

Artificial pheromone for path selection by a foraging swarm of robots

Alexandre Campo · Álvaro Gutiérrez ·
Shervin Nouyan · Carlo Pinciroli ·
Valentin Longchamp · Simon Garnier · Marco Dorigo

Abstract Foraging robots involved in a search and retrieval task may create paths to navigate faster in their environment. In this context, a swarm of robots that has found several resources and created different paths may benefit strongly from path selection. Path selection enhances the foraging behavior by allowing the swarm to focus on the most profitable resource with the possibility for unused robots to stop participating in the path maintenance and to switch to another task. In order to achieve path selection, we implement virtual ants that lay artificial pheromone inside a network of robots. Virtual ants are local messages transmitted by robots; they travel along chains of robots and deposit artificial pheromone on the robots that are literally forming the chain and indicat-

ing the path. The concentration of artificial pheromone on the robots allows them to decide whether they are part of a selected path. We parameterize the mechanism with a mathematical model and provide an experimental validation using a swarm of 20 real robots. We show that our mechanism favors the selection of the closest resource is able to select a new path if a selected resource becomes unavailable and selects a newly detected and better resource when possible. As robots use very simple messages and behaviors, the system would be particularly well suited for swarms of microrobots with minimal abilities.

Keywords Swarm robotics · Path selection · Artificial pheromone

A. Campo · S. Nouyan · C. Pinciroli · M. Dorigo
IRIDIA, CoDE, Université Libre de Bruxelles,
50, Av. F. Roosevelt, CP 194/6, 1050 Brussels, Belgium
e-mail: acampo@ulb.ac.be

S. Nouyan
e-mail: snouyan@ulb.ac.be

C. Pinciroli
e-mail: cpinciro@ulb.ac.be

M. Dorigo
e-mail: mdorigo@ulb.ac.be

Á. Gutiérrez
ETSI Telecomunicación, B-317, Universidad Politécnica
de Madrid, Av. Complutense 30, 28040 Madrid, Spain
e-mail: aguti@ctsit.upm.es

V. Longchamp
STI IMT LSRO, Ecole Polytechnique Fédérale de Lausanne,
1015 Lausanne, Switzerland
e-mail: valentin.longchamp@epfl.ch

S. Garnier
Department of Ecology and Evolutionary Biology,
Princeton University, Princeton, NJ 08544, USA
e-mail: sgarnier@princeton.edu

1 Introduction

Collective foraging is a task in which a group of robots have to search for resources and exploit them (Cao et al. 1997; Baleh 1999; Shell and Matarić 2006). Robot foraging is the generalization of many real-world tasks such as search and rescue, mining, agriculture, or exploration of unknown or hostile environments. In some cases, foraging robots may use the resource immediately when it is found; for instance, in the case of batteries charging, the robots should find an electric plug and recharge their batteries immediately (Silverman et al. 2002; Wawerla and Vaughan 2008; Zebrowski and Vaughan 2005). However, in other scenarios, robots have to retrieve items from the resource and deliver them to a specific place, which we call central place (Parker 1998; Groß et al. 2008).

In the literature, a number of strategies making use of path creation techniques have been proposed for collective foraging (Nouyan et al. 2008, 2006; Werger and Matarić 1996).

Path creation allows robots to keep track of the locations of the resources and of the central place. It differs from other techniques such as map creation or dead reckoning that rely on sensors accuracy or extended computational capabilities (Burgard et al. 1998; Gutmann et al. 1999; Dudek and Mackenzie 1993). Instead, the robots literally form a chain from one place to another marking in this way a path. The main limitation of path creation lies in the number of robots used to mark the path and that can therefore not be used for other tasks. When there are several resources in the environment, the number of robots used on the different paths can be reduced by carrying out a path selection.

Selecting which path to maintain to reach a specific resource means allocating less robots to marking, as the robots used in the no longer maintained paths are free to perform different tasks. For instance, the exploitation of a resource may require a minimum number of robots to be successful. This is the case, for example, when retrieval of items involves cooperative transport (Groß and Dorigo 2009; Campo et al. 2006); if robots try to exploit several resources at the same time, they may end up in a deadlock situation in which there are too few robots assigned to each resource and none of them can be exploited. An effective collective foraging behavior should therefore implement the selection of the path to the most profitable resource, with the possibility for unused robots to stop participating in path maintenance and switch to another task.

A number of strategies that improve foraging by implementing path selection have been applied with remarkable success. Payton et al. (2001) proposed a bio-inspired mechanism making use of virtual pheromone to form a distributed computing mesh made of robots. This mechanism creates gradients of virtual pheromone from the central place to the resource and vice-versa. In order to do so, robots that are part of the mesh relay hop-count messages to each other. The robot at the central place emits a message with a given hop-count number. Upon reception of a message, the hop-count number is decreased and the message is relayed to neighboring robots. The robot at the resource also emits messages to create a second pheromone gradient from the resource to the central place. In this way, robots that are not part of the mesh are able to move to the two locations following the shortest path. However, robots relaying messages are unable to distinguish if they are useful to mark out the shortest path or not. Therefore, they cannot switch to another task if necessary and are doomed to stay in place. Moreover, when there are two identical resources in the environment, the mechanism does not allow a collective choice to take place. Most likely the robots would go indifferently to one or the other resource.


Szymanski et al. (2006) proposed a novel mechanism to solve the aforementioned problem of task allocation. The robots at the central place or at the resource emit messages that propagate through the swarm as waves. The messages are

composed of two elements: the number of hops done and the estimate number of hops left to reach the destination. Upon reception of messages, robots update their smallest hop distance to each location. After a while, robots at the central place and at the resource have received at least one message that travelled through the shortest path. Therefore, they send messages with the minimal number of hops to reach the destination. Robots of the swarm also know their hop distance to each location and can compare it with the information of the messages relayed to determine if they are on the shortest path or not. This distributed behavior allows the robots to find the shortest path from one place to the other. More importantly, the robots know when they are on the shortest path and they can switch task and participate in foraging instead of remaining along a useless path. The main limitation of this mechanism is that it cannot select among two identical resources. In addition, it seems very sensitive to transmission errors. An error in a single message is enough to prevent the whole swarm from finding the shortest path.

Schmickl and Crailsheim (2008) introduce a foraging mechanism inspired by the trophallactic behavior of bees. Robots forage in a closed arena, with one resource available and a central place. When they reach the resource, robots receive a load of virtual nectar that they gradually distribute to encountered neighbors. The peer-to-peer transmission of nectar acts as a diffusion process which creates a gradient of virtual nectar inside the population of robots, between the resource and the central place. When foraging, robots simply orient themselves using this gradient. The results of the algorithm are quite promising, but its domain of application differs from chain-based systems as it may only work with a closed arena and a number of robots large enough to cover the surface of the arena to establish a nectar gradient that can be followed by the robots. Furthermore, the system has been tested in simulation exclusively and is based on assumptions such as bi-directional infrared communication during robot motion with spatially directed sensors. Therefore, a real robot implementation seems necessary to fully understand the benefits of this approach.

Garnier et al. (2007) took a completely different direction. Instead of creating another variation of the Bellman-Ford algorithm (Bellman 1958), they simply transposed the pheromone laying behavior of ants on Alice robots (Caprari et al. 2002). As it is particularly difficult to implement pheromone sensors and actuators on robots, the authors substituted pheromone with light projected on the ground thanks to a video projector as proposed by Sugawara et al. (2004). The mechanism proves to be successful and the robots collectively select a single path to the resource. Because of the strong similarity with ants' behavior, we can safely assume that robots could also select the shortest path to the most profitable resource, for instance, the closest one. However, using this mechanism, outside laboratories seems difficult for the moment as there

Fig. 1 **a** Setup with resources at identical distances. **b** Setup with resources at different distances. *CP* central place, *R1* and *R2* resources 1 and 2, respectively



is no cheap and miniaturized technology that would allow robots to deposit and perceive pheromone reliably (Russell 1995, 1999; Stella et al. 1995; Svennebring and Koenig 2003; Purnamadajaja and Russell 2004).

In this study, we combine the realistic and effective implementation of a network of robots with the desirable properties of the pheromone laying mechanism. Path selection and task allocation (marking out the path or foraging) are achieved by implementing virtual ants that lay artificial pheromone inside a network of robots. What we call virtual ants are local messages transmitted by robots; they travel along chains of robots and deposit artificial pheromone on the robots forming the chain and indicating the path. The concentration of artificial pheromone on the robots allows them to decide whether they are part of a selected path. Our mechanism involves only three types of messages and very simple computations by the robots. These features make the mechanism suitable for robots with severely limited capabilities such as microrobots.

We study the properties of our mechanism with an analytical model and use that study to choose effective parameters. Using our mechanism, a swarm of robots is able to select a single path leading to a resource and to favor the path to the most profitable resource. The theoretical results are validated with experiments using a swarm of 20 real robots. In addition, we demonstrate the plasticity of the mechanism with real robots and study its robustness to communication errors.

2 Methods

2.1 The task of selecting between two paths

We have devised a simple task in which robots are offered two resources. We focus on the path selection process and assume that robots have already established chains between the central place and each resource using, for example, a mechanism such as the one presented by Nouyan et al. (2006, 2008). Once chains are formed, robots do not move anymore and simply exchange messages.

Figure 1 shows the two different setups employed. In the first setup, the two resources are situated at equal distance from the central place, while in the second setup, one resource is closer to the central place than the other. We always use




Fig. 2 **a** The *e-puck* robot. **b** A swarm of 20 *e-pucks* arranged in a long and a short path for the setup with resources at different distances

a total of 20 robots that are positioned at 10 cm interval to form chains.

The task allows us to test whether the robots are able to (i) perform a collective choice when offered two resources at identical distance, (ii) modulate their choice depending on the distances of the available resources from the central place, and (iii) exhibit plasticity after sudden environmental changes.

2.2 E-puck robots

E-pucks (see Fig. 2) are modular, robust, and inexpensive robots designed for research and educational purposes (Mondada et al. 2009). They are wheeled cylindrical robots, 7 cm of diameter, equipped with a variety of sensors, and whose mobility is ensured by a differential drive system. We have equipped each robot with an *E-puck Range & Bearing* board (Gutiérrez et al. 2008, 2009). The board allows the robots to communicate locally, obtaining at the same time both the range and the bearing of the emitter without the need for any centralized control or external reference. The range and bearing system allows robots to send and receive messages up to 6 m and in 12 different directions thanks to the 12 sensor/actuator pairs disposed around the body of the robot. Robots are capable of sending the same message in all directions, or specifically in one direction by using a single sensor/actuator pair. The power used to emit infrared signals can be adjusted in real time, which allows robots to decide at which maximum distance their messages can be perceived.

In this study, we have limited the emission range to 13 cm and used eight sensor/actuator pairs to limit the size of our

experimental setup, to reduce the probability of message collisions and to make sure that a message is specifically sent to a given robot and cannot be perceived by another robot in the neighborhood.

2.3 The simulator

Our simulation platform is a fast multi-robot simulator for the *e-puck* robot. The simulator uses a quad-tree to speed up algorithms that handle the detection of collisions and the detection of local infrared communication. Motion of objects is handled by a custom rigid body physics engine, specialized to simulate only the dynamics in environments containing flat terrain and walls. This restriction allows simplifications of the physics computations and thereby reduces the computational resources necessary for running simulations.

In the simulator, an *e-puck* is modelled as a cylindrical body that holds eight infrared sensors distributed around the body and a range and bearing communication system. A differential drive system made up of two wheels is fixed to the body of the simulated robot. All the simulated sensors and actuators are modelled after measures taken on real robots. In order to simulate effectively the sensory input of real robots, we added noise to the simulated sensory data. In the range and bearing communication system, noise is added to the range (± 2.5 cm) and bearing ($\pm 20^\circ$) values. By default, messages transmitted with the range and bearing system were delivered with no errors. For the study of robustness to communication errors (see Sect. 4.5), we implemented a probability that a message became corrupted. Corrupted messages could be either lost or their content could be altered into any other possible content.

2.4 The robots' controller

The robots' controller is designed to produce a selection mechanism at the collective level. The main part of the selection mechanism consists of virtual ants travelling along the chains of robots from the central place to the resource and vice-versa. Virtual ants are implemented as two messages transmitted from robot to robot (called VA_{CP} when going to the central place and VA_R when going to a resource).

The virtual ants deposit artificial pheromone on the robots forming the chain and indicating the path. As virtual ants take less time to traverse a shorter path, in a given amount of time they will make more trips on short paths than on long ones. Shorter paths will therefore receive larger amounts of pheromone than longer paths. The amplification process will increase further this difference: relay robots preferentially send virtual ants in directions that have already higher amounts of pheromone. Amplification determines a collective choice even when resources are at identical distances. In fact, even though on average these paths receive the same

amount of pheromone, a slight and random difference in pheromone concentrations can be amplified and lead to a final collective choice.

Only robots that perceive either the central place or a resource have the possibility to emit new virtual ants. These ants are then relayed by any robot that perceives them until they reach their destination. There is no fixed population of virtual ants: they are created at a rate determined by the emitters' concentration of pheromone and disappear if they are not relayed, typically when reaching the extremity of a chain.

2.4.1 Establishing the communication network

Before emitting or relaying virtual ants, robots need to know in which direction to send them. In a first stage, the robots organize themselves in a network with a shortest path tree topology that allows ants emitted from a resource to travel straight to the central place along the shortest path.

In order to do so, the robot perceiving the central place probabilistically emits wavefront signals. These signals propagate through the swarm and inform each robot of the swarm about which of its immediate neighbors is closer to the central place (see Fig. 3a). In practice, wavefront signals are implemented by a particular message (called SPT) which is first emitted at the central place. Upon detection of an SPT message, a robot interrupts any current action and repeats the message, sending it in all directions. In order to ensure that the messages are propagated only forward, robots that have just sent an SPT message enter an inhibited state of $h = 0.35$ s during which they do not receive or send any message.

As described in Fig. 3b, this method creates a communication network with a shortest path tree topology and is general enough to work with any arrangement of the robots, as long as the swarm forms a single connected network. We have also implemented a regular update of the communication network because it is the most sensitive element of the system. If a communication error happens during the construction of the shortest path tree, or if a robot has moved, our assumption that virtual ants are able to travel between the central place and the resources may not be verified anymore. To cope with these unlikely but possible events, the communication network is updated by sending new wavefront signals at a rate of $u = 0.01$ updates per second.

2.4.2 The motion of virtual ants within the network

Since virtual ants are messages in the system, the emission and relay of these messages within the communication network are achieved by the robots themselves.

Figure 4 provides a detailed view of a robot relaying a virtual ant. Robots have a total of eight pairs of




Fig. 3 The robots organize themselves in a network with a shortest path tree topology. **a** The robot perceiving the central place emits a wavefront signal that is in turn retransmitted by every robot perceiving

it, thus propagating the signal through the entire swarm. **b** Each robot remembers the direction from which the wavefront was perceived first, creating the desired topology at the collective level

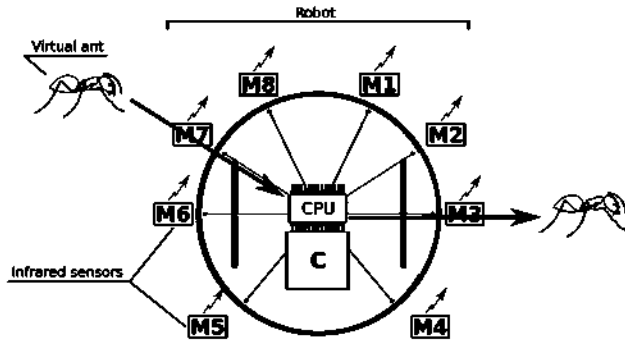


Fig. 4 Detailed schema of a robot relaying a virtual ant. During this process, the robot receives the virtual ant on sensor/actuator $j = 7$, updates the memory register M_j associated to sensor/actuator j as well as the concentration of pheromone C , and finally transmits the virtual ant in a new direction through another sensor/actuator $i = 3$

sensors/actuators, positioned around their body, that receive or send virtual ants. The physical location of the sensors/actuators allows the robots to infer the incoming direction of virtual ants and also to send them in specific directions.

Because of infrared reflections on the ground or on the body of other robots, a robot might perceive an ant he just relayed ahead of himself. In order to avoid that a robot relays several times the same ant by mistake, we use two different messages for ants going to the central place (VA_{CP}) or ants going to the resources (VA_R). For instance, a virtual ant VA_{CP} perceived by a robot in the direction of the central place will not be relayed again by this same robot because it is already heading towards its destination.

There are eight memory registers M_1, \dots, M_8 associated to the eight pairs of infrared sensors/actuators. When a virtual ant is perceived by the sensor/actuator i , memory M_i is updated by adding an amount $l \in [0, 1]$ of pheromone:

$$M_i = M_i + l. \quad (1)$$

Therefore, Eq. 1 implements the same fundamental principle as described by Goss et al. (1989) to explain path

selection in real ants. In addition, the relaying robot updates his concentration of pheromone C according to Eq. 2:

$$C = C + (1 - C)l. \quad (2)$$

When updated, the concentration of pheromone C converges asymptotically to 1. Pheromone concentration is used to tell, after convergence, whether a robot is on the selected path: this is the case if a robot's C value exceeds a given threshold. C is also the rate per second at which robots may emit new ants. Note that only robots perceiving the central place or the resources may emit new ants (that is create them), other robots are relaying these ants.

When a relaying robot perceives with sensor/actuator j an ant going to a resource, it has to decide in which direction to send it. The robot calculates the probability $P(i, j)$ to send this ant with actuator i using Eq. 3:

$$P(i, j) = \frac{M_i^s}{\sum_{k \neq j} M_k^s}, \quad (3)$$

where $s > 1$ is a factor that introduces non-linearity in the system by amplifying small differences among possible directions. For instance, if we consider that $s = 2$ and there are only two sensors/actuators and one of them has received 10% more pheromones than the other, the probability to relay a virtual ant with that sensor/actuator is higher by 21%: $110^2 / (110^2 + 100^2) = 1.21 \cdot 100^2 / (110^2 + 100^2)$.

Similarly to its natural counterpart, artificial pheromone evaporates at a constant rate f per second:

$$C(t + \Delta t) = C(t) - \Delta t \cdot C(t) \cdot f, \quad (4)$$

$$M_i(t + \Delta t) = M_i(t) - \Delta t \cdot M_i(t) \cdot f \quad \forall i \in [1, 8], \quad (5)$$

where t is the time and Δt is the controller time step. This evaporation ensures that pheromone concentration decreases on a path that is not selected. The selection of a path is maintained only if virtual ants keep on travelling along the path and depositing artificial pheromone. If a sudden change in the environment makes a path unavailable, the swarm slowly

Table 1 Summary of variables and parameters of the experiment

Variables	Description	Initial
α	Rate of ants emitted at central place	0.5 ants/s
β_i	Rate of ants emitted at R_i	0.5 ants/s
Parameters	Description	Default
n	Number of resources in the experiment	2
s	Amplification factor	2
l	Pheromone deposited (learn)	0.88
f	Evaporation rate (forget)	0.029
d_1	Distance of R1 to central place	0.8 m
d_2	Distance of R2 to central place	1.6 m
h	Duration of inhibition period	0.35 s
v	Speed of a virtual ant	5 m/s
ϵ	Minimum pheromone concentration	0.016 ants/s
u	Update rate of the network topology	0.01 /s

forgets it. However, ignored paths may need to be reactivated. For this purpose, we define a minimum concentration of pheromone ϵ for every robot that guarantees a minimal traffic of virtual ants.

An experiment typically unfolds with robots waiting for an initial wavefront signal that will create the first shortest path tree. Once the wavefront signal is received, robots at the resources emit virtual ants toward the central place. When receiving virtual ants, the robot at the central place starts to send virtual ants as well. The virtual ants are emitted at a rate that is equal to the pheromone concentration of the emitting robot. Initially, the pheromone concentration on all the robots is set to 0.5. We consider that the selection process is completed when the pheromone concentration of the non-selected path has decreased below a predefined threshold.

3 Analytical model

3.1 Model description

The previous definitions of the robots' task and controller involve several free parameters, such as the distances d_1 and d_2 of the two resources to the central place, as well as the evaporation rate of pheromone f (forgetting process) and the amount of pheromone l deposited by a virtual ant on a robot during relaying (learning process).

In order to understand the impact of these parameters on the collective behavior of the robots, we devise a mathematical model and study its steady states along with the probabilities to reach them. Our approach is not aimed at finding the best parameters for a specific case. Rather, we try to obtain a

better understanding of the collective behavior and to extract general rules to parameterize effectively the robots in the subsequent experiments.

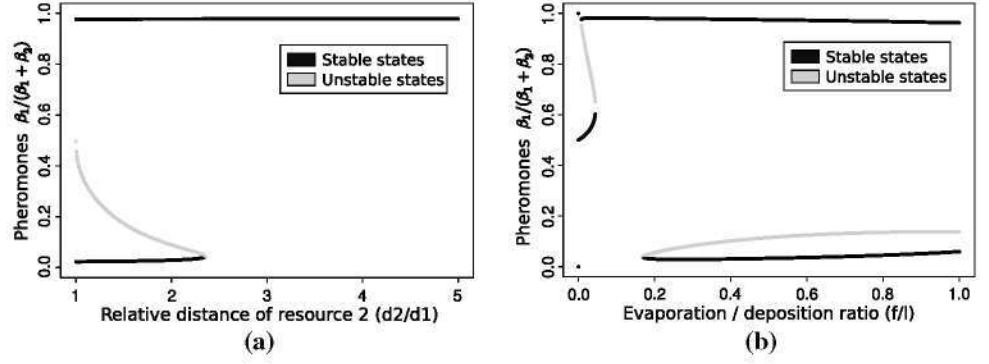
We model the concentrations of pheromone at the central place (CP) and at the resources (R1 and R2) over time with the variables α , β_1 , and β_2 , respectively (see Table 1). As said in Sect. 2, pheromone concentration is used as the rate of emission of virtual ants (ants/s). Hence, the number of ants emitted at R_i over a time interval T is given by $m_i = \beta_i T$. We assume that the transmission of an ant lasts exactly h seconds, so the number m_i of ants that can be emitted at R_i over a period T must satisfy the constraint $m_i h \leq T$.

When a robot is relaying a virtual ant, it cannot receive another ant at the same time because of hardware limitations. If another ant is sent to the robot during transmission, we say that a *collision* happens. The first ant is transmitted, while the other one is lost. When $n > 1$ resources are present in the environment, the possible paths to the central place intersect. Robots at intersections may receive ants coming from each resource and are likely to experience collisions more frequently than others. For this reason, we include collisions at intersections in our model to take their effect into account. We call x_i the number of ants per second emitted at R_i not experiencing a collision. At any moment, the probability p_i that an ant emitted at R_i collides at the intersection point with another ant emitted at R_j ($j \neq i$) is given by:

$$p_i = \frac{(x_j T)h}{T} = x_j h.$$

Hence, the rate of ants emitted at R_i and not experiencing collisions is: $x_i = \beta_i(1 - p_i) = \beta_i(1 - x_j h)$. We obtain a system of equations that can be solved using a simple Gaussian elimination:

Fig. 5 Equilibria of the analytical model when pheromone concentrations are stable. **a** Impact of distances ratio d_2/d_1 . **b** Impact of evaporation/deposition ratio f/l



$$\begin{cases} x_i = \beta_i(1 - x_j h), \\ x_j = \beta_j(1 - x_i h), \end{cases} \Leftrightarrow \begin{cases} x_i = \frac{\beta_i - \beta_i \beta_j h_j}{\beta_i \beta_j h_i h_j - 1}, \\ x_j = \frac{\beta_j - \beta_i \beta_j h_i}{\beta_i \beta_j h_i h_j - 1}. \end{cases}$$

Following the description of the robots' controller given in Sect. 2.4, the dynamics of α and β_i are modelled by the following differential equations:

$$\begin{cases} \frac{d\alpha}{dt} = -f(\alpha - \epsilon) + l(1 - \alpha) \sum_i \varphi_i, \\ \frac{d\beta_i}{dt} = -f(\beta_i - \epsilon) + l(1 - \beta_i) \alpha \frac{\varphi_i^s}{\sum_k \varphi_k^s}. \end{cases}$$

where $\varphi_i = x_i v / d_i$ is the number of ants travelling along the path from R_i to the central place.

3.2 Parameter study

We study the stable states of the model when pheromone concentrations reach a steady state. We focus on the influence of the relative distances of resources, as well as the deposition and evaporation coefficients. The values of parameters h and v in Table 1 are set to the maximum allowed by the hardware communication system at the time of the study. The amplification factor s is set to 2 (it must be strictly higher than 1 to introduce non-linearities and trigger amplification). Table 1 gives the default values of the parameters.

We define the relative concentration of pheromone $k = \beta_1 / (\beta_1 + \beta_2)$, a coefficient to study path selection. If $k \approx 0$, the concentration of pheromone at R1 is negligible while it is high at R2. In this case, we say that the path to R2 is selected. Conversely, if $k \approx 1$, the path to R1 is selected. If $k \approx 0.5$, the swarm is unable to select a single path.

Figure 5a shows coefficient k when the relative distance to R2 increases. When resources are at the same distance, we observe two stable states with one unstable state at $k = 0.5$. The swarm selects indifferently one of the two paths. As soon

as R2 is farther than R1, the stable state $k \approx 1$ in which the path to R1 is selected has a larger basin of attraction and the probability to converge there is increased. When the ratio of distances $d_2/d_1 > 2.5$, the swarm always selects the path to R1.

Figure 5b shows the impact of the evaporation and deposition ratio (f/l) on the model's stable states when the ratio of distances $d_2/d_1 = 2$. We have set the deposition coefficient to $l = 0.88$ to achieve a relatively fast collective behavior. We observe that low values of f generate a third stable state in which $k = 0.5$ and the swarm is unable to make a selection. When $f/l < 0.2$, the model may have up to three stable states, and depending on noise and initial conditions, it may not be able to select a path. As soon as $f/l > 0.2$, the model has two stable states in which the swarm selects the path to either R1 or R2. The unstable state is very close to 0 and we observe a larger basin of attraction of stable state $k = 1$ in which the path to R1 is selected.

Based on the previous observations, we have set parameters d_1, d_2, l, f with the values reported in Table 1. Unless stated otherwise, these settings are used to perform the experiments reported in the following section.

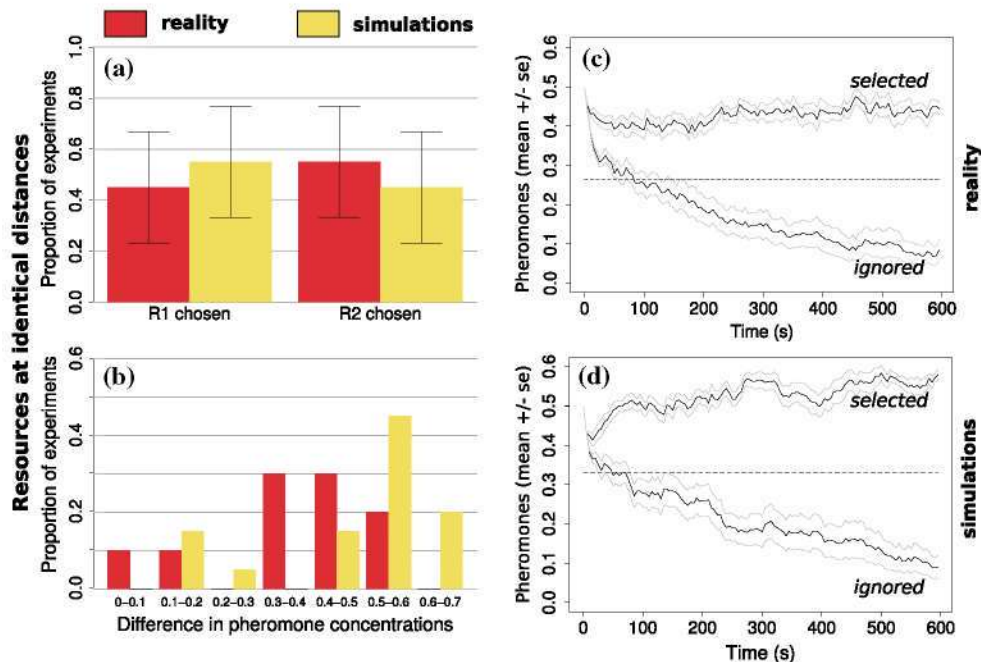
4 Results

In the following, we report results obtained from experiments with real and simulated robots. For each setup with resources at identical or different distances, we perform 20 repetitions. Experiments concern the collective choice of paths, the plasticity of the selection process, and the robustness to communication errors.

4.1 Resources at identical distances

When resources are located at the same distance from the central place, the swarm arbitrarily selects one of the paths (binomial test, $p > 0.35$; see Fig. 6). The swarm's collective choices in reality and in simulations do not differ statistically (proportions test, $p > 0.75$). At the beginning of the

Fig. 6 Results with real and simulated robots for setups with resources at identical distances. **a** Proportion of repetitions that ended with the choice of R1 or R2. **b** Distribution of differences in pheromone concentrations ($|\beta_1 - \beta_2|$), accounting for the discrimination of the swarm among the paths. **c, d** Average pheromone concentrations over time at the selected and ignored paths



experiments, the concentrations of pheromone on the two paths are similar, and virtual ants are routed towards either resource. The amplification of random fluctuations in the emission of ants triggers the swarm’s collective choice. One of the two paths receives more and more virtual ants laying artificial pheromone, while the pheromone concentration on the other path decreases quickly to the allowed minimum. The selection achieved by the swarm is stable, both in reality and in simulations. In simulations, most of the experiments (85%) end with a clear selection, where the difference in pheromone concentrations is high enough to define a threshold (>0.25) for discriminating among paths. With real robots, we observe that 80% of the experiments end with a clear discrimination of the paths. The remaining experiments exhibit differences that are too small to allow discrimination. This slight difference between simulations and reality probably arises because of two related factors that are reducing the strength of the amplification process: the non-modelled collisions of messages travelling in opposite directions and the vanishing of virtual ants.

4.2 Resources at different distances

When resources are located at different distances, the swarm selects the path to the closest resource more often (binomial test, $p < 0.001$). As shown in Fig. 7, that path is chosen in 95% of the cases with real robots. The concentration of pheromone is maintained for the chosen path, while it quickly drops for the other one. We observe a good agreement between reality and simulations, both for the evolution of pheromone concentrations in time and the differences in

concentrations. However, with real robots, concentrations of pheromone are not as high as expected. Again, this behavior can be explained by the non-modelled collisions of virtual ants. This phenomenon is also responsible for the better capability of the swarm to choose the path to the closest resource in reality, although this difference is not statistically significant (proportions test, $p > 0.18$). Collisions of virtual ants are more likely to occur on a longer path because virtual ants spend more time travelling from one end to the other. Hence, the path to the farthest resource is more affected by the negative feedback of collisions, which results in the swarm discriminating more accurately among the two resources.

4.3 Resources of different qualities

The system proposed is not bound to select paths only based on their length. It is also possible for the robots to modulate their collective choice as a function of the quality of the resources detected. We make the assumption that robots can evaluate the quality of the resources when they detect them and that they can translate quality into a number $Q \in [0, 1]$. For this specific experiment, we modify the behavior of the robots so that if they detect a resource, they emit virtual ants at the rate $r = C \cdot Q$, where C is the pheromone concentration on the robot and Q is the quality of the resource detected.

We run two simulation experiments using resources at identical distances. In one experiment, resource R1 has a quality set to $Q_{R1} = 1$, while resource R2 has a quality $Q_{R2} = 0.8$. In the second experiment, the qualities of the resources are inverted. Figure 8a and b summarizes the collective choice of the robots in these two experiments. When

Fig. 7 Results with real and simulated robots for setups with resources at different distances. **a** Proportion of repetitions that ended with the choice of R1 or R2. **b** Distribution of differences in pheromone concentrations ($|\beta_1 - \beta_2|$), accounting for the discrimination of the swarm among the paths. **c, d** Average pheromone concentrations over time at the selected and ignored paths

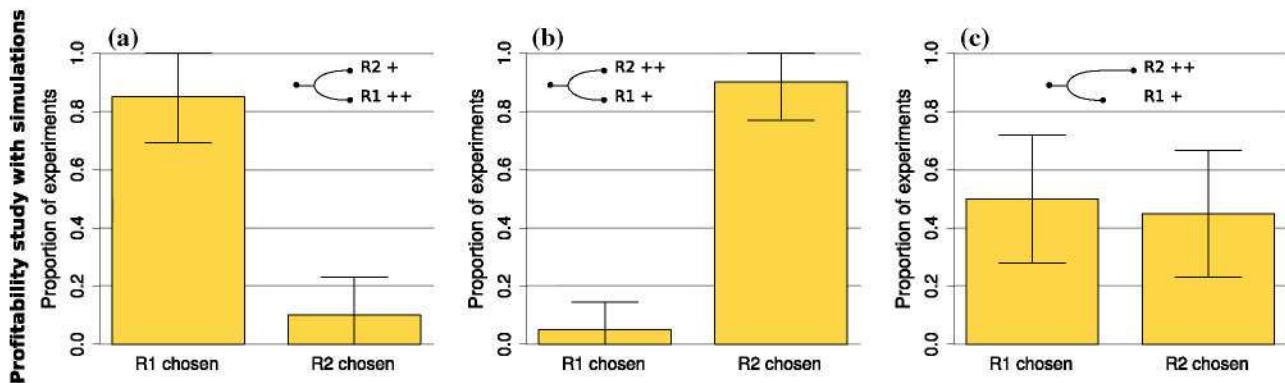
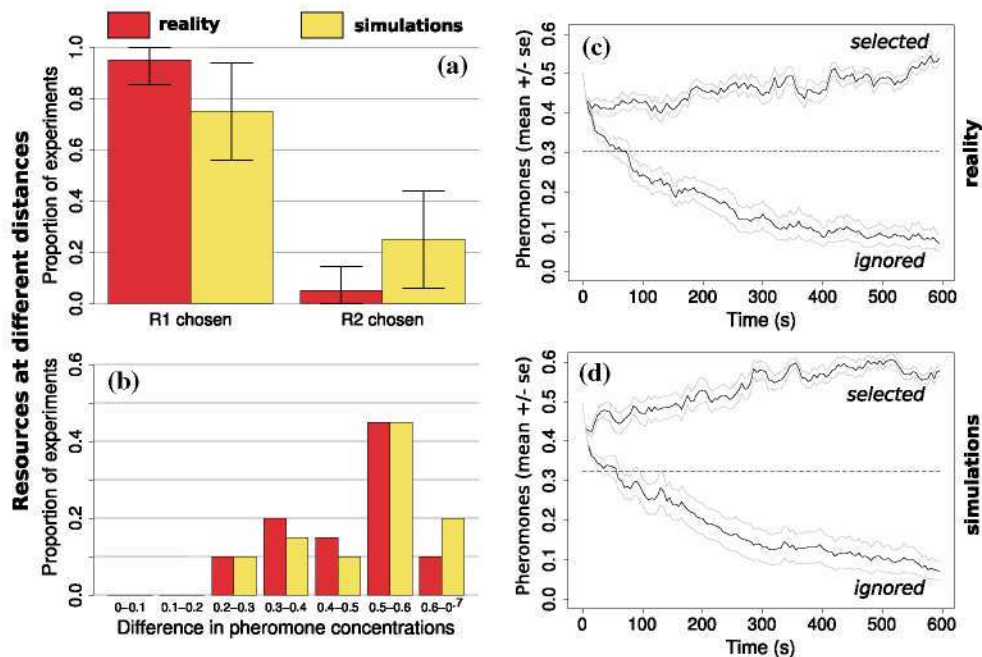


Fig. 8 Selection of the resources based on their quality. We report the proportion of trials that end up with the selection of R1 or R2 in three different sets of simulation experiments. **a, b** Resources are situated at identical distances from the central place. **a** R1 has a higher quality than R2, **b** R1 has a lower quality than R2. Robots choose the resource with

best quality. **c** Resources are situated at different distances and their respective qualities are tuned to compensate for the effects of distance. The choice of the robots is randomly in favor of R1 that is closer but with a lower quality, or R2 that is farther but with a higher quality

R1 has a better quality than R2, it is selected in 85% of the trials. In the symmetrical situation where R2 has a better quality than R1, the robots select R2 in 85% of the cases. This contrasts with the previous experiments in which resource quality is not defined and only distance is taken into account by the robots and shows that the collective choice of the robots can be steered by the quality of the resources detected.

It is worth mentioning that the collective choice of the robots depends on the resources profitabilities which are function of their distances and qualities. In order to show this duality, we run simulation experiments with resources at different distances in which the quality of the resources are specifically tuned to compensate for the effects of distance ($Q_{R1} = 0.92$; $Q_{R2} = 1.00$). In Fig. 8c, we observe that the

robots do not distinguish between the two offered resources and they select randomly R1 or R2 (binomial test, $p > 0.74$). This is because R1 is closer but with a lower quality, while R2 is farther with a higher quality.

4.4 Plasticity of the selection process

In order to test the response of the swarm when there is a sudden environmental change, that is, when a selected resource becomes unavailable or when a new resource appears in the environment, we run three sets of experiments with real robots. In all these experiments, we use the setup with resources located at different distances.

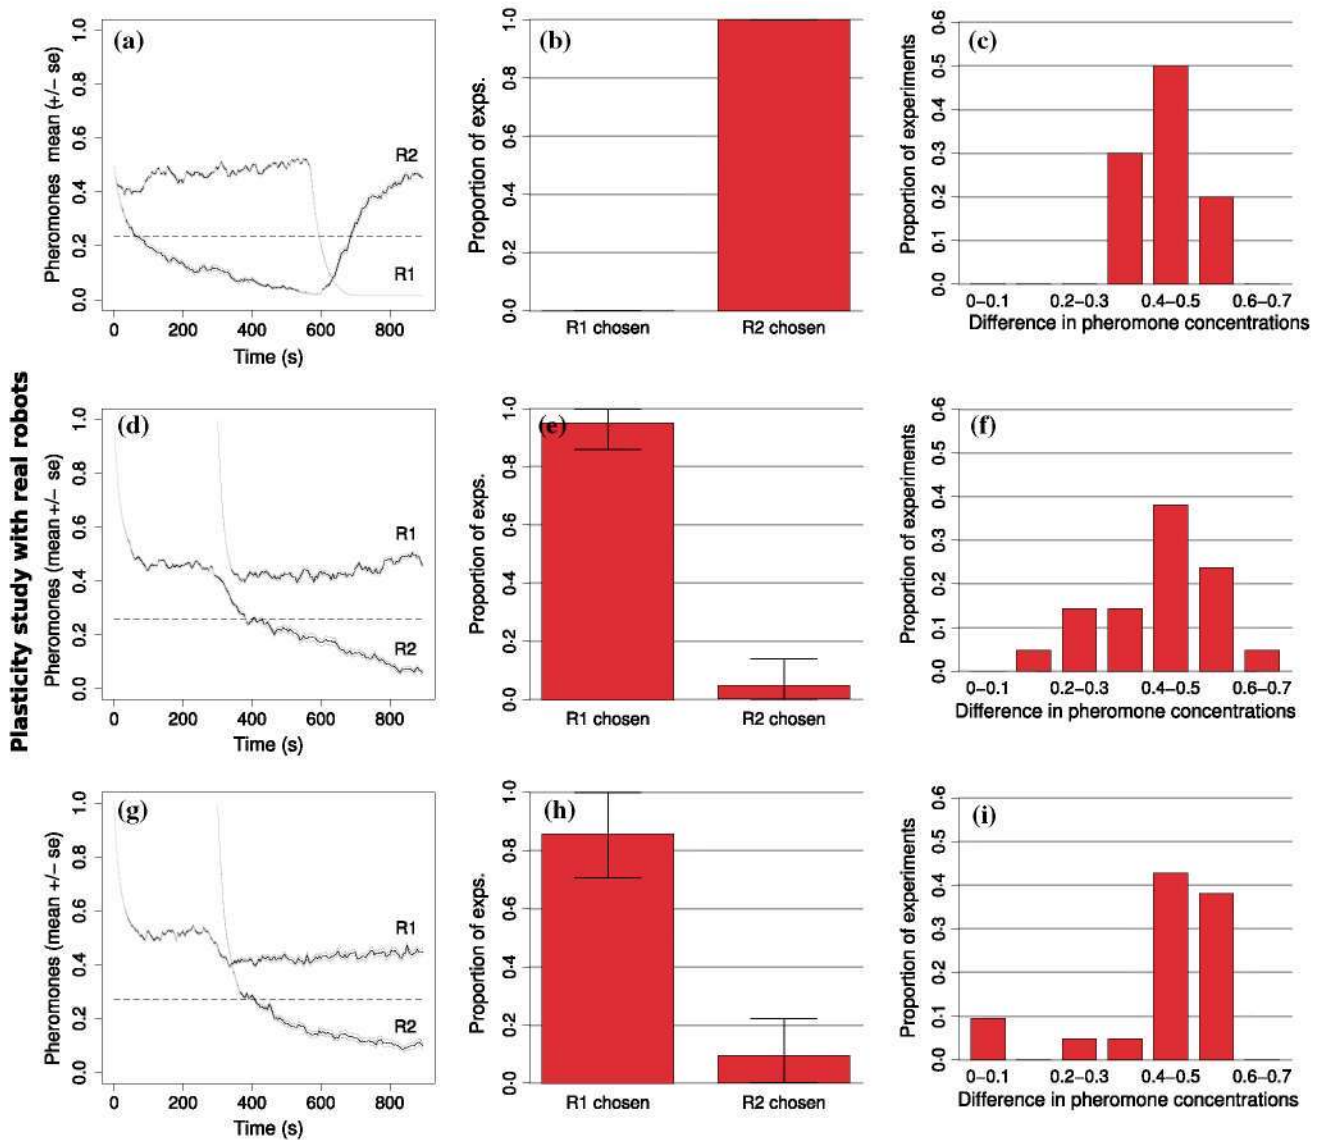


Fig. 9 Plasticity of the swarm selection in experiments with real robots. Three different experiments are carried out using a setup with resources at different distances. *First column* Average pheromone concentrations over time on the paths. *Second column* Proportions of trials in which R1 or R2 is selected at the end of experiments. *Third column* Difference in pheromone concentrations ($|\beta_1 - \beta_2|$) at the end of experiments. **A, B, C** Initially R1 is selected. After 600 s, the path to R1 is


blocked. The robots adapt their choice towards the remaining resource R2. **D, E, F** Initially R2 is selected as it is the only resource available. After 300 s, R1 becomes available. The robots detect the new and better resource R1 and select it. **G, H, I** Initially R1 is selected as it is the only resource available. After 300 s, R2 becomes available. The robots detect the new resource R2 but maintain their initial choice in favor of the better resource R1

In the first set of experiments, the selected path to the closest resource is suddenly blocked (see Fig. 9a–c). After 600 s, we introduce an obstacle along the selected path to the closest resource R1. After the obstacle is introduced, the swarm reorganizes itself and selects the only path left. This is the outcome of two factors: first, the pheromone concentration in the shortest path drops quickly as virtual ants are blocked by the obstacle. Second, the longest path has maintained a minimum concentration of pheromone ϵ that allows it to have a low activity. As soon as the swarm has forgotten about the

unavailable resource R1, it strongly amplifies its activity in the second path.

In the second set of experiments, we provide only one resource (R2) and later we introduce a new resource closer to the central place (R1) (see Fig. 9d–f). The resource R1 becomes available in the environment 300 s after the start of the experiment. The robot that detects it immediately starts to emit virtual ants towards the central place. During 60 s on average (until $t = 360$ s), a competition between the two paths takes place. The new resource is closer and the detecting

Fig. 10 Robustness of the selection process in simulation experiments. We report the proportion of trials ($\pm 95\%$ confidence intervals) that end up with the selection of R1 or R2 as a function of the probability error per message transmitted. The performance of the robots degrades gracefully until the probability is greater than 0.188, in which case no path is selected



robot starts with a higher concentration of pheromone (we have set it to 1) but on the rest of the path, the concentration of pheromone is still minimal. Hence, many virtual ants are emitted at R1 to the central place, but none are sent back to R1. It takes several virtual ants emitted at R1 to level up the pheromone concentration along the whole new path so that virtual ants start to travel from the central place to R1. At the same time, the path initially selected experiences a decrease in the pheromone concentration due to collisions of virtual ants emitted on the short path. As expected, the short path receives virtual ants faster than the long path and, consequently, reacts faster. The difference in reaction time of the short path is amplified and the robots select the short path in 95% of the trials.

The previous experiments are not sufficient to claim the plasticity of the system. Indeed, the robots could systematically select the path to a new resource appearing in the environment, even if the new resource is less profitable than the one currently selected. In order to validate the plasticity of our system, we run a new experiment in which we provide one resource (R1) and introduce a new and more distant resource (R2) after 300 s (see Fig. 9g–i). Again, during a period of 60 s on average (until $t = 360$ s), the paths to the two resources compete. The selected short path has a small drop in the pheromone concentration incurred by the sudden activity on the long path. However, the pheromone concentration on the long path never grows because of the inherent disadvantage of profitability. The robots maintain their initial choice in 90% of the cases.

In summary, plasticity allows robots to modify their collective choice in favor of a path to a newly detected resource if and only if that resource is more profitable than the one already selected.

4.5 Robustness to communication errors

We test the robustness of our system to communication errors by introducing a probability that a message sent by a robot

is corrupted. When corrupted, messages can be either lost or transformed into any other type of message with equal probability. In total, we use only three types of messages: virtual ants (VA_{CP} and VA_R), and wavefront signals used to build the topology of the network (SPT).

In this specific experiment, we rely exclusively on simulation experiments and we use a setup with resources located at different distances. We also must adapt the parameters of the system to make pheromone deposition and evaporation slower. Indeed, corrupted messages reduce the number of virtual ants travelling along the paths. An immediate consequence is that pheromone concentration drops faster when there are communication errors. In order to cope with this effect, we set the amount of pheromone deposited per virtual ant to $l = 0.06$ and the rate of pheromone evaporation to $f = 0.01$. All other parameters are kept identical to the default configuration reported in Table 1.

The probability that a message gets corrupted is studied in the range $[0, 0.25]$. In Fig. 10, we report for each error probability tested the proportion of trials that end up with the selection of the short path and of the long path. When there are no communication errors, the robots select the short path in 95% of the cases. When the error probability is in the range $[0.002, 0.056]$, we observe an improvement in the capability of the robots to select the short path, in the best case always selecting the short path. This is because errors are more likely to happen on the long path. The performance of the system gracefully degrades as the probability of communication errors increases and the robots either select the short path or do not make any choice. In any case, the long path is never selected in more than 5% of the trials. When the error probability is equal to 0.188, the robots are not able anymore to perform a collective choice in 92% of the trials, and if they do a collective choice, it is indifferently in favor of the short or of the long path. Finally, when the communication error probability is greater than 0.188, the robots are not anymore able to select any path reliably.

5 Conclusions

In this article, we described an approach to path selection inspired by the trail laying behavior of ants during a foraging task. We devised and studied an analytical model to gain high level understanding of the robots' collective behavior at steady state and to be able to parameterize it appropriately. We experimentally validated the mechanism with a swarm of 20 real robots and studied properties of plasticity and robustness to communication errors.

Artificial pheromone allowed robots to perform a collective choice between two identical resources available in the environment. Moreover, the robots' choice was modulated by the profitability of the resources. When offered two resources located at different distances, robots selected the path to the closest resource most of the times. Robots could also distinguish the resources based on their quality, in addition to their distance to the central place. If a sudden change occurred in the environment and a resource became unavailable, robots quickly reoriented their choice in favor of a path to another available resource. Finally, if a new resource appeared, robots changed their choice if and only if that resource was more profitable than the one already selected.

Since each robot of the swarm stored a local concentration of pheromone, we defined a concentration threshold above which robots were considered part of a selected path. Robots that were not needed to maintain paths could focus on another task. Hence, the swarm was able to focus on a resource and reallocate robots towards its exploitation once exploration and selection were completed.

The collisions of virtual ants with those going in opposite directions were not modelled neither analytically nor in simulations. Experimental results showed that they generated a negative feedback reducing the strength of amplification and slightly slowing down the selection process of the robots. Interestingly, because collisions are more likely to occur along paths to farther resources, the discrimination ability of the swarm was enhanced and robots better distinguished differences between the two paths.

Our system showed the ability to cope with communication errors without any specific additional mechanism to correct errors or check the validity of the messages transmitted. The performance of the robots degraded gracefully as the probability of communication errors increased. In comparison, the performance of systems based on dynamic programming could be disrupted by a single communication error. In an experiment involving 20 robots each transmitting one message per time step, a communication error probability of 0.05 per message means on average one communication error per time step in the whole swarm. We showed that in these conditions, robots using artificial pheromones are still performing at the top of their capabilities. However, in the case a robot fails completely, an entire path could be

disabled. Since robots can reorganize and select an alternative when their initial collective choice is no longer available, it would be interesting to have at hand a mechanism maintaining redundant paths to resources.

Tuning parameters is critical to obtain the different properties of the system. In our experiments, the parameter settings used to study the robustness of the system differed from the settings used to study the plasticity of the system. The rates of pheromone deposition and evaporation had to be lowered so that the system would not forget about a path because of non-systematic communication errors. In contrast, the plasticity of the system relied on the ability of the system to forget an inaccessible path thanks to pheromone evaporation. Therefore, the properties of robustness and plasticity may not coexist in all situations and appropriate parameters should be set depending on the desired behavior of the system and the environment in which it is used. In order to increase the efficiency of the system in a specific application, it may be necessary to optimize the parameters with tools such as genetic algorithms or simulated annealing.

On a long-term perspective, the outlined mechanism is an interesting candidate for implementations with miniaturized robots. Control is totally distributed and does not rely on complex computations or sophisticated communications. Robots need only to send three different types of messages (SPT, VA_{CP} , VA_R), and do not need explicit knowledge of distances. The communication requirements are rather low: during experiments with real robots, we measured that a robot communicated on average 0.24 bytes per second. Notably, the implemented system does not require to assign identifiers to each robot, thus improving scalability of the collective behavior.

The need for robots to perform a collective choice and select only one single resource was one of our main concerns in this work. However, if the number of robots available for foraging is large, it would be more interesting to have the robots exploit several resources at the same time and avoid negative interferences due to overcrowding (Goldberg and Mataric 1997). We suggest a simple modification of our selection system that may allow to take advantage of that situation. The modification consists in setting the parameter s to 1, thereby removing amplification from the system and distributing pheromone concentrations on the paths proportionally to resource profitability. In that case using a threshold to allocate robots on paths would no longer be feasible, rather a probability of dismantling a path could be used. Paths with the lowest pheromone concentrations would gradually disappear and the paths to the most interesting resources would remain exploited.

Perspectives for future work include the integration of the present selection mechanism with other works that focussed on path creation and resource exploitation (Nouyan et al. 2008, 2006; Groß et al. 2008). These previous works have

studied thoroughly scenarios of foraging in which chains of robots are used, and very convincing experiments with real robots have been carried out. Plugging path selection into these controllers would allow the study of complex foraging situations with paths to multiple resources.

More specifically, we are considering to investigate scenarios with multiple identical resources in which robots have to focus on resources one after the other to avoid deadlocks triggered by a shortage of robots available for exploitation. We are also interested in studying the ability of robots to deal with several resources in parallel, and self-organized task allocation mechanisms to let the robots decide when to explore the environment, when to exploit resources and when to participate in path maintenance and selection.

Acknowledgements Alexandre Campo and Marco Dorigo acknowledge support from the Belgian F.R.S.-FNRS. Carlo Pinciroli acknowledges support from COMP2SYS, a Marie Curie Early Stage Research Training Site funded by the European Commission under contract MEST-CT-2004-505079. Valentin Longchamp acknowledges support from the Centre Suisse d'Electronique et de Microtechnique via a PhD research grant. Simon Garnier is supported by a research grant from the French Ministry of Education, Research and Technology.

References

- Balch TR (1999) Reward and diversity in multirobot foraging. In: Proceedings of IJCAI-99 workshop on agents learning about, from and with other agents, Morgan Kaufmann, San Mateo, CA, pp 553–561
- Bellman R (1958) On a routing problem. *Q Appl Math* 16(1):87–90
- Burgard W, Derr A, Fox D, Cremers AB (1998) Integrating global position estimation and position tracking for mobile robots: the dynamic Markov localization approach. In: Proceedings of the IEEE international conference on intelligent robots and systems. IEEE Press, Piscataway, NJ, pp 1–6
- Campo A, Nouyan S, Birattani M, Groß R, Dorigo M (2006) Negotiation of goal direction for cooperative transport. In: Proceedings of the 5th international workshop on ant colony optimization and swarm intelligence. Lecture Notes in Computer Science, vol 4150. Springer, Berlin, pp 191–202
- Cao YU, Fukunaga AS, Kahng A (1997) Cooperative mobile robotics: antecedents and directions. *Auton Robots* 4(1):7–27
- Caprari G, Estier T, Siegwart R (2002) Fascination of down scaling—alice the sugar cube robot. *J Micro-Mechatron* 1(3):177–189
- Dudek G, Mackenzie P (1993) Model-based map construction for robot localization. In: Proceedings of vision interface. CIPPR Society Press, Toronto, pp 97–102
- Garnier S, Tache F, Combe M, Grimal A, Theraulaz G (2007) Alice in pheromone land: an experimental setup for the study of ant-like robots. In: Proceedings of IEEE swarm intelligence symposium. IEEE Press, Piscataway, NJ, pp 1–5
- Goldberg D, Mataric MJ (1997) Interference as a tool for designing and evaluating multi-robot controllers. In: The fourteenth national conference on artificial intelligence. AAAI Press, Menlo Park, CA, pp 637–642
- Goss S, Aron S, Deneubourg JL, Pasteels JM (1989) Self-organized shortcuts in the argentine ant. *Naturwissenschaften* 76(12):579–581
- Groß R, Dorigo M (2009) Towards group transport by swarms of robots. *Int J Bio-Inspired Comput* 1(1–2):1–13
- Groß R, Nouyan S, Bonani M, Mondada F, Dorigo M (2008) Division of labour in self-organised groups. In: Proceedings of the 10th international conference on simulation of adaptive behavior. Lecture Notes in Artificial Intelligence, vol 5040. Springer, Berlin, pp 426–436
- Gutiérrez A, Campo A, Dorigo M, Amor D, Magdalena L, Monasterio-Huelin F (2008) An open localization and local communication embodied sensor. *Sensors* 8(11):7545–7563
- Gutiérrez A, Campo A, Dorigo M, Donate J, Monasterio-Huelin F, Magdalena L (2009) Open e-puck range and bearing miniaturized board for local communication in swarm robotics. In: Proceedings of the IEEE international conference on robotics and automation. IEEE Press, Piscataway, NJ, pp 3111–3116
- Gutmam JS, Weigel T, Nebel B (1999) Fast, accurate and robust self-localization in polygonal environments. In: Proceedings of the IEEE international conference on intelligent robots and systems. IEEE Press, Piscataway, NJ, pp 1412–1419
- Mondada F, Bonani M, Raemy X, Pugh J, Cianci C, Klapcoz A, Magnenat S, Zufferey JC, Floreano D, Martinoli A (2009) The e-puck, a robot designed for education in engineering. In: 9th conference on autonomous robot systems and competitions. IPCB, Instituto Politécnico de Castelo Branco, Castelo Branco, pp 59–65
- Nouyan S, Groß R, Bonani M, Mondada F, Dorigo M (2006) Group transport along a robot chain in a self-organised robot colony. In: Proceedings of the 9th international conference on intelligent autonomous systems. IOS Press, Amsterdam, pp 433–442
- Nouyan S, Campo A, Dorigo M (2008) Path formation in a robot swarm: self-organized strategies to find your way home. *Swarm Intell* 2(1):1–23
- Parker LE (1998) Alliance: an architecture for fault tolerant multirobot cooperation. *IEEE Trans Robot Autom* 14(2):220–240
- Payton D, Daily M, Estowski R, Howard M, Lee C (2001) Pheromone robotics. *Auton Robots* 11(3):319–324
- Purnamadajaja AH, Russell RA (2004) Pheromone communicate: implementation of necrophoric bee behaviour in a robot swarm. In: Proceedings of the IEEE conference on robotics, automation and mechatronics. IEEE Press, Piscataway, NJ, pp 638–643
- Russell RA (1995) Laying and sensing odor markings as a strategy for assisting mobile robots navigation tasks. *IEEE Robot Autom Mag* 2(3):3–9
- Russell RA (1999) Ant trails—an example for robots to follow? In: Proceedings of the IEEE International conference on robotics and automation. IEEE Press, Piscataway, NJ, pp 2698–2703
- Schmickl T, Crailsheim K (2008) Trophallaxis within a robotic swarm: Bio-inspired communication among robots in a swarm. *Auton Robots* 25(1):171–188
- Shell DA, Mataric MJ (2006) On foraging strategies for large-scale multi-robot systems. In: Proceeding of the IEEE/RSJ international conference on intelligent robots and systems. IEEE Press, Piscataway, NJ, pp 2717–2723
- Silverman MC, Nies D, Jung B, Sukhatme GS (2002) Staying alive: a docking station for autonomous robot recharging. In: Proceedings of the IEEE international conference on robotics and automation. IEEE Press, Piscataway, NJ, pp 1050–1055
- Stella E, Musio F, Vasaneli L, Distanto A (1995) Goal-oriented mobile robot navigation using an odour sensor. In: Proceedings of the Intelligent vehicles symposium. IEEE Press, Piscataway, NJ, pp 147–151
- Sugawara K, Kazama T, Watanabe T (2004) Foraging behavior of interacting robots with virtual pheromone. In: Proceedings of IEEE/RSJ international conference on intelligent robots and systems. IEEE Press, Piscataway, NJ, pp 3074–3079

- Svennebring J, Koenig S (2003) Trail-laying robots for robust terrain coverage. In: Proceedings of the IEEE international conference on robotics and automation. IEEE Press, Piscataway, NJ, pp 75–82
- Szymanski M, Breitling T, Seyfried J, Wörn H (2006) Distributed shortest-path finding by a micro-robot swarm. In: Proceedings of the 5th international workshop on ant colony optimization and swarm intelligence. Lecture Notes in Computer Science, vol 4150. Springer, Berlin, pp 404–411
- Wawerla J, Vaughan RT (2008) Optimal robot recharging strategies for time discounted labour. In: Proceedings of the 11th international conference on artificial life. MIT Press, Cambridge, MA, pp 670–677
- Werger B, Mataric MJ (1996) Robotic food chains: externalization of state and program for minimal-agent foraging. In: Proceedings of the 4th international conference on simulation of adaptive behavior. MIT Press/Bradford Books, Cambridge, MA, pp 625–634
- Zebrowski P, Vaughan RT (2005) Recharging robot teams: a tanker approach. In: Proceedings of the 12th international conference on advanced robotics. IEEE Press, Piscataway, NJ, pp 803–810