 19 robots with these rates of recovery. However, the initial distribution is different from this optimal one so that the decision-making rules lead them toward it.

Next, for the simulations, we have a deterministic model (D), the semi-stochastic with (SS-TAM) and without TAMs' help (SS). In the models with information sharing available (i.e., working with Assumptions A6 and A7), we tested three different networks in robot-robot interaction: Fully connected and two regular networks with degree 3 and 1. We also tested a variation with switching-links strategy in the information-sharing structure to observe the effects of changing neighbors. In particular, after a robot transition, they could abandon previous neighbors and connect to some of those in the new region. For this variation, we added a prefix (-S). The following sections describes the results in two groups of different sizes.

5.1 Simulation Results

First, we settle the regions rate such that the optimal distribution point is 2-5-9 (i.e., 2 in the red region, 5 for the gray region, and 9 for the blue region). This indicates that we need 16 robots for this configuration. With this number of robots, we tried different initial conditions avoiding the optimal one. For reasons of space, we opted for showing only the results with the initial distribution 5-5-6, almost the same number of robots in each region. Note that only three robots should move from red to blue region, but system decentralization and robot autonomy generated more travels. Figure 4 shows the results of 30 simulations of 10,000 steps for all the models with and without switching-links strategy.

Note that the deterministic model delivered the best performance; t-test showed that it has no competitor. SS-TAM-S with a network of degree 3 got the second place (losing by an average of 18 objects, with $p = 0.001$). Although the box-plots show that the variation in the results with the switching-links strategy decreases, the t-tests indicate there is no significant deterioration in introducing this strategy in any model with any network structure. On the other hand, it means that robots could work with local-range communication.

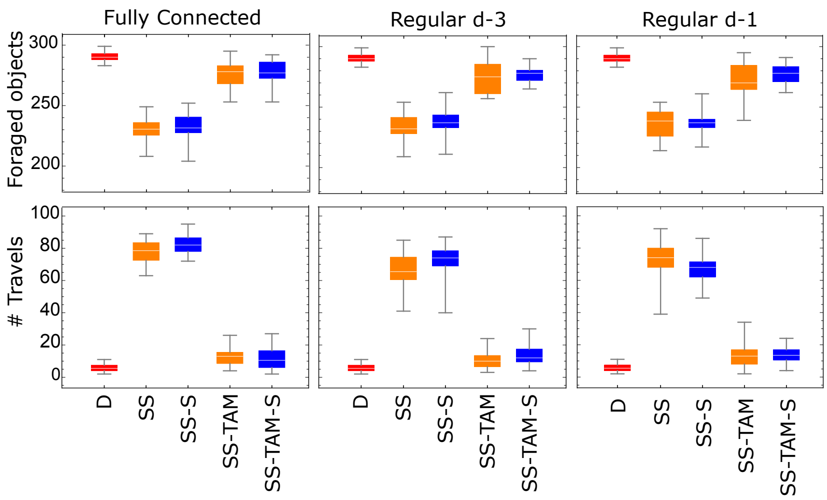

Figure 4: Performance and travels of robots while foraging in the environment with different network structures. Desired distribution 2-5-9.

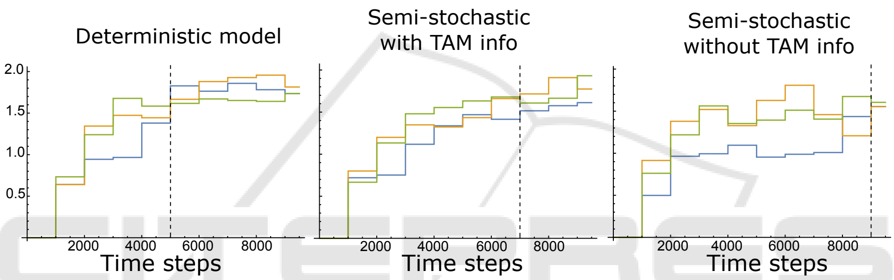

Figure 5: Evolution of the performance with a period of update in TAMs of 1,000 steps. The dashed lines represent the moment when the marginal utilities reached a value near to the optimal distribution.

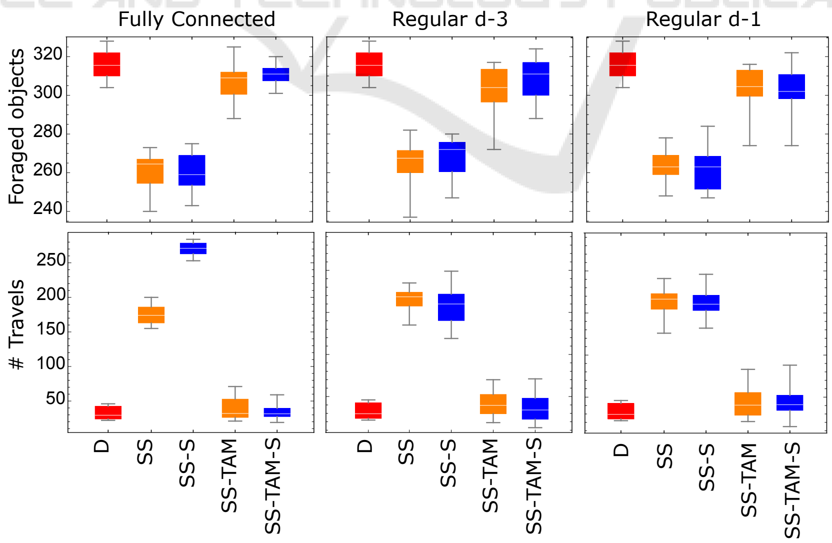

Figure 6: Performance and travels of robots while foraging in the environment with different network structures. Desired distribution 5-11-16.

Note also that robots with lower amount of travels obtained the best performances. In the SS model, which did not include the TAMs' information, robots required more travels because they only have the information shared among neighbors. In other words, without Assumption A6, robots invested more time traveling than foraging due to lack of information. In contrast to SS, robots foraging with SS-TAM obtain better results due to more accurate information received from TAMs. Moreover, they have more autonomy than those foraging with the deterministic model. Although they were considering only an estimation of the number of partners, they reached a near-optimal performance.

Next, to see which mechanism was near to the optimal distribution point $\Delta_q^{\star}$, we computed the Euclidean distance of the last distribution. We allowed simulations to run about 4 periods of 1,000 steps after the deterministic model reached a value near the equilibrium. Table 1 shows the results of the Euclidean distance.

Table 1: Values of the Euclidean distance of the implemented models and the computed optimal distribution.

	Structure		
Model	Fully	Reg. d-3	Reg. d-1
D	1.31	1.31	1.31
SS	1.75	1.72	3.42
SS-S	2.23	2.39	2.60
SS-TAM	1.71	1.04	1.74
SS-TAM-S	1.53	0.80	1.48

The SS and SS-TAM models (with and without TAM information, respectively) delivered a value near to the optimal distribution. To know the reason behind their difference in performance, we observed the evolution of the marginal utilities. The SS-TAM model was faster in reaching such value near the optimal distribution. Since the only difference between SS and SS-TAM is the information available in the TAMs of each region, we can affirm that TAMs were fundamental for a better performance in SS-TAM. Figure 5 shows the settling time for the evolution of marginal utilities of a simulation with each model.

It is clear that the faster model to reach the balance is the deterministic one. However, we reached a near-optimal solution where robots could keep their autonomy in making decisions to transition between regions without restriction. Although robots following the SS model without TAMs' help reached a good distribution, they required a large time due to diversity in the estimations of other robots.

We increased the number of robots by moving the optimal distribution point downward. We settled the regions rate in 20%, 40%, and 60% and the optimal distribution point became 5-11-16 (i.e., 5 in the red region, 11 for the gray region, and 16 for the blue region). This indicates that we need 32 robots for this configuration. With this group of robots, we tried different initial conditions avoiding the optimal robot distribution. In this occasion, we opted for 1-15-16 as initial distribution. The following paragraphs details the effects of working with a large group.

Again, the deterministic model delivered the best performance; t-test showed it has surpassed all other models. Only the SS-TAM-S model with a fully connected network was near, with a difference of 13 objects in the average performance and $p = 0.002$. Moreover, t-test results indicated that there is no significant difference between the models with and without the switching-link strategy. Nevertheless, it was observed a lower variability among the simulations (i.e., the box-plots seem shorter).

We arrived to similar conclusions of those obtained with the small group. The semi-stochastic model without TAMs' help worsened its performance because a greater number of robots increases the diversity in the estimations of each robot too. In the travel plots, we could confirm that SS has the greatest amount of travels and the lowest number of foraged objects. In other words, TAMs information was key to stimulate robots to focus on foraging and avoid too much travels between regions.

6 CONCLUSIONS

We adjusted a MAS strategy to work in a MRTA scenario successfully. The optimal solution provided by Bertsekas worked because the performance is non-linear; the robots exhibited a diminishing return curve. However, we could only reach near-optimal solutions. We offered a decentralized solution that allowed robots to keep their autonomy; however, it was necessary to add environmental aids. In particular, the information provided by TAMs was critical for robots to reach a better performance. TAMs helped robots to have a better reference for decentralized decisions (autonomy).

The different structures of robot-robot communication show that robots could forage with a minimal condition on information sharing, not necessarily a fully connected network. However, they would sacrifice a small portion of their performance. Therefore, it depends on the environment dimensions and the hardware embedded in the robots to decide between fully communication and highest performance, or scarcely connection and a still-good performance.

All decision-making models stimulated robots to work for the group goals (distributed control) through the diminishing return utility functions. Nevertheless, the time to reduce the distance to the optimal point indicates that it is important to achieve such point as faster as robots can. Although we have provided a good alternative for the deterministic solution, in future works, we want to explore alternatives to improve the speed to reach such point without depending on TAMs or losing autonomy.

Moreover, the phenomenon of diminishing returns is present in many scenarios where the incorporation of an additional worker to a job may improve the performance, but each additional worker increases (gradually) in smaller portions the performance. We have seen this kind of shapes in many previous works of MRTA because robots are sharing limited resources. Therefore, there is a great opportunity to adjust this same set of rules for those systems. It would be easily adjusted for each environment and its conditions.

ACKNOWLEDGEMENTS

GMBO is grateful to Fapemig, CNPq and CAPES for their financial support. Juan M Nogales also is grateful with the OEA for the scholarship.

REFERENCES

Ahuja, R. K., Ergun, Ö., Orlin, J. B., and Punnen, A. P. (2002). A survey of very large-scale neighborhood search techniques. *Discrete Applied Mathematics*, 123(1):75–102.

Akbarimajd, A. and Simzan, G. (2014). Application of artificial capital market in task allocation in multi-robot foraging. *Int. Journal of Computational Intelligence Systems*, 7(3):401–417.

Anderson, C. and Bartholdi, J. J. (2000). Centralized versus decentralized control in manufacturing: lessons from social insects. *Complexity and complex systems in industry*, pages 92–105.

Antoun, A., Valentini, G., Hocquard, E., Wiandt, B., Trianni, V., and Dorigo, M. (2016). Kilogrid: A modular virtualization environment for the kilobot robot. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 3809–3814, Daejeon, South Korea. IEEE.

Bertsekas, D. P. (1999). *Nonlinear programming*. Athena scientific Belmont.

Bobadilla, L., Martinez, F., Gobst, E., Gossman, K., and LaValle, S. M. (2012). Controlling wild mobile robots using virtual gates and discrete transitions. In *American Control Conf.*, pages 743–749. IEEE.

Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm intelligence: from natural to artificial systems*. Number 1. Oxford university press.

Bonabeau, E., Sobkowski, A., Theraulaz, G., Deneubourg, J.-L., et al. (1997). Adaptive task allocation inspired by a model of division of labor in social insects. In *Proc. Biocomputing and emergent computation*, pages 36–45, Skovde, Sweden.

Brutschy, A., Garattoni, L., Brambilla, M., Francesca, G., Pini, G., Dorigo, M., and Birattari, M. (2015). The TAM: abstracting complex tasks in swarm robotics research. *Swarm Intelligence*, 9(1):1–22.

Buchanan, E., Pomfret, A., and Timmis, J. (2016). Dynamic task partitioning for foraging robot swarms. *Proc. of the 10th Int. Conf. on Swarm Intelligence*, pages 113–124.

Cao, Y. U., Fukunaga, A. S., and Kahng, A. (1997). Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4(1):7–27.

Chaimowicz, L., Campos, M. F. M., and Kumar, V. (2002). Dynamic role assignment for cooperative robots. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, volume 1, pages 293–298 vol.1.

Dasgupta, P. (2011). *Multi-Robot Task Allocation for Performing Cooperative Foraging Tasks in an Initially Unknown Environment*, pages 5–20. Springer Berlin Heidelberg, Berlin, Heidelberg.

De Marco, R. and Farina, W. (2001). Changes in food source profitability affect the trophallactic and dance behavior of forager honeybees (apis mellifera). *Behavioral Ecology and Sociobiology*, 50(5):441–449.

Färe, R. (1980). *Laws of diminishing returns*. Number 176. Springer Verlag.

Gerkey, B. P. and Matari, M. J. (2004). A formal analysis and taxonomy of task allocation in multi-robot systems. *The Int. Journal of Robotics Research*, 23(9):939–954.

Jeanne, R. L. and Nordheim, E. V. (1996). Productivity in a social wasp: per capita output increases with swarm size. *Behavioral Ecology*, 7(1):43–48.

Khan, M. U., Li, S., Wang, Q., and Shao, Z. (2016). Cps oriented control design for networked surveillance robots with multiple physical constraints. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 35(5):778–791.

Krause, A. and Guestrin, C. (2007). Near-optimal observation selection using submodular functions. In *Proc. of the 22nd Conf. on Artificial Intelligence*, volume 7, pages 1650–1654.

Kube, C. and Bonabeau, E. (2000). Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1):85 – 101.

Lerman, K. and Galstyan, A. (2002). Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots*, 13(2):127–141.

Liu, Y. and Nejat, G. (2013). Robotic urban search and rescue: A survey from the control perspective. *Journal of Intelligent & Robotic Systems*, 72(2):147–165.

Marden, J. R., Arslan, G., and Shamma, J. S. (2009). Cooperative control and potential games. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(6):1393–1407.

Mataric, M. J. (1994). Reward functions for accelerated learning. In *Proc. of the Eleventh Int. Conf. on Machine Learning*, pages 181–189. Morgan Kaufmann.

Nogales, J. M. and Finke, J. (2013). Optimal distribution of heterogeneous agents under delays. In *2013 American Control Conf.*, pages 3212–3217.

Phan, C. and Liu, H. H. T. (2008). A cooperative uav/ugv platform for wildfire detection and fighting. In *Proc. of the 7th Int. Asia Simulation Conf. on System Simulation and Scientific Computing*, pages 494–498.

Pini, G., Brutschy, A., Scheidler, A., Dorigo, M., and Birattari, M. (2014). Task partitioning in a robot swarm: Object retrieval as a sequence of subtasks with direct object transfer. *Artificial life*, 20(3):291–317.

Pitonakova, L., Crowder, R., and Bullock, S. (2016). Task allocation in foraging robot swarms: The role of information sharing. *Proc. of the Artificial Life Conf.*, page 306.

Raty, T. D. (2010). Survey on contemporary remote surveillance systems for public safety. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 40(5):493–515.

Şahin, E. (2005). *Swarm Robotics: From Sources of Inspiration to Domains of Application*. Springer Berlin Heidelberg, Santa Monica, CA.

Schmickl, T., Thenius, R., and Crailsheim, K. (2012). Swarm-intelligent foraging in honeybees: benefits and costs of task-partitioning and environmental fluctuations. *Neural Computing and Applications*, 21(2):251–268.

Seeley, T. D., Visscher, P. K., and Passino, K. M. (2006). Group decision making in honey bee swarms. *American Scientist*, 94(3):220–229.

Stone, P. and Veloso, M. (1999). Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence*, 110(2):241 – 273.

Sung, C., Ayanian, N., and Rus, D. (2013). Improving the performance of multi-robot systems by task switching. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 2999–3006.

Viguria, A., Maza, I., and Ollero, A. (2007). Set: An algorithm for distributed multirobot task allocation with dynamic negotiation based on task subsets. In *Proc. of IEEE Int. Conf. on Robotics and Automation*, pages 3339–3344.

Wawerla, J. and Vaughan, R. T. (2010). A fast and frugal method for team-task allocation in a multi-robot transportation system. In *Proc. of the IEEE Conf. on Robotics and Automation*, pages 1432–1437. IEEE.

Yan, Z., Jouandeau, N., and Ali-Chrif, A. (2012). Multi-robot heuristic goods transportation. In *6th IEEE Int. Conf. Intelligent Systems*, pages 409–414.

Yan, Z., Jouandeau, N., and Cherif, A. A. (2013). A survey and analysis of multi-robot coordination. *Int. Journal of Advanced Robotic Systems*, 10(12):399.

Zahadat, P., Hahshold, S., Thenius, R., Crailsheim, K., and Schmickl, T. (2015). From honeybees to robots and back: division of labour based on partitioning social inhibition. *Bioinspiration & biomimmics*, 10(6):66005.