

An Ant Based Algorithm for Task Allocation in Large-Scale and Dynamic Multiagent Scenarios

Fernando dos Santos and Ana L. C. Bazzan
PPGC – Universidade Federal do Rio Grande do Sul
Caixa Postal 15064, CEP 91501-970, Porto Alegre, RS, Brasil
{fsantos,bazzan}@inf.ufrgs.br

ABSTRACT

This paper addresses the problem of multiagent task allocation in extreme teams. An extreme team is composed by a large number of agents with overlapping functionality operating in dynamic environments with possible inter-task constraints. We present eXtreme-Ants, an approximate algorithm for task allocation in extreme teams. The algorithm is inspired by the division of labor in social insects and in the process of recruitment for cooperative transport observed in ant colonies. Division of labor offers fast and efficient decision-making, while the recruitment ensures the allocation of tasks that require simultaneous execution. We compare eXtreme-Ants with two other algorithms for task allocation in extreme teams and we show that it achieves balanced efficiency regarding quality of the solution, communication, and computational effort.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multiagent systems

General Terms

Algorithms

Keywords

Artificial Intelligence, Multiagent Systems, Swarm Intelligence

1. INTRODUCTION

This paper deals with the issue of how to efficiently allocate tasks among agents in large-scale and dynamic multiagent environments. A large-scale environment means thousands of agents that must coordinate themselves to perform the available tasks. Scerri *et al.* [9] call these scenarios *extreme teams*. Task allocation in extreme teams is associated with four features: (i) dynamic environments, in which task

can appear and disappear; (ii) agents are able to perform multiple tasks given resources that are available to them; (iii) agents have overlapping functionality to perform the tasks but with differing levels of capability; and (iv) inter-task constraints can be present, imposing simultaneous execution requirements.

Extreme teams can be formalized as an extended generalized assignment problem (E-GAP) [9]. The E-GAP model captures precisely the characteristics of extreme teams and defines the solution as the allocation that maximizes the reward, measured by the capabilities of the agents that take part in the allocation. Efficient multiagent techniques to deal with E-GAP are a prerequisite to build teams of robots to act in extreme situations.

Besides the reward, the communication channel must be used in an efficient way in order to avoid blocking/overloading the channel. Moreover, the computational effort employed by the agents to decide which tasks to perform must be as low as possible, enabling them to act in environments where the time available to make a decision is highly restricted.

Social insects (e.g., ants) have the characteristics of extreme teams. Thus, we can conclude that Nature (despite the simplicity of the insects) has provided these insects with the capability to effectively act in these teams.

To perform tasks that are related to the survival of the colony, social insects adopt a division of labor among workers. Theraulaz *et al.* [10] present a mathematical model to replicate some mechanics of division of labor. Their model is based on individual response thresholds and tasks stimuli. Moreover, it is not required that individuals have complete information about the environment and there is no need of team leaders.

Simultaneous execution of tasks also exists in social insects, as for instance in some species of ants. The task in question is the transportation of large prey. Instead of seizing and transporting a large prey individually, some species form groups of ants to cooperatively transport a prey. These groups are formed via a process called recruitment [4]. In this sense, the large prey can be seen as a set of interdependent subtasks, where each one is simultaneously executed by an ant.

We propose a multiagent approximate task allocation algorithm, called eXtreme-Ants, which is inspired by the division of labor in social insects and in the process of recruitment present in ants. Agents running eXtreme-Ants are efficient to act in extreme teams, with low computational effort and communication. We evaluate eXtreme-Ants empirically

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'09, July 8–12, 2009, Montréal, Québec, Canada.
Copyright 2009 ACM 978-1-60558-325-9/09/07 ...\$5.00.

in a domain-independent simulator and compare it with two other algorithms that are also GAP-based: Swarm-GAP [2] and LA-DCOP [9]. The algorithm eXtreme-Ants achieves rewards close to the ones achieved by LA-DCOP, but with lower communication and computational effort. Regarding Swarm-GAP, eXtreme-Ants yields better rewards, particularly in the presence of inter-task constraints that impose simultaneous execution.

The remaining of this paper is organized as follows. Section 2 discusses the GAP and its extension (E-GAP), the model of division of labor in social insects, and the process of recruitment used by ants to cooperatively transport large prey. Section 3 details two other algorithms for dealing with GAP-based task allocation, Swarm-GAP and LA-DCOP. Section 4 presents the proposed algorithm. Section 5 discusses the empirical evaluation via a series of experiments. Section 6 points out the conclusion and future directions.

2. BACKGROUND

2.1 GAP and EGAP models

The generalized assignment problem (GAP) is a model used to formalize the multiagent task allocation problem [5]. A GAP is composed by a set \mathcal{J} of tasks to be performed by a set \mathcal{I} of agents. Each agent $i \in \mathcal{I}$ has a capability to perform each task $j \in \mathcal{J}$ denoted by $Cap(i, j) \rightarrow [0, 1]$. Each agent also has a limited amount of resource $i.res$ and uses $Res(i, j)$ when performs task j . An allocation matrix M is used to represent the allocation, where m_{ij} is given by Equation 1.

$$m_{ij} = \begin{cases} 1 & \text{if } i \text{ performs } j \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

The goal is to find M that maximizes the allocation reward, which is given by the agents' capabilities, as shown in Equation 2.

$$M = \underset{M'}{\operatorname{argmax}} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} Cap(i, j) \times m'_{ij} \quad (2)$$

The allocation M must respect all agents' resources limitations (Equation 3) and each task must be allocated to at most one agent (Equation 4).

$$\forall i \in \mathcal{I}, \sum_{j \in \mathcal{J}} Res(i, j) \times m_{ij} \leq i.res \quad (3)$$

$$\forall j \in \mathcal{J}, \sum_{i \in \mathcal{I}} m_{ij} \leq 1 \quad (4)$$

The GAP model was extended by Scerri *et al* [9] to incorporate two features related to extreme teams: scenario dynamics and inter-task constraints. This extended model was called extended generalized assignment problem (E-GAP).

Inter-task constraints are interdependencies among tasks. We focus on AND constraints here, but the formalization can be extended to other constraint types as well. In the case of an AND constraint, the agents only receive the reward if all constrained tasks are *simultaneously* executed. The AND constrained tasks can be viewed as a decomposition of a large task into interdependent subtasks. The execution of only some subtasks does not lead to the successful execution of the large task, wasting the agents' resources and not producing the desired effect in the system. Moreover, in the case of physical robots, they can be damaged attempting to

perform an effort greater than their capabilities (e.g., trying to remove a large piece of collapsed building from a blocked road).

In order to formalize AND constrained tasks, the E-GAP model defines a set $\bowtie = \{\alpha_1, \dots, \alpha_p\}$ containing p sets α_k of AND constrained tasks in the form $\alpha_k = \{j_1 \wedge \dots \wedge j_q\}$. Each AND constrained task j belongs to at most one set α_k . The number of tasks that are being performed in a set α_k is given by Equation 5.

$$x_k = \sum_{i \in \mathcal{I}} \sum_{j \in \alpha_k} m_{ij} \quad (5)$$

Let $v_{ij} = Cap(i, j) \times m_{ij}$. Given the constraints of \bowtie , the reward $Val(i, j, \bowtie)$ of an agent i performing a task j is given by Equation 6.

$$Val(i, j, \bowtie) = \begin{cases} v_{ij} & \text{if } \forall \alpha_k \in \bowtie, j \notin \alpha_k \\ v_{ij} & \text{if } \exists \alpha_k \in \bowtie \text{ with } j \in \alpha_k \wedge x_k = |\alpha_k| \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

To represent the dynamics in the scenario all E-GAP variables are indexed by a time step t . The goal is to find a sequence of allocations \vec{M} , one for each time step t , as shown in Equation 7. A delay cost function $DC^t(j^t)$ can be used to define the cost of not performing a task j at time step t .

$$f(\vec{M}) = \sum_t \sum_{i^t \in \mathcal{I}^t} \sum_{j^t \in \mathcal{J}^t} (Val^t(i^t, j^t, \bowtie^t) \times m_{ij}^t) - \sum_t \sum_{j^t \in \mathcal{J}^t} (1 - \sum_{i^t \in \mathcal{I}^t} m_{ij}^t) \times DC^t(j^t) \quad (7)$$

Furthermore, the agents' resource limitations must be respected at each time step t (Equation 8) and each task must be allocated to at most one agent (Equation 9).

$$\forall t, \forall i^t \in \mathcal{I}^t, \sum_{j^t \in \mathcal{J}^t} Res^t(i^t, j^t) \times m_{ij}^t \leq i^t.res \quad (8)$$

$$\forall t, \forall j^t \in \mathcal{J}^t, \sum_{i^t \in \mathcal{I}^t} m_{ij}^t \leq 1 \quad (9)$$

2.2 Division of Labor in Social Insects

An effective division of labor is responsible for the ecological success of insect societies. A social insect colony with hundreds of thousands of members operates without explicit coordination. An individual cannot assess the needs of the colony; it has just a fairly simple local information. From individual workers aggregation, the complex colony behavior emerges. The key point is the plasticity of the individuals, in other words, the existence of a behavioral flexibility. This flexibility allows the individuals to engage in different tasks responding to changing conditions in the colony [1].

This behavior is the basis of the theoretical model described by Theraulaz *et al.* [10]. In this model, interactions among members of the colony and individual perception of local needs result in a dynamic distribution of tasks. The model is based on individuals' internal response threshold related to tasks stimuli. Assuming the existence of \mathcal{J} tasks to be performed, each task $j \in \mathcal{J}$ has an associated stimulus s_j . The stimulus is related to the demand for the task execution, and can be a number of encounters, a chemical concentration, or any quantitative cue sensed by individuals. Given a set of \mathcal{I} individuals which can perform the

tasks in \mathcal{J} , each individual $i \in \mathcal{I}$ has an internal response threshold θ_{ij} , which is related to the likelihood of reacting to the stimulus associated with task j . The threshold can be seen as a genetic characteristic (also called polymorphism, which is responsible for the existence of differences in the morphologies of insects belonging the same society), or as a temporal polyethism (in which individuals of the same age tend to perform identical sets of tasks), or simply as individual variability.

In the model of Theraulaz *et al.* [10] the individual internal threshold θ_{ij} and the task stimulus s_j are used to compute the probability (tendency) $T_{ij}(s_j)$ of the individual i to perform task j , as shown in Equation 10.

$$T_{ij}(s_j) = \frac{s_j^2}{s_j^2 + \theta_{ij}^2} \quad (10)$$

The use of the tendency replicates the plasticity of individuals because any individual is able to perform any task if the corresponding task stimulus is high enough to overcome the individual's internal threshold. This flexibility enables the survival of the colony in an eventual absence of specialized individuals, since other individuals start to perform the tasks as soon as the stimulus exceeds their thresholds.

2.3 Recruitment for Cooperative Transport

In some species of ants the transportation of large prey in a cooperative way involves two or more ants that cannot do the transport alone [8]. The main purpose of the cooperative transport is to maximize the trade-off between the gained energy (food) and the energy spent to take the prey to the nest. Further, this process speeds up the transport.

The group involved in the cooperative transport is formed by a process called *recruitment*. When a single scout ant discovers a prey, it firstly attempts to seize and transport it individually. After unsuccessful attempts, the recruitment process starts. In order to recruit nestmates, ants employ a mechanism called long-range recruitment (LRR). Some species also employ a second mechanism called short-range recruitment (SRR). In both mechanisms the ants use communication through the environment (stigmergy) [4].

In the SRR the scout that discovers the prey releases a secretion. Shortly thereafter, nestmates in the vicinity are attracted to the prey location by the secretion odor. When the prey cannot be moved by the scout and the ants recruited via SRR, one of the ants begins the LRR. Hölldobler *et al.* report that SRR is sufficient to summon enough ants to transport the prey in the majority of the cases [4].

In the LRR the scout ant that discovers the prey returns to the nest to recruit nestmates. In the course towards the nest, the scout lays a pheromone trail. Nestmates encountered in the course are stimulated by the scout via direct contact (e.g. antennation, trophallaxis), thus establishing a chain of communication among nestmates. After, the stimulated nestmate also begins to lay a pheromone trail even though it had not yet experienced the prey stimulus itself. When the scout arrives at the nest, nestmates are attracted by the pheromone and run to the prey site.

After the recruited ants arrive where the prey is, they begin the cooperative transport. The number of ants engaged in the transport is regulated at the prey site and depends on its characteristics, such as weight, size, rotational forces, and difficulty to move [3]. When the number of ants present

is not enough to move the prey, more ants are recruited by one of the aforementioned processes, until the prey is successfully transported. Although in a more economic approach the scout ant should recruit the exact number or nestmates, it was suggested that the scout cannot make a fine assessment of the number of ants required to retrieve the prey [8]. Therefore, the most effective strategy may be to recruit a constant number of ants followed by a regulation of the group size during the transport.

In summary, the recruitment for cooperative transport consists in three steps:

1. The scout ant that has discovered the prey starts the recruitment, inviting nestmates with pheromones;
2. The ants that accept to join the recruitment move to the prey site;
3. The size of the transportation group is regulated according to the prey characteristics.

3. RELATED WORK

The research regarding multiagent task allocation has shown significant advances in the last few years (auction, contracting, coalition formation, organizations, etc). A complete review of the subject is outside the scope of this paper. We just mention that auctions are normally centralized mechanisms, in which agents put bids to an auctioneer, depending on their capabilities and resources. After receiving all bids, the auctioneer makes the allocation of the tasks among the bidders. Centralized auctioneers can have severe bottlenecks. Further, auctions require high amounts of communication [11].

Here we concentrate in the line of research that deals with coordination for task allocation. Within this research line, one approach is the framework of distributed constraint optimization problem (DCOP). A DCOP consists of a set of variables that can assume values from a discrete domain. Each variable is assigned to one agent which has the control over its value. The goal of the agents is to choose values for the variables to optimize a global objective function. This function can be described as an aggregation over a set of cost functions related to pairs of variables. A DCOP can be represented by a constraint graph, where vertices are variables and edges are cost functions between variables. An E-GAP can be converted to a DCOP taking agents as variables and tasks as the domain of values. However, this modeling leads to dense constraint graphs, since one cost function is required between each pair of variables to ensure the mutual exclusion constraint of the E-GAP model. Despite the existence of complete algorithms for DCOP such as Adopt [6] and DPOP [7], these are not efficient to deal with the problem of multiagent task allocation. Adopt and DPOP demand high communication and space respectively.

In order to deal with the particular characteristics of extreme teams, Scerri *et al.* present an approximate algorithm called Low-communication Approximate DCOP (LA-DCOP), which uses tokens to represent tasks and further minimize communication [9]. An agent decides whether or not to perform a task based both on its capability and on a threshold associated to the task. To deal with inter-task constraints, LA-DCOP uses a differentiated kind of token, called potential token. If an agent in LA-DCOP is able to

perform more than one task, it must select those that maximize its capability given its resources. This selection is a maximization problem, which can be reduced to a binary knapsack problem (BKP), proved to be NP-Complete. The computational complexity of LA-DCOP thus depends of the complexity of its function to deal with the BKP.

Swarm-GAP is another approximate algorithm which can deal with extreme teams [2]. An agent in Swarm-GAP decides whether or not to perform a task based on the model of division of labor used by social insects colonies. This algorithm also uses tokens to represent tasks. To deal with inter-task constraints, the agents in Swarm-GAP just increase the tendency to allocate a constrained task by a factor called execution coefficient. The execution coefficient is computed using the rate between the number of constrained tasks which are allocated and the total number of constrained tasks.

4. eXtreme-Ants

4.1 Basic Ideas

eXtreme-Ants is an approximate algorithm that solves E-GAPs. Agents running eXtreme-Ants use the model of division of labor in social insects (Equation 10) to decide whether or not to perform the tasks. The notation used hereafter to represent agents, tasks, and all other terms is the one from the EGAP model (Section 2.1) and from the model of division of labor (Section 2.2). The internal threshold θ_{ij} of an agent i towards a task j is defined via the concept of polymorphism and corresponds to the inverse of the capability $Cap(i, j)$, as shown in Equation 11. If an agent is not capable regarding a particular task, then its internal threshold is set to infinity, avoiding the allocation of the task to the agent. This makes sense if we consider the capability as a kind of morphism. For example, a fire brigade agent is more capable of fighting fires than rescuing civilians. Thus it have low thresholds related to fire fighting tasks and high thresholds to rescue civilians.

$$\theta_{ij} = \begin{cases} 1 - Cap(i, j) & \text{if } Cap(i, j) > 0 \\ \infty & \text{otherwise} \end{cases} \quad (11)$$

Each task $j \in \mathcal{J}$ has an associated stimulus s_j . The stimulus controls the execution of the tasks by the agents. Low stimuli mean that the tasks will only be performed by agents with low thresholds (thus, more capable). High stimuli increase the chance of tasks being performed, even by agents with high thresholds (less capable).

Since in the E-GAP each task must be allocated to at most one agent, eXtreme-Ants uses tokens to represent this mutual exclusion constraint. A token contains a list of tasks it represents. An agent that holds a token has the exclusive right to perform the tasks contained in the token. If the agent does not perform all tasks, it passes the token to another agent. This way, eXtreme-Ants avoid conflicts in the allocation and reduces the communication.

To deal with AND constraints among tasks, agents in eXtreme-Ants reproduce the recruitment process of ants. When an agent detects that it is not capable of performing all AND constrained tasks perceived, it recruits other agents to form a group committed with the execution.

Algorithm 1: eXtreme-Ants for agent i

```

1 when perceived set of tasks  $\mathcal{J}$ 
2    $token := newToken()$ ;
3   add each  $j \in \mathcal{J}$  to  $token.tasks$ ;
4   evaluateToken( $token$ );

5 when perceived set of AND constrained tasks  $\alpha_k$ 
6   if has enough resources to perform all tasks of  $\alpha_k$ 
7     then
8       foreach  $j \in \alpha_k$  do
9         if  $roulette() < T_{ij}$  then
10           accept to perform task  $j$ ;
11       if all tasks of  $\alpha_k$  were accepted then
12         perform the tasks and decrease  $i.res$ ;
13       else
14         discard accepts of all tasks in  $\alpha_k$  (lines 7-9);
15         performRecruitment( $\alpha_k$ );
16       else
17         performRecruitment( $\alpha_k$ );

17 when received token
18   evaluateToken( $token$ );

19 procedure evaluateToken( $token$ )
20   /* decides whether or not to perform the tasks */
21   foreach  $j \in token.tasks$  do
22     if  $roulette() < T_{ij}$  and  $i.res \geq Res(i, j)$  then
23       accept task  $j$  and decrease  $i.res$ ;
24        $token.tasks := token.tasks - j$ ;
25   if there are non accepted tasks in  $token.tasks$  then
26     send  $token$  to a teammate;
```

4.2 Algorithm Details

Algorithms 1 and 2 present the details of our approach. Each agent a_i reacts to two events: perception of a set of tasks¹ (which can be AND constrained), and receipt of messages. In the following we detail the algorithm operation.

When an agent perceives a set \mathcal{J} of tasks (line 1) it creates a token to store the perceived tasks. The agent then decides whether or not to perform the tasks contained in the token, given its tendency and the available resources (lines 19-26). When a task is performed by an agent i , the available resources at i is decreased by the amount required. If some tasks contained in the token were not performed, the agent sends the token to a randomly selected agent. As in [9], to avoid agents passing token back and forth, each token maintains a list of visited agents. The token can revisit an agent only after all were visited.

When an agent perceives a set of AND constrained tasks α_k (line 5), it acts as a scout ant. Firstly it attempts to perform all the constrained tasks. If it fails, it begins a recruitment process. We develop a protocol that reproduces the three steps of the recruitment process of ants via the use of messages. There are five kinds of messages used in the recruitment protocol of eXtreme-Ants:

request: to invite an agent to join the recruitment for a task $j \in \alpha_k$ and to commit to it;

¹We use the term “set of tasks” because our approach does not constraint the number of tasks an agent can perceive on each time step.

Algorithm 2: eXtreme-Ants for agent i (cont.)

```
27 procedure performsRecruitment(AND set  $\alpha_k$ )
28   repeat
29      $j :=$  pick a task from  $\alpha_k$ ;
30     send("request",  $j$ ) to a teammate;
31   until the maximum number of requests sent for all
      $j \in \alpha_k$  is reached or the recruitment for  $\alpha_k$  is
     finished or aborted.

32 when received ("request",  $j$ ) from agent  $i_s$ 
33   /* decides whether or not to commit */
34   if roulette() <  $T_{ij}$  and  $i.res \geq Res(i, j)$  then
35     commit to  $j$ ;
36     send("committed",  $j$ ) to  $i_s$ ;
37   else
38     if request timeout reached then
39       send("timeout",  $j$ ) to  $i_s$ ;
40     else
41       forward("request",  $j$ ) to a teammate;

42 when received ("committed",  $j$ ) from  $i_c$ 
43    $\alpha_k :=$  AND group which contains  $j$ ;
44   if recruitment for  $\alpha_k$  is finished or aborted then
45     send("release",  $j$ ) to  $i_c$ ; return
46   if at least one agent committed with each  $j \in \alpha_k$ 
     then
47     recruitment for  $\alpha_k$  is finished;
48     /* forms the group of engaged agents */
49     foreach  $j \in \alpha_k$  do
50       pick a committed agent  $i_p$  with probability
       proportional to  $Cap(i_p, j)$ ;
51       send("engage",  $j$ ) to  $i_p$ ;
52       send("release",  $j$ ) to non selected agents;

53 when received ("engage",  $j$ )
54   accept task  $j$  and decrease  $i.res$ ;

55 when received ("release",  $j$ )
56   decommit to  $j$ ;

57 when received ("timeout",  $j$ )
58    $\alpha_k :=$  AND group which contains  $j$ ;
59   if the number of received timeouts for each  $j \in \alpha_k$  is
     equal the number of requests sent then
60     recruitment for  $\alpha_k$  is aborted;
61     foreach  $j \in \alpha_k$  do
62       send("release",  $j$ ) to committed agents;
```

committed: to inform that the agent joins the recruitment for a task $j \in \alpha_k$ and commits to it;

engage: to inform that the agent was indeed selected to perform the task $j \in \alpha_k$;

release: to inform that the agent was not selected to perform the task $j \in \alpha_k$ and must decommit with it.

timeout: to inform that a request for a task $j \in \alpha_k$ has reached its timeout.

For the first step of the recruitment, the scout agent sends a certain number of **request** messages for each task $j \in \alpha_k$ (lines 27-31). The requests are sent to randomly selected agents. These requests take the role of pheromone

released in the air (SRR) or released on the way to the nest (LRR). As it occurs with the scout ant, which recruits a fixed number of nestmates independently of prey characteristics, eXtreme-Ants fixes a maximum number of requests that must be sent for each AND constrained task. This maximum number must be experimentally determined to maximize the total reward.

In the second step, the agents must decide whether or not to join the recruitment. When an agent receives a request originated by a scout agent i_s for a task j (lines 32-41), it uses the tendency (Equation 10) to decide if it accepts the request and then joins the recruitment, avoiding double commitment. If the request is accepted, the agent commits to the task, reserving the amount of resources required. A **committed** message is sent to the scout to inform the commitment. If the request is not accepted, the agent forwards it to another agent, reproducing the chain of communication present in the LRR.

In the third step, the size of the group of agents engaged in the simultaneous execution of the AND constrained tasks must be regulated. In eXtreme-Ants the regulation is done by the scout agent. When the scout receives enough commitments for each constrained task $j \in \alpha_k$ (line 46), it forms the group of agents that will execute the tasks simultaneously. Following the E-GAP definition, just one agent must be selected among those committed for each constrained task. The scout then performs a probabilistic selection, picking an agent i_p with probability proportional to its capability $Cap(i_p, j)$. The scout then informs i_p that it was the selected one and thus must engage in the execution of j (via an **engage** message, line 51). All other non selected agents are released (via a **release** message, line 52). Agents that commit to an already allocated task are also released to avoid deadlocks. At this moment the recruitment is finished. As the result, a group of agents is formed, in which each agent is allocated to a task $j \in \alpha_k$, enabling the simultaneous execution of all AND constrained tasks in α_k .

After the group of engaged agents is formed, the requests not yet accepted by other agents become obsolete. To avoid agents passing obsolete requests back and forth, eXtreme-Ants uses a timeout mechanism. The timeout is given by the number of agents that a recruitment request is allowed to visit. When the timeout of a recruitment request is detected (line 38), the scout is notified via a **timeout** message. When the scout receives a timeout notification (lines 57-62) it checks if a timeout was received for all recruitment requests for each AND constrained task $j \in \alpha_k$. If it is the case, it aborts the recruitment process and releases the committed agents.

It is important to note that due to the algorithm asynchronism, the scout agent can perform other actions during the recruitment. These actions comprise the perception and allocation of another tasks, and even a recruitment for other AND constrained task groups. Although eXtreme-Ants reproduce the inter-agent communication via messages, it can be easily modified to use some kind of indirect communication (e.g., pheromones) when the environment allows it.

5. EXPERIMENTS AND RESULTS

We compare eXtreme-Ants to Swarm-GAP [2] and LA-DCOP [9]. We have evaluated eXtreme-Ants in a domain-independent simulator that allows experimentation with large number of agents, performing experiments similar

to Swarm-GAP and LA-DCOP which have also used such a simulator.

Basically, each experiment consists of 2000 tasks, grouped in five classes, where each class determines the task characteristics. The number of agents varies from 500 to 4000. This means that the load (ratio between tasks and agents) is 4 in the first case and 0.5 in the latter. Each agent has a 60% probability of having a non-zero capability for each class. In this case the agent has a randomly assigned capability ranging from 0 to 1. Regarding the AND constraints among tasks, 60% of the tasks are related in groups of five tasks. The simulated communication channel is reliable (every sent message is received) and fully connected (each agent is connected to every other agent). Each experiment consists of 1000 time steps. The total number of tasks is kept constant. At each time step, each task has a probability of 10% of being replaced by a task potentially requiring a different capability. The tasks are persistent, which means that non allocated tasks are kept in the next time step. At each time step, each token or message is allowed to move from one agent to another only once. Despite that each task can have a particular stimulus value, we adopt the same value for all tasks. Each datapoint in the plots we show here represents the average over 20 runs.

As defined by the E-GAP, the goal is to maximize the total reward, which is the sum of the reward at each time step (Equation 2). We have tested each algorithm with threshold (for LA-DCOP) in the interval $[0.0, 0.9]$ and stimulus (for eXtreme-Ants and Swarm-GAP) in $[0.1, 0.9]$. From these values we have selected those that have yielded the maximum total reward in each scenario. These values are shown in Table 1 and will be used in the comparisons.

Table 1: Stimulus and threshold values that yield the maximum total reward for each algorithm.

Agents	eXtreme-Ants (Stimulus)	Swarm-GAP (Stimulus)	LA-DCOP (Threshold)
500	0.3	0.2	0.0
1000	0.3	0.3	0.4
1500	0.2	0.2	0.6
2000	0.2	0.2	0.6
2500	0.2	0.2	0.6
3000	0.2	0.2	0.6
3500	0.2	0.2	0.7
4000	0.2	0.2	0.7

The first experiment compares the total reward achieved by each algorithm using the stimulus and threshold values previously selected (Table 1). Additionally, in the case of eXtreme-Ants the number of recruitment requests is five for each AND constrained task, and the timeout is 20.

Figure 1 shows the total rewards achieved by each algorithm. On average, eXtreme-Ants yields rewards that are 25% higher than those of Swarm-GAP and 19% lower than those of LA-DCOP (t-test, 99% confidence).

When an agent performs a task, it uses an amount of its resources. Thus, the agents must avoid to waste their resources performing tasks that do not yield any reward (e.g., tasks that belong to an AND constrained set, but are not simultaneously performed). The second experiment, shown in Figure 2, compares the percentage of resources used by each agent to perform the tasks. As we can see, all algorithms use almost the same percentage of resources. There is no signi-




Figure 1: Total reward as a function of the number of agents.

ficative difference between eXtreme-Ants and Swarm-GAP in the cases with 3000, 3500, and 4000 agents, and between eXtreme-Ants and LA-DCOP in the cases with 500 and 1000 agents (t-test, 99% confidence).




Figure 2: Percentage of resources used by each agent to perform tasks as a function of the number of agents.

From these two experiments, we can see that despite the fact that agents in Swarm-GAP use almost the same amount of resources, the rewards achieved are worse than those achieved by eXtreme-Ants and LA-DCOP. This is due to the way Swarm-GAP deals with AND constrained tasks. The use of an execution coefficient (see Section 3) does not ensure the simultaneous execution of the AND constrained tasks. Thus, the agents use their resources to perform tasks, but this does not translate into a reward. Both eXtreme-Ants and LA-DCOP outperform Swarm-GAP regarding the total reward. This is due to the existence of an explicit coordination mechanism to deal with AND constrained tasks, ensuring their simultaneous allocation.

LA-DCOP yields higher rewards than eXtreme-Ants because in the former each agent maximizes its capability when performing the tasks, taking into account the available resources. On the other hand, agents in eXtreme-Ants make a simple one-shot decision to perform tasks. The maximization leads to a better exploitation of the agents' resources. However, as we show next, there is a trade-off between the achieved reward and the communication/computational effort.

In the next experiment, shown in Figure 3, we compare the amount of communication used in each algorithm. Communication is measured as the sum of messages sent by the agent over all time steps, regardless of message type (e.g., token, recruitment request, etc.). The results are statistically significant at 99% confidence level (t-test). On average, agents in eXtreme-Ants sent 121% fewer messages than those in LA-DCOP and 80% more messages than those in Swarm-GAP. The smallest difference to LA-DCOP occurs with 3000 agents. Even in this case LA-DCOP sends 66% more messages than eXtreme-Ants.




Figure 3: Total number of messages sent as a function of the number of agents.

As mentioned, in Swarm-GAP the absence of an explicit coordination mechanism to ensure the simultaneous allocation leads to a small number of messages. Thus it sends fewer messages than eXtreme-Ants and LA-DCOP. This however has a great impact in the reward achieved by Swarm-GAP.

The last experiment aims at evaluating the computational effort of the agents in each algorithm. This is important because low computational effort means that the agents are more efficient to act in environments in which the available time to make a decision is restricted. We define the computational effort as the number of evaluated tasks by each agent at each time step. This number is computed as follows. Each time an agent decides whether or not to perform a task, an internal counter is incremented. In the case of eXtreme-Ants and Swarm-GAP, each probabilistic decision causes just one increment in the counter. On the other hand, since an agent in LA-DCOP solves a BKP to decide which tasks to perform, the increment in the counter is related to the number of retained tasks. To solve a BKP our imple-

mented version of LA-DCOP uses a greedy approach, which sorts the tasks by the agent's capability and then selects the tasks to perform constrained by the agent's resources. If n is the number of retained tasks, the sort causes an increment of $n \log(n)$ in the counter, followed by an increment of at most n to select the tasks to perform.

Figure 4 shows the average computational effort of each agent. The external plot emphasizes the area that concentrates the majority of the points. The internal plot shows the full area.




Figure 4: Computational effort (number of evaluated tasks by each agent) as a function of the number of agents.

On average, the computational effort of Swarm-GAP is 55% lower than that of eXtreme-Ants. Since in Swarm-GAP there is no explicit coordination mechanism to deal with AND constrained tasks, the agents do not have to make additional evaluations regarding the simultaneous allocation of constrained tasks, reducing the computational effort of Swarm-GAP. However, the absence of such mechanism affects the total reward, as shown previously. The higher computational effort of both eXtreme-Ants and LA-DCOP are due to the presence of an explicit coordination mechanism to deal with constrained tasks.

eXtreme-Ants outperforms LA-DCOP regarding computational effort in all cases, being these efforts on average 151% lower than those of LA-DCOP. The most significant result is for the case of 500 agents, in which the computational effort of LA-DCOP is 493% higher than that of eXtreme-Ants. All these figures are statistically significant at 99% confidence t-test.

As shown in the experiments, the probabilistic allocation of eXtreme-Ants, based on the model of division of labor, reduces the amount of communication and the computational effort when compared to LA-DCOP. The reduction in the computational effort is due to the simple one-shot decision, which does not require any local maximization. The low computational effort causes the reduction in the number of messages sent, since in LA-DCOP the tasks which are not selected in the local maximization are sent to other agents.

In comparison to Swarm-GAP, in both eXtreme-Ants and LA-DCOP, the presence of an efficient coordination mecha-

nism to deal with inter-task constraints leads to better total rewards.

Finally, we emphasize that the choice of one particular algorithm shall consider the constraints of the scenario. When the total reward is a key point and there are no constraints in the communication and in the time the agents have to make a decision, LA-DCOP is a good choice. On the other hand, in scenarios with such constraints eXtreme-Ants is more appropriate. It yields lower rewards but the decision is faster and there is a better use of the communication channel.

6. CONCLUSIONS

In this paper we presented an approximate algorithm for task allocation in extreme teams, called eXtreme-Ants. The algorithm is inspired by the division of labor in social insects and in the process of recruitment present in ants that transport prey cooperatively.

The experimental results show that the use of the model of division of labor to decide whether or not to perform given tasks allows the agents to make efficient coordinated actions. Since the decision-making is probabilistic, it is fast and requires a reduced communication and computational effort, enabling the agents to act in environments where the available time to make a decision is highly restricted. Moreover, the incorporated recruitment process provides efficient allocation of constrained tasks that requires simultaneous execution. This avoids that agents waste their resources and leads to better total rewards. The efficiency of eXtreme-Ants regarding communication and computational effort suggests that techniques that are inspired by social insects can be considered for multiagent task allocation.

We intend to work in the direction of changing the stimuli values dynamically, indicating different priorities in the execution of the tasks. More than one kind of resource for an agent can also be considered. Besides, these resources can change over time, as for instance, the battery charge of a robot. We also intend to evaluate the performance when unreliable communication channels are present and where failures may occur, and to apply this approach in the RoboCup Rescue simulator.

7. ACKNOWLEDGMENTS

This research is partially supported by the Air Force Office of Scientific Research (AFOSR) (grant number FA9550-06-1-0517) and by the Brazilian National Council for Scientific and Technological Development (CNPq).

8. REFERENCES

- [1] E. Bonabeau, G. Theraulaz, and M. Dorigo. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, USA, 1999.
- [2] P. R. Ferreira, Jr., F. Boffo, and A. L. C. Bazzan. Using swarm-GAP for distributed task allocation in complex scenarios. In N. Jamali, P. Scerri, and T. Sugawara, editors, *Massively Multiagent Systems*, number 5043 in Lecture Notes in Artificial Intelligence, pages 107–121. Springer, Berlin, 2008.
- [3] B. Hölldobler. Territorial behavior in the green tree ant (*Oecophylla smaragdina*). *Biotropica*, 15(4):241–250, 1983.
- [4] B. Hölldobler, R. C. Stanton, and H. Markl. Recruitment and food-retrieving behavior in *Novomessor* (formicidae, hymenoptera). *Behavioral Ecology and Sociobiology*, 4(2):163–181, 1978.
- [5] S. Martello and P. Toth. *Knapsack Problems: Algorithms and Computer Implementations*. John Wiley & Sons, New York, NY, USA, 1990.
- [6] P. J. Modi, W.-M. Shen, M. Tambe, and M. Yokoo. An asynchronous complete method for distributed constraint optimization. In *Proc. of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 161–168, New York, USA, 2003. ACM Press.
- [7] A. Petcu and B. Faltings. A scalable method for multiagent constraint optimization. In L. P. Kaelbling and A. Saffioti, editors, *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pages 266–271, Edinburgh, Scotland, August 2005. Professional Book Center.
- [8] S. K. Robson and J. F. A. Traniello. Resource assessment, recruitment behavior, and organization of cooperative prey retrieval in the ant *Formica schaufussi* (hymenoptera: Formicidae). *Journal of Insect Behavior*, 11(1):1–22, 1998.
- [9] P. Scerri, A. Farinelli, S. Okamoto, and M. Tambe. Allocating tasks in extreme teams. In F. Dignum, V. Dignum, S. Koenig, S. Kraus, M. P. Singh, and M. Wooldridge, editors, *Proc. of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 727–734, New York, USA, 2005. ACM Press.
- [10] G. Theraulaz, E. Bonabeau, and J. Deneubourg. Response threshold reinforcement and division of labour in insect societies. In *Royal Society of London Series B - Biological Sciences*, volume 265, pages 327–332, 2 1998.
- [11] Y. Xu, P. Scerri, K. Sycara, and M. Lewis. Comparing market and token-based coordination. In *Proc. of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, pages 1113–1115, New York, NY, USA, 2006. ACM.