

Task Partitioning in Swarms of Robots: Reducing Performance Losses Due to Interference at Shared Resources

Giovanni Pini, Arne Brutschy, Mauro Birattari, and Marco Dorigo

IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium
{gpini, arne.brutschy, mbiro, mdorigo}@ulb.ac.be

Abstract. The performance of large groups of robots is often limited by a commonly shared resource. This effect, termed interference, can have a large impact on robotic swarms. This article studies the issue of interference in a swarm of robots working on a harvesting task. The environment of the robots is spatially constrained, i.e., there is a commonly shared resource, the nest, which limits the group's performance when used without any arbitration mechanism. The article studies the use of task partitioning for reducing concurrent accesses to the resource, and therefore limiting the impact of interference on the group's performance. In our study, we spatially partition the environment into two subparts, thereby partitioning the corresponding harvesting task as well. We employ a simple method to allocate individuals to the partitions. The approach is empirically studied both in an environment with a narrow nest area and an environment without this constraint. The results of the task partitioning strategy are analyzed and compared to the case in which task partitioning is not employed.

1 Introduction

In collective robotics, interference is a critical problem limiting the growth of a group: the time each robot spends in non-task-relevant behaviors such as obstacle avoidance increases when the density of individuals rises—see e.g., [1]. The performance on tasks that suffer from physical interference can typically be improved by spatial partitioning; for example, by keeping each robot in its own “working area”. A known approach that uses this rationale is the so called bucket-brigade [2,3]. In this approach, robots hand over objects to robots working in the following area, until the objects reach their destination. As tasks usually cannot be partitioned arbitrarily, this approach effectively limits the number of robots that can be employed in the task. A possible solution to this problem, treating working areas as non-exclusive, raises other problems: How should individuals be allocated to tasks? How can such an allocation help in limiting the amount of interference?

In this paper, we study how task partitioning can help in reducing sources of interference. Additionally, we show a simple way to achieve self-organized allocation to such a task partition when using a robotic swarm.

We use the foraging problem, one of the canonical testbeds for collective robotics, as a base for our studies. In our experiments, a swarm of homogeneous robots has to

harvest prey objects from a source area and transport them to a home area. In this study, we limit ourselves to a harvesting task that is pre-partitioned by the designer into two subtasks with a sequential interdependency. We study a simple threshold-based model of self-organized allocation and focus on two questions: Under which environmental conditions is it advantageous to partition the task? Can this partition reduce interference between group members? These questions are studied in two experiments using a simulated robot swarm.

The paper is organized as follows. We first review related works in Sect. 2. In Sect. 3 we explain the task partitioning and the allocation method employed in this study. Section 4 gives the methods used in the experiments by describing the environments, the simulated robots, and the controller. In Sect. 5 the results of the experiments are given and discussed. Section 6 draws some conclusions and discusses future work.

2 Related Work

Interference has long been acknowledged as being one of the key issues in multi-robot cooperation [4]. Lerman and Galstyan devised a mathematical model that allows a quantification of the interference and its effect on group performance [1]. Probably, the most thorough study was published by Goldberg, who identified several types of multi-robot interactions [5]. Goldberg notes that one of the most common types of interference is physical interference in a central area, for example the nest. This kind of interference results from resource conflicts, in this case physical space, and can be arbitrated by either making sure that robots stay in different areas all the time or by employing a scheduling mechanism to ensure that robots use the same space only at different times.

A simple method for reducing interference by using the first arbitration method mentioned is the so-called bucket-brigade: robots are forced to stay in exclusive working areas and to pass objects to the following robot as soon as they cross the boundaries of their area [2,3]. Recently, this has been extended to work with adaptive working areas by Lein and Vaughan [6]. To the best of our knowledge, current works concerned with bucket brigading only studied the influence of interference due to obstacle avoidance. Other sources of interference (e.g., object manipulation) were never studied, although they might have a critical impact on the performance of any task partitioning approach. To quote Shell and Matarić [3]: *“If the cost of picking up or dropping pucks is significant [...], then bucket brigading may not be suitable.”*

Task allocation for multi-robot systems is a wide field, which can be divided in intentional and self-organized task allocation. Intentional task allocation relies on negotiation and explicit communication to create global allocations, whereas in self-organized task allocation global allocations result from local, stochastic decisions. A formal analysis and taxonomy that covers intentional task allocation has been proposed by Gerkey and Matarić [7]. Kalra and Martinoli recently compared the two best-known approaches of intentional and self-organized task allocation [8].

The field of self-organized task allocation is in its early stages, as most studies tackle simple problems without task interdependencies. An exception is the work of Scheidler et al, in which helper components of a computing system are allocated to tasks with sequential interdependencies by using a threshold based method [9]. Studies in

self-organized task allocation are mostly based on threshold-based approaches, taking inspiration from division of labor in social insects. Krieger and Billeter were among the first to propose threshold-based approaches in multi-robot task allocation [10]. Labella et al. used threshold-based task allocation in a multi-foraging task [11]. Similarly, Campo and Dorigo used a notion of the group's internal energy to allocate individuals to a multi-foraging task [12]. Finally, Liu et al. studied a multi-foraging task while focusing on the influence of the use of different social cues on the overall group performance [13].

3 Task Partitioning and Allocation

In this work, we study a collective foraging task. By spatially partitioning the environment, the global foraging task is automatically partitioned into two subtasks: 1) harvesting prey objects from a harvesting area (*source*) and 2) transporting them to a home area (*nest*). Robots working on the first subtask harvest prey objects from the source and pass them to the robots working on the second subtask, which store the objects in the nest. These subtasks have a sequential interdependency in the sense that they have to be performed one after the other in order to complete the global task once: delivering a prey object to the home area. Robots can decide to switch from one subtask to the other, thus creating a task allocation problem: individual robots have to be allocated to subtasks and different allocations yield different performance. As a prey object has to be passed directly from one robot to the other, a robot usually has to wait some time before passing a prey object to or receiving a prey object from a robot working on the other subtask. This waiting time can therefore give an indication of the allocation quality for the respective subtask: if the waiting time is very long, there might not be enough robots allocated to the other subtask. Thus, the robots can use this waiting time to decide whether to switch subtask or not. Ideally, the waiting time should be the same for the two subtasks in order for the system to reach a stable state and deliver optimal performance.

Our robots exploit a simple threshold-based model to decide when to switch task: when the waiting time t_w is higher than a threshold Θ , a robot switches its subtask. The robot's waiting time is a function of the average time the robots working in the other subtask need to complete their task. The task-completion time of a robot depends on two factors: 1) round-trip-time (i.e., distance to travel) and 2) time lost due to interference. Thus, the robot's threshold Θ is a function of the round-trip-time and the interference of the robots in the other subtask. Therefore, the optimal task switching threshold depends on the task (i.e., time to harvest a prey object) and the environment (i.e., distance between the source and the nest). As the parameters of the environment are not pre-programmed into the robots, determining the optimal threshold can be a complex problem. In this work, we limit ourselves to a simple method for setting this threshold: at the start of the experiment, each robot draws a random threshold that is used as its task switching threshold throughout the experiment.

In the following, we study the properties of this simple self-organized task allocation strategy, compare this strategy to a strategy without task partitioning, and analyze how it can help to reduce interference. We refer to the two strategies as *partitioned* and *non-partitioned*, respectively.

4 Methods

This section describes the environments in which the experiments are carried out, the simulated robots, and the robot’s controller. Additionally, we describe how we run the experiments and we introduce some metrics that we use to evaluate the properties of the system.

4.1 Environments

We study task allocation in two different environments. In these two environments, the nest is marked by a light source that can be perceived by all robots, thus providing directional information. The environment is spatially partitioned in two parts: the source is located on the left and the nest is located on the right side of the arena. We refer to the two sides of the arena as *harvest area* and *store area*, respectively. The *exchange zone* is located between these two areas. Robots working on the left side, called *harvesters*, gather prey objects in the source and move them to the exchange zone, where they pass them to the robots working on the other side. These are referred to as *stomers*: their role is to transport prey objects to the nest and store them there. The nest, the source, and the exchange zone can be detected through environmental cues (ground color).

At time $t = 0$, the robots are randomly placed in the harvest area. The experiments run for $t_{\max} = 18,000$ time steps (a simulated time of one hour, with a time step length of 200 ms). The experiments are run in two different arenas (see Fig. 1). The first arena (Fig. 1a) is 4.125 m long with a width of 1.6 m at the source and exchange zone, whereas the nest is 0.4 m wide. The exchange zone is located 3.125 m away from the source. This arena is characterized by the presence of an area, critical for the task, in which high interference between robots can be expected (the nest). Thus, this arena is referred to as the *narrow-nest* environment.

The second arena (Fig. 1b) has a rectangular shape: it is 3.75 m long and 1.6 m wide. Here as well the exchange zone is located 3.125 m from the source. The arena shape does not suggest the presence of any zone where interference can be higher than in other places. This arena is referred to as the *wide-nest* environment.

The area of both arenas is 6 m^2 , 5 m^2 for the harvest area and 1 m^2 for the store area. The overall area is the same in the two arenas, so that the same group size results in the same robot density. Thus, results are comparable across the two environments.

4.2 Simulation

The experiments are carried out in a custom simulation environment that models geometries and functional properties of simple objects and robots. Our robots’ model is purely kinematic. Prey objects are simulated as an attribute a robot can possess and not as physical entities. Although the experiments are conducted in simulation only, the simulated robots have a real counterpart: the swarm-bot robotic platform [14]. The platform consists of a number of mobile autonomous robots called s-bots, which have been used for several studies, mainly in swarm intelligence and collective robotics—see for instance Groß et al. [15] and Nouyan et al. [16]. The simulated s-bots are of round

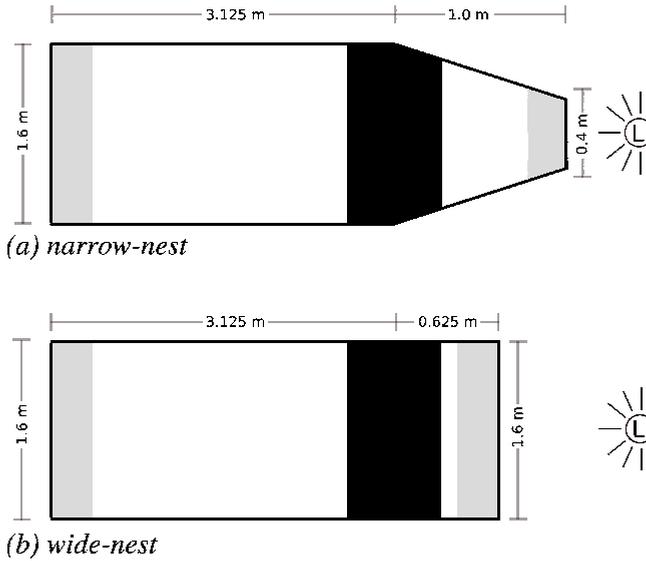


Fig. 1. Depiction of (a) the *narrow-nest* environment used in the first experiment and (b) the *wide-nest* environment used in the second experiment. The gray stripes are the source (left), and the nest (right), each 0.25 m deep. The black stripe is the exchange zone, that is 0.5 m deep. The light source is marked with “L”.

shape, with a diameter of 0.116 m. Each of them is equipped with 16 infrared proximity sensors, used to perceive obstacles up to a distance of 0.15 m. Eight ambient light sensors can be used to perceive light gradients up to a distance of 5.0 m. The robots are equipped with 4 ground sensors used to perceive nest, source and exchange zone. A 8 LEDs ring is used to signal when a prey object is carried. An omni-directional camera allows the perception of LEDs in a circle of radius 0.6 m surrounding the robot. A uniform noise of 10% is added to all sensor readings at each simulation step. The robots can move at a maximum speed of 0.1 m/s by means of a differential drive system.

4.3 Controller

All the robots share the same, hand coded, finite state machine controller depicted in Fig. 2. The controller consists of two parts, each corresponding to a possible subtask a robot can perform. Gray states refer to the **harvest** subtask, white states to the **store** subtask. Since all the robots start in the harvest area, their controller is initially set to perform anti-phototaxis. In this way they will reach the source, where they can start retrieving prey objects. The behavior of each robot is a function of the task it is performing. Harvesters not carrying a prey object move towards the source, where they can find prey. Harvesters carrying a prey object, move to the exchange zone and wait for a free storer. Upon arrival of such a storer, the harvester passes the prey object to it. Storer carrying a prey object move towards the nest, where they can deposit the object. Storer not carrying a prey object head to the exchange zone and search for a harvester with a prey

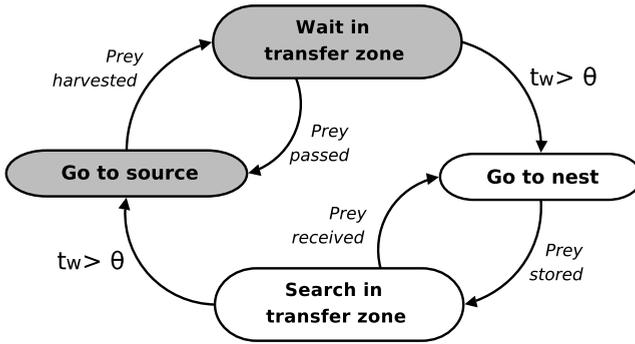


Fig. 2. Simplified state diagram of the controller of the robots. Gray states belong to the harvest task, white states to the store task. The obstacle avoidance state has been omitted for clarity, as it is applicable in all states of the robot. t_w is the time spent in the exchange zone and Θ is the threshold.

object. Robots can detect other robots carrying a prey on the basis of the color of their LED ring. While moving, each robot avoids obstacles (walls and other robots).

Task switches can occur: a harvester carrying a prey object can decide to become a storer, and a storer not carrying a prey object can decide to become a harvester. As mentioned before, robots switch task depending on an internal threshold Θ , representing the maximum amount of control cycles they can spend in the transfer zone waiting to pass (harvesters) or receive (storers) a prey object. If a robot remains in the transfer zone longer than its threshold without passing or receiving prey objects ($t_w > \Theta$), it switches its task. The optimal threshold value is not trivial to determine. In the work presented here, we use a simple method to set the threshold Θ : at the beginning of the experiment, each robot draws a random threshold, sampled uniformly in the interval $[0, 1000]$. We chose this method because it is independent of the environment and does not rely on complex approximation techniques. The threshold value does not change during the experiment. In case of the non-partitioned strategy, the threshold is set to $\Theta = 0$, causing the robots to switch subtask immediately as soon as they reach the exchange zone.

4.4 Experiments

The goal of the experiments is to investigate whether task partitioning can reduce interference in task-critical zones, and how to allocate a robotic swarm to partitions. As pointed out by Lerman and Galstyan [1], interference is related to the number of individuals in the system. Additionally, the physical interference between robots is also a function of the environment the robots act in. The higher the group size, the higher the density, resulting in a higher amount of physical interference. Thus, in order to study interference in our experiments, we increase the size of the group in each of the two environments shown in Fig. 1, while using both strategies (non-partitioned and partitioned). We study the performance of the system when the group size N ranges in the interval $[1, 40]$. We run 50 repetitions for each value of N and each experimental setting.

4.5 Metrics

In order to quantify the influence of interference, we measure the *group performance* P by the number of prey objects collected by the swarm at the end of the experiment ($t_{\max} = 1$ hour). From the group performance measure we can derive the *individual efficiency* as follows:

$$I_{\text{eff}} = P/N, \quad (1)$$

where N is the size of the group. Individual efficiency can help to understand the effect of interference on the performance.

In order to measure the influence of environmental features on the interference, we define an interference measure taking inspiration from Rosenfeld et al. [17]. In their work, interference is measured as the time spent performing actions not strictly related to the task, but rather lost due to negative interactions with the environment (e.g., obstacle avoidance maneuvers). By registering the number of collisions for each area of the arena, we can draw conclusions about where physical interferences happen most often. We measure interference through the state of the controller: in our case a robot is experiencing interference each time its controller perceives an obstacle.

In case of a partitioned task, there is another source of inefficiency that adds to interference: the time lost in the exchange zone. We define the *strategy cost* C as the sum of time lost because of physical interference and time lost in the exchange zone:

$$C = T_{\text{int}} + T_{\text{part}}, \quad (2)$$

where T_{int} is the amount of time steps during which the controller perceives an obstacle, and T_{part} is the total amount of time steps spent in prey passing maneuvers. By using this metric, the cost of the non-partitioned strategy is purely due to interference ($T_{\text{part}} = 0$), while in case of the partitioned strategy, prey passing costs add to interference costs. In a way, passing a prey object produces another kind of interference in the system. The strategy cost captures this effect, thus allowing for a comparison of strategies.

5 Results and Discussion

The graphs in Fig. 3a and 4a show the performance P for different group sizes in the narrow-nest and wide-nest environment respectively.

Figure 3b and 4b show the individual efficiency I_{eff} of the robots in the narrow-nest and wide-nest environment respectively. Black curves are the average computed over the 50 repetitions of each setting, gray curves indicate the 95% confidence interval on the expected value. The performance graph in Fig. 3a shows that the partitioned strategy improves performance in the narrow-nest environment. The graph shows that the non-partitioned strategy performs better than the partitioned strategy for small group sizes (up to $N = 13$ robots). However, increasing the group size makes the non-partitioned strategy collapse: the number of gathered prey objects drops dramatically for groups larger than 13. The individual efficiency graph (Fig. 3b) can explain the behavior of the system. The robots employing the partitioned strategy are less efficient, for small group sizes, than those performing the non-partitioned strategy. However, the addition of more individuals affects the efficiency of the non-partitioned group in a

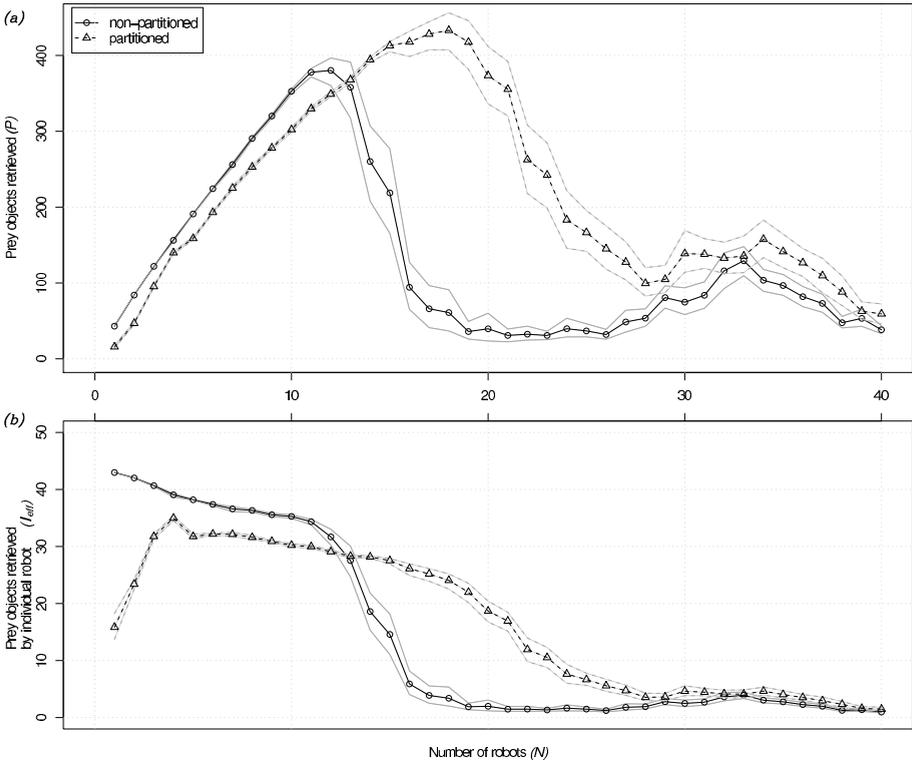


Fig. 3. (a) Performance P and (b) individual efficiency I_{eff} for increasing number of robots in the narrow-nest environment. The black continuous line refers to the case of no task partitioning, the black dashed line to the case of partitioning. Gray lines indicate the 95% confidence interval on the expected value.

more dramatic way. At a certain point, the drop in efficiency becomes very steep for the non-partitioned strategy. On the other hand, the partitioned strategy scales better: individual efficiency drops smoothly. This explains why a group using the partitioned strategy performs better: it can benefit from the work of more individuals and therefore collects more prey objects. These considerations do not hold in the wide-nest environment. Figure 4b shows that the individual efficiency drops smoothly for both strategies. In addition, for group sizes $N < 33$ it is always higher in the non-partitioned strategy. Therefore, the non-partitioned strategy performs better than the partitioned strategy, as shown by the graph in Fig. 4a. In both the environments, independently of the strategy used to accomplish the task, the system collapses when the area is saturated by the swarm. Figure 5 shows the effect of an increasing number of robots on the strategy cost in the narrow-nest environment. The graph compares the cost C of each of the two strategies for different group sizes.

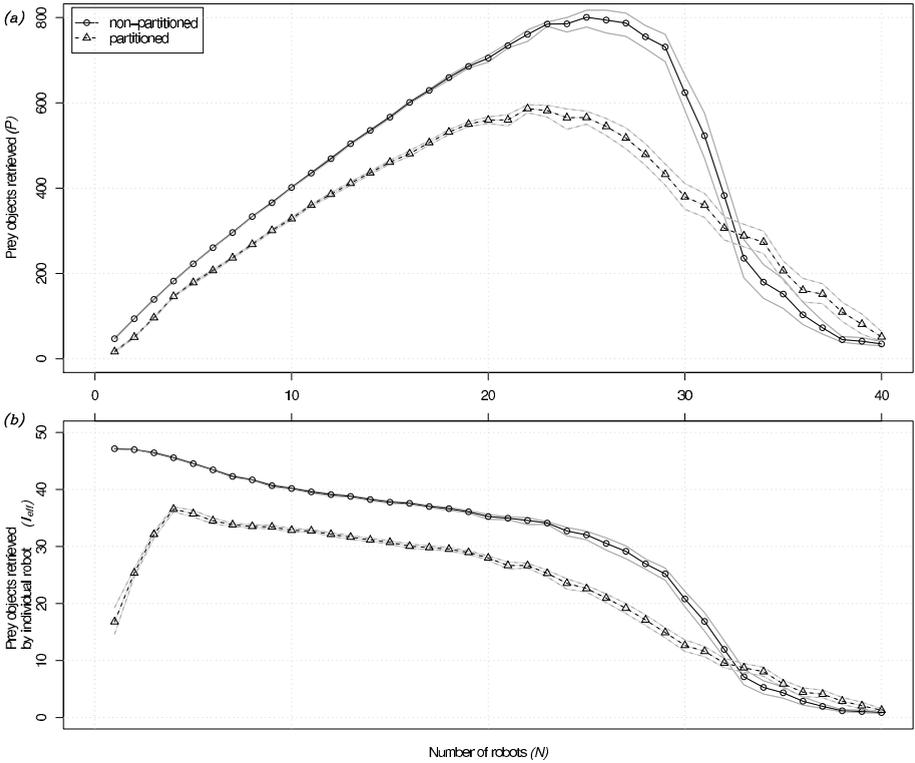


Fig. 4. (a) Performance P and (b) individual efficiency I_{eff} for increasing number of robots in the wide-nest environment. The black continuous line refers to the case of no task partitioning, the black dashed line to the case of partitioning. Gray lines indicate the 95% confidence interval on the expected value.

In case of the partitioned strategy (Fig. 5a), the graph shows each component of the cost (T_{int} and T_{part}). Clearly, task partitioning has the effect of reducing the cost due to interference but has the disadvantage of increasing the cost due to time lost. The probability of two or more robots encountering each other increases with the robot density. Although this determines a higher interference cost (i.e., T_{int}), it decreases the cost due to lower waiting time (i.e., T_{part}) in the case of the partitioned strategy. Partitioning performs better when the gain from interference reduction is greater than the loss of performance due to partitioning inefficiencies. These considerations hold in the narrow-nest environment, where the likelihood of physical interference in a task-critical zone is very high. In the wide-nest environment, interference in the nest is as likely as interference in the exchange zone. Thus, it is not beneficial to pay the cost of waiting and the non-partitioned strategy performs better for any group size. The mechanism by which partitioning reduces interference costs can be deduced by comparing the interference graphs in Fig. 6. The graphs show the number of times that physical interference (as defined in Sect. 4.5) was registered in each region of the narrow-nest environment. The total area was discretized in squares of 1 cm^2 . Figure 6 shows the results obtained with

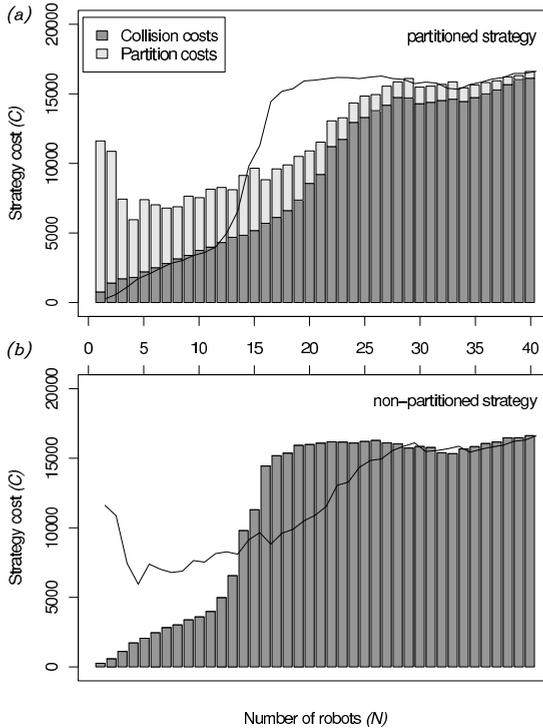


Fig. 5. Cost of interference in the narrow-nest environment. Bars represent the cost C , sum of interference time T_{int} and partition time T_{part} (i.e., waiting times). For easy reference, the outline of the bars of the respective other graph has been added to each graph. (a) Costs for the partitioned strategy, where interference cost stem from waiting times and collisions. (b) Cost in case of the non-partitioned strategy, where only physical interference through collisions exists.

18 robots, in the case of the non-partitioned strategy (Fig. 6a) and in the case of the partitioned strategy (Fig. 6b). The graphs show that the use of the non-partitioned strategy leads to high interference in the nest, which becomes congested. Partitioning the task reduces the robot density in the nest, thus spreading the interference more uniformly across the arena. In addition, the overall interference diminishes because the exchange zone is wider: the robots have more freedom of movement and collide less often. Although the graphs show only data collected with 18 robots, experiments with different group sizes produced similar results.

6 Conclusions and Future Work

Interference can be an issue when working with swarms of robots. In this work, we used task partitioning and allocation to reduce interference between robots sharing the same physical space. We manually partitioned the environment and employed a simple self-organized strategy for allocating individuals to subtasks. Results show that a partitioning strategy improves performance in a constrained environment. Additionally, we

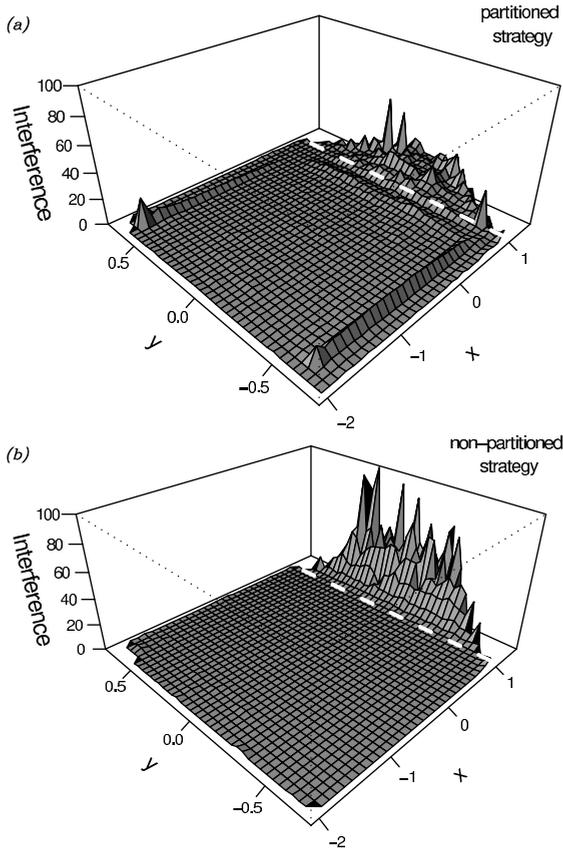


Fig. 6. Mean interference values registered for (a) the partitioned strategy and (b) the non-partitioned strategy, both in the narrow-nest environment. Shown values are observation means of 50 repetitions with $N = 18$ robots. Coordinates on the x - and y -axis are given in meters. The arena is stretched along the y -axis for better visualization. The dashed white line marks the location of the exchange zone.

identified cases in which partitioning is not advantageous and a non-partitioned strategy should be used. The proposed strategy is fairly simple and far from being an optimal solution, nevertheless we improved the performance of the swarm when interference was costly. We believe that our approach can be generalized to a wider range of applications. It can be applied to those situations in which system's performance drops because of concurrent access to a shared resource. If the accesses can be distributed more efficiently by partitioning the resource, the impact of interference can be limited. Future work will concern the identification of the optimal allocation in the studied environments as well as the development and study of a strategy that can find this optimal allocation in a self-organized and adaptive way. In addition, the interference metric proposed in Sect. 4.5 could be used by the robots to decide whether to partition the task.

In this way, we could achieve even better performance, since partitioning would be employed only when strictly needed. Finally, the goal is to validate the system using the real robots.

Acknowledgements. This work was supported by the SWARMANOID project, funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant IST-022888 and by the VIRTUAL SWARMANOID project funded by the Fund for Scientific Research F.R.S.–FNRS of Belgium’s French Community. The information provided is the sole responsibility of the authors and does not reflect the European Commission’s opinion. The European Commission is not responsible for any use that might be made of data appearing in this publication. Marco Dorigo and Mauro Birattari acknowledge support from the Fund for Scientific Research F.R.S.–FNRS of Belgium’s French Community, of which they are a research director and a research associate, respectively.

References

1. Lerman, K., Galstyan, A.: Mathematical model of foraging in a group of robots: Effect of interference (2002)
2. Fontán, M.S., Matarić, M.J.: A study of territoriality: The role of critical mass in adaptive task division (1996)
3. Shell, D., Matarić, M.J.: On foraging strategies for large-scale multi-robot systems (2006)
4. Goldberg, D., Matarić, M.J.: Maximizing reward in a non-stationary mobile robot environment. *Autonomous Agents and Multi-Agent Systems* 6, 287–316 (2003)
5. Goldberg, D.: Evaluating the dynamics of agent-environment interaction (2001)
6. Lein, A., Vaughan, R.: Adaptive multi-robot bucket brigade foraging (2008)
7. Gerkey, B.P., Matarić, M.J.: A formal analysis and taxonomy of task allocation in multi-robot systems (2004)
8. Kalra, N., Martinoli, A.: A comparative study of market-based and threshold-based task allocation (2006)
9. Scheidler, A., Merkle, D., Middendorf, M.: Stability and performance of ant queue inspired task partitioning methods (2008)
10. Krieger, M.J.B., Billeter, J.B.: The call of duty: Self-organised task allocation in a population of up to twelve mobile robots (2000)
11. Labella, T.H., Dorigo, M., Deneubourg, J.L.: Division of labor in a group of robots inspired by ants’ foraging behavior (2006)
12. Campo, A., Dorigo, M.: Efficient multi-foraging in swarm robotics (2007)
13. Liu, W., Winfield, A., Sa, J., Chen, J., Dou, L.: Towards energy optimization: Emergent task allocation in a swarm of foraging robots (2007)
14. Mondada, F., Pettinaro, G.C., Guignard, A., Kwee, I.V., Floreano, D., Deneubourg, J.L., Nolfi, S., Gambardella, L.M., Dorigo, M.: SWARM-BOT: A new distributed robotic concept (2004)
15. Groß, R., Bonani, M., Mondada, F., Dorigo, M.: Autonomous self-assembly in swarm-bots (2006)
16. Nouyan, S., Campo, A., Dorigo, M.: Path formation in a robot swarm. Self-organized strategies to find your way home (2008)
17. Rosenfeld, A., Kaminka, G.A., Kraus, S.: A study of scalability properties in robotic teams (2005)