Foraging-inspired Self-organisation for Terrain Exploration with Failure-prone Agents

Arles Rodríguez, Jonatan Gómez
ALIFE Research Group
Universidad Nacional de Colombia
Bogotá, Colombia
aerodriguezp@unal.edu.co, jgomezpe@unal.edu.co

Ada Diaconescu
Telecom ParisTech
LTCI CNRS
Paris, France
ada.diaconescu@telecom-paristech.fr

***Abstract*—Mobile ad-hoc sensor systems are employed increasingly for distributed tasks in unreliable conditions, such as terrain exploration and measuring. Here, self-organising solutions can help ensure reliability, availability and scalability, while making use of unreliable components (or agents) with limited resources. These enable agents to act independently, and to exchange and combine their partial solutions into a (more) complete result, which can be transmitted to users before all agents fail. In previous work, we have studied how foraging-inspired self-organisation can help mobile agents achieve a collaborative task – terrain exploration and information gathering. Obtained results revealed two key aspects impacting success rates: the strategy for exploring as much uncharted terrain as fast as possible, and the strategy for encountering the highest number of agents that hold complementary information. In this paper, we explore further techniques for these aspects and introduce passive pheromone evaporation. Our results show that a hybrid approach improving exploration efficiency features higher success rates than basic stigmergy models, random exploration and Lévy walks.**

***Keywords*-self-organising systems, information gathering and sharing, terrain exploration, foraging, mobile sensors, multi-agent systems, decentralized control, reliability, failures.**

I. Introduction

An important challenge in mobile sensor systems is to develop algorithms that allow components (e.g. processes, agents, robots or nodes) to collect and aggregate information by using range-limited communication and focusing on mobility [1]. To tackle this challenge, some constraints must be considered: each component has partial knowledge of the environment, acts locally and may fail. Coordination and cooperation between components are essential for reaching global objectives while the way in which components communicate and move are key to the system's success [2]. Scalability is also required as systems must support increasing component numbers, which must take local decisions and communicate in a peer-to-peer fashion as broadcast is not available [3]. Application areas of interest include: communication coverage by mobile robots; robot-based terrain exploration, such as for altitude charting; decentralised computations in sensor networks, such as average functions for measurements; or, automatic record replication for data synchronisation and self-healing in distributed environments, such as computer clusters or clouds. Each of these application domains features different resource constraints, failure risks, and performance targets.

Communication coverage with mobile robots has been studied via a combination of sensing and locomotion, so as to form a network over a targeted space [1]. Some approaches were proposed based on random motions and initial robot locations within a terrain [4], [1]. Robot foraging and communication coverage share some of the same principles: make agents move towards unexplored areas and avoid other agents in order to maximize the explored space; distribute sensor nodes throughout the terrain; and, define an exploration strategy from local perceptions [5]. Some exploration approaches are inspired from foraging in animals, like fruit flies or spider monkeys where patterns are super-diffusive and appear to obey a search for resources that matches a Lévy walk pattern [6], [7], [26]. Lévy walks were also applied to model opportunistic mobile networks assuming that humans move following this model [8], [9]. Additionally, stigmergy was adopted for inter-robot communication via the environment, requiring only that a robot passes close enough to a location where 'communication' was placed to be affected by it [10], [11]. There are two main ways to define pheromones in robots: *simulated*, in which the environment is *not* modified, and it is mapped instead as a logical grid with different pheromone concentrations (robots share references of places, or use a map defined as a shared memory) [10], [12], [13]; or *situated*, in which the environment *is* modified and robots use specialized sensors to detect this – e.g. chemical sensors for alcohol-based trails [14], infrared sensors [15], or light sensors for pheromones implemented as light projections onto the ground [16].

Exploration with unreliable agents was also proposed in [17], by using a trust measure between agents that sense the world in a defective way. The approach uses a motion strategy that is either based on agent reliability, or on a random motion in case agents do not have sufficient information. Similar to our work, [17] also established a difference between direct information (collected by each agent) and indirect information (obtained from others) and aimed to achieve information coherence across agents as a result. However, rather than dealing with unreliable information, this paper aims to evaluate common exploration techniques, identify the causes behind their performance differences, and explore solutions that combine their most favourable characteristics.

In previous work, we have studied how basic exploration techniques and foraging-inspired self-organisation help mobile agents achieve such collaborative tasks [18]. Obtained results allowed us to identify two key aspects impacting agent performance and success rates: *exploration of new paths*, which addresses the problem of covering unexplored territory; and, *finding and exchanging* with agents that detain new information, which speeds-up global data collection by merging results of local explorations. Results also showed that increasing agent population *densities* produced better results, even when relying on random walks.

In this paper, we explore some alternatives for the two aspects above, their most promising combinations and the addition of passive pheromone evaporation. We focus on decentralised techniques inspired from foraging behaviours in animals: Random motions and Lévy walks (for the exploration aspect) and, stigmergy-based strategies (for the agent search and exchange aspect, and for avoiding explored territory). The main contributions of this paper consist in proposing various exploration approaches and analysing them with respect to speed, failure resistance, message exchanges and scalability, which impact their applicability to targeted domains. Speed relates to the rapidity of data collection; failure resistance to the ability of covering the entire terrain in the presence of increasing agent failures; the number of messages relates to communication overheads, to be minimized in sensor networks for improving battery life [19]; and, scalability shows the impact of terrain sizes on the other performance results (at the same agent density – $density = number\_of\_agents/world\_size$).

The remaining of this paper is organized as follows. Section II defines the problem, Section III provides some basic exploration solutions and Section IV discusses experiments and results for these solutions. This allows to identify the best candidates to consider for more advanced hybrid solutions. Section V explores such hybrid alternatives and also considers the impact of pheromone evaporation. Results are analysed in Section VI, discussed in Section VII and conclusions drawn in Section VIII.

II. Problem Definition

Each agent is designed to explore a simulated world, to sense some desired data throughout the environment, and to share its collected data with agents that it encounters. Agents fail with a certain probability and the objective is to allow at least one agent to collect all the information before failing. The simulated world is a bi-dimensional non-toroidal space defined as a matrix of properties $Props^{width \times height}$. $Props$ is a collection $Props = \{\tau_w, data\}$, with τ_w the amount of pheromone in each world location (x, y): $\tau_w(x, y) \in \mathbb{R} \wedge \tau_w(x, y) \in [0, 1]$; and $data$ the information of interest in this world, like altitudes or distances to a target, with values that do not change: $data \in \mathbb{R} \wedge data \in [0, 1]$.

Agents are endowed with different perception capabilities: *Percept* $= \{$*pheromone, data, socialStatus, neighbour, msg, loc, proximity_sensor*$\}$, where *pheromone* is a vector $\mathbb{R}^n$ with values in $[0, 1]$ representing the amount of pheromone in an agent's vicinity (Moore neighbourhood with $r = 1$ and centre in the agent's location [23]); *data* is the information of interest in the agent's current location; *socialStatus*$= \{$*seeker*, *carrier*$\}$ indicates the agent's role; *neighbour* returns the id of an adjacent agent (randomly selected from its Moore neighbourhood); *msg* stores messages received from adjacent agents; *loc* returns the agent's location as $(row,\ column)$; and, *proximity_sensor* indicates the presence of obstacles (bordering walls in the experiments).

An agent can perform the following actions *Actions*$=$ $\{$*none, down, left, right, up, upleft, upright, downright, downleft, Fail, Collect, Send, Recv*$\}$. The first nine actions represent possible movements in a direction *dir* of the bi-dimensional world. *Fail* stops the agent's thread to simulate failure. *Collect* senses *data* from the agent's current location and stores it in its local memory. *Send* and *Recv* enable information exchanges as shown above.

Figure 1 depicts an example of stigmergy-based agent exploration in a simulated world. Agents start from random locations and communicate via *Trophallaxis*: when two agents occupy adjacent locations, they exchange their information (Green circles in Fig. 1). Like in nature, a trophallactic cascade scheme is formed which may be more efficient than direct transfer via the environment [20]. Inter-agent communication is implemented so that whenever an agent s senses a neighbour r, it gets its id and exchanges its local information I_s with r's information I_r via $Send$ and $Receive$ primitives: agent s performs $Send(r, msg)$ where $msg = \{I_s\}$. Hence, agent s sends its current information I_s to r. Then, agent r performs $Recv(msg)$, where $msg = \{I_s\}$, and completes its information $I_r = I_r \cup I_s$. Similarly, agent r also detects and sends its information to agent s, which updates its local information $I_s = I_s \cup I_r$. This exchange is inspired by traditional Asynchronous Distributed Systems with FIFO communication channels, where messages received first are processed first by each agent [2], [21], [22].

Algorithm 1 provides an example of agent program. Computation is executed in rounds, where in each round the agent senses its environment, computes an action and effects it onto its environment. This cycle is repeated until the agent completes the exploration or fails [24], [25]. Failure is triggered with a certain probability (p_f), defined for all agents. For instance, $p_f = 0.1$ means that each agent has a probability of failure of 1 in 10 rounds.

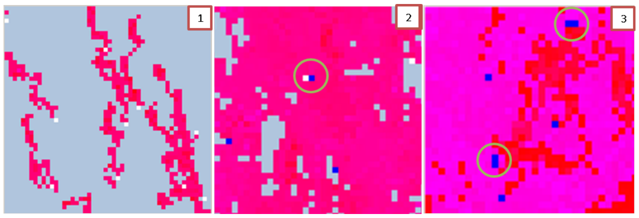

Figure 1. World exploration using stigmergy with 10 agents: white points are *seekers* and blue points are *carriers*; unexplored terrain is initially light blue; pheromone traces are red and fade into blue as concentrations decrease.

Algorithm 1 Example of agent program

Percept p
Action action
round $\leftarrow 0$
while $Agent.status \neq Fail$ **do**
  $\lambda \leftarrow U[0,1]$ ▷ uniform random number
  if $\lambda < p_f$ **then**
    Agent.status $\leftarrow Fail$
    break
  end if
  p $\leftarrow$ environment.sense()
  Action $\leftarrow$ computeAction(p)
  environment.act(Agent, Action)
  round $\leftarrow$ round+1
end while

III. Basic Approaches

A. Overview

The agents' success in completing their collective task relies critically on each agent's motion process. The motion strategy impacts both the exploration of uncharted terrain and the encounter of other agents – both essential for data collection. In this section, four motion processes are defined based on two main criteria (Fig 2): two purely exploratory processes – random and Lévy walks – and two stigmergy-based processes that use pheromones to improve collective exploration and agent encounters. The input of all algorithms is the agent's perception *Percept* and the output is the agent's direction of movement *dir* (Cf. Section II).

When following a *random walk*, at each round, an agent chooses its movement direction randomly from the set of possible directions. The following subsections describe the other motion mechanisms (summarised in Fig. 2).

B. Lévy Walks

In this approach, exploration starts in a random direction, which is then maintained during a random time that is inversely proportional to the distance travelled. This process has been used for dispersion and mixing of swarms in robotic communication coverage [6]. In this paper we adapt the motion mechanism in [6] as shown in Algorithm 2: *randomDir()* returns a random direction *dir*, α is a uniform random number that represents the increment rate of an

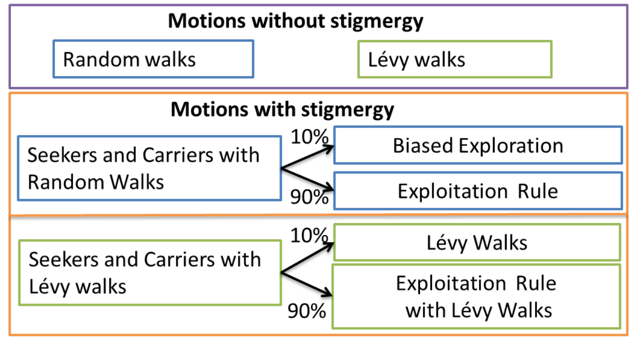

Figure 2. Overview of Basic Solution Approaches

accumulator and T is the accumulator's threshold ($T = 1$ in the tests). The 'repeat' section is executed at each round, moving the agent in the current direction dir, until T is crossed, an obstacle is encountered (as detected by *neighbour_sensor()* or *proximity_sensor()*) or agent fails.

Algorithm 2 Reactive Lévy walk for an agent [6]

while $Agent.status \neq Fail$ **do**
  dir $\leftarrow$ randomDir()
  accumulator $\leftarrow 0$
  $\alpha \leftarrow U[0,1]$ ▷ uniform random number
  repeat ▷ repeat is executed at each round
    move(dir)
    acumulator $\leftarrow$ acumulator $+\alpha$
  until $accumulator \geq T \vee neighbour\_sensor() \vee proximity\_sensor() \vee Agent.status = Fail$
end while

C. Seekers and Carriers with Random Walks

This approach is based on the Ant Colony System algorithm (ACS) [27], yet is adapted to use stigmergy for enabling agents to avoid explored paths and to find other agents (instead of looking for food points as in the original version). Each agent i has a socialStatus determined by its collected amount of information: *seeker*, if i is looking for other agents and hence explores locations with more pheromone, and *carrier* if i is looking for uncharted terrain and hence pursues locations with less pheromone. Initially, all locations have a pheromone value $\tau_w = 0.5$. More details on this algorithm are presented in [18]. As in ACS [27], a random variable $q \in [0,1]$ dictates when to apply an *exploitation rule* or *biased exploration* (Eq. 1):

$$dir = \begin{cases} \text{exploitation rule} & \text{if } q \leq 0.9 \\ \text{biased exploration} & \text{otherwise} \end{cases} \quad (1)$$

The *exploitation rule* determines a direction depending on the agent's socialStatus. A seeker will choose the direction with the *maximum* amount of pheromone in its vicinity, looking for carriers; and a carrier will choose the direction with the *minimum* amount of pheromone in its vicinity, looking for uncharted terrain. If more than one direction has the same *minimum* or *maximum* value, a random direction is picked.

Biased exploration is a random-proportional rule [27] which gives an agent i a probability of choosing a direction $p_d(x,y)$ depending on the amount of pheromone τ_w in its vicinity $neighbourhood(i)$ (Eq. 2). $neighbourhood(i)$ includes the locations in the Moore neighbourhood of i with $r = 1$. This prevents agents from getting trapped in a confined area (e.g. carriers surrounded by pheromone traces). For seekers $\tau'_w(x,y) = \tau_w(x,y)$ and for carriers $\tau'_w(x,y) = 1 - \tau_w(x,y)$.

$$p_d(x,y) = \frac{\tau'_w(x,y)}{\sum_{(k,l)\in \text{neighbourhood(i)}} \tau_w(k,l)} \tag{2}$$

Whenever an agent i moves, at each round t, it updates its internal pheromone value $\tau_{a_t}(i)$ (as in Eq 3) and the pheromone amount in its location $\tau_w(x,y)$ (as in Eq 4).

$$\tau_{a_t}(i) = (\tau_{a_{t-1}}(i) + 0.01 * (0.5 - \tau_{a_{t-1}}(i))) \tag{3}$$

$$\tau_{w_t}(x,y) = \tau_{w_{t-1}}(x,y) + 0.01*(\tau_{a_{t-1}}(i) - \tau_{w_{t-1}}(x,y)) \tag{4}$$

Equation 3 is based on the local update rule of ACS [27]. If an agent i is, or turns into, a seeker and finds or receives new information, then its pheromone value is updated to $\tau_{a_t}(i) = 0$. Based on Eq. 3 this value increases at each round until a certain point. Based on Eq. 4, this value reduces the pheromone amount in explored locations by seekers.

If an agent i is, or becomes, a carrier and finds or receives new information, then its internal pheromone value is updated to $\tau_{a_t}(i) = 1$. In this case, Eq. 3 decreases this pheromone value at each round; and Eq. 4 increases the amount of pheromone in the locations explored by carriers.

When information is exchanged, agent r that receives information I_s from agent s calculates the difference $dif = I_s \setminus I_r$. Two scenarios can occur:

- if $dif = \emptyset$, meaning that r had at least the same information as s, then agent r (re)turns into a carrier and sets its pheromone value to one $\tau_{a_t}(r) = 1$;
- if $dif \neq \emptyset$, then r (re)turns into a seeker and sets its pheromone value to zero $\tau_{a_t}(r) = 0$.

The aim is to have some agents (carriers) focused on exploration and further data collection, and the others (seekers) focus on finding carrier agents and getting their data. This should help distribute data collected among agents, before the ones having more data fail. The criterion for designating agent roles is based on the relative data-collection success at the time of each encounter. Additionally, each time an agent i receives new information from neighbours or its environment, it resets its internal pheromone value depending of its *socialStatus* (seekers $\tau_{a_t}(i) = 0$, carriers $\tau_{a_t}(i) = 1$) to stimulate trails that produce new information. This mechanism of pheromone resetting is defined for all the stigmergy-based approaches in this paper.

D. Seekers and Carriers with Lévy Walks

In previous experiments presented in [18], the average amount of information collected by agents was higher for Lévy walks compared to random walks and to seekers and carriers exploration. Lévy walks represent a super-diffusive pattern that appears advantageous to exploration. Therefore, we assumed that the initial seekers and carriers approach could be improved by replacing its random exploration parts by Lévy walks. The new approach is the same for selecting between an exploitation rule and a biased exploration rule (Eq. 1). Yet, biased exploration is replaced from the random-proportional rule of Eq. 2 to Lévy walks (Alg. 2) as shown in Eq 5. Also, the exploitation rule is updated so that if more than one direction contains the same maximum (for seekers) or minimum (for carriers) pheromone amount, and if the last direction given by the Lévy walk algorithm is within these options, then this direction is chosen (otherwise a random one is used as before).

$$dir = \begin{cases} \text{exploitation rule with Lévy walk} & \text{if } q \leq 0.9 \\ \text{Lévy walk} & \text{otherwise} \end{cases} \tag{5}$$

IV. Experiments and Results

Experiments for each basic approach in Section III were performed considering different failure probabilities $pf = 0$, 1×10^{-4}, 3×10^{-4}, 5×10^{-4}, 7×10^{-4} and 9×10^{-4}. These p_f values were selected to explore the range in-between no failures and failure rates where all algorithms fail. Additionally, different world sizes were considered while maintaining the same agent density – i.e. agent populations of 10, 20, 30, 40 and 50 for world sizes of 50×50, 71×71, 87×87, 100×100, and 112×112 , respectively. An agent's size is 1×1. Each experiment stops when an agent collects all the information or if all the agents fail. Presented results are averages from experiments performed 30 times.

Table I presents the success rates of each algorithm (averaged over the 30 executions). An experiment is successful if an agent manages to collect all the information. Random exploration performs well for low failure rates (until $pf = 10^{-4}$), yet declines quickly as failure rates increase. Better success rates are observed for Lévy walks and for both seekers and carriers approaches, with successful experiments occurring for 10 and 20 agents up to $pf = 5 \times 10^{-4}$ and with higher success rates for $pf = 10^{-4}$ compared to random walks. It can also be noted that among these three algorithms, for $pf = 3 \times 10^{-4}$ and 5×10^{-4}, the best success rate varies with the world scale. Also, the overall success of all algorithms declines as the scale increases. This may be due to the initial agent locations within the world, set randomly; further experiments are required in future work to establish this. Figure 3 presents a box-plot with the *collected information percentages* averaged across all agents, at the end of each experiment. This represents

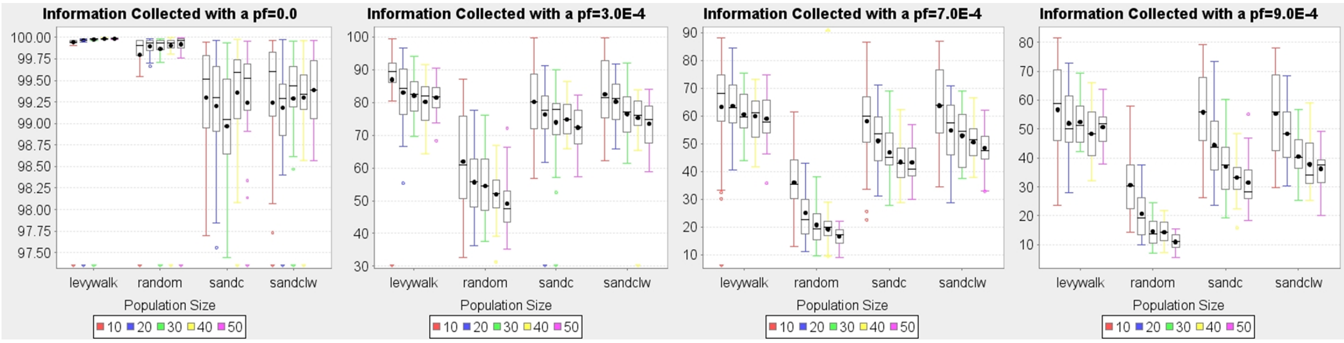

Figure 3. Box-plot of *average information collected* by agents for random motions (random), Lévy walks (levywalk), Seekers and Carriers (sandc) and Seekers and Carriers with Lévy walks (sandclw)

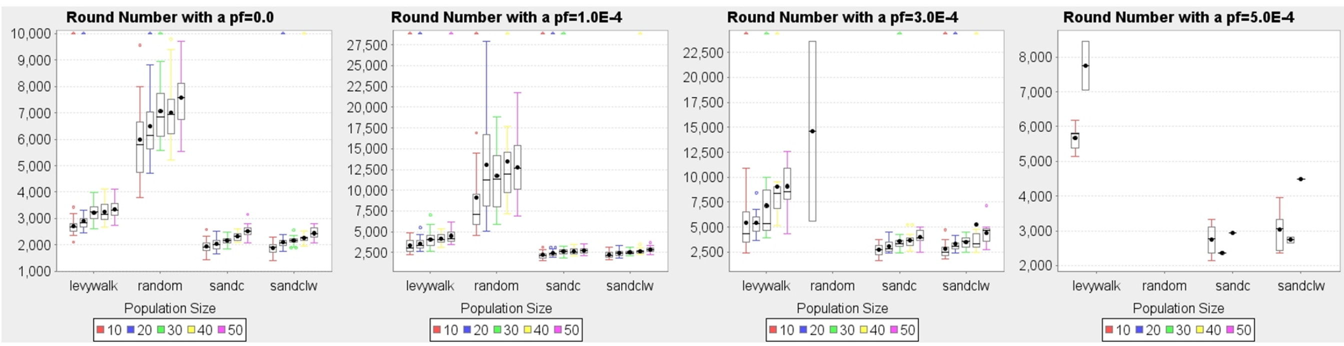

Figure 4. Box-Plot of *round number* for random motions (random), Lévy walks (levywalk), Seekers and Carriers (sandc) and Seekers and Carriers with Lévy walks (sandclw)

Population (World size)	p_f	Success Rate			
		Random	Lévy Walks	Seeks and Carriers	SandC With Lw
10 (50×50)	9.0E-4	0%	0%	3.33%	0%
	7.0E-4	0%	0%	0%	0%
	5.0E-4	0%	16.67%	13.33%	20%
	3.0E-4	6.67%	83.33%	53.33%	58.97%
	1.0E-4	86.67%	100%	96.67%	100%
	0.0	100%	100%	100%	100%
20 (71×71)	9.0E-4	0%	0%	0%	0%
	7.0E-4	0%	0%	0%	0%
	5.0E-4	0%	6.67%	3.33%	6.67%
	3.0E-4	0%	50%	50%	57.89%
	1.0E-4	93.33%	100%	100%	100%
	0.0	100%	100%	100%	100%
30 (87×87)	9.0E-4	0%	0%	0%	0%
	7.0E-4	0%	0%	0%	0%
	5.0E-4	0%	0%	3.33%	3.33%
	3.0E-4	0%	36.67%	43.33%	39.47%
	1.0E-4	83.33%	100%	100%	100%
	0.0	100%	100%	100%	100%
40 (100×100)	9.0E-4	0%	0%	0%	0%
	7.0E-4	0%	0%	0%	0%
	5.0E-4	0%	0%	0%	0%
	3.0E-4	0%	23.33%	43.33%	31.58%
	1.0E-4	73.33%	100%	100%	100%
	0.0	100%	100%	100%	100%
50 (112×112)	9.0E-4	0%	0%	0%	0%
	7.0E-4	0%	0%	0%	0%
	5.0E-4	0%	0%	0%	0%
	3.0E-4	0%	56.67%	26.67%	36.84%
	1.0E-4	80%	100%	100%	100%
	0.0	100%	100%	100%	100%

Table I
INFORMATION COLLECTED AND SUCCESS RATES

the data distribution degree among agents and can provide an indicator of behaviour when failures increase (as the most successful agents may fail before completion). When introducing failures, performance degrades, unsurprisingly, as p_f increases. We also show different agent populations for each algorithm to assess scalability. In the absence of failures, all agents in the four techniques manage to collect almost 100% of the data by the end of the experiments. Lévy walks perform particularly well, since they promote exploration equally for all agents. When p_f is increased, data collected by agents with random motions decreases faster than in the other methods. For the three other techniques (also with the best success rates in Table I), data collected by *seekers and carriers with Lévy walks* is slightly higher compared to *seekers and carriers with random walks*.

Figure 4 presents the average *round number* of the agent(s) that complete data collection in the successful experiments for more than one technique (with $pf = 0$, 10^{-4}, 3×10^{-4} and 5×10^{-4}). It is observed that higher p_f values cause an increase in the round number necessary to gather the complete information. When an algorithm failed to complete for a p_f value across all experiments, its round number was not represented in the graph. In all cases, there is little difference between the round numbers achieved by the two *seekers and carriers* approaches. Yet, both these approaches finish the task quite faster than Lévy walks (e.g. 1000 rounds faster for $pf = 0$ and 10^{-4}). This may be due to the fact that seekers and carries approaches promote agent encounters and exchanges. When p_f increases, this exchange aspect seems to outrun the average agent information collection (as this latter aspect was slightly better for Lévy walks, Cf. Fig. 3). Round numbers also increase with the world size; to be explored in future work.

Figure 5 presents a box-plot of *messages sent* for the different techniques. It can be observed than a higher amount of messages is sent for random motions without failures ($pf = 0$). This may be explained by the fact that random motion only promotes exploration around the initial location [6], so some agents keep exploring the same locations and communicate among themselves, exchanging the same information repeatedly. The Lévy walks technique exchanged the smallest number of messages, as it focused on exploration, but took a higher number of rounds in gathering the information (than seekers and carriers). In each technique, the number of messages decreased with the p_f increase, as there were less agents to exchange with. The number of messages also increased with the world scales.

In terms of speed and success rates, *seekers and carriers with Lévy walks* appear better than the other techniques. However, *Lévy walks* feature higher amounts of collected information across agents and smaller message numbers compared to the two seekers and carriers methods. In the following section we aim to take the best features of each technique and to add passive pheromone evaporation, in order to improve the exploration capabilities of the agents.

V. Hybrid Approaches and Evaporation

A. Lévy Walkers and Carriers

This approach redefines the seekers and carriers technique by replacing seekers with Lévy walks, as follows: if an agent i is a *carrier*, its behaviour is the same as defined for seekers and carriers (subsection III-C). Hence, carriers will increase pheromone amounts in visited locations and choose paths with minimum pheromone amounts to avoid visited places. Motion is determined by the rules in Eq. 1 and Eq. 2. Pheromone is deposited as defined in Eq. 3 and Eq. 4. If an agent i is a *seeker*, its motion process is a Lévy walk (Alg. 2) and it does not deposit pheromone.

Since only one agent type (carriers) impact pheromone amounts, the initial amounts in all world locations is set to $\tau_w = 0$. The aim of this hybrid approach is to combine the best information collection techniques: Lévy walks by enhanced exploration, and carriers by avoiding already explored territory. Initially, all agents are Levy walkers. As they encounter other agents, the ones that collected more information become carriers. This should allow them to explore more uncharted territory by avoiding pheromone locations. Yet, if pheromone-based exploration does not produce more information, then carriers swap back into Lévy walkers.

B. Lévy Walkers and Carriers with Lévy Walks

This approach is similar to the previous one (*Lévy walkers and carriers*) except that for the carriers the rules in Eq. 1 are replaced by Eq. 5. This means that a *carrier* will choose the direction with the minimum amount of pheromone, but if several such directions exist, and one of them is the same as the current output of the Lévy walk algorithm, then this direction will be selected (if it does not match the direction of the Lévy walk, a random direction is chosen among the minimum ones). Their biased exploration is also a Lévy walk, rather than random. Carriers deposit pheromone as defined by Eq. 3 and Eq. 4 and the initial pheromone amounts for all world locations is set to ($\tau_w = 0$). Seekers are replaced by Lévy walkers (Alg. 2) and they do not deposit pheromone.

C. Carriers and Pheromone Evaporation

This approach considers having only carriers exploring the world. It was observed in previous experiments that exploration with pheromone is fast but some agents collect more information than others (Cf. Fig 3). Moreover, it is possible that the agents that collected the most information fail before completing the task, and their pheromone traces are left behind preventing other agents from recollecting the information. A solution to this issue is to also enable passive pheromone evaporation (by the environment, rather than by agents). In this approach, each agent determines its movements by using stigmergy inspired by the Ant Colony Optimization algorithm (ACO) [28] and the exploration and exploitation rules of ACS. Each agent explores and updates local pheromone just like a *carrier* (Eq 1). There are *no* socialStatus changes in this approach. If an agent finds new information from neighbours or from new locations, the pheromone of the agent is set to $\tau_t(i) = 1$ (as for carriers in the aforementioned approaches). In the same way, all the world locations are initialised to $\tau_w = 0$.

Evaporation is applied to all world locations, independently of agent movements, by using the definition in [28]. Namely, it is defined via (Eq. 6), with evaporation rate $\rho = 0.01$:

$$\tau_w(x,y) = (1-\rho)\tau_{w_{t-1}}(x,y), \text{ for } \forall(x,y) \in \{\text{world locations}\} \quad (6)$$

As [29] indicates, natural pheromone evaporation makes paths less dominant, rewarding exploration while also addressing cases where path-producing agents fail.

D. Carriers with Lévy Walks and Pheromone Evaporation

This model is similar to the previous one (*carriers and pheromone evaporation*), except that the rules in Eq. 1 are replaced by Eq. 5.

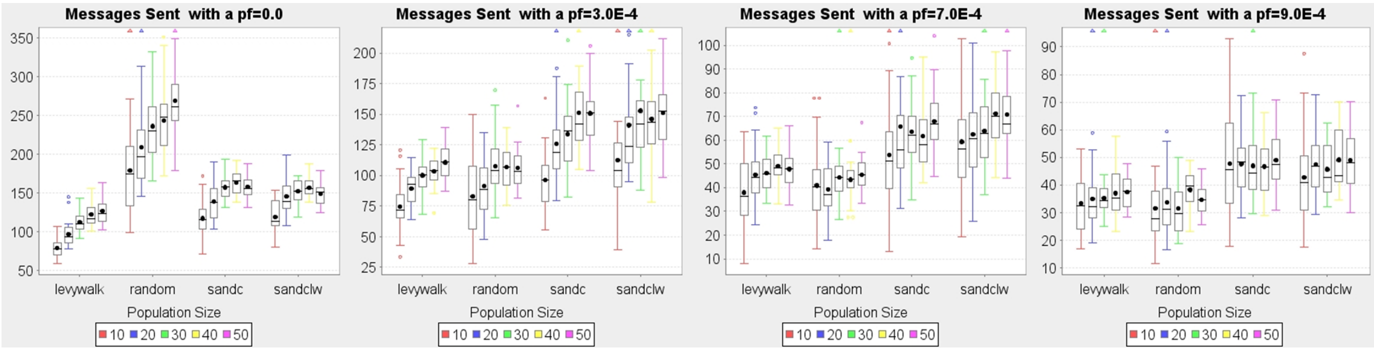

Figure 5. Box-Plot of *messages sent* for random motions (random), Lévy walks (levywalk), Seekers and Carriers (sandc) and Seekers and Carriers with Lévy Walks (sandclw)

VI. Experiments and Results with Hybrid Approaches and Evaporation

Several experiments were performed on these hybrid models, in the same way as for the basic models (Cf. Section IV): $pf = 0$, 10^{-4}, 3×10^{-4}, 5×10^{-4}, 7×10^{-4} and 9×10^{-4}; agent populations of 10, 20, 30, 40 and 50 and world sizes of 50×50, 71×71, 87×87, 100×100, and 112×112 respectively. Each experiment was performed 30 times and stopped when one agent collected all the information or if all the agents failed. The values of p_f were selected because most of our experiments were successful for $p_f = 10^{-4}$ and failed for 10^{-3}; so p_f values between these two extremes were the most relevant to explore.

As can be observed in Table II, success rates are even better for the hybrid algorithms than for *seekers and carriers with Lévy walks*, which was the best option from the basic approaches (Cf. column "SandC with LW" in both Tables I and II). This indicates that using seekers to increase agent encounters did not improve success rates. Success rate results also indicate that passive pheromone evaporation plays an important role in failure resistance, since both approaches using it were better than the others. Namely, *Carriers and Pheromone Evaporation* and *Carriers with Lévy walks and Pheromone Evaporation* achieve some successes for all p_f values, even for $pf = 9 \times 10^{-4}$, for 10, 20 and 30 agents; improved success rates were also observed for the other population sizes compared to the other approaches.

Figure 6 depicts the average data collected by agents at the end of each experiment. When agents do not fail, they seem to collect more data when using carriers and Lévy walks (the two approaches without evaporation) than when evaporation is used. However, this advantage diminishes as the failure rate p_f increases, with all hybrid approaches performing about the same by the time p_f reaches 9×10^{-4}. Moreover, Figure 6 shows that it could be beneficial to replace the random-probabilistic rule by Lévy walks, since *Lévy Walkers and Carriers with Lévy Walks* collects slightly more information than *Lévy Walkers and Carriers* and has somewhat higher success rates (except for 10 and 20 agents with $pf = 5 \times 10^{-4}$; more study will be performed here in future work).

Pop	p_f	Success Rates				
		SandC with Lw	Lw and C	Lw and C-Lw	C and Evap	C-Lw and Evap
10	9.0E-4	0%	6.67%	0%	40%	26.67%
	7.0E-4	0%	3.33%	16.67%	50%	53.33%
	5.0E-4	20%	46.88%	40%	86.67%	90%
	3.0E-4	58.97%	70%	83.33%	100%	96.67%
	1.0E-4	100%	96.67%	100%	100%	100%
	0.0	100%	100%	100%	100%	100%
20	9.0E-4	0%	0%	0%	10%	20%
	7.0E-4	0%	3.33%	3.33%	40%	40%
	5.0E-4	6.67%	40.63%	36.67%	83.33%	83.33%
	3.0E-4	57.89%	80%	83.33%	100%	100%
	1.0E-4	100%	100%	100%	100%	100%
	0.0	100%	100%	100%	100%	100%
30	9.0E-4	0%	0%	0%	3.33%	10%
	7.0E-4	0%	3.33%	0%	16.67%	46.67%
	5.0E-4	3.33%	18.75%	20%	76.67%	90%
	3.0E-4	39.47%	73.33%	73.33%	100%	100%
	1.0E-4	100%	100%	100%	100%	100%
	0.0	100%	100%	100%	100%	100%
40	9.0E-4	0%	0%	0%	0%	6.67%
	7.0E-4	0%	0%	0%	10%	23.33%
	5.0E-4	0%	6.25%	33.33%	63.33%	93.33%
	3.0E-4	31.58%	90%	96.67%	100%	100%
	1.0E-4	100%	100%	100%	100%	100%
	0.0	100%	100%	100%	100%	100%
50	9.0E-4	0%	0%	0%	0%	0%
	7.0E-4	0%	0%	0%	26.67%	16.67%
	5.0E-4	0%	3.23%	6.67%	86.67%	86.67%
	3.0E-4	36.84%	83.33%	86.67%	100%	100%
	1.0E-4	100%	100%	100%	100%	100%
	0.0	100%	100%	100%	100%	100%

Table II
Success rates for Seekers and Carriers with Lévy Walks (SandC with Lw), Lévy Walkers and Carriers (Lw and C), Lévy Walkers and Carriers with Lévy Walks (Lw and C-Lw), Carriers and Pheromone Evaporation (C and Evap) and, Carriers with Lévy Walks and Pheromone Evaporation (C-Lw and Evap)

Figure 7 presents the *round number* for the best agent(s) and it is observed that hybrid methods are faster than the basic methods in Section III. It can also be noted how fast information gathering can be linked to higher success rates. Here again, the two approaches using pheromone evapora-

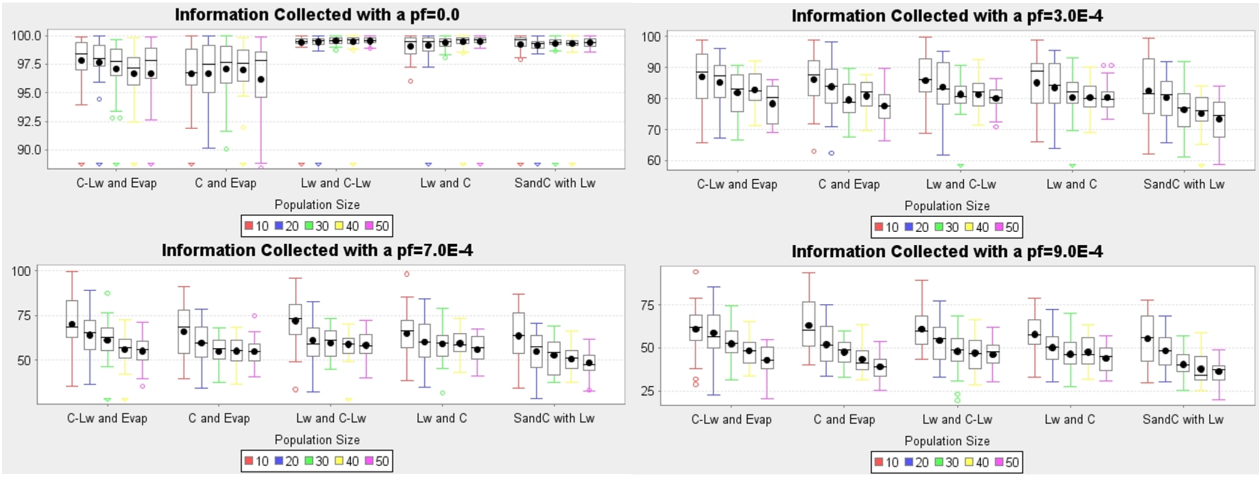

Figure 6. Box-plot of information collected for Seekers and Carriers with Lévy Walks (SandC with Lw), Lévy Walkers and Carriers (Lw and C), Lévy Walkers and Carriers with Lévy Walks (Lw and C. w Lw), Carriers and Pheromone Evaporation (C and Evap) and, Carriers with Lévy Walks and Pheromone Evaporation (C-LW and Evap)

tion are faster than the others (with *Carriers with Lévy walks and Pheromone Evaporation* appearing the fastest in most cases). It can also be noted how some techniques no longer occur in the box-plot for higher p_f because they do not manage to get any success before all agents fail.

Figure 8 presents the average number of messages sent. When agents do *not* fail, the hybrid approaches exchange less messages than the basic ones. As p_f increases, messages sent by all approaches decrease (since agents fail), and the difference between approaches also diminishes. For $pf = 7$ and 9×10^{-4} the two approaches using evaporation are the ones sending most messages. This can also account for their lower round numbers and higher success rates.

VII. Discussion

The above experiments aimed to determine the impact of two main factors on the performance metrics (speed, failure resistance, messages sent and scalability), as compared to the basic approaches. These changes were compared to the best of the basic approaches (*seekers and carriers with Lévy walks*).

The first factor consisted in replacing seekers with Lévy walkers (V-A and V-B). The aim was to favour exploration rather than agent encounters, considering that agents will still meet (e.g. as they aim to explore the same uncharted areas). On the contrary, seekers may be inefficient in data collection since they follow already explored paths in search of carrier agents. Hence, agent encounters are more productive between exploring agents, since the data they exchange is more complementary. Indeed, this technique managed to improve success rates at increasing failure rates (by lowering the round number), while maintaining the average data collected about the same (compared to *seekers and carriers with Lévy walks*).

The second factor added passive pheromone evaporation (V-C and V-D). The aim was to address the problem of traces left by failed agents, which prevented other agents from re-exploring them. This technique succeeded in further improving success rates (via even lower round numbers), while average data collection was not significantly impacted. Nonetheless, the applicability of this approach depends on the possibility and cost of implementing this feature in the targeted environment. Also, the approaches using evaporation appear to exchange more messages than when no evaporation is used, which should be taken into account in terms of energy consumption.

The proposed techniques are applicable in different contexts depending on the domain-specific constrains. For instance, if a targeted domain allows for pheromone modelling, either situated or simulated, then hybrid approaches that optimise exploration can be considered – using Lévy walks to improve exploration, pheromone trails to enhance agent information sharing and to avoid already charted territory, and pheromone evaporation to "forget" such trails in case the agents that posed them had failed in the meantime. Otherwise, when pheromone marking and/or evaporation are unavailable, a basic Lévy walk may be the best option. Also, in cases where the agent density is much higher and/or the agent failure rate is much lower, and where the time to completion is not a critical parameter, a basic Lévy walk or even a Random walk may do. Finally, in cases with no pheromone availability and low failure rates, but that require to have many agents collect the data (rather than a single one), the Lévy walks are the best.

More information on the experiments, including complete measurement tables and source code are available at: http://alife.unal.edu.co/~aerodriguezp/termites/.

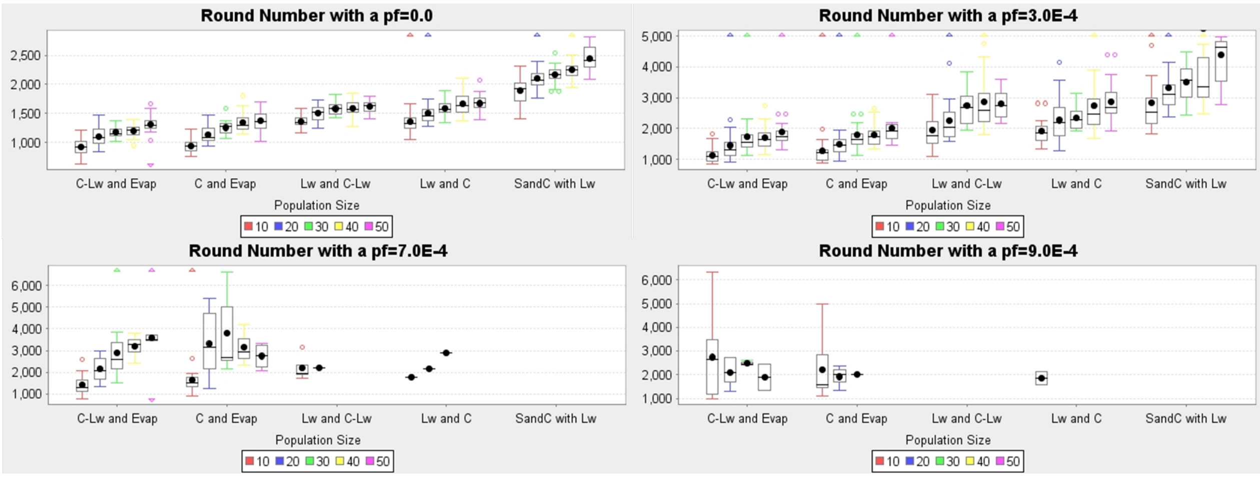

Figure 7. Box-Plot of round number for Seekers and Carriers with Lévy Walks (SandC with Lw), Lévy Walkers and Carriers (Lw and C), Lévy Walkers and Carriers with Lévy Walks (Lw and C. w Lw), Carriers and Pheromone Evaporation (C and Evap) and, Carriers with Lévy Walks and Pheromone Evaporation (C-LW and Evap)

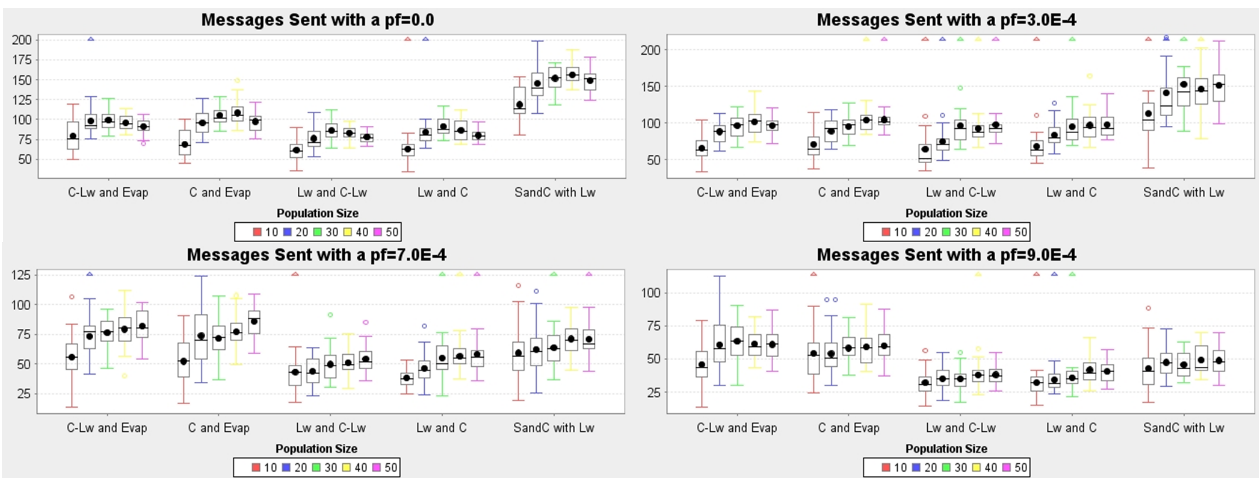

Figure 8. Box-Plot of messages sent for Seekers and Carriers with Lévy Walks (SandC with Lw), Lévy Walkers and Carriers (Lw and C), Lévy Walkers and Carriers with Lévy Walks (Lw and C. w Lw), Carriers and Pheromone Evaporation (C and Evap) and, Carriers with Lévy Walks and Pheromone Evaporation (C-LW and Evap)

VIII. Conclusions

In this paper, we proposed different solutions for the problem of terrain exploration and data gathering via a set of unreliable agents. Since none of the agents can complete the task on its own, because of limited action time before failure, agents must self-organise and share the partial data they collect to obtain the global result faster. Here, completion speed was an essential characteristic for the solution success rates. This, in turn, depended on the agents' ability to explore uncharted territory, and to meet agents with complementary information.

The best exploration characteristics were achieved for solutions following Lévy walks (better than random walks) and using pheromones to mark (and avoid) already charted terrain. Passive pheromone evaporation was highly beneficial for removing traces left by failed agents and allowing paths to be re-explored by the remaining agents. Attempting to introduce agents with seeker behaviours, which purposefully searched for exploratory agents by following their pheromone paths, did not produce any improvements. On the contrary, seeker techniques seemed to only increase encounters between agents with similar information. Instead, allowing agents to meet passively during their explorations produced much better results overall. Hence, hybrid approaches that managed to improve exploration while still managing to encounter agents that held complementary information provided the best results.

For applications in real environments, the number of message exchanges must be considered, since impacting energy consumption and battery life. The approaches with the best success rates also performed well in this respect. Scalability is also essential and results obtained indicated linearity of solutions with increasing terrain sizes at the same population density. Future work will concentrate on the impact of the agents' initial locations and on alternative techniques for seekers, agent role changes and pheromone use.

Experimental results showed that mechanisms that favour exploration are more resistant to failure than those that focus on increasing communication between agents. Hybrid models that combine Lévy walks and adapted Ant Colony algorithms can be promising providing a higher success rate in the presence of failures via faster exploration compared to other algorithms. Further work will explore additional combinations, like Lévy Walkers that leave pheromone traces (like the carriers) for improving exploration, and applying passive evaporation.

Acknowledgments

This work was supported by the *"Convocatoria del Programa Nacional de Proyectos para el Fortalecimiento de la Investigación, la Creación y la Innovación en Posgrados de La Universidad Nacional De Colombia 2013-2015"* code 23418.

References

[1] A. F. T. Winfield and J. Nembrini, "Emergent Swarm Morphology Control of Wireless Networked Mobile Robots," in *Morphogenetic Engineering, Toward Programmable Complex Systems*. Springer, 2012, ch. 10, pp. 239–271.

[2] M. Raynal, *Distributed Algorithms for Message-Passing Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013.

[3] A. Tanenbaum and M. V. Steen, *Distributed systems: principles and paradigms*. Prentice-Hall, 2006.

[4] J. Beal, N. Correll, L. Urbina, and J. Bachrach, "Behavior modes for randomized robotic coverage," *Second International Conference on Robot Communication and Coordination*, 2009.

[5] S. Lopes, B. Frisch, A. Boeing, K. Vinsen, and T. Bräunl, "Autonomous exploration of unknown terrain for groups of mobile robots," *IEEE Intelligent Vehicles Symposium, Proceedings*, no. Iv, pp. 157–162, 2011.

[6] J. Beal, "Superdiffusive dispersion and mixing of swarms with reactive levy walks," *International Conference on Self-Adaptive and Self-Organizing Systems, SASO*, pp. 141–148, 2013.

[7] S. Benhamou, "How Many Animals Really Do the Lévy Walk?" *Ecology*, vol. 88, no. 8, pp. 1962–1969, 2007.

[8] I. Rhee, M. Shin, S. Hong, K. Lee and S. Chong, "On the Levy-Walk Nature of Human Mobility," *IEEE/ACM Transactions on Networking*, vol. 19, no. 3, pp. 630–643, 2011.

[9] B. Saha, S. Misra, and S. Pal, "Utility-based Exploration for Performance Enhancement in Opportunistic Mobile Networks," *IEEE Transactions on Computers*, vol. 9340, no. c, pp. 1–1, 2015.

[10] A. Abraham, C. Grosan, and V. Ramos, *Stigmergic Optimization*. Spinger, 2010.

[11] T. Balch and R. C. Arkin, "Communication in reactive multiagent robotic systems," *Auton. Robots*, vol. 1, no. 1, pp. 27–52, 1994.

[12] R. T. Vaughan, K. Stø y, G. S. Sukhatme, and M. J. Matarić, "LOST: Localization-space trails for robot teams," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 796–812, 2002.

[13] D. Payton, R. Estkowski, and M. Howard, "Pheromone Robotics and the Logic of Virtual Pheromones," *Swarm Robotics*, vol. 3342, pp. 45–57, 2005.

[14] T. Sharpe and B. Webb, "Simulated and situated models of chemical trail following in ants," in *Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior on From Animals to Animats 5*. Cambridge, MA, USA: MIT Press, 1998, pp. 195–204.

[15] D. Payton, M. Daily, R. Estowski, M. Howard, and C. Lee, "Pheromone robotics," *Autonomous Robots*, vol. 11, no. 3, pp. 319–324, 2001.

[16] S. Garnier, F. Tâche, M. Combe, A. Grimal, and G. Theraulaz, "Alice in pheromone land: An experimental setup for the study of ant-like robots," *Proceedings of the IEEE Swarm Intelligence Symposium, SIS 2007*, no. Sis, pp. 37–44, 2007.

[17] Q. a. Nguyen Vu, S. Hassas, F. Armetta, B. Gaudou, and R. Canal, "Combining trust and self-organization for robust maintaining of information coherence in disturbed MAS," *Proceedings 5th IEEE International Conference on Self-Adaptive and Self-Organizing Systems, SASO 2011*, pp. 178–187, 2011.

[18] A. Rodriguez, J. Gomez, and A. Diaconescu, "Towards Failure-Resistant Mobile Distributed Systems Inspired by Swarm Intelligence and Trophallaxis," in *Proceedings of the European Conference on the Synthesis and Simulation of Living Systems ECAL*, 2015.

[19] M. HENI and R. BOUALLEGUE, "Power Control in Reactive Routing Protocol for Mobile Ad Hoc Network," pp. 53–68, 2012.

[20] M. E. Suárez and B. L. Thorne, "Rate, Amount, and Distribution Pattern of Alimentary Fluid Transfer via Trophallaxis in Three Species of Termites (Isoptera: Rhinotermitidae, Termopsidae)," *Annals of the Entomological Society of America*, vol. 93, no. Shapiro 1990, pp. 145–155, 2000.

[21] T. Chandra and S. Toueg, "Unreliable failure detectors for reliable distributed systems," *Journal of the ACM (JACM)*, vol. 43, no. 2, 1996.

[22] A. Kshemkalyani and M. Singhal, *Distributed computing: principles, algorithms, and systems*. Cambridge University Press, 2008.

[23] L. Gray, "at Wolfram's New Kind of Science," *Notices of the AMS*, vol. 50, pp. 200–211, 2002.

[24] P. G. Balaji and D. Srinivasan, "An introduction to multi-agent systems," *Studies in Computational Intelligence*, vol. 310, pp. 1–27, 2010.

[25] S. Russell and P. Norvig, *Inteligencia Artificial. Un enfoque moderno. 2da Edición*, 2004.

[26] S. Benhamou, "Free-flight odor tracking in Drosophila is consistent with an optimal intermittent scale-free search," *PLoS ONE*, vol. 2, no. 4, pp. 2351–2352, 2007.

[27] M. Dorigo and L. M. Gambardella, "Ant Colony System: A cooperative learning approach to the traveling salesman problem," *IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION*, 1997.

[28] M. Dorigo and T. Stutzle, *Ant colony optimization*. MIT Press, 2004, vol. 1.

[29] J. E. Bell and P. R. McMullen, "Ant colony optimization techniques for the vehicle routing problem," *Advanced Engineering Informatics*, vol. 18, pp. 41–48, 2004.